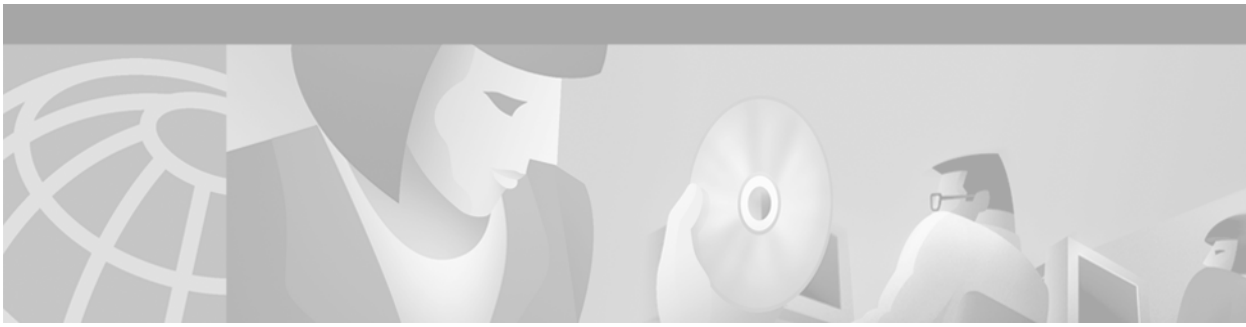

Configuring Integrated IS-IS

This chapter describes how to configure Integrated Intermediate System-to-Intermediate System (IS-IS). For a complete description of the integrated IS-IS commands listed in this chapter, refer to the "Integrated IS-IS Commands" chapter of the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols* publication. To locate documentation of other commands that appear in this chapter, use the command reference master index, or search online.

IS-IS is an International Organization for Standardization (ISO) dynamic routing specification. IS-IS is described in ISO 10589. The Cisco implementation of IS-IS allows you to configure IS-IS as an IP routing protocol.

To identify the hardware platform or software image information associated with a feature, use the Feature Navigator on Cisco.com to search for information about the feature or refer to the software release notes for a specific release. For more information, see the "Identifying Supported Platforms" section in the "Using Cisco IOS Software" chapter in this book.

IS-IS Configuration Task List

To configure IS-IS, perform the tasks described in the following sections. The tasks in the first section are required; the tasks in the remaining sections are optional:

- Enabling IS-IS and Assigning Areas (Required)
- Enabling IP Routing for an Area on an Interface (Optional)
- Monitoring IS-IS (Optional)

In addition, you can filter routing information and specify route redistribution. For more information about these features, see the "Filter Routing Information" and "Redistribute Routing Information" sections, respectively, in the "Configuring IP Routing Protocol-Independent Features" chapter of this document.

Enabling IS-IS and Assigning Areas

Unlike other routing protocols, enabling IS-IS requires that you create an IS-IS routing process and assign it to a specific interface, rather than to a network. You can specify more than one IS-IS routing process per Cisco router, using the multiarea IS-IS configuration syntax. You then configure the parameters for each instance of the IS-IS routing process.

Note Multiarea IS-IS is supported only for ISO CLNS.

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level 1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (interarea routing).

Some networks use legacy equipment that supports only Level 1 routing. These devices are typically organized into many small areas that cannot be aggregated due to performance limitations. Cisco routers are used to interconnect each area to the Level 2 backbone.

A single Cisco router can participate in routing in up to 29 areas, and can perform Level 2 routing in the backbone. In general, each routing process corresponds to an area. By default, the first instance of the routing process configured performs both Level 1 and Level 2 routing. You can configure additional router instances, which are automatically treated as Level 1 areas. You must configure the parameters for each instance of the IS-IS routing process individually.

For IS-IS multiarea routing, you can configure only one process to perform Level 2 routing, although you can define up to 29 Level 1 areas for each Cisco router. If Level 2 routing is configured on any process, all additional processes are automatically configured as Level 1. You can configure this process to perform Level 1 routing at the same time. If Level 2 routing is not desired for a router instance, remove the Level 2 capability using the **is-type** router configuration command. Use the **is-type** router configuration command also to configure a different router instance as a Level 2 router.

Network entity titles (NETs) define the area addresses for the IS-IS area and the system ID of the router. Refer to the "Configuring ISO CLNS" chapter in the *Cisco IOS Apollo Domain, Banyan VINES, ISO CLNS, and XNS Configuration Guide* for a more detailed discussion of NETs.

To enable IS-IS and specify the area for each instance of the IS-IS routing process, use the following commands in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# router isis [area tag]`	Enables IS-IS routing for the specified routing process, and places the router in router configuration mode. Use the *area tag* arguments to identify the area to which this IS-IS router instance is assigned. A value for *tag* is required if you are configuring multiple IS-IS areas. The first IS-IS instance configured is Level 1-2 by default. Later instances are automatically Level 1. You can change the level of routing to be performed by a particular routing process using the **is-type** router configuration command.
Step 2	`Router(config)# net network-entity-title`	Configures NETs for the routing process. Specify a NET for each routing process if you are configuring multiarea IS-IS. You can specify a name for a NET and for an address. **Note** Multiarea IS-IS is supported only for ISO CLNS.

See the "IS-IS Configuration Examples" section at the end of this chapter for examples of configuring IS-IS as an IP routing protocol.

Enabling IP Routing for an Area on an Interface

To enable IP routing and specify the area for each instance of the IS-IS routing process, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface interface-type interface-number`	Enters interface configuration mode.
Step 2	`Router(config-if)# ip router isis [area tag]`	Configures an IS-IS routing process for ISO Connectionless Network Service (CLNS) on an interface and attaches an area designator to the routing process.
Step 3	`Router(config-if)# ip address ip-address-mask`	Defines the IP address for the interface. An IP address is required on all interfaces in an area enabled for IS-IS if any one interface is configured for IS-IS routing.

See the "IS-IS Configuration Examples" section at the end of this chapter for examples of configuring IS-IS as an IP routing protocol.

IS-IS Interface Parameters Configuration Task List

The Cisco IS-IS implementation allows you to alter certain interface-specific IS-IS parameters. Most interface configuration commands can be configured independently from other attached routers. The **isis password** interface configuration command should configure the same password on all routers on a network. The settings of other commands (**isis hello-interval**, **isis hello-multiplier**, **isis retransmit-interval**, **isis retransmit-throttle-interval**, **isis csnp-interval**, and so on) can be different on different routers or interfaces. However, if you decide to change certain values from the defaults, it makes sense to configure them on multiple routers and interfaces.

To alter IS-IS parameters, perform the optional tasks described in the following sections:

- Configuring IS-IS Link-State Metrics (Optional)
- Setting the Advertised Hello Interval (Optional)
- Setting the Advertised CSNP Interval (Optional)
- Setting the Retransmission Interval (Optional)
- Setting the LSP Transmissions Interval (Optional)
- Setting the Retransmission Throttle Interval (Optional)
- Setting the Hello Multiplier (Optional)
- Specifying Designated Router Election (Optional)
- Specifying the Interface Circuit Type (Optional)
- Assigning a Password for an Interface (Optional)
- Limiting LSP Flooding (Optional)

Configuring IS-IS Link-State Metrics

You can configure a cost for a specified interface. You can configure the *default-metric* value for Level 1 or Level 2 routing. To configure the metric for the specified interface, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis metric default-metric [level-1 | level-2]`	Configures the metric (or cost) for the specified interface.

Setting the Advertised Hello Interval

You can specify the length of time (in seconds) between hello packets that the Cisco IOS software sends on the interface.

To specify the length of time between hello packets for the specified interface, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis hello-interval {seconds | minimal} [level-1 | level-2]`	Specifies the length of time (in seconds) between hello packets the Cisco IOS software sends on the specified interface.

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because only a single type of hello packet is sent on serial links, it is independent of Level 1 or Level 2.) Specify an optional level for X.25, Switched Multimegabit Data Service (SMDS), and Frame Relay multiaccess networks. X25, SMDS, ATM, and Frame Relay networks should be configured with point-to-point subinterfaces.

Setting the Advertised CSNP Interval

Complete sequence number protocol data units (CSNPs) are sent by the designated router to maintain database synchronization. You can configure the IS-IS CSNP interval for the interface.

To configure the CSNP interval for the specified interface, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis csnp-interval seconds {level-1 | level-2}`	Configures the IS-IS CSNP interval for the specified interface.

This feature does not apply to serial point-to-point interfaces. It applies to WAN connections if the WAN is viewed as a multiaccess meshed network.

Setting the Retransmission Interval

You can configure the number of seconds between retransmission of IS-IS link-state packets (LSPs) for point-to-point links. To set the retransmission level, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis retransmit-interval` *seconds*	Configures the number of seconds between retransmission of IS-IS LSPs for point-to-point links.

The value you specify should be an integer greater than the expected round-trip delay between any two routers on the attached network. The setting of this parameter should be conservative, or needless retransmission will result. The value should be larger for serial lines.

Setting the LSP Transmissions Interval

To configure the delay between successive IS-IS LSP transmissions, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis lsp-interval` *milliseconds*	Configures the delay between successive IS-IS LSP transmissions.

Setting the Retransmission Throttle Interval

You can configure the maximum rate at which IS-IS LSPs will be re-sent on point-to-point links, in terms of the number of milliseconds between packets. This configuration is different from the retransmission interval, which is the amount of time between successive retransmissions of the same LSP.

The retransmission throttle interval is typically not necessary, except in cases of very large networks with high point-to-point neighbor counts. To set the retransmission throttle interval, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis retransmit-throttle-interval` *milliseconds*	Configures the IS-IS LSP retransmission throttle interval.

Setting the Hello Multiplier

To specify the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down, use the following command in interface configuration command. The default value is 3.

Command	Purpose
`Router(config-if)# isis hello-multiplier multiplier [level-1 | level-2]`	Sets the hello multiplier.

Specifying Designated Router Election

You can configure the priority to use for designated router election. Priorities can be configured for Level 1 and Level 2 individually.

To specify the designated router election, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis priority number-value [level-1 | level-2]`	Configures the priority to use for designated router election.

Specifying the Interface Circuit Type

You can specify adjacency levels on a specified interface. This parameter is also referred to as the *interface circuit type*.

To specify the interface circuit type, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis circuit-type [level-1 | level-1-2 | level-2-only]`	Configures the type of adjacency desired for neighbors on the specified interface (the interface circuit type).

Assigning a Password for an Interface

You can assign different passwords for different routing levels. Specifying Level 1 or Level 2 configures the password for only Level 1 or Level 2 routing, respectively. If you do not specify a level, the default is Level 1. By default, authentication is disabled.

To configure a password for the specified level, use the following command in interface configuration mode:

Command	Purpose
`Router(config-if)# isis password password [level-1 | level-2]`	Configures the authentication password for a specified interface.

Limiting LSP Flooding

Limiting LSP flooding is important to IS-IS networks in general, and is not limited to configuring multiarea IS-IS networks. In a network with a high degree of redundancy, such as a fully meshed set of point-to-point links over a nonbroadcast multiaccess (NBMA) transport, flooding of LSPs can limit network scalability. You can reduce LSP flooding in two ways:

- Blocking Flooding on Specific Interfaces

  The advantage of full blocking over mesh groups is that it is easier to configure and understand, and fewer LSPs are flooded. Blocking flooding on all links permits the best scaling performance, but results in a less robust network structure. Permitting flooding on all links results in poor scaling performance.

- Configuring Mesh Groups

  The advantage of mesh groups over full blocking is that mesh groups allow LSPs to be flooded over one hop to all routers on the mesh, while full blocking allows some routers to receive LSPs over multiple hops. This relatively small delay in flooding can have an impact on convergence times, but the delay is negligible compared to overall convergence times.

Blocking Flooding on Specific Interfaces

You can completely block flooding (full blocking) on specific interfaces, so that new LSPs will not be flooded out over those interfaces. However, if flooding is blocked on a large number of links, and all remaining links go down, routers cannot synchronize their link-state databases even though there is connectivity to the rest of the network. When the link-state database is no longer updated, routing loops usually result.

To use CSNPs on selected point-to-point links to synchronize the link-state database, configure a CSNP interval using the **isis csnp-interval** interface configuration command on selected point-to-point links over which normal flooding is blocked. You should use CSNPs for this purpose only as a last resort.

Configuring Mesh Groups

Configuring mesh groups (a set of interfaces on a router) can help to limit redundant flooding. All routers reachable over the interfaces in a particular mesh group are assumed to be densely connected (each router has many links to other routers), where many links can fail without isolating one or more routers from the network.

Normally, when a new LSP is received on an interface, it is flooded out over all other interfaces on the router. When the new LSP is received over an interface that is part of a mesh group, the new LSP will not be flooded out over the other interfaces that are part of that same mesh group.

Mesh groups rely on a full mesh of links between a group of routers. If one or more links in the full mesh go down, the full mesh is broken, and some routers might miss new LSPs, even though there is connectivity to the rest of the network. When you configure mesh groups to optimize or limit LSP flooding, be sure to select alternative paths over which to flood in case interfaces in the mesh group go down.

To minimize the possibility of incomplete flooding, you should allow unrestricted flooding over at least a minimal set of links in the mesh. Selecting the smallest set of logical links that covers all physical paths results in very low flooding, but less robustness. Ideally you should select only enough links to ensure that LSP flooding is not detrimental to scaling performance, but enough links to ensure that under most failure scenarios no router will be logically disconnected from the rest of the network.

Miscellaneous IS-IS Parameters Configuration Task List

The following tasks differ from the preceding interface-specific IS-IS tasks because they configure IS-IS itself, rather than the interface.

To configure optional IS-IS parameters as described in the following sections:

- Generating a Default Route (Required)
- Specifying the System Type (Optional)
- Configuring IS-IS Authentication Passwords (Optional)
- Summarizing Address Ranges (Optional)
- Setting the Overload Bit (Optional)
- Changing the Routing Level for an Area (Optional)
- Tuning LSP Interval and Lifetime (Optional)
- Customizing IS-IS Throttling of LSP Generation, SPF Calculation, and PRC (Optional)
- Modifying the Output of show Commands (Optional)

Generating a Default Route

You can force a default route into an IS-IS routing domain. Whenever you specifically configure redistribution of routes into an IS-IS routing domain, the Cisco IOS software does not, by default, redistribute the *default route* into the IS-IS routing domain. The following command generates a default route into IS-IS, which can be controlled by a route map. You can use the route map to identify the level into which the default route is to be announced, and you can specify other filtering options configurable under a route map. You can use a route map to conditionally advertise the default route, depending on the existence of another route in the routing table of the router.

To generate a default route, use the following command in router configuration mode:

Command	Purpose
`Router(config-router)# default-information originate [route-map map-name]`	Forces a default route into the IS-IS routing domain.

See also the discussion of redistribution of routes in the "Configuring IP Routing Protocol-Independent Features" chapter of this document.

Specifying the System Type

You can configure the router to act as a Level 1 (intra-area) router, as both a Level 1 router and a Level 2 (interarea) router, or as an interarea router only.

To specify router level support, use the following command in router configuration mode:

Command	Purpose
`Router(config-router)# is-type {level-1 | level-1-2 | level-2-only}`	Configures the system type (area or backbone router).

Configuring IS-IS Authentication Passwords

You can assign passwords to areas and domains.

The area authentication password is inserted in Level 1 (station router level) LSPs, and the routing domain authentication password is inserted in Level 2 (area router level) LSPs.

To configure either area or domain authentication passwords, use the following commands in router configuration mode, as needed:

Command	Purpose
`Router(config-router)# area-password password`	Configures the area authentication password.
`Router(config-router)# domain-password password`	Configures the routing domain authentication password.

Summarizing Address Ranges

You can create aggregate addresses that are represented in the routing table by a summary address. This process is called *route summarization*. One summary address can include multiple groups of addresses for a given level. Routes learned from other routing protocols also can be summarized. The metric used to advertise the summary is the smallest metric of all the more-specific routes.

To create a summary of addresses for a given level, use the following command in router configuration mode:

Command	Purpose
`Router(config-router)# summary-address address mask {level-1 | level-1-2 | level-2}`	Creates a summary of addresses for a given level.

Setting the Overload Bit

You can configure the router to set the overload bit (also known as the hippity bit) in its nonpseudonode LSPs. Normally the setting of the overload bit is allowed only when a router runs into problems. For example, when a router is experiencing a memory shortage, the link-state database may not be complete, resulting in an incomplete or inaccurate routing table. By setting the overload bit in their LSPs, other routers can ignore the unreliable router in their shortest path first (SPF) calculations until the router has recovered from its problems.

The result will be that no paths through this router are seen by other routers in the IS-IS area. However, IP and CLNS prefixes directly connected to this router will be still be reachable.

This command can be useful when you want to connect a router to an IS-IS network, but do not want real traffic flowing through it under any circumstances. Examples are as follows:

- A test router in the lab, connected to a production network.
- A router configured as an LSP flooding server, for example, on an NBMA network, in combination with the mesh-group feature.
- A router that is aggregating virtual circuits (VCs) used only for network management. In this case, the network management stations must be on a network directly connected to the router with the **set-overload-bit** router configuration command configured.

Unless you specify the **on-startup** keyword, this command sets the overload bit immediately and it remains set until the **no set-overload-bit** command is specified. If you specify the **on-startup** keyword, you must indicate whether it is set for a specified number of *seconds* or until BGP has converged. If BGP does not signal IS-IS that it has converged, IS-IS will turn off the overload bit after 10 minutes.

In addition to setting the overload bit, you might also want to suppress certain types of IP prefix advertisements from LSPs. For example, allowing IP prefix propagation between Level1 and Level 2 effectively makes a node a transit node for IP traffic, which may be undesirable. The **suppress** keyword used with the **interlevel** or **external** keyword (or both) accomplishes that suppression while the overload bit is set.

To set the overload bit, use the following command in router configuration mode:

Command	Purpose
`Router(config-router)# set-overload-bit [on-startup {seconds | wait-for-bgp}] [suppress {[interlevel] [external]}]`	Sets the overload bit.

Changing the Routing Level for an Area

You can change the routing level configured for an area using the **is-type** router configuration command. If the router instance has been configured for a Level 1-2 area (the default for the first instance of the IS-IS routing process in a Cisco router), you can remove Level 2 (interarea) routing for the area using the **is-type** command and change the routing level to Level 1 (intra-area). You can also configure Level 2 routing for an area using the **is-type** command, but the instance of the IS-IS router configured for Level 2 on the Cisco router must be the only instance configured for Level 2.

To change the routing level for an IS-IS routing process in a given area, use the following command in router configuration mode:

Command	Purpose
`Router (config)# is-type {level-1 | level-1-2 | level-2-only}`	Configures the routing level for an instance of the IS-IS routing process.

Tuning LSP Interval and Lifetime

By default, the router sends a periodic LSP refresh every 15 minutes. LSPs remain in a database for 20 minutes by default. If they are not refreshed by that time, they are deleted. You can change the LSP refresh interval or the LSP lifetime. The LSP interval should be less than the LSP lifetime or else LSPs will time out before they are refreshed. The software will adjust the LSP refresh interval if necessary to prevent the LSPs from timing out.

To change the LSP refresh interval or lifetime, use the appropriate command in router configuration mode:

Command	Purpose
`Router (config-router)# lsp-refresh-interval` *seconds*	Sets the LSP refresh interval.
`Router (config-router)# max-lsp-lifetime` *seconds*	Sets the maximum time that link-state packets (LSPs) can remain in a router's database without being refreshed.

Customizing IS-IS Throttling of LSP Generation, SPF Calculation, and PRC

Partial Route Computation (PRC)

PRC is the software's process of calculating routes without performing an SPF calculation. This is possible when the topology of the routing system itself has not changed, but a change is detected in the information announced by a particular IS or when it is necessary to attempt to reinstall such routes in the RIB.

Benefits of Throttling IS-IS LSP Generation, SPF Calculation, and PRC

IS-IS throttles three main events: link-state PDU (LSP) generation, Shortest Part First (SPF) computation, and partial route computation (PRC). Throttling slows down the frequency of these events during times of network instability. Although throttling these events slows down network convergence, not throttling could result in a network not functioning. If network topology is unstable, throttling slows down the scheduling of these intervals until the topology becomes stable.

The throttling of LSP generation prevents flapping links from cause many LSPs to be flooded through the network. The throttling of SPF computation and PRC prevents the router from crashing from the demand of too many calculations.

How Throttling of IS-IS LSP Generation, SPF Calculation, and PRC Works

IS-IS throttling of LSP generation, SPF calculations, and PRC occurs by default. You can customize the throttling of these events with the **lsp-gen-interval, spf-interval**, and **prc-interval** commands, respectively.

The arguments in each command behave similarly. For each command:

- The first argument indicates the maximum number of seconds between LSP generations or calculations.
- The second argument indicates the initial wait time (in milliseconds) before running the first LSP generation or calculation.
- The third argument indicates the minimum amount of time to wait (in milliseconds) between the first and second LSP generation or calculation. (In addition to this wait time, there might be some other system overhead between LSP generations or calculations.)

Each subsequent wait interval is twice as long as the previous one until the wait interval reaches the maximum wait time specified, upon which the wait interval remains constant. After the network calms down and there are no triggers for 2 times the maximum interval, fast behavior is restored (the initial wait time).

Other commands are available to control the delay between successive LSPs, the retransmission of the same LSA, and the retransmission of LSPs on a point-to-point interface.

Perform this task to customize throttling of LSP generation, SPF calculation, PRC, or any combination of the three, beginning in router configuration mode:

Command	Purpose
`Router(config-router)# lsp-gen-interval [level-1 | level-2] lsp-max-wait [lsp-initial-wait lsp-second-wait]`	Sets IS-IS LSP generation throttling timers. • The default *lsp-max-wait* interval is 5 seconds. • The default *lsp-initial-wait* interval is 50 milliseconds. • The default *lsp-second-wait* interval is 5000 milliseconds.
`Router(config-router)# spf-interval [level-1 | level-2] spf-max-wait [spf-initial-wait spf-second-wait]`	Sets IS-IS SPF throttling timers. • The default *spf-max-wait* interval is 10 seconds. • The default *spf-initial-wait* interval is 5500 milliseconds. • The default *spf-second-wait* interval is 5500 milliseconds.
`Router(config-router)# prc-interval prc-max-wait [prc-initial-wait prc-second-wait]`	Sets IS-IS partial route computation throttling timers. • The default *prc-max-wait* interval is 10 seconds. • The default *prc-initial-wait* interval is 2000 milliseconds. • The default *prc-second-wait* interval is 5000 milliseconds.

Modifying the Output of show Commands

To customize display output when the IS-IS multiarea feature is used, making the display easier to read, use the following command in EXEC mode:

Command	Purpose
`Router# isis display delimiter [return count | character count]`	Specifies the delimiter to be used to separate displays of information about individual IS-IS areas.

For example, the following command causes information about individual areas to be separated by 14 dashes (-) in the display:

```
isis display delimiter - 14
```

The output for a configuration with two Level 1 areas and one Level 2 area configured is as follows:

```
dtp-5# show clns neighbors
--------------
Area L2BB:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0009 Tu529       172.21.39.9     Up     25        L1L2 IS-IS
--------------
```

```
Area A3253-01:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0053 Et1         0060.3e58.ccdb  Up     22        L1   IS-IS
0000.0000.0003 Et1         0000.0c03.6944  Up     20        L1   IS-IS
--------------
Area A3253-02:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 Et2         0000.0c03.6bc5  Up     27        L1   IS-IS
0000.0000.0053 Et2         0060.3e58.ccde  Up     24        L1   IS-IS
```

Monitoring IS-IS

To monitor the IS-IS tables and databases, use the following commands in EXEC mode, as needed:

Command	Purpose
Router# **show isis database** [**level-1**] [**level-2**] [**l1**] [**l2**] [**detail**] [**lspid**]	Displays the IS-IS link-state database.
Router# **show isis** *area-tag* **routes**	Displays the IS-IS Level 1 routing table.
Router# **show isis spf-log**	Displays how often and why the router has run a full SPF calculation.
Router# **show isis** *area-tag* **topology**	Displays a list of all connected routers in all areas.

IS-IS Configuration Examples

This section includes the following examples:

- Enabling IS-IS Configuration Example
- Multiarea IS-IS Configuration for CLNS Network Example
- IS-IS Throttle Timers Example

Enabling IS-IS Configuration Example

The following example shows how to configure three routers to run IS-IS as an IP routing protocol. Figure 51 illustrates the example configuration.

Router A Configuration

```
router isis
 net 49.0001.0000.0000.000a.00
interface ethernet 0
 ip router isis
interface serial 0
 ip router isis
```

Router B Configuration

```
router isis
 net 49.0001.0000.0000.000b.00
interface ethernet 0
 ip router isis
interface ethernet 1
 ip router isis
```

```
interface serial 0
 ip router isis
```

Router C Configuration

```
router isis
 net 49.0001.0000.0000.000c.00
interface ethernet 1
 ip router isis
interface ethernet 2
 ip router isis
```

Figure 51 IS-IS Routing

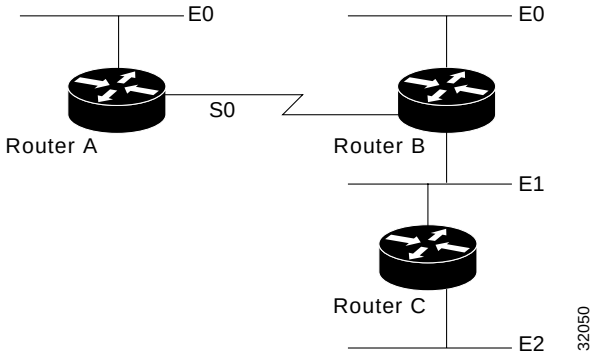

Multiarea IS-IS Configuration for CLNS Network Example

The following example shows a multiarea IS-IS configuration with two Level 1 areas and one Level 1-2 area. Figure 52 illustrates this configuration.

```
clns routing

.
.
.

interface Tunnel529
 ip address 10.0.0.5 255.255.255.0
 ip router isis BB
 clns router isis BB

interface Ethernet1
 ip address 10.1.1.5 255.255.255.0
 ip router isis A3253-01
 clns router isis A3253-01
!
interface Ethernet2
 ip address 10.2.2.5 255.255.255.0
 ip router isis A3253-02
 clns router isis A3253-02

.
.
.
```

```
router isis BB                          ! Defaults to "is-type level-1-2"
 net 49.2222.0000.0000.0005.00
!
router isis A3253-01
 net 49.0553.0001.0000.0000.0005.00
 is-type level-1
!
router isis A3253-02
 net 49.0553.0002.0000.0000.0005.00
 is-type level-1
```

Figure 52 Multiarea IS-IS Configuration with Three Level 1 Areas and One Level 2 Area

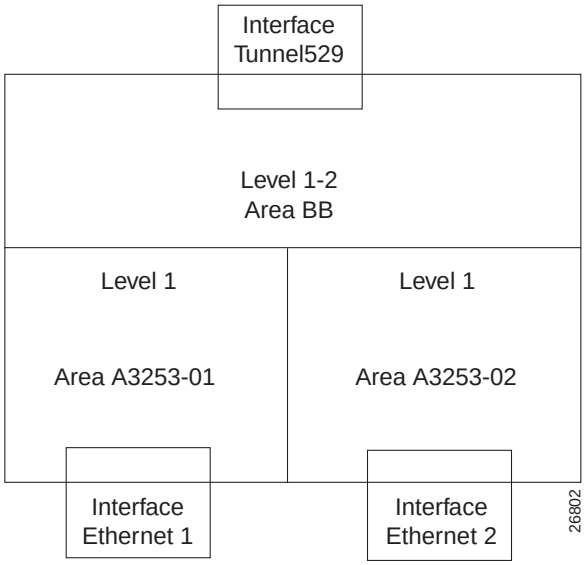

IS-IS Throttle Timers Example

This example shows a system configured with IS-IS throttling of LSP generations, SPF calculations and PRC:

```
router isis
 spf-interval 5 10 20
 prc-interval 5 10 20
 lsp-gen-interval 2 50 100
```