

Chapter 1

ISO CLNS Routing Protocols

1

This chapter describes routing protocol options for the International Organization for Standardization (ISO) Connectionless Network Services (CLNS) protocol. The chapter “Switching ISO CLNS” contains all the information you need for configuring ISO CLNS switching, as well as configuration requirements not related specifically to ISO CLNS routing protocols support. This chapter focuses on ISO CLNS routing protocols. Topics in this chapter include:

- An introduction to the Cisco-supported ISO CLNS routing protocols
- Overview of the sequence for configuring and starting routing processes
- Details of commands used for configuring static and dynamic routing
- Examples of configuring the supported ISO CLNS protocols—static, ISO-IGRP, and IS-IS routing
- Information concerning monitoring and debugging ISO CLNS routing processes on Cisco routers
- Summaries of ISO CLNS routing configuration commands

This chapter also explains how to optimize and fine-tune your CLNS-based internet work. Configuration examples are grouped in a separate section, with configuration command lists and illustrations.

Cisco-Supported Routing Protocols

The basic function of a router is to forward packets to the proper destination. All routers do this by looking up the destination address in a table. The tables may be built dynamically or statically. If you are configuring all of the entries in the table, you are using *static* routing. If you have a routing process building the tables, you are using *dynamic* routing. It is possible, and sometimes necessary, to use both static and dynamic routing simultaneously.

Cisco supports two dynamic routing protocols for CLNP networks: *ISO-IGRP* and *IS-IS*. Both routing protocols support the concept of *areas*. Within an area, all routers know how to reach all of the Station IDs. Between areas, routers know how to reach the proper area.

IS-IS supports two levels of routing: *station routing* (within an area) and *area routing* (between areas). ISO-IGRP supports three levels of routing: *station routing*, *area routing*, and *interdomain* routing.

Using Names to Represent a Full NSAP or NET

Various CLNS configuration commands allow the use of a name to represent a full NSAP or NET. When names are used in this way, the following system behaviors occur:

- Although using names as proxies for addresses is allowed with CLNS commands, the name is never written out to NVRAM.
- The **clns host** command is generated after the **router** subcommands when the configuration file is parsed. As a result, the NVRAM version of the configuration cannot be edited to specifically change the address defined in the original **clns host** command. This affects all commands that accept names.
 - **net** (router subcommand)
 - **clns is-neighbor** (interface subcommand)
 - **clns es-neighbor** (interface subcommand)
 - **clns route** (global configuration command)
- In the following command example, the **net** command specification assumes that a **clns host** command already has been entered (with the name cisco representing a valid NET).

```
router iso-igrp
net cisco
```

Configuring ISO CLNS Routing

Follow these steps to configure your router for ISO CLNS routing:

- Step 1:** Enable CLNS routing with the **clns routing** global configuration command (described in the chapter “Switching ISO CLNS”).
- Step 2:** Create dynamic routing processes with the **router** global configuration command and either the **iso-igrp** or **isis** keywords. You do not need to explicitly specify a routing process to use static routing facilities.
- Step 3:** Specify which interface should be actively routing ISO CLNS with the **clns router** interface subcommand running each routing process defined in step 2.

If you do not intend to use a dynamic routing protocol on a specific interface, but wish to switch packets over the link, use the **clns enable** interface subcommand (described in the chapter “Switching ISO CLNS”).

Steps 1 through 3 enable CLNS routing. Optional commands are also available for customizing the routing environment, and you may need to follow steps 4 through 6 to complete configuration:

- Step 4:** Redistribute routing information.
- Step 5:** Map NSAP addresses to media addresses.

Step 6: Adjust ES-IS parameters as necessary (described in the chapter “Switching ISO CLNS”).

Tasks related to static or dynamic routing are described in the following sections or in the chapter “Switching ISO CLNS.” Configuration information is followed by descriptions of the EXEC commands related to the ISO CLNS routing protocols supported by Cisco. Summaries of the global configuration commands and interface subcommands described in these sections appear at the end of this chapter.

Configuring CLNS Static Routing

Static routing is used when it is not possible or desirable to use dynamic routing.

For example, if you are using routers that do not support the same dynamic routing protocol, then you must use static routing. In addition, if your network includes WAN links that involve paying for connect time or per packet, you might not want to run a routing protocol over that link. Finally, if you want routers to advertise connectivity to external networks, but you are not running an interdomain routing protocol, you *must* use static routes.

Note: An interface that is configured for static routing cannot reroute *around* failed links.

Routes are entered by specifying pairs (NSAP-prefix, next-hop-NET). NET is a Network Entity Title, similar in function to an NSAP. In the routing table, the best match means the longest NSAP-prefix entry that matches the beginning of the destination NSAP. In the following static routing table, the next-hop-NETs are listed for completeness, but are not necessary to understand the routing algorithm. Table 1-2 offers examples of how the routing entries in Table 1-1 can be applied to various NSAPs.

Table 1-1 Sample Routing Table Entries

Entry #1	NSAP Prefix	Next Hop NET
1	47.0005.000c.0001	47.0005.000c.0001.0000.1234.00
2	47.0004	47.0005.000c.0002.0000.0231.00
3	47.0005.0003	47.0005.000c.0001.0000.1234.00
4	47.0005.000c	47.0005.000c.0004.0000.0011.00
5	47.0005	47.0005.000c.0002.0000.0231.00

Table 1-2 Hierarchical Routing Examples

Datagram Destination NSAP	Table Entry Number Used
47.0005.000c.0001.0000.3456.01	1
47.0005.000c.0001.6789.2345.01	1
47.0004.1234.1234.1234.1234.01	2
47.0005.0003.4321.4321.4321.01	3
47.0005.000c.0004.5678.5678.01	4
47.0005.000c.0005.3456.3456.01	5
47.0005.0023.9876.9876.9876.01	6

Octet boundaries must be used for the internal boundaries of NSAPs and NETs.

You enter a specific static route by using the **clns route** global configuration command (applied globally or to a specific interface). The syntax for the global configuration version of this command follows:

```
clns route nsap-prefix {next-hop-NET|name}  
no clns route nsap-prefix
```

NSAPs that start with *nsap-prefix* are forwarded to *next-hop-NET* or the *name* of the next hop node.

The following variation of the **clns route** global configuration command uses the **discard** keyword to explicitly tell a router to discard packets with the specified *nsap-prefix* (full syntax listed).

```
clns route nsap-prefix discard  
no clns route nsap-prefix
```

The **clns route** global configuration command can also be applied to a specific interface. This version has the following syntax:

```
clns route nsap-prefix interface-type unit [SNPA-address]  
no clns route nsap-prefix
```

The optional argument *SNPA-address* is required for multiaccess networks.

Note: You must include the argument *SNPA-address* when the *interface-type unit* specifies a multiaccess network. Otherwise you receive an error message that complains about a bad SNPA.

The **clns route** command is not literally *applied* to a specific interface. There are two forms of the **clns route** command:

- Configured with the destination NSAP and next-hop NET

- Configured with the destination NSAP, next-hop interface, and SNPA (referred to as an interface-static route).

Examples:

The following example sets a static route for a router with the global version of the **clns route** command:

```
clns route 47.0004.000c 47.0005.0001.0000.0001.0000.00
```

This example sets static routes for specific interfaces:

```
clns route 39.0001 serial 0
clns route 38.0002 ethernet 1 0000.0c00.1550
!SNPA address specified
clns route 39.0003 serial 1 4085551021
!X.121 address specified
```

Configuring ISO-IGRP CLNS Dynamic Routing

Use the commands described in this section to configure CLNS dynamic routing. In the previous chapter, you turned on CLNS processing in the router. To specify an ISO-IGRP routing process, you must enable the ISO-IGRP routing process and identify the address for the router. Use the **router** global configuration command. This command has the following syntax:

```
router iso-igrp tag
no router iso-igrp tag
```

The keyword **iso-igrp** specifies dynamic routing using the ISO-IGRP protocol. You can specify up to ten ISO-IGRP processes. The argument *tag* defines a meaningful name for routing process. For example, you could define routing process named *Finance* for the Finance department, and another routing process named *Marketing* for the Marketing department. Creating a name for a routing process means that you use names when configuring routing.

Note: The argument *tag* is optional; if it is not specified a null tag is assumed. The *tag* argument must be unique among all CLNS router processes for a given router.

The **no router iso-igrp** global configuration command with the appropriate tag disables ISO-IGRP routing for the system.

Example:

In the following example, a router is specified in *Manufacturing*. (The command must be typed on one line.)

```
router iso-igrp Manufacturing
```

Configuring Network Entity Titles

The **net** router subcommand configures a Network Entity Title (NET) for the routing process. The syntax for this command is as follows:

```
net network-entity-title
```

The argument *network-entity-title* is the NET; NETs define the area addresses for the ISO-IGRP area, as described in the chapter “Switching ISO CLNS.” Only one NET per router is allowed for ISO-IGRP ISO-IGRP routing. There is no default value for this command. The **net** router subcommand allows you to specify a name for an NET, as well as an address.

Example 1:

The following example illustrates specifying an NET:

```
router iso-igrp Finance
net 47.0004.004d.0001.0000.0c11.1111.00
```

Example 2:

The following illustrates using a name for a NET:

```
clns host NAME 39.0001.0000.0c00.1111.00
!
router iso-igrp Marketing
net NAME
!
```

Specifying Router Level Support

Each interface supported by ISO-IGRP must also have a **clns router** interface subcommand entry to enable the routing process on the interface. Following is the command syntax for this command:

```
clns router iso-igrp tag [level 2]
no clns router iso-igrp tag
```

The argument *tag* is the tag defined for the routing process the **router iso-igrp** global configuration command above.

If you want this interface to advertise Level 2 information only, use the **level 2** keyword. This option reduces the amount of router-to-router traffic by telling the router to send out only Level 2 routing updates on certain interfaces. The Level 1 information will not be passed on the interfaces for which the Level 2 option set.

Use the **no clns router iso-igrp** command with the appropriate *tag* to disable the CLNS routing protocol on the interface.

Example:

In the following example, the interface will advertise Level 2 information only on interface serial 0:

```
interface serial 0
  clns router iso-igrp marketing level 2
```

Configuring ISO-IGRP Metrics

Two ISO-IGRP metrics can be configured. These are the bandwidth and delay associated with an interface. Refer to the chapter “Adjusting Interface Characteristics” for details about the **bandwidth** and **delay** interface subcommands used to set these metrics.

Redistributing Routes into an ISO-IGRP Domain

A router may be configured to do interdomain dynamic routing by putting the router into two domains and configuring the router to redistribute the routing information between the domains. Routers configured this way are referred to as *border* routers.

Note: It is not necessary to use redistribution between areas.

If you have a router that is in two routing domains, you may want to redistribute routing information between the two domains. The **redistribute** router configuration subcommand configures what routes are redistributed into the ISO-IGRP domain.

redistribute *router-name tag*
no redistribute *router-name tag*

The **redistribute** router configuration subcommand configures routes redistributed into the ISO-IGRP domain. The argument *router-name* specifies the source of routes that will be redistributed into ISO-IGRP. The *tag* relates to the *tag* specified with the **router iso-igrp** global configuration subcommand.

The *router-name* keywords supported are **iso-igrp**, **isis**, and **static**. Static routes are only redistributed into ISO-IGRP when a **redistribute static** command is entered. The default is to not redistribute static routes into ISO-IGRP.

Example:

The following example illustrates redistribution of IS-IS routes of France and ISO-IGRP routes of Germany into the ISO-IGRP area tagged Backbone:

```
router iso-igrp Backbone
 redistribute isis France
 redistribute iso-igrp Germany
```

Redistributing Static Routes

The **redistribute static** router subcommand causes the specified routing process to advertise static CLNS routes. Its full syntax is as follows:

redistribute static
no redistribute static

Use the **no redistribute static** command to stop the routing process from advertising static routes.

Example:

In the following example, the router will advertise any static routes it knows about in the Chicago domain.

```
router iso-igrp Chicago
 redistribute static
```

Note: Only the router that injects the static route needs to have a **redistribute** configuration command defined.

Specifying Preferred Routes

When multiple routing processes are running in the same router for CLNS, it is possible for the same route to be advertised by more than one routing process. The router will always pick the route whose routing protocol has the lowest administrative distance. Use the **distance** router subcommand to assign administrative distance for a particular routing protocol. The command syntax is:

distance *value*
no distance *value*

By default the following administrative distances are assigned:

- Static routes—10
- ISO-IGRP routes—100

- IS-IS routes—110

The lower the *value* of the distance, the more preferred the route.

Example:

To give preference to IS-IS routes over ISO-IGRP, the following configuration commands are required:

```
router isis
distance 90
```

Configuring IS-IS CLNS Dynamic Routing

This section describes the various global and interface commands associated with configuring ISO CLNS IS-IS routing support on Cisco routers. Each separate command description includes a simple example of the command. Examples set in the context of application configurations are provided in the subsequent section in this chapter, “CLNS Static, ISO-IGRP, and IS-IS Routing Examples.”

Note: ISO-IGRP and IS-IS should not be configured for same area. This means that the **net** subcommand should *not* specify an NSAP address where all bytes up to (but not including the System ID) are the same when enabling both ISO-IGRP and IS-IS routing.

Enabling IS-IS Routing

The **router isis** global configuration command configures the IS-IS routing process. The command syntax is as follows:

```
router isis [tag]
no router isis [tag]
```

Associated with the process is a *tag* that is used later as a reference to this process. You can specify *only one* IS-IS process per router. The argument *tag* defines a meaningful name for a routing process. For example, you could define a routing process named *Pieinthesky*. Creating a name for a routing process means that you use names when configuring routing. The argument *tag* is optional; if it is not specified, a null tag is assumed.

Note: The *tag* argument must be unique among all CLNS router processes for a given router.

The **no router isis** global configuration command with the appropriate tag disables IS-IS routing for the system.

Example

The following example illustrates starting IS-IS routing with the optional tag argument:

```
router isis Pieinthesky
```

Configuring Network Entity Titles

The **net** router subcommand configures a Network Entity Title (NET) for the routing process. The syntax for this command is as follows:

```
net network-entity-title  
no net network-entity-title
```

NETs define the area addresses for the IS-IS area, as described in the chapter “Switching ISO CLNS.” For IS-IS, multiple NETs per router are allowed, with a maximum of three. There is no default value for this command. The **net** router subcommand allows you to specify a name for an NET, as well as an address.

Note: Although IS-IS allows you to configure multiple NETs, ISO-IGRP allows only one NET per router.

The **no net** command must be specified with the NET; it removes a specific NET.

Example 1:

The following example illustrates specifying a single NET:

```
router isis Pieinthesky  
net 47.0004.004d.0001.0000.0c11.1111.00
```

Example 2:

The following illustrates using a name for an NET:

```
clns host NAME 39.0001.0000.0c00.1111.00  
!  
router isis  
net NAME
```

Specifying Router Level Support

The **is-type** global configuration command configures the level at which the router should operate. The command syntax is as follows:

```
is-type [level-1 | level-1-2 | level-2-only]
no is-type [level-1 | level-1-2 | level-2-only]
```

The **is-type** command includes an optional argument with three valid alternative keywords:

- **level-1**
- **level-1-2**
- **level-2-only**

If **level-1** is specified, the router acts as a station router. If **level-1-2** is specified, the router acts as both a station router and an area router. If **level-2-only** is specified, the router acts as an area router only. The default value is **level-1-2**.

The command **no is-type** resets the parameter to the default.

Example

The following example specifies a router as capable of being used as an area router only.

```
is-type level-2-only
```

Redistributing Routes into an IS-IS Domain

If you have a router that is in two routing domains, you may want to redistribute routing information between the two domains.

The **redistribute** router configuration subcommand configures what routes are redistributed into the IS-IS domain.

```
redistribute router-name tag
no redistribute router-name tag
```

The **redistribute** global configuration command causes the routes learned by routing process tag to be advertised in this routing process. The argument *router-name* specifies the type of routing protocol that is redistributed into IS-IS. Supported *router-name* keywords are **static** and **iso-igrp**. Static routes are always redistributed into IS-IS, unless a **no redistribute static** is performed. Redistribution only occurs for level-2 routing.

Example:

The following example illustrates redistribution of ISO-IGRP routes of Michigan and ISO-IGRP routes of Ohio into the IS-IS area tagged Illinois.

```
router isis Illinois
redistribute iso-igrp Michigan
redistribute iso-igrp Ohio
```

Configuring IS-IS for an Interface

The interface subcommand **clns router isis** enables IS-IS routing for OSI on a specific interface. The syntax for this command is as follows:

```
clns router isis tag  
no clns router isis tag
```

Use the same text for the argument *tag* as specified in the global configuration command **router isis** described earlier in this chapter.

The **no router isis** global configuration command with the appropriate area tag disables IS-IS routing for the system. If no *tag* is specified, the null tag is used.

Configuring IS-IS Link State Metrics

The interface subcommand **isis metric** configures the metric (or cost) for the specified interface. The syntax for this command is as follows:

```
isis metric default-metric delay-metric expense-metric error_metric [level-1 | level-2]  
no isis metric [level-1 | level-2]
```

Only the metric *default-metric* is supported by Cisco routers. The default value for this metric is ten. You can configure this metric for Level 1 and/or Level 2 routing.

The **no isis metric** command resets the *default-metric* value to ten. Specification of the **level-1** or **level-2** optional keywords resets metric only for Level 1 or Level 2 routing, respectively.

Specifying Designated Router Election

Use the **isis priority** interface subcommand to configure the priority to use for Designated Router election. Priorities can be configured for Level 1 and Level 2 individually.

```
isis priority value [level-1 | level-2]  
no isis priority [level-1 | level2]
```

The default *value* is 64. The **no isis priority** command resets priority to 64. Specification of the **level-1** or **level-2** optional keywords resets priority only for Level 1 or Level 2 routing, respectively.

Specifying Interface Circuit Type

Use the **isis circuit-type** interface subcommand to configure the type of *adjacency* desired for this interface. The command syntax is as follows:

```
isis circuit-type [level-1 | level-1-2 | level-2-only]  
no isis circuit-type
```

If **level-1** is specified, a Level 1 adjacency may be established if there is at least one area address in common between this system and its neighbors.

If **level-1-2** is specified, a Level 1 and 2 adjacency is established if the neighbor is also configured as **level-1-2** and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established.

If **level-2-only** is specified, a Level 2 adjacency is established if and only if the neighbor is configured exclusively to be a Level 2 router.

The default value for this command is **level-1-2**. The **no isis circuit-type** command resets the circuit type to Level 1 and Level 2.

Configuring IS-IS Password Authentication

The **isis password** interface subcommand configures the authentication password for an interface. The command syntax for this command is as follows:

```
isis password password [level-1 | level-2]
no isis password [level-1 | level-2]
```

Different passwords can be assigned for different routing levels using the optional **level-1** and **level-2** keyword arguments. By default authentication is disabled.

The **no isis password** command disables authentication for IS-IS. Specifying **level-1** or **level-2** optional keywords disables the password only for Level 1 or Level 2 routing, respectively. If no keyword is specified, the default is **level-1**.

CLNS Static, ISO-IGRP, and IS-IS Routing Examples

This section provides configuration examples of both intra- and interdomain static and dynamic routing using static, ISO-IGRP, and IS-IS routing techniques.

Basic Static Routing

Configuring FDDI, Ethernets, Token Rings, and serial lines for CLNS can be as simple as just enabling CLNS on the interfaces. This is all that is ever required on serial lines using HDLC encapsulation. If all systems on an Ethernet or Token Ring support ISO 9542 ES-IS, then nothing else is required. In this case, an Ethernet and a serial line can be configured as in the following example.

```
clns net 47.0004.004D.0055.0000.0C00.BF3B.00
clns routing
interface ethernet 0
clns enable
interface serial 0
clns enable
clns route 47.0004.004d.0099 serial 0
clns route 47.0005 serial 0
```

Static Routing


The following is a more complete example of CLNS static routing on a system with two Ethernet interfaces. After configuring routing, you define an NET and enable CLNS on the Ethernet 0 and Ethernet 1 interfaces. You must then define an IS-neighbor and define a static route with the **clns route** command, as shown. In this situation, there is an IS on Ethernet 1 that does not support ES-IS.

```
clns host foo
clns net foo
clns routing
interface Ethernet 0
clns enable
interface Ethernet 1
clns enable
clns is-neighbor foo 0000.0C00.62e7
clns route 47.0004.000c foo
```

Static Intradomain Routing

Figure 1-1 and the configurations that follow demonstrate how to use static routing inside of a domain. Imagine a company with two branch offices in Detroit and Chicago, connected with an X.25 link. These offices are both in the domain named Sales.

Figure 1-1 CLNS X.25 Intradomain Routing



The following is one way to configure the router in Chicago.

```
clns host chicago 47.0004.0050.0002.0000.0c00.243b.00
clns host detroit 47.0004.0050.0001.0000.0c00.1e12.00
clns routing
router iso-igrp sales
net chicago
interface ethernet 0
clns router iso-igrp sales
interface serial 0
```

```

encapsulation x25
x25 address 31342174523156
x25 nvc 4
clns router iso-igrp sales
clns is-neighbor detroit 31343136931281 broadcast

```

This configuration will bring up an X.25 virtual circuit between the router in Chicago and the router in Detroit. Routing updates will be sent across this link. This implies that the virtual circuit could be up continuously.

If this is undesirable, use the following configuration instead.

```

clns host chicago 47.0004.0050.0002.0000.0c00.243b.00
clns host detroit 47.0004.0050.0001.0000.0c00.1e12.00
router iso-igrp sales
net 47.0004.0060.0002.0000.0c00.243b.00
!
interface ethernet 0
clns router iso-igrp sales
!
interface serial 0
encapsulation x25
x25 address 31342174523156
x25 nvc 4
clns enable
clns is-neighbor detroit 31343136931281
!
clns route 47.0004.0050.0001 detroit

```

If the Chicago office should grow to contain multiple routers, it would be appropriate for each of those routers to know how to get to Detroit. Add the following command to redistribute information between routers in Chicago:


```

router iso-igrp sales
redistribute static

```

Static Interdomain Routing

Figure 1-2 and the example configurations that follow illustrate how to configure two routers that distribute information across domains. In this example, Castor and Pollux communicate across a serial link.



Example Configuration for Castor:

```
router iso-igrp orion
net 47.0006.0200.0100.0102.0304.0506.00
!
clns host pollux 47.0006.0200.0200.1112.1314.1516.00
!
interface ethernet 0
clns router iso-igrp orion
!
interface serial 1
clns enable
!
clns route 39.0001 pollux
```

Example Configuration for Pollux:

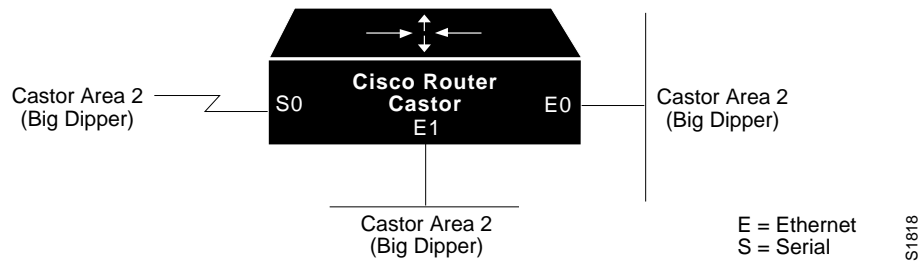
```
router iso-igrp pleiades
net 47.0006.0200.0200.1112.1314.1516.00
!
clns host castor 47.0006.0200.0100.0001.0102.0304.0506.00
!
interface ethernet 0
clns router iso-igrp pleiades
!
interface serial 0
clns enable
!
clns route 47.0006.0200.0100 castor
```

CLNS routing updates will not be sent on the serial link; however, CLNS packets will be sent and received over the serial link.

Routing Within the Same Area

Figure 1-2 and the example configuration that follows illustrate how to configure dynamic routing within a routing domain. The router may exist in one or more areas within the domain. The router named Castor exists in a single area.

Figure 1-2 CLNS Dynamic Routing—Within a Single Area



```

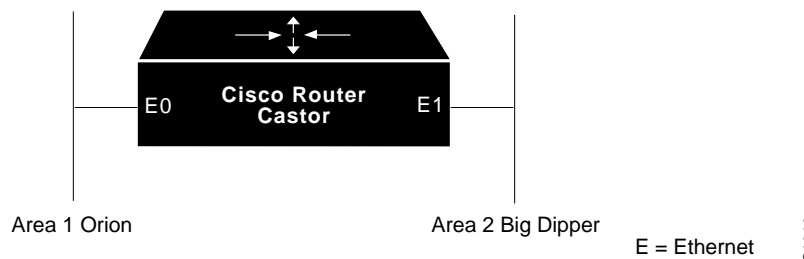
clns routing
router iso-igrp bigdipper
net 47.0004.004D.0002.0000.0C00.0506.00
interface Ethernet 0
clns router iso-igrp bigdipper
interface Ethernet 1
clns router iso-igrp bigdipper
interface Serial 0
clns router iso-igrp bigdipper

```

Dynamic Routing in More Than One Area

Figure 1-3 and the example configuration that follows illustrate how to configure a router named Castor that exists in two areas.

Figure 1-3 CLNS Dynamic Routing Within Two Areas



```

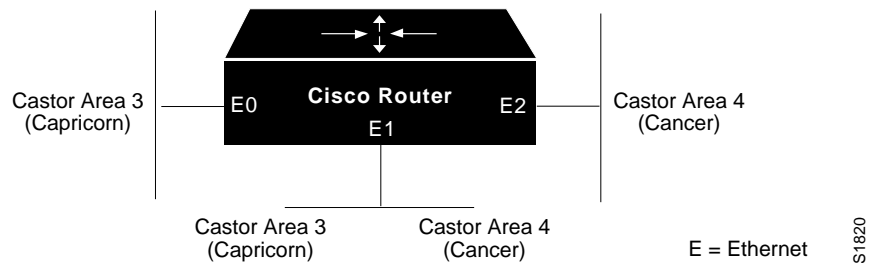
clns routing
router iso-igrp orion
net 47.0004.004D.0001.212223242526.00
router iso-igrp bigdipper
net 47.0004.004D.0002.212223242526.00
interface ethernet 0
clns router iso-igrp orion
interface ethernet 1
clns router iso-igrp bigdipper

```

Dynamic Routing in Overlapping Areas

Figure 1-4 and the example that follows illustrate how to configure a router with overlapping areas.

Figure 1-4 CLNS Dynamic Routing Within Overlapping Areas




```
clns routing
router iso-igrp capricorn
net 47.0004.004D.0003.0000.0C00.0508.00
router iso-igrp cancer
net 47.0004.004D.0004.0000.0C00.0506.00
interface ethernet 0
clns router iso-igrp capricorn
interface ethernet 1
clns router iso-igrp capricorn
clns router iso-igrp cancer
interface ethernet 2
clns router iso-igrp cancer
```

Dynamic Interdomain Routing

Figure 1-5 and the examples that follow illustrate how to configure three domains that want to be transparently connected.

Figure 1-5 CLNS Interdomain Dynamic Routing



S1821

Example Configuration for Router X:

```

clns routing
router iso-igrp A
net 47.0006.0200.0002.0102.0104.0506.00
redistribute iso-igrp B
router iso-igrp B
net 47.0007.0200.0003.0102.0104.0506.00
redistribute iso-igrp A
interface serial 0
clns router iso-igrp A
interface serial 1
clns router iso-igrp B

```

Example Configuration for Router Y:

```

clns routing
router iso-igrp B
net 47.0007.0200.0004.0102.0104.0506.00
redistribute iso-igrp C
router iso-igrp C
net 47.0008.0200.0005.0102.0104.0506.00
redistribute iso-igrp B
interface serial 0
clns router iso-igrp B
interface serial 1
clns router iso-igrp C

```

Router X will inject a prefix route for domain A into domain B. Domain B will inject this prefix route and a prefix route for domain B into domain C.

You can also configure a border router between domain A and domain C.

IS-IS Routing Configuration

The examples that follow illustrate the basic syntax and configuration command sequence for IS-IS routing.

Example 1: Level 1 and Level 2 Routing

The following example illustrates using the IS-IS protocol to configure a single area address for Level 1 and Level 2 routing.

```
clns routing
router isis
net 47.0004.004d.0001.0000.0c00.1111.00
interface Ethernet 0
clns router isis
interface Ethernet 1
clns router isis
interface Serial 0
clns router isis
```

Example 2: Level 2 Routing Only

The following example illustrates a similar configuration, featuring a single area address being used for specification of Level 1 and Level 2 routing. However, in this case, interface Serial 0 is configured for Level 2 routing only.

```
clns routing
router isis
net 47.0004.004d.0001.0000.0c00.1111.00
interface Ethernet 0
clns router isis
interface Ethernet 1
clns router isis
interface Serial 0
clns router isis
isis circuit-type level-2-only
```

Example 3: OSI Configuration

The following example illustrates an OSI configuration example. In this example, IS-IS runs with two area addresses, metrics tailored, and different circuit types specified for each interface.

```
clns routing
router isis area1
net 47.0004.004d.0001.0000.0c11.1111.00
net 47.0004.004d.0011.0000.0c11.1111.00
is-type level-1-2
interface Ethernet 0
clns router isis area1
isis metric 5 level-1
isis circuit-type level-1
interface Ethernet 1
clns router isis area1
isis metric 2 level-2
```

```
isis circuit-type level-2-only
interface serial 0
clns router isis area1
isis circuit-type level-1-2
isis priority 3 level-1
isis priority 1 level-2
```

Example 4: ISO CLNS Dynamic Route Redistribution

The following example illustrates route redistribution between IS-IS and ISO-IGRP domains. In this case, the IS-IS domain is on interface Ethernet 0; the ISO-IGRP domain is on interface Serial 0. The IS-IS routing process is assigned a null tag; the ISO-IGRP routing process is assigned a tag of *remote-domain*.

```
router isis
net 39.0001.0001.0000.0c00.1111.00
redistribute iso-igrp remote-domain

router iso-igrp remote-domain
net 39.0002.0001.0000.0c00.1111.00
redistribute isis

interface Ethernet 0
clns router isis

interface Serial 0
clns router iso-igrp remote
```

Monitoring a CLNS Network

Use the EXEC commands described in this section to obtain displays of activity on the ISO CLNS network.

Displaying CLNS Routes

Use the EXEC command **show clns route** to display all of the destinations to which this router knows how to route packets. Enter this command at the EXEC prompt:

```
show clns route [nsap]
```

The **show clns route** command shows the IS-IS Level 2 routing table as well as static and ISO-IGRP learned prefix routes. IS-IS area addresses and prefix routes are stored in this table.

Destinations are sorted by category. The optional argument *nsap* specifies the CLNS address.

Sample output follows:

```
ISO-IGRP Routing Table for Domain 39.0003, Area 0020
System Id      Next-Hop      SNPA          Interface     Metric  State
0000.0C00.224D 0000.0000.0000 --             --          0          Up

ISO-IGRP Routing Table for Domain 39.0003
Area Id        Next-Hop      SNPA          Interface     Metric  State
0020           0000.0000.0000 --             --          0          Up

CLNS Prefix Routing Table
39.0003 [100/0]
  via 39.0004.0030.0000.0C00.224D.00, ISO-IGRP, Up
39.0004.0040 [110/10]
  via 0000.0C00.2D55, IS-IS, Up, Serial2
39.0004.0030 [110/0]
  via 0000.0C00.224D, IS-IS, Up
39.0004.0030.0000.0C00.224D.00, Local NET Entry
39.0003.0020.0000.0C00.224D.00, Local NET Entry
```

The first two display fields show ISO-IGRP information:

- **System Id**—The identification value of the system listed in Level 1 forwarding table
- **Area Id**—The identification value of the area listed in Level 1 forwarding table
- **Next-Hop**—System ID of best cost next-hop to listed address
- **SNPA**—SNPA of next-hop system
- **Interface**—Interface through which next-hop system is known by router
- **Metric**—ISO-IGRP metric for the route
- **State**—Up (active) or Down (non-operational)

The third portion shows CLNS Prefix information:

- **Destination prefix**—Example: 39.0003.
- **Administrative distance/metric**—Example: [100/0].
Here, a value of 0 for metric indicates the route is internally reachable.
- **Next-hop address**—Either an NET (if a static route) or System id, if route obtained via IS-IS or ISO-IGRP.
- **Link status**—Up (active) or Down (nonoperational)
- **Interface type**—Only appears if the specific interface through which the destination is reachable is unambiguously known to router; Local NET Entry indicates destination is on a directly connected network.

Neighbors are not included in the **show clns route** display.

Output for the **show clns route nsap** command is the same as **show clns route**, but only lists a single entry.

Displaying the IS-IS Level 1 Routing Table

The **show isis route EXEC** command displays the IS-IS Level 1 routing table. The command syntax is as follows:

show isis routes

Following is sample output:

```
IS-IS Level-1 Routing Table - Version 34
System Id      Next-Hop      SNPA          Interface    Metric  State
0000.0C00.0C35 0000.0C00.0C35 0000.0c00.0c36 Ethernet1    20     Up
0800.2B16.24EA 0800.2B16.24EA aa00.0400.2d05 Ethernet0    10     Up
0800.2B14.060E 0800.2B14.060E aa00.0400.9205 Ethernet0    10     Up
0800.2B14.0528 0800.2B14.0528 aa00.0400.9105 Ethernet0    10     Up
0000.0C00.40AF 0000.0000.0000 --            --          0         Up
0000.0C00.62E6 0000.0C00.62E6 0000.0c00.62e7 Ethernet1    10     Up
AA00.0400.2D05 0800.2B16.24EA aa00.0400.2d05 Ethernet0    10     Up
```

This display presents the Level 1 forwarding table for IS-IS learned routes. The following fields are provided in this display:

- **System id**—The identification value of the system listed in Level 1 forwarding table
- **Next-Hop**—System ID of best cost next-hop to listed address
- **SNPA**—SNPA of next-hop system
- **Interface**—Interface through which next-hop system is known by router
- **Metric**—IS-IS metric for the route
- **State**—Up (active) or Down (nonoperational)

Displaying Protocol-Specific Information

The **show clns protocol** command lists the protocol-specific information for each ISO-IGRP routing process in this router. Enter this command at the EXEC prompt:

show clns protocol [domain] area-tag]

There will always be at least two routing processes, a Level 1 and a Level 2, and there may be more. The optional argument *domain* specifies a particular ISO-IGRP routing domain; the optional argument *area-tag* specifies a particular IS-IS area.

Following is sample output:

```
ISO-IGRP Level 1 Router: remote
Routing for domain: 39.0003 area: 0020
Sending Updates every 45 seconds. Next due in 11 seconds
Invalid after 135 seconds,
Hold down for 145 seconds
Sending Router Hellos every 17 seconds. Next due in 9 seconds
Invalid after 51 seconds,
IGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
Interfaces in domain/area:
    TokenRing1
--More--
```

```

ISO-IGRP Level 2 Router: DOMAIN_remote
  Routing for domain: 39.0003
  Redistribute:
    isis (Null Tag)
  Sending Updates every 45 seconds. Next due in 2 seconds
  Invalid after 135 seconds,
  Hold down for 145 seconds
  Sending Router Hellos every 17 seconds. Next due in 0 seconds
  Invalid after 51 seconds,
  ISO-IGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  Interfaces in domain/area:
    TokenRing1
--More--
IS-IS Router: <Null Tag>
  System Id: 0000.0C00.224D.00  IS-Type: level-1-2
  Manual area address(es):
    39.0004.0030
  Routing for area address(es):
    39.0004.0030
  Interfaces supported by IS-IS:
    Serial2
  Next global update in 530 seconds
  Redistributing:
    static
    iso-igrp (remote)

```

In the ISO-IGRP portion of the display:

- The first line provides the domain address and area number for Level 1 routing processes. For Level 2 routing processes, this command lists the domain address.
- The next set of fields indicate some of the protocol timers. The field labeled `Sending updates` displays when the next routing updates will be sent.
- The `Invalid` field indicates how long routing updates are to be believed.
- The `Hold Down` field indicates how long a route will be held down before new information is to be believed.
- The `Sending Router Hellos` field indicates how often the routers will send Hello packets to each other and when the next is due.
- The field labeled `Invalid` indicates how long a neighbor entry will be remembered.
- The `ISO-IGRP metric weight` displays lists the weights applied to the various components of the metric. These fields are followed by the list of interfaces that are in this area.

In the IS-IS portion of the display:

- The first portion specifies the relevant IS-IS area tag.
- The second line lists the System ID and IS Type.
- The next two information fields display area addresses.
- Area addresses are followed by a list of interfaces on the router supporting IS-IS.
- Several information fields are provided that indicate the next expected IS-IS update and any configuration of route redistribution.

Displaying the IS-IS Database

The **show isis database** EXEC command displays the IS-IS link state database. If **detail** is specified, the contents of each LSP is displayed. Otherwise, a summary display is provided. The command syntax is as follows:

```
show isis database [level-1 | level-2 | l1 | l2 | detail]
```

Note: The notations **l1** and **l2** are abbreviations for the options **level-1** and **level-2**, respectively. Each of the options shown in brackets for this command can be entered in a arbitrary string within the same command entry. For example, the following are both valid command specifications and provided the same display: **show isis data detail l2** and **show isis data l2 detail**.

Following is sample output when the command is specified with no options or as **show isis data l1 l2**:

```
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.0C35.00-00  0x0000000C  0x5696        792           0/0/0
0000.0C00.40AF.00-00* 0x00000009  0x8452        1077          1/0/0
0000.0C00.62E6.00-00  0x0000000A  0x38E7        383           0/0/0
0000.0C00.62E6.03-00  0x00000006  0x82BC        384           0/0/0
0800.2B16.24EA.00-00  0x00001D9F  0x8864        1188          1/0/0
0800.2B16.24EA.01-00  0x00001E36  0x0935        1198          1/0/0

IS-IS Level-2 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.0C35.03-00  0x00000005  0x04C8        792           0/0/0
0000.0C00.3E51.00-00  0x00000007  0xAF96        758           0/0/0
0000.0C00.40AF.00-00* 0x0000000A  0x3AA9        1077          0/0/0
```

This display presents IS-IS related information for the link state database. The following fields are provided in this display:

- **LSPID**—The link state PDU id. The first six octets form the System id. The next octet is the pseudo id. When this value is zero, the LSP describes links from the system. When it is nonzero, the LSP is a pseudo-node LSP. The designated router for an interface is the only system that originates pseudonode LSPs. The last octet is the LSP number. If there is more data than can fit in a single LSP, additional LSPs are sent with increasing LSP numbers. An asterisk (*) indicates that the LSP was originated by the local system.
- **LSP Seq Num**—The sequence number for the LSP that allows other systems can determine they have received the latest information from the source.
- **LSP Checksum** —The checksum of the entire LSP packet.
- **LSP Holdtime**—The amount of time the LSP remains valid, in seconds.

- ATT— The attach bit. This indicates that the router is also a Level 2 router, and it can reach other areas.
- P—The P bit. Detects if the IS is area partition repair capable.
- OL—The Overload bit. Determines if the IS is congested.

A sample output of the **show isis database detail** command follows:

```
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.0C35.00-00  0x0000000C   0x5696        325            0/0/0
  Area Address: 47.0004.004D.0001
  Area Address: 39.0001
  Metric: 10   IS 0000.0C00.62E6.03
  Metric: 0    ES 0000.0C00.0C35
--More--
0000.0C00.40AF.00-00* 0x00000009   0x8452        608            1/0/0
  Area Address: 47.0004.004D.0001
  Metric: 10   IS 0800.2B16.24EA.01
  Metric: 10   IS 0000.0C00.62E6.03
  Metric: 0    ES 0000.0C00.40AF
--More--
0000.0C00.62E6.00-00  0x0000000C   0x34E9        813            0/0/0
  Area Address: 47.0004.004D.0001
  Metric: 10   IS 0000.0C00.62E6.03
  Metric: 0    ES 0000.0C00.62E6
--More--
0000.0C00.62E6.03-00  0x00000007   0x80BD        812            0/0/0
  Metric: 10   IS 0000.0C00.40AF.00
  Metric: 0    IS 0000.0C00.62E6.00
  Metric: 10   IS 0000.0C00.0C35.00
--More--
0800.2B16.24EA.00-00  0x00001D9F   0x8864        715            1/0/0
  Area Address: 47.0004.004D.0001
  Metric: 10   IS 0800.2B16.24EA.01
  Metric: 0    ES 0800.2B16.24EA
  Metric: 0    ES AA00.0400.2D05
--More--
0800.2B16.24EA.01-00  0x00001E36   0x0935        724            1/0/0
  Metric: 0    IS 0800.2B16.24EA.00
  Metric: 0    IS 0800.2B16.24EA.00
  Metric: 0    IS 0000.0C00.40AF.00
  Metric: 0    ES 0800.2B14.060E
  Metric: 0    ES AA00.0400.9105
  Metric: 0    ES 0800.2B14.0528
  Metric: 0    ES AA00.0400.9205

IS-IS Level-2 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.0C35.03-00  0x00000005   0x04C8        317            0/0/0
  Metric: 0    IS 0000.0C00.0C35.00
--More--
0000.0C00.3E51.00-00  0x00000009   0xAB98        1182           0/0/0
  Area Address: 39.0004
  Metric: 10   IS 0000.0C00.40AF.00
  Metric: 10   IS 0000.0C00.3E51.05
--More--
```

```
0000.0C00.40AF.00-00* 0x0000000A 0x3AA9 599 0/0/0
Area Address: 47.0004.004D.0001
Area Address: 39.0001
Metric: 10 IS 0800.2B16.24EA.01
Metric: 10 IS 0000.0C00.3E51.00
```

In addition to the information displayed in **show isis database**, this **show** command displays the contents of each LSP.

Debugging ISO CLNS Dynamic Routing Protocols

Use the EXEC commands described in this section to troubleshoot and monitor the ISO CLNS routing protocols. For each **debug** command, there is a corresponding **undebug** command that turns the message logging off.

debug clns-igrp-packets

Causes all ISO-IGRP routing activity to be displayed.

debug isis-adj-packets

Logs information for adjacency-related functions such as Hello packets sent and received and IS-IS adjacencies going up and down.

debug isis-update-packets

Logs incoming and outgoing sequence number packets and link state packets. It also logs flooding-related functions.

debug isis-spf-events

Logs events associated with the Dijkstra algorithm.

ISO CLNS Global Command Summary (Routing Protocols)

This section provides an alphabetical summary of the ISO CLNS global configuration commands.

[no] clns route *nsap-prefix interface-type unit [SNPA-address]*

A variation of the **clns route** global configuration command that creates an interface-static route.

[no] clns route *nsap-prefix* {*next-hop-NET*|*name*}

Enters a specific static route. NSAPs that start with *nsap-prefix* are forwarded to *next-hop-NET* or the *name* of the next hop.

[no] clns route *nsap-prefix* [**discard**]

A variation of the **clns route** command that uses the **discard** keyword to explicitly tell a router to discard packets with NSAPs that start with the specified *nsap-prefix*.

[no] is-type [**level-1** | **level-1-2** | **level-2-only**]

Configures the level at which the router should operate. If **level-1** is specified, the router acts as a station router. If **level-1-2** is specified, the router acts as both a station router and a router that routes between areas. If **level-2-only** is specified, the router acts as an inter-area router only. The default value is **level-1-2**. The **no is-type** command resets routing level to Level 1 and 2.

[no] router [**isis** | **iso-igrp**] *tag*

Identifies the area the router will work in and lets the router know that it will be routing dynamically rather than statically. The keywords **isis** and **iso-igrp** specify dynamic routing using the IS-IS or ISO-IGRP protocols. The argument *tag* defines a meaningful name for a routing process.

ISO CLNS Router Subcommand Summary

This section provides an alphabetical list of the ISO CLNS router subcommands.

[no] distance *value*

Assigns administrative distance for a particular routing protocol.

Default values:

- Static routes—10
- ISO-IGRP routes—100
- IS-IS routes—110

[no] net *network-entity-title*

Configures a Network Entity Title (NET) for the routing process.

[no] redistribute *router-name tag*

Redistributes routing information throughout a routing domain. The argument *router-name* can be **isis**, **iso-igrp**, or **static**. The argument *tag* is the defined tag for the routing process from which the routing information is to be redistributed. The **no** form disables route redistribution.

[no] redistribute static

Causes the router to inject any static CLNS routes into the domain. The **no** form stops redistribution.

ISO CLNS Interface Subcommand Summary (Routing Protocols)

This section provides an alphabetical list of the ISO CLNS interface subcommands.

[no] clns router iso-igrp tag [level2]

This command specifies ISO-IGRP routing. The argument *tag* is the tag defined for the NET using the **router** global configuration command. The optional **level2** keyword allows the interface to advertise Level 2 information. The **no router iso-igrp** global configuration command with the appropriate tag disables ISO-IGRP routing for the system.

[no] clns router isis area-tag

Enables IS-IS routing for OSI on a specific interface. Use the same text described earlier for the argument *area-tag* as specified in the global configuration command **router isis**. The **no router isis** global configuration command with the appropriate area tag disables IS-IS routing for the system.

[no] isis circuit-type [level-1 | level-1-2 | level-2-only]

Configures the type of adjacency desired for this interface. If **level-1** is specified, at most a Level 1 adjacency may be established if there is at least one area address in common between this system and its neighbors. If **level-1-2** is specified, a Level 1 and 2 adjacency is established if the neighbor is configured as **level-1-2** and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. If **level-2-only** is specified, a Level 2 adjacency is established if and only if the neighbor is configured exclusively to be a Level 2 router. The default value for this command is **level-1-2**. The **no isis circuit-type** command resets the circuit type to Level 1 and Level 2.

[no] isis metric *default-metric delay-metric expense-metric error_metric* [**level-1** | **level-2**]

Configures the metric (or cost) for the specified interface. Only the metric *default-metric* is supported by Cisco routers. The default value for this metric is ten. You can configure this metric for Level 1 and/or Level 2 routing. The **no isis metric** command resets the *default-metric* value to ten. Specifying the **level-1** or **level-2** optional keywords resets metric only for Level 1 or Level 2 routing, respectively.

[no] isis password *password* [**level-1** | **level-2**]

Configures the authentication password for an interface. Different passwords can be assigned for different routing levels using the optional **level-1** and **level-2** keyword arguments. By default authentication is disabled. The **no isis password** command disables authentication for IS-IS. Specification of the **level-1** or **level-2** optional keywords disables the password only for Level 1 or Level 2 routing, respectively.

[no] isis priority *value* [**level-1** | **level-2**]

Configures the priority to use for designated router election. Priorities can be configured for Level 1 and Level 2 individually. The default *value* is 64. The **no isis priority** command resets priority to 64. The higher the number, the higher the priority. Specification of the **level-1** or **level-2** optional keywords resets priority only for Level 1 or Level 2 routing, respectively.