

Configuring ISO CLNS

The ISO CLNS protocol is a standard for the network layer of the OSI model. This chapter describes how to configure ISO CLNS. For a complete description of the ISO CLNS commands in this chapter, refer to the “ISO CLNS Commands” chapter of the *Network Protocols Command Reference, Part 3*. To locate documentation of other commands that appear in this chapter, use the command reference master index or search online.

Note Cisco access servers currently support only ES-IS, but not IS-IS.

ISO CLNS Configuration Task List

To configure ISO CLNS, you must configure the routing processes, associate addresses with the routing processes, and customize the routing processes for your particular network.

You must perform some combination of the tasks in the following sections to configure the ISO CLNS protocol:

- Understand Addresses
- Understand Routing Processes
- Configure ISO IGRP Dynamic Routing
- Configure IS-IS Dynamic Routing
- Configure CLNS Static Routing
- Configure Miscellaneous Features
- Configure CLNS over WANs
- Enhance ISO CLNS Performance
- Monitor and Maintain the ISO CLNS Network
- Configure TARP on ISO CLNS

See the “ISO CLNS Configuration Examples” section at the end of this chapter for configuration examples.

Understand Addresses


Addresses in the ISO network architecture are referred to as NSAP addresses and network entity titles (NETs). Each node in an OSI network has one or more NETs. In addition, each node has many NSAP addresses. Each NSAP address differs from one of the NETs for that node in only the last byte. This byte is called the *N-selector*. Its function is similar to the port number in other protocol suites.

Our implementation supports all NSAP address formats that are defined by ISO 8348/Ad2; however, we provide ISO IGRP or IS-IS dynamic routing only for NSAP addresses that conform to the address constraints defined in the ISO standard for IS-IS (ISO 10589).

An NSAP address consists of the following two major fields, as shown in Figure 15:

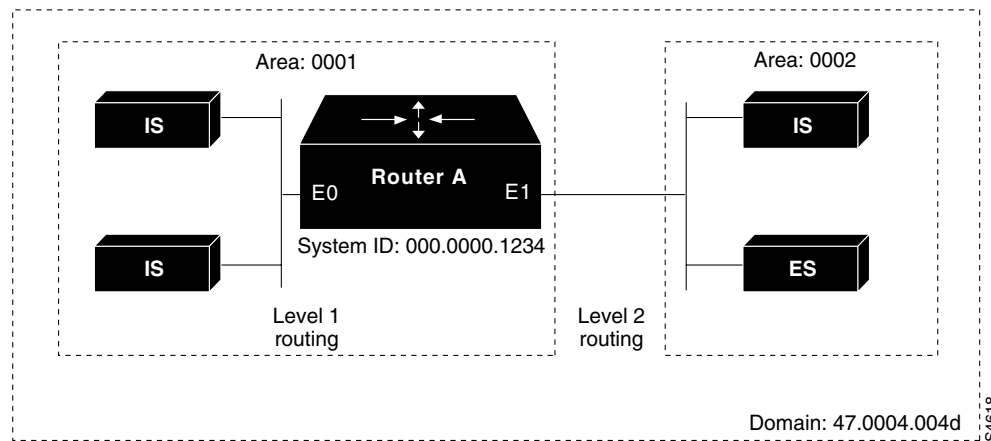
- The initial domain part (IDP) is made up of 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI). The length of the IDI and the encoding format for the domain specific part (DSP) are based on the value of the AFI.
- The DSP is made up of a high-order DSP, an area identifier, a system identifier, and a 1-byte N-selector (labeled S).

Figure 15 NSAP Address Fields



Assign addresses or NETs for your domains and areas. The domain address uniquely identifies the routing domain. All routers within a given domain are given the same domain address. Within each routing domain, you can set up one or more areas, as shown in Figure 16. Determine which routers are to be assigned to which areas. The area address uniquely identifies the routing area and the system ID identifies each node.

Figure 16 Sample Domain and Area Addresses



The key difference between the ISO IGRP and IS-IS NSAP addressing schemes is in the definition of area addresses. Both use the system ID for Level 1 routing (routing within an area). However, they differ in the way addresses are specified for area routing. An ISO IGRP NSAP address includes three separate fields for routing: the *domain*, *area*, and *system ID*. An IS-IS address includes two fields: a single continuous *area* field (comprising the domain and area fields) and the *system ID*.

ISO IGRP NSAP Address

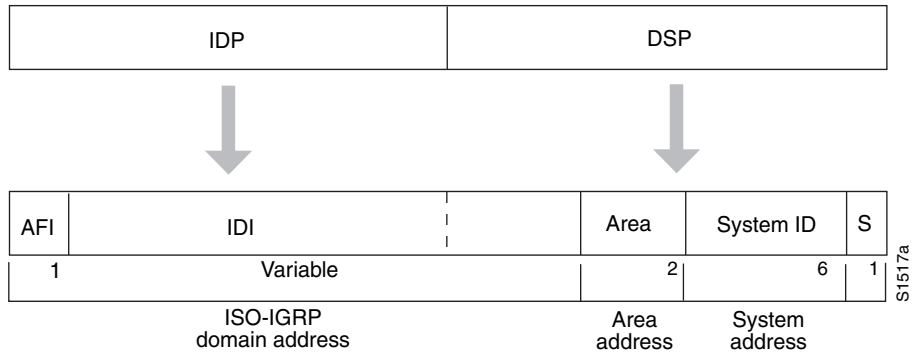
The ISO IGRP NSAP address is divided into three parts: a domain part, an area address, and a system ID. Domain routing is performed on the domain part of the address. Area routing for a given domain uses the area address. System routing for a given area uses the system ID part. The NSAP address is laid out as follows:

- The domain part is of variable length and comes before the area address.
- The area address is the 2 bytes before the system ID.
- The system ID is the 6 bytes before the N-selector.
- The N-selector (S) is the last byte of the NSAP address.

Cisco's ISO IGRP routing implementation interprets the bytes from the AFI up to (but not including) the area field in the DSP as a *domain identifier*. The area field specifies the *area*, and the system ID specifies the *system*.

Figure 17 illustrates the ISO IGRP NSAP addressing structure. The maximum address size is 20 bytes.

Figure 17 ISO IGRP NSAP Addressing Structure



IS-IS NSAP Address


An IS-IS NSAP address is divided into two parts: an area address and a system ID. Level 2 routing (routing between areas) uses the area address. Level 1 routing (routing within an area) uses the system ID address. The NSAP address is laid out as follows:

- The area address is the NSAP address, not including the system ID and N-selector.
- The system ID is found between the area address and the N-selector byte.
- The N-selector (S) is the last byte of the NSAP address.

The IS-IS routing protocol interprets the bytes from the AFI up to (but not including) the system ID field in the DSP as an *area identifier*. The system ID specifies the *system*.

Figure 18 illustrates the IS-IS NSAP addressing structure. The maximum address size is 20 bytes.

Figure 18 IS-IS NSAP Addressing Structure



Addressing Rules

All NSAP addresses must obey the following constraints:

- No two nodes can have addresses with the same NET; that is, addresses that match all but the N-selector (S) field in the DSP.
- No two nodes residing within the same area can have addresses in which the system ID fields are the same.

- ISO IGRP requires at least 10 bytes of length: 1 byte for domain, 2 bytes for area, 6 bytes for system ID, and 1 byte later in this chapter for N-selector.
- ISO IGRP and IS-IS should not be configured for the same area. Do *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.
- A router can have one or more area addresses. The concept of multiple area addresses is described in the “Assign Multiple Area Addresses to IS-IS Areas” section later in this chapter.
- Cisco’s implementation of IS-IS requires at least 8 bytes: one byte for area, 6 bytes for system ID, and 1 byte for N-selector.

Addressing Examples

The following are examples of OSI network and GOSIP NSAP addresses using the ISO IGRP implementation.

The following is the OSI network NSAP address format:

```
|          Domain|Area|          System ID| S|
47.0004.004D.0003.0000.0C00.62E6.00
```

The following is an example of the GOSIP NSAP address structure. This structure is mandatory for addresses allocated from the International Code Designator (ICD) 0005 addressing domain. Refer to the GOSIP document, *U.S. Government Open Systems Interconnection Profile (GOSIP)*, Draft Version 2.0, April 1989, for more information.

```
|          Domain|          Area|System ID| S|
47.0005.80.ffff00.0000.ffff.0004.0000.0c00.62e6.00
| | | | | |
AFI IDI DFI AAI Resv RD
```

Sample Routing Table

You enter static routes by specifying NSAP prefix and next-hop NET pairs (by using the **clns route** command). The NSAP prefix can be any portion of the NSAP address. NETs are similar in function to NSAP addresses.

If an incoming packet has a destination NSAP address that does not match any existing NSAP addresses in the routing table, the Cisco IOS software will try to match the NSAP address with an NSAP prefix to route the packet. In the routing table, the best match means the longest NSAP prefix entry that matches the beginning of the destination NSAP address.

Table 2 shows a sample static routing table in which the next-hop NETs are listed for completeness, but are not necessary to understand the routing algorithm. Table 3 offers examples of how the longest matching NSAP prefix can be matched with routing table entries in Table 2.

Table 2 Sample Routing Table Entries

Entry	NSAP Address 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 3 Hierarchical Routing Examples

Datagram Destination NSAP Address	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
47.0005.0001.0005.3456.3456.01	5

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

Understand Routing Processes

The basic function of a router is to forward packets: receive a packet in one interface and send it out another (or the same) interface to the proper destination. All routers do this by looking up the destination address in a table. The tables can be built either dynamically or statically. If you are configuring all the entries in the table yourself, you are using *static* routing. If you use a routing process to build the tables, you are using *dynamic* routing. It is possible, and sometimes necessary, to use both static and dynamic routing simultaneously.

When you configure only ISO CLNS and not routing protocols, the Cisco IOS software only makes forwarding decisions. It does not perform other routing-related functions. In such a configuration, the software compiles a table of adjacency data, but does not advertise this information. The only information that is inserted into the routing table is the NSAP and NET addresses of this router, static routes, and adjacency information.

You can route ISO CLNS on some interfaces and transparently bridge it on other interfaces simultaneously. To do this, you must enable concurrent routing and bridging by using the **bridge crb** command. For more information on bridging, refer to the “Configuring Transparent Bridging” chapter in the *Bridging and IBM Networking Configuration Guide*.

Dynamic Routing

Cisco supports the following two dynamic routing protocols for ISO CLNP networks:

- ISO IGRP
- IS-IS

When dynamically routing, you can choose either ISO IGRP or IS-IS, or you can enable both routing protocols at the same time. Both routing protocols support the concept of *areas*. Within an area, all routers know how to reach all the system IDs. Between areas, routers know how to reach the proper area.

ISO IGRP supports three levels of routing: *system routing*, *area routing*, and *interdomain routing*. Routing across domains (interdomain routing) can be done either statically or dynamically with ISO IGRP. IS-IS supports two levels of routing: *station routing* (within an area) and *area routing* (between areas).

Intermediate Systems (IS) and End Systems (ES)

Some ISs keep track of how to communicate with all the ESs in their areas and thereby function as Level 1 routers (also referred to as *local routers*). Other ISs keep track of how to communicate with other areas in the domain, functioning as Level 2 routers (sometimes referred to as *area routers*). Cisco routers are always Level 1 and Level 2 routers when routing ISO IGRP; they can be configured to be Level 1 only, Level 2 only, or both Level 1 and Level 2 routers when routing IS-IS.

ESs communicate with ISs using the ES-IS protocol. Level 1 and Level 2 ISs communicate with each other using either ISO IS-IS or Cisco's ISO IGRP protocol.

Static Routing

Static routing is used when it is not possible or desirable to use dynamic routing. The following are some instances of when you would use static routing:

- If your network includes WAN links that involve paying for connect time or for per-packet charges, you would use static routing, rather than paying to run a routing protocol and all its routing update packets over that link.
- If you want routers to advertise connectivity to external networks, but you are not running an interdomain routing protocol, you *must* use static routes.
- If you must interoperate with another vendor's equipment that does not support any of the dynamic routing protocols that Cisco supports, you must use static routing.
- For operation over X.25, Frame Relay, or SMDS networks, static routing is generally preferable.

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

Routing Decisions

A CLNP packet sent to any of the defined NSAP addresses or NETs will be received by the router. The Cisco IOS software uses the following algorithm to select which NET to use when it sends a packet:

- If no dynamic routing protocol is running, use the NET defined for the outgoing interface, if it exists; otherwise, use the NET defined for the router.
- If ISO IGRP is running, use the NET of the ISO IGRP routing process that is running on the interface.
- If IS-IS is running, use the NET of the IS-IS routing process that is running on the interface.

Configure ISO IGRP Dynamic Routing

The ISO IGRP is a dynamic distance-vector routing protocol designed by Cisco for routing an autonomous system that contains large, arbitrarily complex networks with diverse bandwidth and delay characteristics.

To configure ISO IGRP, complete the tasks in the following sections. Only enabling ISO IGRP is required; the remaining task is optional, although you might be required to perform it, depending upon your specific application:

- Enable ISO IGRP
- Configure ISO IGRP Parameters

In addition, you can also configure the following miscellaneous features, described later in this chapter:

- Filter routing information—See the “Create Packet-Forwarding Filters and Establish Adjacencies” section.
- Redistribute routing information from one routing process to another—See the “Redistribute Routing Information” section.
- Configure administrative distances—See the “Specify Preferred Routes” section.

Enable ISO IGRP

To configure ISO IGRP dynamic routing, you must enable the ISO IGRP routing process, identify the address for the router, and specify the interfaces that are to route ISO IGRP. Optionally, you can set a level for your routing updates when you configure the interfaces. CLNS routing is enabled by default on routers when you configure ISO IGRP. You can specify up to ten ISO IGRP routing processes.

To configure ISO IGRP dynamic routing on the router, perform the following tasks in global configuration mode:

Task	Command
Step 1 Enable the ISO IGRP routing process and enter router configuration mode.	router iso-igrp <i>[tag]</i>
Step 2 Configure the NET or address for the routing process.	net <i>network-entity-title</i>

Although IS-IS allows you to configure multiple NETs, ISO IGRP allows only one NET per routing process.

You can assign a meaningful name for the routing process by using the *tag* option. You can also specify a name for a NET in addition to an address. For information on how to assign a name, see the “Specify Shortcut NSAP Addresses” section later in this chapter.

You can configure an interface to advertise Level 2 information only. This option reduces the amount of router-to-router traffic by telling the Cisco IOS software to send out only Level 2 routing updates on certain interfaces. Level 1 information is not passed on the interfaces for which the Level 2 option is set.

To configure ISO IGRP dynamic routing on the interface, perform the following task in interface configuration mode:

Task	Command
Enable ISO IGRP on specified interfaces; also set the level type for routing updates.	clns router iso-igrp <i>tag [level 2]</i>

See the “Dynamic Routing in Overlapping Areas Example,” “Dynamic Interdomain Routing Example”, and “ISO CLNS over X.25 Example” sections at the end of this chapter for examples of configuring dynamic routing.

Configure ISO IGRP Parameters

Cisco's ISO IGRP implementation allows you to customize certain ISO IGRP parameters. You can perform the optional tasks discussed in the following sections:

- Adjust ISO IGRP Metrics.
- Adjust ISO IGRP Timers.
- Enable or Disable Split Horizon.

Adjust ISO IGRP Metrics

You have the option of altering the default behavior of ISO IGRP routing and metric computations. This allows, for example, the tuning of system behavior to allow for transmissions via satellite. Although ISO IGRP metric defaults were carefully selected to provide excellent operation in most networks, you can adjust the metric.

Note Adjusting the ISO IGRP metric can dramatically affect network performance, so ensure that all metric adjustments are made carefully. Because of the complexity of this task, it is not recommended unless it is done with guidance from an experienced system designer.

You can use different metrics for the ISO IGRP routing protocol on CLNS. To configure the metric constants used in the ISO IGRP composite metric calculation of reliability and load, perform the following task in router configuration mode:

Task	Command
Adjust the ISO IGRP metric.	metric weights <i>qos k1 k2 k3 k4 k5</i>

Two additional ISO IGRP metrics can be configured: the bandwidth and delay associated with an interface. Refer to the “Interface Commands” chapter of the *Configuration Fundamentals Command Reference* for details about the **bandwidth** and **delay** interface configuration commands used to set these metrics.

Note Using the **bandwidth** and **delay** commands to change the values of the ISO IGRP metrics also changes the values of IP IGRP metrics.

Adjust ISO IGRP Timers

The basic timing parameters for ISO IGRP are adjustable. Because the ISO IGRP routing protocol executes a distributed, asynchronous routing algorithm, it is important that these timers be the same for all routers in the network.

To adjust ISO IGRP timing parameters, perform the following task in router configuration mode:

Task	Command
Adjust the ISO IGRP timers (in seconds).	timers basic <i>update-interval holddown-interval invalid-interval</i>

Enable or Disable Split Horizon

Split horizon blocks information about routes from being advertised out the interface from which that information originated. This feature usually optimizes communication among multiple routers, particularly when links are broken.

To either enable or disable split horizon for ISO IGRP updates, perform the following tasks in interface configuration mode:

Task	Command
Enable split horizon for ISO IGRP updates.	cls split-horizon
Disable split horizon for ISO IGRP updates.	no cls split-horizon

The default for all LAN interfaces is for split horizon to be enabled; the default for WAN interfaces on X.25, Frame Relay, or SMDS networks is for split horizon to be disabled.

Configure IS-IS Dynamic Routing

IS-IS is a dynamic routing specification described in ISO 10589. Cisco's implementation of IS-IS allows you to configure IS-IS as an ISO CLNS routing protocol.

To configure IS-IS, complete the tasks in the following sections. Only enabling IS-IS is required; the remainder of the tasks are optional, although you might be required to perform them depending upon your specific application.

- Enable IS-IS
- Assign Multiple Area Addresses to IS-IS Areas
- Configure IS-IS Parameters
- Configure IS-IS Interface Parameters

In addition, you can also configure the following miscellaneous features described later in this chapter:

- Filter routing information—See the “Create Packet-Forwarding Filters and Establish Adjacencies” section.
- Redistribute routing information from one routing process to another—See the “Redistribute Roting Information” section.
- Configure administrative distances—See the “Specify Preferred Routes” section.

Enable IS-IS

To configure IS-IS dynamic routing, you must enable the IS-IS routing process, identify the address for the router, and specify the interfaces that are to route IS-IS. CLNS routing is enabled by default when you configure IS-IS dynamic routing. You can specify *only one* IS-IS process per router.

To configure IS-IS dynamic routing on the router, perform the following tasks in global configuration mode:

Task	Command
Step 1 Enable IS-IS routing and enter router configuration mode.	router isis [<i>tag</i>]
Step 2 Configure the NET for the routing process.	net <i>network-entity-title</i>

You can assign a meaningful name for the routing process by using the *tag* option. You can also specify a name for a NET in addition to an address. For information on how to assign a name, see the “Specify Shortcut NSAP Addresses” section later in this chapter.

To configure IS-IS dynamic routing on an interface, perform the following task in interface configuration mode:

Task	Command
Specify the interfaces that should be actively routing IS-IS.	<code>clns router isis [tag]</code>

Note For IS-IS, multiple NETs per router are allowed, with a maximum of three. However, only one IS-IS process is allowed, whether you run it in integrated mode, ISO CLNS only, or IP only.

See the “IS-IS Routing Configuration Examples” section at the end of this chapter for examples of configuring IS-IS routing.

Assign Multiple Area Addresses to IS-IS Areas

IS-IS routing supports the assignment of multiple area addresses on the same router. This concept is referred to as *multihoming*. Multihoming provides a mechanism for smoothly migrating network addresses, as follows:

- Splitting up an area—Nodes within a given area can accumulate to a point that they are difficult to manage, cause excessive traffic, or threaten to exceed the usable address space for an area. Multiple area addresses can be assigned so that you can smoothly partition a network into separate areas without disrupting service.
- Merging areas—Use transitional area addresses to merge as many as three separate areas into a single area that share a common area address.
- Transition to a different address—You may need to change an area address for a particular group of nodes. Use multiple area addresses to allow incoming traffic intended for an old area address to continue being routed to associated nodes.

You must statically assign the multiple area addresses on the router. Cisco currently supports assignment of up to three area addresses on a router. The number of areas allowed in a domain is unlimited.

All the addresses must have the same system ID. For example, you can assign one address (*area1* plus system ID), and two additional addresses in different areas (*area2* plus system ID and *area3* plus system ID) where the system ID is the same.

A router can dynamically learn about any adjacent router. As part of this process, the routers inform each other of their area addresses. If two routers share at least one area address, the set of area addresses of the two routers are merged. The merged set cannot contain more than three addresses. If there are more than three, the three addresses with the lowest numerical values are kept, and all others are dropped.

To configure multiple area addresses in IS-IS areas, perform the following tasks beginning in global configuration mode:

Task	Command
Step 1 Enable IS-IS routing and enter router configuration mode.	router isis <i>[tag]</i>
Step 2 Configure NETs for the routing process. The router can have up to three NETs. Enter each command separately.	net <i>network-entity-title</i>

See the “NETs Configuration Examples” section at the end of this chapter for examples of configuring NETs and multiple area addresses.

Configure IS-IS Parameters

Cisco’s IS-IS implementation allows you to customize certain IS-IS parameters. You can perform the optional tasks discussed in the following sections:

- Specify Router-Level Support.
- Configure IS-IS Authentication Passwords.
- Ignore IS-IS Link-State Packet (LSP) Errors.
- Log Adjacency State Changes.
- Change IS-IS LSP MTU Size.

Specify Router-Level Support

It is seldom necessary to configure the IS type because the IS-IS protocol will automatically establish this. However, 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 configure the IS-IS level, perform the following task in router configuration mode:

Task	Command
Configure the IS-IS level at which the router is to operate.	is-type { level-1 level-1-2 level-2-only }

Configure IS-IS Authentication Passwords

You can assign authentication passwords to areas and domains. An area password is inserted in Level 1 (station router) link-state PDUs (LSPs), complete sequence number PDUs (CSNPs), and partial sequence number PDUs (PSNPs). A routing domain authentication password is inserted in Level 2 (area router) LSP, CSNP, and PSNP.

To configure area or domain passwords, perform the following tasks in router configuration mode:

Task	Command
Configure the area authentication password.	area-password <i>password</i>
Configure the routing domain authentication password.	domain-password <i>password</i>

Ignore IS-IS Link-State Packet (LSP) Errors

You can configure the router to ignore IS-IS LSPs that are received with internal checksum errors, rather than purging the LSPs. LSPs are used by the receiving routers to maintain their routing tables.

The IS-IS protocol definition requires that a received LSP with an incorrect data-link checksum be purged by the receiver, which causes the initiator of the LSP to regenerate it. However, if a network has a link that causes data corruption while still delivering LSPs with correct data-link checksums, a continuous cycle of purging and regenerating large numbers of LSPs can occur, rendering the network nonfunctional.

To allow the router to ignore LSPs with an internal checksum error, perform the following task in router configuration mode:

Task	Command
Ignore LSPs with internal checksum errors rather than purging the LSPs.	ignore-lsp-errors

Log Adjacency State Changes

You can configure IS-IS to generate a log message when an IS-IS adjacency changes state (up or down). This may be useful when monitoring large networks. Messages are logged using the system error message facility. Messages are of the following form:

```
%CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0034 (Serial0) Up, new adjacency
%CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0034 (Serial0) Down, hold time expired
```

To generate log messages when an IS-IS adjacency changes state, perform the following task in router configuration mode:

Task	Command
Log IS-IS adjacency state changes.	log-adjacency-changes

Change IS-IS LSP MTU Size

Under normal conditions, the default maximum transmission unit (MTU) size should be sufficient. However, if the MTU of a link is lowered to less than 1500 bytes, the LSP MTU must be lowered accordingly on each router in the network. If this is not done, routing will become unpredictable.

The MTU size must be less than or equal to the smallest MTU of any link in the network. The default size is 1497 bytes.



Caution The CLNS MTU of a link (which is the applicable value for IS-IS, even if it is being used to route IP) may differ from the IP MTU. To be certain about a link MTU as it pertains to IS-IS, use the **show clns interface** command to display the value.

To change the MTU size of IS-IS link state packets, perform the following task in router configuration mode:

Task	Command
Specify the maximum LSP packet size, in bytes.	lsp-mtu size

Note This rule applies for all routers in a network. If any link in the network has a reduced MTU, all routers must be changed, not just the routers directly connected to the link.

Configure IS-IS Interface Parameters

Cisco's IS-IS implementation allows you to customize certain interface-specific IS-IS parameters. You can perform the optional tasks discussed in the following sections:

- Adjust IS-IS Link-State Metrics.
- Set the Advertised Hello Interval and Hello Multiplier.
- Set the Advertised CSNP Interval.
- Set the Retransmission Interval.
- Specify Designated Router Election.
- Specify the Interface Circuit Type.
- Configure IS-IS Password Authentication.

You are not required to alter any of these parameters, but some interface parameters must be consistent across all routers in the network. Therefore, be sure that if you do configure any of these parameters, the configurations for all routers on the network have compatible values.

Adjust IS-IS Link-State Metrics

You can configure a cost for a specified interface. The default metric is used as a value for the IS-IS metric. This is the value assigned when there is no quality of service (QOS) routing performed. The only metric that is supported by the Cisco IOS software and that you can configure is the *default-metric*, which you can configure for Level 1 or Level 2 routing or both.

To configure the link state metric, perform the following task in interface configuration mode:

Task	Command
Configure the metric (or cost) for the specified interface.	isis metric <i>default-metric</i> { level-1 level-2 }

Set the Advertised Hello Interval and Hello Multiplier

You can specify the length of time (in seconds) between hello packets that the Cisco IOS software sends on the interface. You can also change the default hello packet multiplier used on the interface to determine the hold time transmitted in IS-IS hello packets (the default is 3).

The hold time determines how long a neighbor waits for another hello packet before declaring the neighbor down. This time determines how quickly a failed link or neighbor is detected so that routes can be recalculated.

To set the advertised hello interval and multiplier, perform the following tasks in interface configuration mode:

Task	Command
Specify the length of time, in seconds, between hello packets the software sends on the specified interface.	isis hello-interval <i>seconds</i> { level-1 level-2 }
Specify the number used to multiply the hello interval seconds by to determine the total holding time transmitted in the IS-IS hello packet. If not specified, a multiplier of 3 is used.	isis hello-multiplier <i>multiplier</i> [{ level-1 level-2 }]

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because there is only a single type of hello packet sent on serial links, the hello packet is independent of Level 1 or Level 2.) Specify an optional level for X.25, SMDS, and Frame Relay multiaccess networks.

Use the **isis hello-multiplier** command in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval (**isis hello-interval** command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.

Set the Advertised CSNP Interval

CSNPs are sent by the designated router to maintain database synchronization.

You can configure the IS-IS CSNP interval for the interface by performing the following task in interface configuration mode:

Task	Command
Configure the IS-IS CSNP interval for the specified interface.	isis csnp-interval <i>seconds</i> { level-1 level-2 }

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

Set the Retransmission Interval

You can configure the number of seconds between retransmission of LSPs for point-to-point links.

To set the retransmission level, perform the following task in interface configuration mode:

Task	Command
Configure the number of seconds between retransmission of IS-IS LSPs for point-to-point links.	isis retransmit-interval <i>seconds</i>

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

Specify 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. The designated router enables a reduction in the number of adjacencies required on a multiaccess network, which in turn reduces the amount of routing protocol traffic and the size of the topology database.

To configure the priority to use for designated router election, perform the following task in interface configuration mode:

Task	Command
Configure the priority to use for designated router election.	isis priority <i>value</i> { level-1 level-2 }

Specify the Interface Circuit Type

It is normally not necessary to configure this feature because the IS-IS protocol automatically determines area boundaries and keeps Level 1 and Level 2 routing separate. However, you can specify the adjacency levels on a specified interface.

To configure the adjacency for neighbors on the specified interface, perform the following task in interface configuration mode:

Task	Command
Configure the type of adjacency desired for neighbors on the specified interface (specify the interface circuit type).	isis circuit-type { level-1 level-1-2 level-2-only }

If you specify Level 1, a Level 1 adjacency is established if there is at least one area address common to both this node and its neighbors.

If you specify both Level 1 and Level 2 (the default value), a Level 1 and 2 adjacency is established if the neighbor is also configured as both Level 1 and Level 2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established.

If you specify Level 2 only, a Level 2 adjacency is established. If the neighbor router is a Level 1 router, no adjacency is established.

Configure IS-IS Password Authentication

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

To configure an authentication password for an interface, perform the following task in interface configuration mode:

Task	Command
Configure the authentication password for an interface.	isis password <i>password</i> { level-1 level-2 }

Configure CLNS Static Routing

You do not need to explicitly specify a routing process to use static routing facilities. You can enter a specific static route and apply it globally, even if you have configured the router for ISO IGRP or IS-IS dynamic routing.

To configure a static route, complete the tasks in the following sections. Only enabling CLNS is required; the remaining tasks are optional, although you might be required to perform them depending upon your specific application.

- Enable Static Routes
- Configure Variations of the Static Route
- Map NSAP Addresses to Media Addresses

Enable Static Routes

To configure static routing, you must enable CLNS on the router and on the interface. CLNS routing is enabled on the router by default when you configure ISO IGRP or IS-IS routing protocols. NSAP addresses that start with the NSAP prefix you specify are forwarded to the next-hop node.

To configure CLNS on the router, perform the following tasks in global configuration mode:

Task	Command
Step 1 Configure CLNS.	clns routing
Step 2 Assign an NSAP address to the router if the router has not been configured to route CLNS packets dynamically using ISO IGRP or IS-IS.	clns net { <i>net-address</i> <i>name</i> }
Step 3 Enter a specific static route.	clns route <i>nsap-prefix</i> { <i>next-hop-net</i> <i>name</i> }

Note If you have not configured the router to route CLNS packets dynamically using ISO IGRP or IS-IS, you must assign an address to the router.

You also must enable ISO CLNS for each interface you want to pass ISO CLNS packet traffic to end systems, but for which you do not want to perform any dynamic routing on the interface. This is done automatically when you configure IS-IS or ISO IGRP routing on an interface; however, if you do not intend to perform any dynamic routing on an interface, you must manually enable CLNS. You can assign an NSAP address for a specific interface. This allows the Cisco IOS software to advertise different addresses on each interface. This is useful if you are doing static routing and need to control the source NET used by the router on each interface.

To configure CLNS on an interface, perform the following tasks in interface configuration mode:

Task	Command
Step 1 Enable ISO CLNS for each interface.	clns enable
Step 2 Optionally, assign an NSAP address to a specific interface.	clns net { <i>nsap-address</i> <i>name</i> }

See the “Basic Static Routing Examples,” “Static Intradomain Routing Example,” and “Static Interdomain Routing Example” sections at the end of this chapter for examples of configuring static routes.

Configure Variations of the Static Route

You can perform the following tasks that use variations of the **clns route** global configuration command:

- Bind the next hop to a specified interface and media address when you do not know the NSAP address of your neighbor. Note that this version of the **clns route** command is not literally *applied* to a specific interface.
- Discard packets with a specific NSAP prefix that is outside the domain (ISO-IGRP) or area (IS-IS) of the router.
- Specify a default prefix.

To enter a specific static route, discard packets, or configure a default prefix, perform the following tasks in global configuration mode:

Note To discard or filter packets that have an NSAP prefix within the domain (ISO-IGRP) or area (IS-IS) of the router, refer to the “Create Packet-Forwarding Filters and Establish Adjacencies” section of this chapter

Task	Command
Enter a specific static route for a specific interface.	clns route <i>nsap-prefix type number [snpa-address]</i>
Explicitly tell the software to discard packets with the specified NSAP prefix.	clns route <i>nsap-prefix discard</i>
Configure a default prefix rather than specify an NSAP prefix.	clns route default <i>nsap-prefix type number</i>

Map NSAP Addresses to Media Addresses

Conceptually, each ES lives in one area. It discovers the nearest IS by listening to ES-IS packets. Each ES must be able to communicate directly with an IS in its area.

When an ES wants to communicate with another ES, it sends the packet to any IS on the same medium. The IS looks up the destination NSAP address and forwards the packet along the best route. If the destination NSAP address is for an ES in another area, the Level 1 IS sends the packet to the nearest Level 2 IS. The Level 2 IS forwards the packet along the best path for the destination area until it gets to a Level 2 IS that is in the destination area. This IS then forwards the packet along the best path inside the area until it is delivered to the destination ES.

ESs need to know how to get to a Level 1 IS for their area, and Level 1 ISs need to know all of the ESs that are directly reachable through each of their interfaces. To provide this information, the routers support the ES-IS protocol. The router dynamically discovers all ESs running the ES-IS protocol. ESs that are not running the ES-IS protocol must be configured statically.

It is sometimes desirable for a router to have a neighbor configured statically rather than learned through ES-IS, ISO IGRP, or IS-IS.

Note It is necessary to use static mapping only for ESs that do *not* support ES-IS. The Cisco IOS software continues to dynamically discover ESs that *do* support ES-IS.

Note If you have configured interfaces for ISO IGRP or IS-IS, the ES-IS routing software automatically turns on ES-IS for those interfaces.

Perform the following tasks in interface configuration mode, as needed, to enter static mapping information between the NSAP protocol addresses and the subnetwork point of attachment (SNPA) addresses (media) for ESs or ISs:

Task	Command
Configure all end systems that will be used when you manually specify the NSAP-to-SNPA mapping.	clns es-neighbor <i>nsap snpa</i>
Configure all intermediate systems that will be used when you manually specify the NSAP-to-SNPA mapping.	clns is-neighbor <i>nsap snpa</i>

For more information, see the “Configure CLNS over WANs” section later in this chapter.

Note The SNPA is a data link layer address (such as an Ethernet address, X.25 address, or Frame Relay DLCI address) used to configure a CLNS route for an interface.

Configure Miscellaneous Features

Perform the optional tasks in the following sections to configure miscellaneous features of an ISO CLNS network:

- Specify Shortcut NSAP Addresses
- Use the IP Domain Name System to Discover ISO CLNS Addresses
- Create Packet-Forwarding Filters and Establish Adjacencies
- Redistribute Routing Information
- Specify Preferred Routes
- Configure ES-IS Hello Packet Parameters
- Configure DECnet OSI or Phase V Cluster Aliases
- Configure Digital-Compatible Mode
- Allow Security Option Packets to Pass

Specify Shortcut NSAP Addresses

You can define a name-to-NSAP address mapping. This name can then be used in place of typing the long set of numbers associated with an NSAP address.

To define a name-to-NSAP address mapping, perform the following task in global configuration mode:

Task	Command
Define a name-to-NSAP address mapping.	clns host <i>name nsap</i>

The assigned NSAP name is displayed, where applicable, in **show** and **debug EXEC** commands. However, some effects and requirements are associated with using names to represent NETs and NSAP addresses.

The **clns host** global configuration command is generated after all other CLNS commands when the configuration file is parsed. As a result, you cannot edit the nonvolatile random access memory (NVRAM) version of the configuration to specifically change the address defined in the original **clns host** command. You must specifically change any commands that refer to the original address. This affects all commands that accept names.

The commands that are affected by these requirements include the following:

- **net** (router configuration command)
- **clns is-neighbor** (interface configuration command)
- **clns es-neighbor** (interface configuration command)
- **clns route** (global configuration command)

Use the IP Domain Name System to Discover ISO CLNS Addresses

If your router has both ISO CLNS and IP enabled, you can use the Domain Naming System (DNS) to query ISO CLNS addresses by using the NSAP address type, as documented in RFC 1348. This feature is useful for the ISO CLNS **ping EXEC** command and when making Telnet connections. This feature is enabled by default.

To enable or disable DNS queries for ISO CLNS addresses, perform the following tasks in global configuration mode:

Task	Command
Enable DNS queries for CLNS addresses.	ip domain-lookup nsap
Disable DNS queries for CLNS addresses.	no ip domain-lookup nsap

Create Packet-Forwarding Filters and Establish Adjacencies

You can build powerful CLNS filter expressions, or access lists. These can be used to control either the forwarding of frames through router interfaces, or the establishment of adjacencies with, or the application of filters to, any combination of ES or IS neighbors, ISO IGRP neighbors, or IS-IS neighbors.

CLNS filter expressions are complex logical combinations of CLNS filter sets. CLNS filter sets are lists of address templates against which CLNS addresses are matched. Address templates are CLNS address *patterns* that are either simple CLNS addresses that match just one address, or match multiple CLNS addresses through the use of wildcard characters, prefixes, and suffixes. Frequently used address templates can be given *aliases* for easier reference.

To establish CLNS filters, perform the following tasks in global configuration mode:

Task	Command
Create aliases for frequently used address templates.	clns template-alias <i>name template</i>
Build filter sets of multiple address template permit and deny conditions.	clns filter-set <i>sname</i> [permit deny] <i>template</i>
Build filter expressions, using one or more filter sets.	clns filter-expr <i>ename term</i>

To apply filter expressions to an interface, perform the following tasks in interface configuration mode:

Task	Command
Apply a filter expression to frames forwarded in or out of an interface.	clns access-group <i>name</i> [in out]
Apply a filter expression to IS-IS adjacencies.	isis adjacency-filter <i>name</i> [match-all]
Apply a filter expression to ISO IGRP adjacencies.	iso-igrp adjacency-filter <i>name</i>
Apply a filter expression to ES or IS adjacencies.	clns adjacency-filter { es is } <i>name</i>

See the “CLNS Filter Example” section at the end of this chapter for examples of configuring CLNS filters.

Redistribute Routing Information

In addition to running multiple routing protocols simultaneously, the Cisco IOS software can redistribute information from one routing process to another.

You can also configure the Cisco IOS software to do interdomain dynamic routing by configuring two routing processes and two NETs (thereby putting the router into two domains) and redistributing the routing information between the domains. Routers configured this way are referred to as *border* routers. If you have a router that is in two routing domains, you might want to redistribute routing information between the two domains.

Note It is not necessary to use redistribution between areas. Redistribution only occurs for Level 2 routing.

To configure the router to redistribute routing information into the ISO IGRP domain, perform the following tasks beginning in global configuration mode:

Task	Command
Specify the routing protocol and tag (if applicable) into which you want to distribute routing information.	router iso-igrp [<i>tag</i>]
Specify one or more ISO IGRP routing protocol and tag (if applicable) you want to redistribute.	redistribute iso-igrp [<i>tag</i>] [route-map <i>map-tag</i>]
Specify the IS-IS routing protocol and tag (if applicable) you want to redistribute.	redistribute isis [<i>tag</i>] [route-map <i>map-tag</i>]
Specify the static routes you want to redistribute.	redistribute static [clns ip]

To configure the router to redistribute routing information into the IS-IS domains, perform the following tasks beginning in global configuration mode:

Task	Command
Specify the routing protocol and tag (if applicable) into which you want to distribute routing information.	router isis [<i>tag</i>]
Specify the IS-IS routing protocol and tag (if applicable) you want to redistribute.	redistribute isis [<i>tag</i>] [route-map <i>map-tag</i>]

Note By default, static routes are redistributed into IS-IS.

You can conditionally control the redistribution of routes between routing domains by defining *route maps* between the two domains. Route maps allow you to use tags in routes to influence route redistribution.

To conditionally control the redistribution of routes between domains, perform the following task in global configuration mode:

Task	Command
Define any route maps needed to control redistribution.	route-map <i>map-tag</i> { permit deny } <i>sequence-number</i>

One or more **match** command and one or more **set** commands typically follow a **route-map** command to define the conditions for redistributing routes from one routing protocol into another. If there are no **match** commands, everything matches. If there are no **set** commands, nothing is done (other than the match).

Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the redistribution *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. When all match criteria are met, all **set** actions are performed.

The **match route-map** configuration command has multiple formats. The **match** commands may be given in any order, and *all* defined match criteria must be satisfied to cause the route to be redistributed according to the *set actions* given with the **set** commands.

To define the match criteria for redistribution of routes from one routing protocol into another, perform at least one of the following tasks in route-map configuration mode:

Task	Command
Match routes that have a network address matching one or more of the specified names (the names can be a standard access list, filter set, or expression).	match clns address <i>name</i> [<i>name...name</i>]
Match routes that have a next hop address matching one or more of the specified names (the names can be a standard access list, filter set, or expression).	match clns next-hop <i>name</i> [<i>name...name</i>]
Match routes that have been advertised by routers matching one or more of the specified names (the names can be a standard access list, filter set, or expression).	match clns route-source <i>name</i> [<i>name...name</i>]

Task	Command
Match routes that have the next hop out matching one or more of the specified interfaces.	match clns interface <i>type number</i> [<i>type number...type number</i>]
Match routes that have the specified metric.	match metric <i>metric-value</i>
Match routes that have the specified route type.	match route-type { level-1 level-2 }

To define set actions for redistribution of routes from one routing protocol into another, perform at least one of the following set tasks in route-map configuration mode:

Task	Command
Set the routing level of the routes to be advertised into a specified are of the routing domain.	set level { level-1 level-2 level-1-2 }
Set the metric value to give the redistributed routes.	set metric <i>metric-value</i>
Set the metric type to give the redistributed routes.	set metric-type { internal external }
Set the tag value to associate with the redistributed routes.	set tag <i>tag-value</i>

See the “Dynamic Interdomain Routing Example” and “TARP Configuration Examples” sections at the end of this chapter for examples of configuring route maps.

Specify 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.

If the router is forwarding packets, dynamic routes will always take priority over static routes, unless the router is routing to a destination outside of its domain and area. The router first will look for an ISO IGRP route within in own area, then for an ISO IGRP route within its own domain, and finally for an IS-IS route within its own area, until it finds a matching route. If a matching route still has not been found, the router will check its prefix table, which contains static routes and routes to destinations outside the area (ISO IGRP), domain (ISO IGRP), and area (IS-IS) routes for that router. When the router is using its prefix table it will choose the route that has the lowest administrative distance.

By default, the following administrative distances are assigned:

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

When you change an administrative distance for a routing process, use the following command in router configuration mode:

Task	Command
Specify preferred routes by setting the lowest administrative distance.	distance <i>value</i> [clns]

Note If you want an ISO IGRP prefix route to override a static route, you must set the administrative distance for the routing process to be lower than 10 (assigned administrative distance for static routes). You cannot change the administrative distance for static routes.

Configure ES-IS Hello Packet Parameters

You can configure ES-IS parameters for communication between end systems and routers. In general, you should leave these parameters at their default values.

When configuring an ES-IS router, be aware of the following:

- ES-IS does not run over X.25 links unless the *broadcast* facility is enabled.
- ES hello packets and IS hello packets are sent without options. Options in received packets are ignored.

ISs and ESs periodically send out hello packets to advertise their availability. The frequency of these hello packets can be configured.

The recipient of a hello packet creates an adjacency entry for the system that sent it. If the next hello packet is not received within the interval specified, the adjacency times out and the adjacent node is considered unreachable.

A default rate has been set for hello packets and packet validity; however, you can change the defaults by performing the following tasks in global configuration mode:

Task	Command
Specify the rate at which ES hello and IS hello packets are sent.	clns configuration-time <i>seconds</i>
Allow the sender of an ES hello or IS hello packet to specify the length of time you consider the information in these packets to be valid.	clns holding-time <i>seconds</i>

A default rate has been set for the ES Configuration Timer (ESCT) option; however, you can change the default by performing the following task in interface configuration mode:

Task	Command
Specify how often the end system should transmit ES hello packet PDUs.	clns esct-time <i>seconds</i>

Configure DECnet OSI or Phase V Cluster Aliases

DECnet Phase V *cluster aliasing* allows multiple systems to advertise the same system ID in end-system hello packets. The Cisco IOS software does this by caching multiple ES adjacencies with the same NSAP address, but different SNPA addresses. When a packet is destined to the common NSAP address, the software splits the packet loads among the different SNPA addresses. A router that supports this capability forwards traffic to each system. You can do this on a per-interface basis.

To configure cluster aliases, perform the following task in interface configuration mode:

Task	Command
Allow multiple systems to advertise the same system ID in end-system hello packets.	clns cluster-alias

If DECnet Phase V cluster aliases are disabled on an interface, ES hello packet information is used to replace any existing adjacency information for the NSAP address. Otherwise, an additional adjacency (with a different SNPA) is created for the same NSAP address.

See the “TARP Configuration Examples” section at the end of this chapter for an example of configuring DECnet OSI cluster aliases.

Configure Digital-Compatible Mode

If you have an old DECnet implementation of ES-IS in which the NSAP address advertised in an IS hello packet does not have the N-selector byte present, you may want to configure the Cisco IOS software to allow IS hello packets sent and received to ignore the N-selector byte. The N-selector byte is the last byte of the NSAP address.

To enable Digital-compatible mode, perform the following task in interface configuration mode:

Task	Command
Allow IS hello packets sent and received to ignore the N-selector byte.	clns dec-compatible

Allow Security Option Packets to Pass

By default, the Cisco IOS software discards any packets with security options set. You can disable this behavior. To allow such packets to pass through, perform the following task in global configuration mode:

Task	Command
Allow the software to accept any packets it sees as set with security options.	clns security pass-through

Note The ISO CLNS routing software ignores the Record Route option, the Source Route option, and the QOS option other than congestion experienced. The Security option causes a packet to be rejected with a bad option indication.

Configure CLNS over WANs

This section provides general information about running ISO CLNS over WANs.

You can use CLNS routing on serial interfaces with HDLC, PPP, LAPB, X.25, Frame Relay, DDR, or SMDS encapsulation. Both incoming and outgoing CLNS packets can be fast switched over PPP.

To use HDLC encapsulation, you must have a router at both ends of the link. If you use X.25 encapsulation, and if IS-IS or ISO IGRP is not used on an interface, you must manually enter the NSAP-to-X.121 address mapping. The LAPB, SMDS, Frame Relay, and X.25 encapsulations interoperate with other vendors.

Both ISO IGRP and IS-IS can be configured over WANs.

X.25 is not a broadcast medium and, therefore, does not broadcast protocols (such as ES-IS) that automatically advertise and record mappings between NSAP/NET (protocol addresses) and SNPA (media addresses). (With X.25, the SNPAs are the X.25 network addresses, or the X.121 addresses. These are usually assigned by the X.25 network provider.) If you use static routing, you must configure the NSAP-to-X.121 address mapping with the **x25 map** command.

Configuring a serial line to use CLNS over X.25 requires configuring the general X.25 information and the CLNS-specific information. First, configure the general X.25 information. Then, enter the CLNS static mapping information.

You can specify X.25 nondefault packet and window sizes, reverse charge information, and so on. The X.25 facilities information that can be specified is exactly the same as in the **x25 map** interface configuration command described in the “Configuring X.25 and LAPB” chapter in the *Wide-Area Networking Configuration Guide*.

See the “ISO CLNS over X.25 Example” section at the end of this chapter for an example of configuring CLNS over X.25.

Enhance ISO CLNS Performance

Generally, you do not need to change the router’s default settings for CLNS packet switching, but there are some modifications you can make when you decide to make changes in your network’s performance. The following sections describe ISO CLNS parameters that you can change:

- Specify the MTU Size
- Disable Checksums
- Disable Fast Switching through the Cache
- Set the Congestion Threshold
- Transmit Error Protocol Data Units (ERPDU)s
- Control Redirect Protocol Data Units (RDPDU)s
- Configure Parameters for Locally Sourced Packets

See the “Performance Parameters Example” section at the end of this chapter for examples of configuring various performance parameters.

Specify the MTU Size

All interfaces have a default maximum packet size. However to reduce fragmentation, you can set the MTU size of the packets sent on the interface. The minimum value is 512; the default and maximum packet size depends on the interface type.

Changing the MTU value with the **mtu** interface configuration command can affect the CLNS MTU value. If the CLNS MTU is at its maximum given the interface MTU, the CLNS MTU will change with the interface MTU. However, the reverse is not true; changing the CLNS MTU value has no effect on the value for the **mtu** interface configuration command.

To set the CLNS MTU packet size for a specified interface, perform the following task in interface configuration mode:

Task	Command
Set the MTU size of the packets sent on the interface.	clns mtu size

Note The CTR card does not support the switching of frames larger than 4472 bytes. Interoperability problems might occur if CTR cards are intermixed with other Token Ring cards on the same network. These problems can be minimized by lowering the CLNS MTU sizes to be the same on all routers on the network.

Disable Checksums

When the ISO CLNS routing software originates a CLNS packet, by default it generates checksums. To disable this function, perform the following task in interface configuration mode:

Task	Command
Disable checksum generation.	no clns checksum

Note Enabling checksum generation has no effect on routing packets (ES-IS, ISO IGRP, and IS-IS) originated by the router; it applies to pings and traceroute packets.

Disable Fast Switching through the Cache

Fast switching through the cache is enabled by default for all supported interfaces. To disable fast switching, perform the following task in interface configuration mode:

Task	Command
Disable fast switching.	no clns route-cache

Note The cache still exists and is used after the **no clns route-cache** interface configuration command is used; the software just does not do fast switching through the cache.

Set the Congestion Threshold

If a router configured for CLNS experiences congestion, it sets the congestion-experienced bit. You can set the congestion threshold on a per-interface basis. By setting this threshold, you cause the system to set the congestion-experienced bit if the output queue has more than the specified number of packets in it.

To set the congestion threshold, perform the following task in interface configuration mode:

Task	Command
Set the congestion threshold.	clns congestion-threshold <i>number</i>

Transmit Error Protocol Data Units (ERPDUs)

When a CLNS packet is received, the routing software looks in the routing table for the next hop. If it does not find one, the packet is discarded and an error protocol data unit (ERPDU) is sent.

You can set an interval between ERPDU. Doing so reduces bandwidth if this feature is disabled. When you set the minimum interval between ERPDU, the Cisco IOS software does not send ERPDU more frequently than one per interface per ten milliseconds.

To transmit ERPDU, perform the following tasks in interface configuration mode:

Task	Command
Send an ERPDU when the routing software detects an error in a data PDU; this is enabled by default.	clns send-erpdo
Set the minimum interval, in milliseconds, between ERPDU.	clns erpdo-interval <i>milliseconds</i>

Control Redirect Protocol Data Units (RDPDUs)

If a packet is sent out the same interface it came in on, a redirect protocol data unit (RDPDU) also can be sent to the sender of the packet. You can control RDPDUs in the following ways:

- By default, CLNS sends RDPDUs when a better route for a given host is known. You can disable this feature. Disabling this feature reduces bandwidth because packets may continue to unnecessarily go through the router.
- You can set the interval times between RDPDUs.

Note SNPA masks are never sent, and RDPDUs are ignored by the Cisco IOS software when the router is acting as an IS.

To control RDPDUs, perform either of the following tasks in interface configuration mode:

Task	Command
Send redirect PDUs when a better route for a given host is known.	clns send-rpdu
Set the minimum interval time, in milliseconds, between RDPDUs.	clns rpdu-interval <i>milliseconds</i>

Configure Parameters for Locally Sourced Packets

To configure parameters for packets originated by a specified router, perform either of the following tasks in global configuration mode:

Task	Command
Specify in seconds the initial lifetime for locally generated packets.	clns packet-lifetime <i>seconds</i>
Specify whether to request ERPDU on packets originated by the router.	clns want-erpdu

You should set the packet lifetime low in an internetwork that has frequent loops.

Note The **clns want-erpdu** global configuration command has no effect on routing packets (ES-IS, ISO IGRP, and IS-IS) originated by the router; it applies to pings and traceroute packets.

Monitor and Maintain the ISO CLNS Network

The following EXEC commands monitor and maintain the ISO CLNS caches, tables, and databases:

Task	Command
Clear and reinitialize the CLNS routing cache.	clear clns cache
Remove ES neighbor information from the adjacency database.	clear clns es-neighbors
Remove IS neighbor information from the adjacency database.	clear clns is-neighbors

Task	Command
Remove CLNS neighbor information from the adjacency database.	clear clns neighbors
Remove dynamically derived CLNS routing information.	clear clns route
Invoke a diagnostic tool for testing connectivity	ping clns { <i>host</i> <i>address</i> }
Display information about the CLNS network.	show clns
Display the entries in the CLNS routing cache.	show clns cache
Display ES neighbor entries, including the associated areas.	show clns es-neighbors [<i>type number</i>] [detail]
Display filter expressions.	show clns filter-expr [<i>name</i>] [detail]
Display filter sets.	show clns filter-set [<i>name</i>]
List the CLNS-specific or ES-IS information about each interface.	show clns interface [<i>type number</i>]
Display IS neighbor entries, according to the area in which they are located.	show clns is-neighbors [<i>type number</i>] [detail]
Display both ES and IS neighbors.	show clns neighbors [<i>type number</i>] [detail]
List the protocol-specific information for each IS-IS or ISO IGRP routing process in this router.	show clns protocol [<i>domain</i> <i>area-tag</i>]
Display all the destinations to which this router knows how to route packets.	show clns route [<i>nsap</i>]
Display information about the CLNS packets this router has seen.	show clns traffic
Display the IS-IS link state database.	show isis database [level-1] [level-2] [l1] [l2] [detail] [<i>lspid</i>]
Display the IS-IS Level 1 routing table.	show isis routes
Display a history of the SPF calculations for IS-IS	show isis spf-log
Display all route maps configured or only the one specified.	show route-map [<i>map-name</i>]
Discover the paths taken to a specified destination by packets in the network.	trace clns <i>destination</i>
Display the routing table in which the specified CLNS destination is found.	which-route { <i>nsap-address</i> <i>clns-name</i> }

Configure TARP on ISO CLNS

Some applications (typically used by telephone companies) running on Synchronous Optical Network (SONET) devices identify these devices by a target identifier (TID). Therefore, it is necessary for the router to cache TID-to-network address mappings. Because these applications usually run over OSI, the network addresses involved in the mapping are OSI NSAPs.

When a device must send a packet to another device it does not know about (that is, it does not have information about the NSAP address corresponding to the remote device's TID), the device needs a way to request this information directly from the device, or from an intermediate device in the network. This functionality is provided by an address resolution protocol called Target Identifier Address Resolution Protocol (TARP).

Requests for information and associated responses are sent as TARP protocol data units (PDUs), which are sent as CLNP data packets. TARP PDUs are distinguished by a unique N-selector in the NSAP address. Following are the five types of TARP PDUs:

- Type 1—Sent when a device has a TID for which it has no matching NSAP. Type 1 PDUs are sent to all Level 1 (IS-IS and ES-IS) neighbors. If no response is received within the specified time limit, a Type 2 PDU is sent. To prevent packet looping, a loop detection buffer is maintained on the router. A Type 1 PDU is sent when you use the **tarp resolve** command.
- Type 2—Sent when a device has a TID for which it has no matching NSAP and no response was received from a Type 1 PDU. Type 2 PDUs are sent to all Level 1 and Level 2 neighbors. A time limit for Type 2 PDUs can also be specified. A Type 2 PDU is sent when you use the **tarp resolve** command and specify the option 2.
- Type 3—Sent as a response to a Type 1, Type 2, or Type 5 PDU. Type 3 PDUs are sent directly to the originator of the request.
- Type 4—Sent as a notification when a change occurs locally (for example, a TID or NSAP change). A Type 4 PDU is sent when you use the **tarp query** command.
- Type 5—Sent when a device needs a TID that corresponds to a specific NSAP. Unlike Type 1 and Type 2 PDUs that are sent to all Level 1 and Level 2 neighbors, a Type 5 PDU is sent only to a particular router.

In addition to the type, TARP PDUs contain the sender's NSAP, the sender's TID, and the target's TID (if the PDU is a Type 1 or Type 2).

TARP Configuration Task List

To configure TARP on the router, complete the tasks in the following sections (only the first task is required, all other tasks are optional):

- Enable TARP and Configure a TARP TID
- Disable TARP Caching
- Disable TARP PDU Origination and Propagation
- Configure Multiple NSAP Addresses
- Configure Static TARP Adjacency and Blacklist Adjacency
- Determine TIDs and NSAPs
- Configure TARP Timers
- Configure Miscellaneous TARP PDU Information
- Monitor and Maintain the TARP Protocol

For several examples of configuring TARP, see the “TARP Configuration Examples” section at the end of this chapter.

Enable TARP and Configure a TARP TID

TARP must be explicitly enabled before the TARP functionality becomes available and the router must have a TID assigned. Also, before TARP packets can be sent out on an interface, each interface must have TARP enabled and the interface must be able to propagate TARP PDUs.

The router will use the CLNS capability to send and receive TARP PDUs. If the router is configured as an IS, the router must be running IS-IS. If the router is configured as an ES, the router must be running ES-IS.

To turn on the TARP functionality, perform the following tasks in global configuration mode:

Task	Command
Turn on the TARP functionality.	tarp run
Assign a TID to the router.	tarp tid <i>tid</i>

To enable TARP on one or more interfaces, perform the following task in interface configuration mode:

Task	Command
Enable TARP on the interface.	tarp enable

Disable TARP Caching

By default, TID-to-NSAP address mappings are stored in the TID cache. Disabling this capability clears the TID cache. Re-enabling this capability restores any previously cleared local entry and all static entries.

To disable TID-to-NSAP address mapping in the TID cache, perform the following task in global configuration mode:

Task	Command
Disable TARP TID-to-NSAP address mapping.	no tarp allow-caching

Disable TARP PDU Origination and Propagation

By default, the router originates TARP PDUs and propagates TARP PDUs to its neighbors, and the interface propagates TARP PDUs to its neighbor. Disabling these capabilities means that the router no longer originates TARP PDUs, and the router and the specific interface no longer propagate TARP PDUs received from other routers.

To disable origination and propagation of TARP PDUs, perform the following tasks in global configuration mode:

Task	Command
Disable TARP PDU origination.	no tarp originate
Disable global propagation of TARP PDUs.	no tarp global-propagate

To disable propagation of TARP PDUs on a specific interface, perform the following task in interface configuration mode:

Task	Command
Disable propagation of TARP PDUs on the interface.	no tarp propagate

Configure Multiple NSAP Addresses

A router may have more than one NSAP address. When a request for an NSAP is sent (Type 1 or Type 2 PDU), the first NSAP address is returned. To receive all NSAP addresses associated with the router, enter a TID-to-NSAP static route in the TID cache for each NSAP address.

To create a TID-to-NSAP static route, perform the following task in global configuration mode:

Task	Command
Enter a TID-to-NSAP static route.	tarp map <i>tid nsap</i>

Configure Static TARP Adjacency and Blacklist Adjacency

In addition to all its IS-IS/ES-IS adjacencies, a TARP router propagates PDUs to all its static TARP adjacencies. If a router is not running TARP, the router discards TARP PDUs rather than propagating the PDUs to all its adjacencies. To allow TARP to bypass routers enroute that may not have TARP running, TARP provides a static TARP adjacency capability. Static adjacencies are maintained in a special queue.

To create a static TARP adjacency, perform the following task in global configuration mode:

Task	Command
Enter a static TARP adjacency.	tarp route-static <i>nsap</i>

To stop TARP from propagating PDUs to an IS-IS/ES-IS adjacency that may not have TARP running, TARP provides a blacklist adjacency capability. The router will not propagate TARP PDUs to blacklisted routers.

To blacklist a router, perform the following task in global configuration mode:

Task	Command
Bypass a router not running TARP.	tarp blacklist-adjacency <i>nsap</i>

Determine TIDs and NSAPs

To determine an NSAP address for a TID or a TID for an NSAP address, perform the following tasks in EXEC mode:

Task	Command
Get the TID associated with a specific NSAP.	tarp query <i>nsap</i>
Get the NSAP associated with a specific TID.	tarp resolve <i>tid</i> [1 2]

To determine the TID, the router first checks the local TID cache. If there is a TID entry in the local TID cache, the requested information is displayed. If there is no TID entry in the local TID cache, a TARP Type 5 PDU is sent out to the specified NSAP address.

To determine the NSAP address, the router first checks the local TID cache. If there is an NSAP entry in the local TID cache, the requested information is displayed. If there is no NSAP entry in the local TID cache, a TARP Type 1 or Type 2 PDU is sent out. By default, a Type 1 PDU is sent to all Level 1 (IS-IS and ES-IS) neighbors. If a response is received, the requested information is displayed. If a response is not received within the response time, a Type 2 PDU is sent to all Level 1 and Level 2 neighbors. Specifying the EXEC command **tarp resolve** *tid* 2 causes only a Type 2 PDU to be sent.

You can configure the length of time that the router will wait for a response (in the form of a Type 3 PDU).

Configure TARP Timers

TARP timers provide default values and typically do not need to be changed.

You can configure the amount of time that the router waits to receive a response from a Type 1 PDU, a Type 2 PDU, and a Type 5 PDU. In addition, you can also configure the PDU's lifetime based on the number of hops.

You can also set timers that control how long dynamically created TARP entries remain in the TID cache, and how long the system ID-to-sequence number mapping entry remains in the loop detection buffer table. The loop detection buffer table prevents TARP PDUs from looping.

To configure TARP PDU timers, control PDU lifetime, and set how long entries remain in cache, perform the following tasks in global configuration mode:

Task	Command
Configure the number of seconds that the router will wait for a response from a TARP Type 1 PDU.	tarp t1-response-timer <i>seconds</i>
Configure the number of seconds that the router will wait for a response from a TARP Type 2 PDU.	tarp t2-response-timer <i>seconds</i>
Configure the number of seconds that the router will wait for a response from a TARP Type 2 PDU after the default timer has expired.	tarp post-t2-response-timer <i>seconds</i>
Configure the number of seconds that the router will wait for a response from a TARP Type 5 PDU.	tarp arp-request-timer <i>seconds</i>
Configure the number of routers that a TARP PDU can traverse before it is discarded.	tarp lifetime <i>hops</i>
Configure the number of seconds a dynamically-created TARP entry remains in the TID cache.	tarp cache-timer <i>seconds</i>
Configure the number of seconds that a system ID-to-sequence number mapping entry remains in the loop detection buffer table.	tarp ldb-timer <i>seconds</i>

Configure Miscellaneous TARP PDU Information

TARP default PDU values typically do not need to be changed.

You can configure the sequence number of the TARP PDU, set the update remote cache bit used to control whether the remote router updates its cache, specify the N-selector used in the PDU to indicate a TARP PDU, and specify the network protocol type used in outgoing PDUs.

To configure miscellaneous PDU information, perform the following tasks in global configuration mode:

Task	Command
Change the sequence number in the next outgoing TARP PDU.	tarp sequence-number <i>number</i>
Set the update remote cache bit in all subsequent outgoing TARP PDUs so that the remote router does or does not update the cache.	tarp urc [0 1]
Specify the N-selector used to identify TARP PDUs.	tarp selector <i>hex-digit</i>
Specify the protocol type used in outgoing TARP PDUs. Only FE (to indicate CLNP) is supported.	tarp protocol <i>hex-digit</i>

Monitor and Maintain the TARP Protocol

Use the following EXEC commands to monitor and maintain the TARP caches, tables, and databases:

Task	Command
Reset the TARP counters that are shown with the show tarp traffic command.	clear tarp counters
Remove all system ID-to-sequence number mapping entries in the TARP loop detection buffer table.	clear tarp ldb-table
Remove all dynamically created TARP TID-to-NSAP address mapping entries in the TID cache.	clear tarp tid-table
Display all global TARP parameters.	show tarp
List all adjacencies that are blacklisted (that is, adjacencies that will not receive propagated TARP PDUs).	show tarp blacklisted-adjacencies
Display information about a specific TARP router stored in the local TID cache.	show tarp host <i>tid</i>
List all interfaces on the router that have TARP enabled.	show tarp interface [<i>type number</i>]
Display the contents of the loop detection buffer table.	show tarp ldb
List all the static entries in the TID cache.	show tarp map
List all static TARP adjacencies.	show tarp static-adjacencies
Display information about the entries in the TID cache.	show tarp tid-cache
Display statistics about TARP PDUs.	show tarp traffic

ISO CLNS Configuration Examples

The following sections provide configuration examples of both intra- and interdomain static and dynamic routing using static, ISO IGRP, and IS-IS routing techniques:

- NETs Configuration Examples
- Dynamic Routing within the Same Area Example
- Dynamic Routing in More Than One Area Example
- Dynamic Routing in Overlapping Areas Example
- Dynamic Interdomain Routing Example
- IS-IS Routing Configuration Examples
- Router in Two Areas Example
- Basic Static Routing Examples
- Static Intradomain Routing Example
- Static Interdomain Routing Example
- CLNS Filter Example
- Route Map Examples
- DECnet Cluster Aliases Example
- ISO CLNS over X.25 Example
- Performance Parameters Example
- TARP Configuration Examples

NETs Configuration Examples

The following are simple examples of configuring NETs for both ISO IGRP and IS-IS.

ISO IGRP

The following example illustrates specifying an NET:

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

The following example illustrates using a name for an NET:

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

The use of this **net** router configuration command configures the system ID, area address, and domain address. Only a single NET per routing process is allowed.

```
router iso-igrp local
 net 49.0001.0000.0c00.1111.00
```

IS-IS

The following example illustrates specifying a single NET:

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

The following example illustrates using a name for an NET:

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

IS-IS Multihoming Example

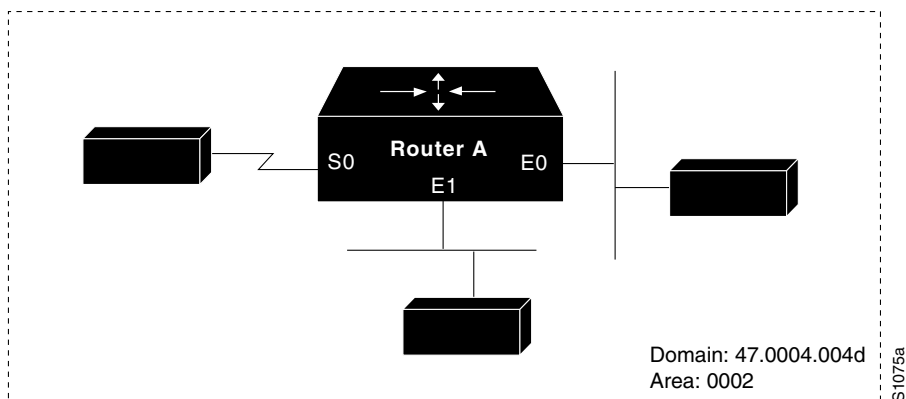
The following example illustrates the assignment of three separate area addresses for a single router using **net** commands. Traffic received that includes an area address of 47.0004.004d.0001, 47.0004.004d.0002, or 47.0004.004d.0003, and that has the same system ID, is forwarded to this router.

```
router isis eng-area1
! | IS-IS Area | System ID | S |
net 47.0004.004d.0001.0000.0c00.1111.00
net 47.0004.004d.0002.0000.0c00.1111.00
net 47.0004.004d.0003.0000.0c00.1111.00
```

Dynamic Routing within the Same Area Example

Figure 19 and the example configuration that follows illustrate how to configure dynamic routing within a routing domain. The router can exist in one or more areas within the domain. The router named Router A exists in a single area.

Figure 19 CLNS Dynamic Routing within a Single Area



```
! define a tag castor for the routing process
router iso-igrp castor
! configure the net for the process in area 2, domain 47.0004.004d
 net 47.0004.004d.0002.0000.0c00.0506.00
! specify iso-igrp routing using the previously specified tag castor
interface ethernet 0
 clns router iso-igrp castor
! specify iso-igrp routing using the previously specified tag castor
interface ethernet 1
 clns router iso-igrp castor
```

```

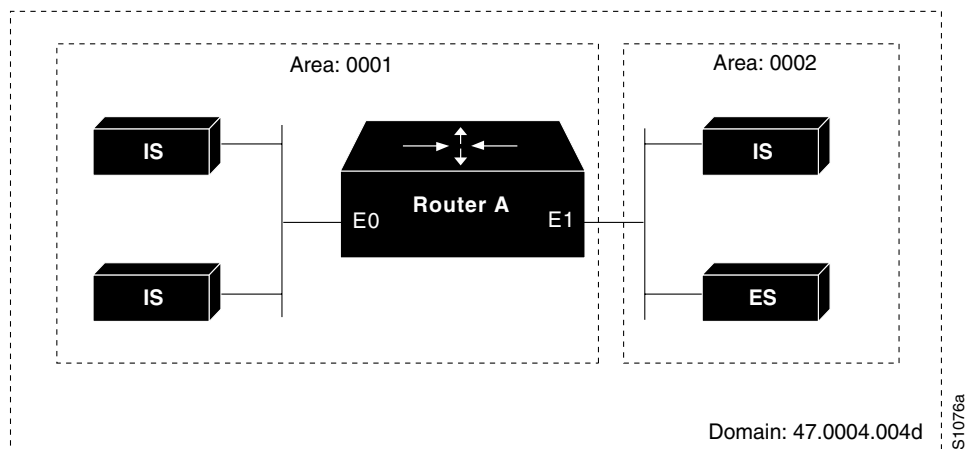
! specify iso-igrp routing using the previously specified tag castor
interface serial 0
  clns router iso-igrp castor

```

Dynamic Routing in More Than One Area Example

Figure 20 and the example configuration that follows illustrate how to configure a router named Router A that exists in two areas.

Figure 20 CLNS Dynamic Routing within Two Areas



```

! define a tag orion for the routing process
router iso-igrp orion
! configure the net for the process in area 1, domain 47.0004.004d
net 47.0004.004d.0001.212223242526.00
! specify iso-igrp routing using the previously specified tag orion
interface ethernet 0
  clns router iso-igrp orion
! specify iso-igrp routing using the previously specified tag orion
interface ethernet 1
  clns router iso-igrp orion

```

Dynamic Routing in Overlapping Areas Example

The example that follows illustrates how to configure a router with overlapping areas:

```

! define a tag capricorn for the routing process
router iso-igrp capricorn
! configure the NET for the process in area 3, domain 47.0004.004d
net 47.0004.004d.0003.0000.0C00.0508.00
! define a tag cancer for the routing process
router iso-igrp cancer
! configure the NET for the process in area 4, domain 47.0004.004d
net 47.0004.004d.0004.0000.0C00.0506.00
! specify iso-igrp routing on interface ethernet 0 using the tag capricorn
interface ethernet 0
  clns router iso-igrp capricorn
! specify iso-igrp routing on interface ethernet 1 using the tags capricorn and cancer
interface ethernet 1
  clns router iso-igrp capricorn
  clns router iso-igrp cancer

```

```


! specify iso-igrp routing on interface ethernet 2 using the tag cancer
interface ethernet 2
  clns router iso-igrp cancer

```

Dynamic Interdomain Routing Example

Figure 21 and the configurations that follow illustrate how to configure three domains that are to be transparently connected.

Figure 21 CLNS Dynamic Interdomain Routing



Router Chicago

The following configuration shows how to configure Router Chicago for dynamic interdomain routing:

```

! define a tag A for the routing process
router iso-igrp A
! configure the NET for the process in area 2, domain 47.0007.0200
  net 47.0007.0200.0002.0102.0104.0506.00
! redistribute iso-igrp routing information throughout domain A
  redistribute iso-igrp B
! define a tag B for the routing process
router iso-igrp B
! configure the NET for the process in area 3, domain 47.0006.0200
  net 47.0006.0200.0003.0102.0104.0506.00
! redistribute iso-igrp routing information throughout domain B
  redistribute iso-igrp A
! specify iso-igrp routing with the tag A
interface ethernet 0
  clns router iso-igrp A
! specify iso-igrp routing with the tag B
interface serial 0
  clns router iso-igrp B

```

Router Detroit

The following configuration shows how to configure Router Detroit for dynamic interdomain routing. Comment lines have been eliminated from this example to avoid redundancy.

```

router iso-igrp B
  net 47.0006.0200.0004.0102.0104.0506.00
  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

```

Chicago injects a prefix route for domain A into domain B. Domain B injects this prefix route and a prefix route for domain B into domain C.

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

IS-IS Routing Configuration Examples

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

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:

```
! route dynamically using the is-is protocol
router isis
! configure the NET for the process in area 47.0004.004d.0001
net 47.0004.004d.0001.0000.0c00.1111.00
! enable is-is routing on ethernet 0
interface ethernet 0
  clns router isis
! enable is-is routing on ethernet 1
interface ethernet 1
  clns router isis
! enable is-is routing on serial 0
interface serial 0
  clns router isis
```

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 interface 0 is configured for Level 2 routing only. Most comment lines have been eliminated from this example to avoid redundancy.

```
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
! configure a level 2 adjacency only for interface serial 0
isis circuit-type level-2-only
```

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. Most comment lines have been eliminated from this example to avoid redundancy.

```
! enable is-is routing in area 1
router isis areal
! Router is in areas 47.0004.004d.0001 and 47.0004.004d.0011
net 47.0004.004d.0001.0000.0c11.1111.00
net 47.0004.004d.0011.0000.0c11.1111.00
! enable the router to operate as a station router and an interarea router
is-type level-1-2
```

```
!  
interface ethernet 0  
  clns router isis area1  
  ! specify a cost of 5 for the level-1 routes  
  isis metric 5 level-1  
  ! establish a level-1 adjacency  
  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  
  ! set the priority for serial 0 to 3 for a level-1 adjacency  
  isis priority 3 level-1  
  isis priority 1 level-2
```

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 Ethernet interface 0; the ISO IGRP domain is on serial interface 0. The IS-IS routing process is assigned a null tag; the ISO IGRP routing process is assigned a tag of *remote-domain*. Most comment lines have been eliminated from this example to avoid redundancy.

```
router isis  
  net 39.0001.0001.0000.0c00.1111.00  
  ! redistribute iso-igrp routing information throughout remote-domain  
  redistribute iso-igrp remote-domain  
!  
router iso-igrp remote-domain  
  net 39.0002.0001.0000.0c00.1111.00  
  ! redistribute is-is routing information  
  redistribute isis  
!  
interface ethernet 0  
  clns router isis  
!  
interface serial 0  
  clns router iso-igrp remote
```


Router in Two Areas Example

The following two examples show how to configure a router in two areas. The first example configures ISO IGRP; the second configures IS-IS.

ISO IGRP

In the following example, the router is in domain 49.0001 and has a system ID of aaaa.aaaa.aaaa. The router is in two areas: 31 and 40 (decimal). Figure 24 illustrates this configuration.

Figure 22 ISO IGRP Configuration



```

router iso-igrp test-proc1
! 001F in the following net is the hex value for area 31
net 49.0001.001F.aaaa.aaaa.aaaa.00
router iso-igrp test-proc2
! 0028 in the following net is the hex value for area 40
net 49.0001.0028.aaaa.aaaa.aaaa.00
!
interface ethernet 1
  clns router iso-igrp test-proc1
!
interface serial 2
  clns router iso-igrp test-proc1
!
interface ethernet 2
  clns router iso-igrp test-proc2

```

IS-IS

To run IS-IS instead of ISO IGRP, use this configuration. The illustration in Figure 24 still applies. Ethernet interface 2 is configured for IS-IS routing and is assigned the tag of test-proc2.

```

router iso-igrp test-proc1
net 49.0002.0002.bbbb.bbbb.bbbb.00
router isis test-proc2
net 49.0001.0002.aaaa.aaaa.aaaa.00
!
interface ethernet 1
  clns router iso-igrp test-proc1
!
interface serial 2
  clns router iso-igrp test-proc1
!
interface ethernet 2
  clns router is-is test-proc2

```

To allow CLNS packets only to blindly pass through an interface without routing updates, you could use a simple configuration. The following example shows such a configuration:

```

clns routing
interface serial 2
! permits serial 2 to pass CLNS packets without having CLNS routing turned on
clns enable

```

Basic Static Routing Examples

Configuring FDDI, Ethernets, Token Rings, and serial lines for CLNS can be as simple as 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.

Example 1

In the following example, an Ethernet and a serial line can be configured as follows:


```

! enable clns packets to be routed
clns routing
! configure the following network entity title for the routing process
clns net 47.0004.004d.0055.0000.0C00.BF3B.00
! pass ISO CLNS traffic on ethernet 0 to end systems without routing
interface ethernet 0
  clns enable
! pass ISO CLNS traffic on serial 0 to end systems without routing
interface serial 0
  clns enable
! create a static route for the interface
clns route 47.0004.004d.0099 serial 0
clns route 47.0005 serial 0
    
```

Example 2

The following is a more complete example of CLNS static routing on a system with two Ethernet interfaces. After configuring routing, you define a NET and enable CLNS on the Ethernet 0 and Ethernet 1 interfaces. You must then define an ES neighbor and define a static route with the **clns route** global configuration command, as shown. In this situation, there is an ES on Ethernet 1 that does not support ES-IS. Figure 23 illustrates this network.

Figure 23 Static Routing



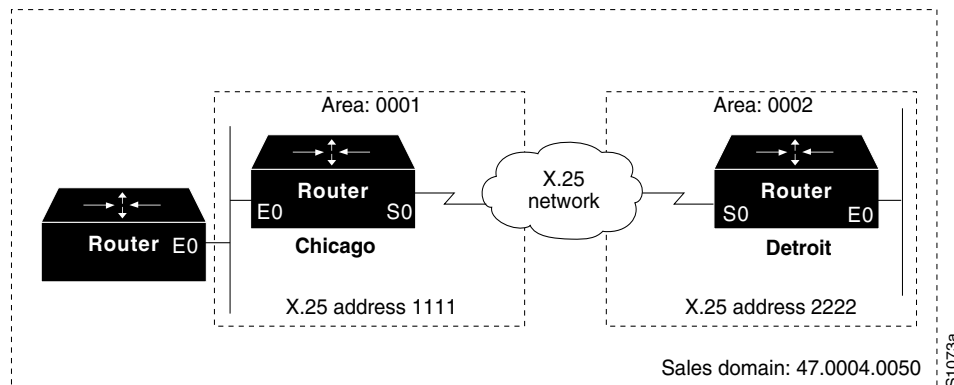
```

clns host sid 39.0001.1111.1111.1111.00
clns host bar 39.0002.2222.2222.2222.00
! assign a static address for the router
clns net sid
! enable CLNS packets to be routed
clns routing
! pass ISO CLNS packet traffic to end systems without routing them
interface ethernet 0
  clns enable
! pass ISO CLNS packet traffic to end systems without routing them
interface ethernet 1
  clns enable
! specify end system for static routing
clns es-neighbor bar 0000.0C00.62e7
! create an interface-static route to bar for packets with the following NSAP address
clns route 47.0004.000c bar
    
```

Static Intradomain Routing Example

Figure 24 and the configurations that follow demonstrate how to use static routing inside of a domain. Imagine a company with branch offices in Detroit and Chicago, connected with an X.25 link. These offices are both in the domain named Sales.

Figure 24 CLNS X.25 Intradomain Routing



The following example shows one way to configure the router in Chicago:

```

! define the name chicago to be used in place of the following NSAP
clns host chicago 47.0004.0050.0001.0000.0c00.243b.00
! define the name detroit to be used in place of the following NSAP
clns host detroit 47.0004.0050.0002.0000.0c00.1e12.00
! enable ISO IGRP routing of CLNS packets
router iso-igrp sales
! configure net chicago, as defined above
net chicago
! specify iso-igrp routing using the previously specified tag sales
interface ethernet 0
  clns router iso-igrp sales
! set the interface up as a DTE with X.25 encapsulation
interface serial 0
  encapsulation x25
  x25 address 1111
  x25 nvc 4
! specify iso-igrp routing using the previously specified tag sales
clns router iso-igrp sales
! define a static mapping between Detroit's nsap and its X.121 address
x25 map clns 2222 broadcast

```

This configuration brings 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 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 Example

Figure 25 and the example configurations that follow illustrate how to configure two routers that distribute information across domains. In this example, Router A (in domain Orion) and Router B (in domain Pleiades) communicate across a serial link.

Figure 25 CLNS Interdomain Static Routing



Router A

The following configuration shows how to configure Router A for static interdomain routing:

```
! define tag orion for net 47.0006.0200.0100.0102.0304.0506.00
router iso-igrp orion
! configure the following network entity title for the routing process
net 47.0006.0200.0100.0102.0304.0506.00
! define the tag bar to be used in place of Router B's NSAP
clns host bar 47.0007.0200.0200.1112.1314.1516.00
! specify iso-igrp routing using the previously specified tag orion
interface ethernet 0
  clns router iso-igrp orion
! pass ISO CLNS traffic to end systems without routing
interface serial 1
  clns enable
! configure a static route to Router B
clns route 39.0001 bar
```

Router B

The following configuration shows how to configure Router B for static interdomain routing:

```
router iso-igrp pleiades
! configure the network entity title for the routing process
net 47.0007.0200.0200.1112.1314.1516.00
! define the name sid to be used in place of Router A's NSAP
clns host sid 47.0006.0200.0100.0001.0102.0304.0506.00
! specify iso-igrp routing using the previously specified tag pleiades
interface ethernet 0
  clns router iso-igrp pleiades
! pass ISO CLNS traffic to end systems without routing
interface serial 0
  clns enable
! pass packets bound for sid in domain 47.0006.0200 through serial 0
clns route 47.0006.0200 sid
```

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

CLNS Filter Example

The following example allows packets if the address starts with either 47.0005 or 47.0023. It implicitly denies any other address.

```
clns filter-set US-OR-NORDUNET permit 47.0005...
clns filter-set US-OR-NORDUNET permit 47.0023...
```

Route Map Examples

The following example redistributes two types of routes into the integrated IS-IS routing table (supporting both IP and CLNS). The first routes are OSPF external IP routes with tag 5, and these are inserted into level-2 IS-IS LSPs with a metric of 5. The second routes are ISO IGRP derived CLNS prefix routes that match CLNS filter expression “osifilter.” These are redistributed into IS-IS as level-2 LSPs with a metric of 30.

```
router isis
 redistribute ospf 109 route-map ipmap
 redistribute iso-igrp nsfnet route-map osimap
!
 route-map ipmap permit
 match route-type external
 match tag 5
 set metric 5
 set level level-2
!
 route-map osimap permit
 match clns address osifilter
 set metric 30
 clns filter-set osifilter permit 47.0005.80FF.FF00
```

Given the following configuration, a RIP learned route for network 160.89.0.0 and an ISO IGRP learned route with prefix 49.0001.0002 will be redistributed into an IS-IS level-2 LSP with a metric of 5:

```
router isis
 redistribute rip route-map ourmap
 redistribute iso-igrp remote route-map ourmap
!
 route-map ourmap permit
 match ip address 1
 match clns address ourprefix
 set metric 5
 set level level-2
!
 access-list 1 permit 160.89.0.0 0.0.255.255
 clns filter-set ourprefix permit 49.0001.0002...
```

DECnet Cluster Aliases Example

The following example enables cluster aliasing for CLNS:

```
clns routing
clns nsap 47.0004.004d.0001.0000.0C00.1111.00
router iso-igrp pleiades
! enable cluster aliasing on interface ethernet 0
interface ethernet 0
 clns cluster-alias
! enable cluster aliasing on interface ethernet 1
interface ethernet 1
 clns cluster-alias
```

The following example denies packets with an address that starts with 39.840F, but allows any other address:

```

clns filter-set NO-ANSI deny 39.840F...
clns filter-set NO-ANSI permit default
    
```

The following example builds a filter that accepts end system adjacencies with only two systems, based only on their system IDs:

```


clns filter-set ourfriends...0000.0c00.1234.**
clns filter-set ourfriends...0000.0c00.125a.**

interface ethernet 0
  clns adjacency-filter es ourfriends
    
```

ISO CLNS over X.25 Example

In the following example, serial interface 1 on Router A acts as a DTE for X.25. It permits broadcasts to pass through. Router B is an IS, which has a CLNS address of 49.0001.bbbb.bbbb.bbbb.00 and an X.121 address of 31102. Router A has a CLNS address of 49.0001.aaaa.aaaa.aaaa.00 and an address of 31101. Figure 26 illustrates this configuration.

Figure 26 Routers Acting as DTEs and DCEs



Router A

```

router iso-igrp test-proc
net 49.0001.aaaa.aaaa.aaaa.00
!
interface serial 1
  clns router iso-igrp test-proc
  ! assume the host is a DTE and encapsulates x.25
  encapsulation x25
  ! define the X.121 address of 31101 for serial 1
  x25 address 31101
  ! set up an entry for the other side of the X.25 link (Router B)
  x25 map clns 31102 broadcast
    
```

Router B

```
router iso-igrp test-proc
 net 49.0001.bbbb.bbbb.bbbb.00
!
interface serial 2
 clns router iso-igrp test-proc
! configure this side as a DCE
 encapsulation x25-dce
! define the X.121 address of 31102 for serial 2
 X25 address 31102
! configure the NSAP of Router A and accept reverse charges
 x25 map clns 31101 broadcast accept-reverse
```

Performance Parameters Example

The following example shows how to set ES hello packet and IS hello packet parameters in a simple ISO IGRP configuration, as well as the MTU for a serial interface:

```
router iso-igrp xavier
 net 49.0001.004d.0002.0000.0C00.0506.00
! send IS/ES hellos every 45 seconds
 clns configuration-time 45
! recipients of the hello packets keep info. in the hellos for 2 minutes
 clns holding-time 120
! specify an mtu of 978 bytes; generally, do not alter the default mtu value
interface serial 2
 clns mtu 978
```

TARP Configuration Examples

The following two sections provide basic and complex examples of TARP configuration.

Basic TARP Configuration Example

The following example enables TARP on the router and interface Ethernet 0. The router is assigned the TID *myname*.


```
clns routing
 tarp run
 tarp tid myname

interface ethernet 0
 tarp enable
```

Complex TARP Configuration Example

Figure 27 and the following example show how to enable TARP on Router A and on interface Ethernet 0, and assign the TID *myname*. A static route is created from Router A (49.0001.1111.1111.1111.00) to Router D (49.0004.1234.1234.1234.00) so that Router D can receive TARP PDUs because Router C is not TARP capable. A blacklist adjacency is also created on Router A for Router B (49.001.7777.7777.7777.00) so that Router A does not send any TARP PDUs to Router B.

Figure 27 Sample TARP Configuration



```
clns routing
tarp run
tarp cache-timer 300
tarp route-static 49.0004.1234.1234.1234.00
tarp blacklist-adjacency 49.0001.7777.7777.7777.00
tarp tid myname

interface ethernet 0
  tarp enable
```