

Configuring ISO CLNS

The International Organization for Standardization (ISO) Connectionless Network Service (CLNS) protocol is a standard for the network layer of the Open Systems Interconnection (OSI) model.

This chapter describes how to configure ISO CLNS. For a complete description of the commands in this chapter, refer to the “ISO CLNS Commands” chapter of the *Router Products Command Reference* publication.

Cisco’s Implementation of ISO CLNS

The Cisco routing software supports packet forwarding and routing for ISO CLNS on networks using a variety of data link layers: Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), and serial.

You can use CLNS routers on serial interfaces with High-Level Data Link Control (HDLC), Point-to-Point Protocol (PPP), Link Access Procedure, Balanced (LAPB), X.25, Switched Multimegabit Data Service (SMDS), or Frame Relay encapsulation. To use HDLC encapsulation, you must have a router at both ends of the link. If you use X.25 encapsulation, you must manually enter the NSAP-to-X.121 mapping. The LAPB, SMDS, Frame Relay, and X.25 encapsulations interoperate with other vendors.

Cisco’s CLNS implementation also is compliant with the Government Open Systems Interconnection Profile (GOSIP) Version 2.

As part of its CLNS support, Cisco routers fully support these ISO and American National Standards Institute (ANSI) standards:

- ISO 9542—Documents the End System-to-Intermediate System (ES-IS) routing exchange protocol.
- ISO 8473—Documents the ISO Connectionless Network Protocol (CLNP).
- ISO 8348/Ad2—Documents network service access point (NSAP) addresses.
- ISO 10589—Documents Intermediate System-to-Intermediate System (IS-IS) Intra-domain Routing Exchange Protocol.

Both the ISO-developed IS-IS and Cisco’s ISO Interior Gateway Routing Protocol (IGRP) dynamic routing protocols are supported for dynamic routing of ISO CLNS. In addition, static routing for ISO CLNS is supported.

ISO CLNS Configuration Task List

To configure ISO CLNS, complete the tasks in the following sections:

- Assign Domain Boundaries, NSAP Addresses, and Area Addresses
- Configure a Routing Process
- Configure ES-IS Hello Packet Parameters
- Create Packet-Forwarding Filters and Establish Adjacencies
- Configure CLNS over WANs
- Configure Miscellaneous Features
- Enhance ISO CLNS Performance
- Monitor and Maintain the ISO CLNS Network

See the end of this chapter for configuration examples.

Assign Domain Boundaries, NSAP Addresses, and Area Addresses

In the following section you will learn how to assign NETs, or addresses, for areas and domains. Addressing background material is provided first, followed by a description of how to assign these addresses. The topic of addressing includes the following sections:

- ISO CLNS Addressing Background
- Configure NETs for Domains and Areas
- Map NSAP Addresses to Media Addresses
- Specify Shortcut NSAP Addresses
- Use the IP Domain Name System to Discover ISO CLNS Addresses

ISO CLNS Addressing Background

Addresses in the ISO network architecture are referred to as network service access point (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 (see Figure 20-1). 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 dynamic routing (ISO-IGRP or IS-IS routing) only for NSAP addresses that conform to the address constraints defined in the ISO standard for IS-IS (ISO 10589).

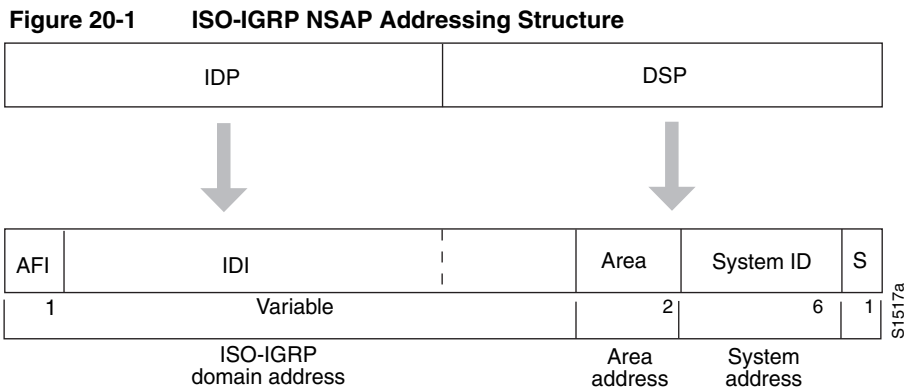
An NSAP address consists of two major fields:

- The IDP is made up of 1-byte AFI and a variable length IDI. The length of the IDI and the encoding format for the DSP are based on the value of the AFI.
- The DSP is made up of a high-order DSP, an area ID, a system ID, and a 1-byte n-selector.

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. However, they differ in the way addresses are specified for area routing. An ISO-IGRP NSAP address includes three separate levels

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 defined for ISO-IGRP and the *system ID*.

Figure 20-1 illustrates the ISO-IGRP NSAP addressing structure.

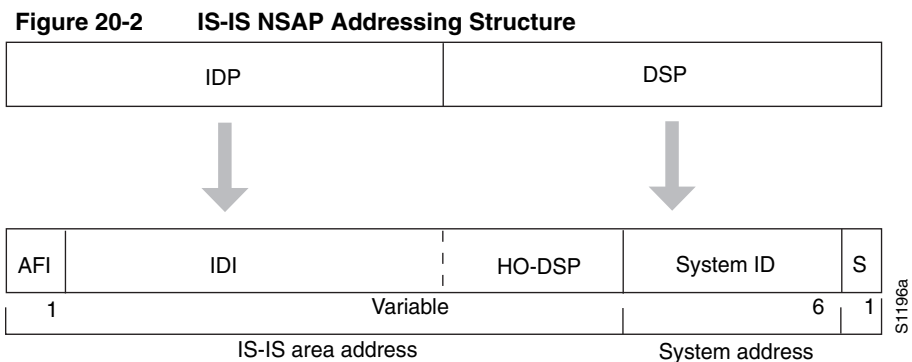


The ISO-IGRP NSAP address is divided into three parts: a domain part, an area address, and a system ID. Domain routing is performed on the domain part of the address. Area routing for a given domain uses the area address. System ID routing for a given area uses the system ID part. The NSAP address is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

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

Our 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 20-2 illustrates the IS-IS NSAP addressing structure.



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

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

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*.

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.
- ISO-IGRP requires at least ten bytes of length; one for domain, two for area, six for system ID, and one for n-selector.
- Our implementation of IS-IS requires at least eight bytes; one for area, six for system ID, and one for n-selector.
- No two nodes residing within the same area can have addresses in which the system ID fields are the same.

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:

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

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

Entering Routes

Routes are entered by specifying pairs (NSAP prefix and next-hop NET). NETs are similar in function to NSAP addresses. In the routing table, the best match means the longest NSAP prefix entry that matches the beginning of the destination NSAP address. In the following sample static routing table, Table 20-1, the next-hop NETs are listed for completeness, but are not necessary to understand the routing algorithm. Table 20-2 offers examples of how the longest matching NSAP prefix can be matched with routing table entries in Table 20-1.

Table 20-1 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 20-2 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.

Configure NETs for Domains and Areas

The first task you have to perform to enable CLNS routing is to assign addresses or NETs for your domains and areas. Assign domain and area addresses the same way, following the addressing rules described at the beginning of this section.

First, establish domains. 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. Determine which routers are to be assigned to which areas. The area address uniquely identifies the routing area.

A router can have one or more area addresses. The concept of multiple area addresses is described in the section that follows, “Assign Multiple Area Addresses to IS-IS Areas.”

Note 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.

Note If you are going to configure a routing process, you might want to read the section “Configure a Routing Process,” which appears after the discussion of addressing.

To configure NETs, perform the following tasks in the order listed:

Step 1 Enter global configuration mode. **configure**¹

Step 2	Enable IS-IS routing, which places you in router configuration mode.	router isis [tag]
	or	
Step 3	Enable ISO-IGRP routing, which places you in router configuration mode.	router iso-igrp [tag]
Step 4	Configure NETs for the routing process.	net network-entity-title

1. This command is documented in the “Interface Commands” chapter of the *Router Products Command Reference*.

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

Assign Multiple Area Addresses to IS-IS Areas

IS-IS routing supports assignment of multiple area addresses on the same router. This concept is referred to as *multihoming*. You must statically assign the multiple area addresses on the router. All of 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. The number of areas allowed in a domain is unlimited.

Multihoming provides a mechanism for smoothly migrating network addresses:

- Splitting up an area—Stations 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 stations. Use multiple area addresses to allow incoming traffic intended for an old area address to continue being routed to associated stations.

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 statically, perform the following task in the order listed:

Step 1	Enter global configuration mode.	configure ¹
Step 2	Enable IS-IS routing, which places you in router configuration mode.	router isis [tag]
Step 3	Configure NETs for the routing process. You can perform this step as many as three times.	net network-entity-title

1. This command is documented in the “Interface Commands” chapter of the *Router Products Command Reference*.

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

Configure a Static NET Address for the Router

You must assign static addresses if you have configured the router to support ISO CLNS but you are not using a routing protocol.

A CLNP packet sent to any of the defined NSAP addresses or NETs will be received by the router. The router uses the following algorithm to select the 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 routing process that is running on this interface.
- If IS-IS is running, use the NET of the IS-IS routing process that is running on this interface.

To assign an address to the router if you are not dynamically routing CLNS packets, perform the following task in global configuration mode:

Task	Command
Assign an address to the router when the router is not configured to dynamically route CLNS packets using ISO-IGRP or IS-IS.	clns net { <i>net-address</i> <i>name</i> }

Map NSAP Addresses to Media Addresses

Conceptually, each end system (ES) lives in one area. It discovers the nearest intermediate system (IS) router 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.

End systems 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. A router dynamically discovers all ESs running the ES-IS protocol. ESs that are not running the ES-IS protocol must be statically configured.

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

Perform the following tasks in interface configuration mode, as needed, to statically enter mapping information between the NSAP protocol addresses and the subnetwork point of attachment (SNPA) addresses for end systems or intermediate systems:

Task	Command
List all end systems that will be used when you manually specify the NSAP-to-SNPA mapping. In this case, the SNPAs are the MAC addresses.	clns es-neighbor <i>nsap snpa</i>
List all intermediate systems that will be used when you manually specify the NSAP-to-SNPA mapping. In this case, the SNPAs are the MAC addresses.	clns is-neighbor <i>nsap snpa</i>

Note It is necessary to use static mapping only for those end systems that do *not* support ES-IS. The router continues to dynamically discover those end systems that *do* support ES-IS.

If there are systems on the Ethernet that do not use ES-IS, or if X.25 is being used and no dynamic routing protocol is running over the X.25 network, you specify the NSAP/NET (protocol address) to SNPA (media address) mappings by performing the following tasks in interface configuration mode.

Task	Command
List all end systems that will be used when you manually specify the NSAP-to-SNPA mapping. In this case, the SNPAs are the X.25 network addresses (X.121 addresses).	clns es-neighbor <i>nsap snpa</i>
or	
List all intermediate systems that will be used when you manually specify the NSAP-to-SNPA mapping. In this case, the SNPAs are the X.25 network addresses (X.121 addresses).	clns is-neighbor <i>nsap snpa</i>
Establish an IP-to-X.121 address map.	x25 map clns <i>snpa x.25-facilities-info</i> ¹

1. This command is documented in the “X.25 and LAPB Commands” chapter of the *Router Products Command Reference* publication.

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

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.

There are some effects and requirements associated with using names to represent NETs and NSAP addresses, however; they include the following:

- Although using names as proxies for addresses is allowed with CLNS commands, they are never written out to NVRAM.

- 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 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:

- **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 Name 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
Allow DNS queries for CLNS addresses.	ip domain-lookup nsap
Disable DNS queries for CLNS addresses.	no ip domain-lookup nsap

Configure a Routing Process

The basic function of a router is to forward packets: to 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 of the entries in the table yourself, you are using *static* routing. If you have a routing process building the tables, you are using *dynamic* routing. It is possible, and sometimes necessary, to use both static and dynamic routing simultaneously.

When you configure only ISO CLNS and not routing protocols, the router only makes forwarding decisions. It does not perform other routing-related functions. In such a configuration, the router 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 network entity title (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. To configure an interface for concurrent routing and bridging, you use the **bridge crb** command to enable concurrent routing and bridging on the router.

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 you are using routers that do not support the same dynamic routing protocol, you must use static routing.
- If your network includes WAN links that involve paying for connect time or per packet, you would use static routing rather than paying to run a routing protocol 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 we support, 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.

Dynamic Routing

We support two dynamic routing protocols for CLNP networks:

- ISO-IGRP
- IS-IS

Both routing protocols support the concept of *areas*. Within an area, all routers know how to reach all of the system IDs. Between areas, routers know how to reach the proper area.

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

Intermediate Systems and End Systems

Some intermediate systems keep track of how to communicate with all of the end systems in their areas and thereby function as Level 1 routers (also referred to as local routers). Other intermediate systems keep track of how to communicate with other areas in the domain, functioning as Level 2 routers (sometimes referred to as area routers). Our 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.

End systems communicate with intermediate systems using the ES-IS protocol. Level 1 and Level 2 intermediate systems communicate with each other using either ISO IS-IS or our ISO-IGRP protocol.

This section describes the tasks associated with each routing protocol. When dynamically routing, you can choose either ISO-IGRP or IS-IS, or you can route both routing protocols at the same time.

Configure CLNS Static Routing

You do not need to explicitly specify a routing process to use static routing facilities. If you choose static routing, the configuration process begins with enabling CLNS routing on the router. CLNS routing is enabled by default when you configure either of the routing protocols.

CLNS static routing is configured when you do the following:

Step 1 Configure CLNS on the router.

Step 2 Assign a static NET address for the router.

Step 3 Enable ISO CLNS for each interface.

Step 4 Enter a specific static route.

Step 5 Configure other, optional variations.

Each of these steps is described in the following sections.

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 CLNS on the Router

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

Task	Command
Configure CLNS on the router.	clns routing

Assign a Static NET Address for the Router

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

A CLNP packet sent to any of the defined NSAP addresses or NETs will be received by the router. The router uses the following algorithm to select the 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 routing process that is running.
- If IS-IS is running, use the NET of the IS-IS routing process that is running.

To assign an address to the router, perform the following task in global configuration mode:

Task	Command
Assign an address to the router when the router is not configured to dynamically route CLNS packets using ISO-IGRP or IS-IS.	clns net { <i>net-address</i> <i>name</i> }

Enable ISO CLNS for Each Interface

You also must enable ISO CLNS for each 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 but intend to pass ISO CLNS packet traffic to end systems, you must enable CLNS yourself.

Enable ISO CLNS when you want to pass ISO CLNS packet traffic to end systems but do not want to perform any dynamic routing on an interface. Perform the following task in interface configuration mode:

Task	Command
Enable ISO CLNS for each interface.	clns enable

Enter a Specific Static Route

You can enter a specific static route and apply it globally even when you are dynamically routing. NSAP addresses that start with the NSAP prefix you specify are forwarded to the next-hop node.

To apply a specific static route globally, perform the following task in global configuration mode:

Task	Command
Enter a specific static route.	clsns route <i>nsap-prefix</i> { <i>next-hop-net</i> <i>name</i> }

Configure Variations of the Clns Route Command

You also can perform the following tasks that use variations of the **clsns 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 **clsns route** command is not literally *applied* to a specific interface.
- Tell a router to discard packets with the specified NSAP prefix.
- Specify a default prefix.

The following list shows how to perform each of these tasks. Perform these tasks in global configuration mode:

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

Configure ISO-IGRP Dynamic Routing

CLNS routing is enabled by default on the router when you configure ISO-IGRP. All you need to do to specify an ISO-IGRP routing process is to 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. You can specify up to ten ISO-IGRP processes.

To configure ISO-IGRP dynamic routing, perform the following tasks in the order listed:

Task	Command
Step 1 Enter global configuration mode.	configure ¹
Step 2 Enable the ISO-IGRP routing process, which places you in router configuration mode.	router iso-igrp [<i>tag</i>]
Step 3 Configure the NET or address for the routing process.	net <i>network-entity-title</i>
Step 4 Enter interface configuration mode.	interface <i>type number</i> ¹
Step 5 Enable ISO-IGRP on specified interfaces; also set the level type for routing updates.	clsns router iso-igrp <i>tag</i> [level 2]

1. This command is documented in the "Interface Commands" chapter of the *Router Products Command Reference* publication.

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

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

The additional tasks that follow allow you to customize ISO-IGRP.

See the “Dynamic Routing within the Same Area Example,” “Dynamic Routing in More Than One Area Example,” and “Dynamic Routing in Overlapping Areas Example” sections at the end of this chapter for examples of configuring dynamic routing.

You can also configure the following ISO-IGRP parameters:

- Adjust ISO-IGRP Metrics
- Adjust ISO-IGRP Timers
- Enable or Disable Split Horizon
- Redistribute Routes into an ISO-IGRP Domain
- Specify Preferred Routes

Adjust ISO-IGRP Metrics

You have the option of altering the default behavior of ISO-IGRP routing and metric computations. This allows, for example, 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. Due to 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. By performing the following task, you can configure the metric constants used in the ISO-IGRP composite metric calculation of reliability and load. Perform this task in router configuration mode:

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

Two additional ISO-IGRP metrics can be configured. These are the bandwidth and delay associated with an interface. Refer to the “Interface Commands” chapter of the *Router Products Command Reference* publication for details about the **bandwidth** and **delay** interface configuration commands used to set these metrics, and to the “Configuring Interfaces” chapter of this manual for configuration information.

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

Redistribute Routes into an ISO-IGRP Domain

You can configure a router to do interdomain dynamic routing by putting it into two domains and configuring it to redistribute 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.

Also, 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. These methods of specifying route redistribution are listed in the following tables.

Static routes by default are redistributed into ISO-IGRP routing domains.

Perform the following tasks in router configuration mode:

Task	Command
Redistribute routes from one routing protocol into other routing domains.	redistribute <i>protocol</i> [<i>tag</i>] [route-map <i>map-tag</i>]
Redistribute static routes by causing the specified routing process to advertise static CLNS routes.	redistribute static

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>

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. The router always picks the route whose routing protocol has the lowest administrative distance. The lower the value of the distance, the more preferred the route.

Default administrative distances are already set. By default, the following administrative distances are assigned:

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

However, if you need to change an administrative distance for a route, perform the following task in router configuration mode:

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

If you want an ISO-IGRP prefix route to override a static route, you must set the distance for the routing process to be lower than 10.

Configure IS-IS Dynamic Routing

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

To configure IS-IS dynamic routing, perform the following required tasks in the order listed:

Task	Command
Step 1 Enter global configuration mode.	configure ¹
Step 2 Enable IS-IS routing, which places you in router configuration mode.	router isis [<i>tag</i>]
Step 3 Configure NETs for the routing process; you can specify a name for a NET as well as an address	net <i>network-entity-title</i>
Step 4 Enter interface configuration mode.	interface <i>type number</i> ¹
Step 5 Specify the interfaces that should be actively routing IS-IS.	clns router isis [<i>tag</i>]

1. This command is documented in the “Interface Commands” chapter of the *Router Products Command Reference* publication.

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.

Configure IS-IS Interface Parameters

Our IS-IS implementation allows you to customize certain interface-specific IS-IS parameters. You can perform the following optional tasks:

- Configure 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 an attached network. Therefore, be sure that if you do configure any of these parameters, the configurations for all routers on the network have compatible values.

Configure 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 router and that you can configure is the *default-metric*, which you can configure for Level 1 and/or Level 2 routing.

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 delay-metric expense-metric error-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 task 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 }

Task	Command
Specify the number that is 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, it 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

Complete sequence number PDUs (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 IS-IS link state PDUs (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 attached network. The setting of this parameter should be conservative, or needless retransmission will result. The value 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.

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

You can specify 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 might be established if there is at least one area address common to both this system and its neighbors.

If you specify both Level 1 and Level 2, 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. This is the default value.

If you specify Level 2, a Level 2 adjacency is established.

Note that it is seldom necessary to configure an interface as Level 1 only or Level 2 only—the protocols will determine the adjacency type automatically.

Configure IS-IS Password Authentication

You can assign different passwords for different routing levels. By default, authentication is disabled. Specifying Level 1 or Level 2 disables 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 IS-IS Parameters

You can configure the following IS-IS parameters:

- Redistribute Routes into an IS-IS Domain
- Specify Preferred Routes
- Specify Router-Level Support
- Configure IS-IS Authentication Passwords
- Ignore IS-IS Link State Packets (LSPs) Errors

- Log Adjacency State Changes
- Set IS-IS LSP MTU Size

Redistribute Routes into an IS-IS Domain

If you have a router that is in two routing domains, you might want to redistribute routing information between the two domains. First, you specify the destination routing protocol, then you define the routing protocol that is to be redistributed into the destination routing protocol.

Redistribution only occurs for Level 2 routing.

You also can conditionally control the redistribution of routes between routing domains by defining route maps between them.

Additionally, you can cause the specified routing process to advertise static CLNS routes. Static routes are always redistributed into IS-IS unless you explicitly disable this feature.

The methods of specifying route redistribution are listed in the following tables.

Perform the following tasks in router configuration mode:

Task	Command
Redistribute routes from one routing protocol into other routing domains.	redistribute <i>protocol</i> [<i>tag</i>] [route-map <i>map-tag</i>]
Redistribute static routes by causing the specified routing process to advertise static CLNS routes.	redistribute static [clns]

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>

See the “Route-Map Examples” section 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. The router always picks the route whose routing protocol has the lowest administrative distance. The lower the value of the distance, the more preferred the route.

By default the following administrative distances are assigned:

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

However, if you need to change an administrative distance for a route, perform the following task in router configuration mode:

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

If you want an IS-IS prefix route to override a static route, you must set the distance for the routing process to be lower than 10.

Specify Router-Level Support

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 }

Note that it is seldom necessary to configure the IS type because the IS-IS protocol will automatically establish this.

Configure IS-IS Authentication Passwords

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

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 Packets (LSPs) Errors

You can allow the router to ignore IS-IS link state packets (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. Because this could render the network nonfunctional, use the **ignore-lsp-errors** to ignore these LSPs rather than purge the LSPs.

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

Set IS-IS LSP MTU Size

Under normal condition, the default 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.



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 set the maximum transmission unit (MTU) size of IS-IS link state packets (LSPs), perform the following task in router configuration mode:

Task	Command
Specify the maximum LSP packet size, in bytes. The size must be less than or equal to the smallest MTU of any link in the network. The default size is 1497 bytes.	lsp-mtu

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 to the routers directly connected to such a link.

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 (ESH) packets and IS hello (ISH) packets are sent without options. Options in received ESH and ISH 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; however, you can change the default by performing the following task in global configuration mode:

Task	Command
Specify the rate at which ESH and ISH packets are sent.	clns configuration-time <i>seconds</i>

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

Task	Command
Allow the sender of an ESH or ISH 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>

Create Packet-Forwarding Filters and Establish Adjacencies

You can build powerful CLNS filter expressions, or access lists, that 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>

Task	Command
Build filter expressions, combining the use of one or more filter sets.	clns filter-expr <i>ename term</i>

Perform the following tasks in interface configuration mode:

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

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

Configure CLNS over WANs

This section provides general information about running ISO CLNS over WANs. For more information, see the relevant chapters describing the specific types of encapsulation that might be used.

You can use CLNS routers on serial interfaces with HDLC, PPP, LAPB, X.25, Frame Relay, DDR, or SMDS encapsulation. 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 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; therefore, ES-IS generally is not used to automatically advertise and record mappings between NSAP/NET (protocol addresses) and subnetwork points of attachment (SNPA) (media addresses). (With X.25, the SNPAs are the X.25 network addresses [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 mapping.

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 static mapping information.

You can specify 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.

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

Configure Miscellaneous Features

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

- Assign Static NSAP Addresses for an Interface
- Configure DECnet OSI or Phase V Cluster Aliases
- Configure Digital-Compatible Mode
- Allow Security-Option Packets to Pass

Assign Static NSAP Addresses for an Interface

You can assign an NSAP address for a specific interface. This allows the router 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 assign an NSAP address for a specified interface, perform the following task in interface configuration mode:

Task	Command
Assign an NSAP address for a specific interface.	clns net { <i>nsap-address</i> <i>name</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 router 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 router 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 on a per-interface basis.	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 “Configuring DECnet Cluster Aliases Example” 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 ISH does not have the N-selector byte present, you may want to configure the router to allow ISHs 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 ISHs sent and received to ignore the N-selector byte.	clns dec-compatible

Allow Security-Option Packets to Pass

By default, the router will discard any packets it sees as set with security options. You can disable this behavior; that is, allow such packets to pass through.

Perform the following task in global configuration mode:

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

Header Options

The ISO CLNS routing software ignores the Record Route option, the Source Route option, and the QOS (quality of service) option other than congestion experienced. The security option causes a packet to be rejected with a bad option indication.

Enhance ISO CLNS Performance

Generally, you do not need to change the 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. This section describes the following 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 "Customizing 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. You can, however, set the maximum transmission unit (MTU) size of the packets sent on the interface to reduce fragmentation. The minimum value is 512; the default and maximum packet size depends on the interface type.

Changing the MTU value of packets sent on the interface can affect the CLNS MTU value. If the CLNS MTU is at its maximum given the interface MTU, then 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 **mtu** value of other packets sent on the interface.

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 <i>size</i>

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 devices on the network.

Disable Checksums

When the ISO CLNS routing software sources a CLNS packet, by default it generates checksums. You can 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) sourced by the system. It applies to pings and trace-route packets.

Disable Fast Switching Through the Cache

Fast switching through the cache is enabled by default for all supported interfaces. You can disable fast switching by performing 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** 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) might be sent.

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

To transmit ERPDUs, 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-erpd
Determine the minimum interval, in milliseconds, between ERPDUs.	clns erpdu-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 PDU (RDPDU) also can be sent to the sender of the packet. You can control RDPDUs in the following ways:

- You can allow CLNS to send RDPDUs when a better route for a given host is known; this is the default behavior. This reduces bandwidth if it is disabled.
- You can set interval times between RDPDUs.

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

Task	Command
Allow CLNS to send redirect PDUs when a better route for a given host is known.	clns send-rdpdu
Determine the minimum interval time, in milliseconds, between RDPDUs.	clns rdpdu-interval <i>milliseconds</i>

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

Configure Parameters for Locally Sourced Packets

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

Task	Command
Globally specify in seconds the initial lifetime for locally generated packets for the specified router.	clns packet-lifetime <i>seconds</i>
Specify whether to request error PDUs on packets sourced by the router.	clns want-erpd

It is a good idea to set the packet lifetime low in an internetwork that has frequent loops.

Note The `clns want-erpdu` command has no effect on routing packets (ES-IS, ISO-IGRP, and IS-IS) sourced by the system. It applies to pings and trace route packets.

Monitor and Maintain the ISO CLNS Network

Use the EXEC commands described in this section to monitor and maintain the ISO CLNS caches, tables, and databases.

Task	Command
Clear and reinitialize the CLNS routing cache.	<code>clear clns cache</code>
Remove ES neighbor information from the adjacency database.	<code>clear clns es-neighbors</code>
Remove IS neighbor information from the adjacency database.	<code>clear clns is-neighbors</code>
Remove CLNS neighbor information from the adjacency database.	<code>clear clns neighbors</code>
Remove all of the dynamically derived CLNS routing information.	<code>clear clns route</code>
Invoke a diagnostic tool for testing connectivity (privileged).	<code>ping clns {host address}</code>
Test network node reachability using a simple ping facility (user).	<code>ping clns {host address}</code>
Display information about the CLNS network.	<code>show clns</code>
Display the entries in the CLNS routing cache.	<code>show clns cache</code>
Display ES neighbor entries including the associated areas.	<code>show clns es-neighbors [type number unit] [detail]</code>
Display filter expressions.	<code>show clns filter-expr [name] [detail]</code>
Display filter sets.	<code>show clns filter-set [name]</code>
List the CLNS-specific or ES-IS information about each interface.	<code>show clns interface [type number unit]</code>
Display IS neighbor entries according to the area in which they are located.	<code>show clns is-neighbors [type number unit] [detail]</code>
Display both ES and IS neighbors.	<code>show clns neighbors [type number unit] [detail]</code>
List the protocol-specific information for each IS-IS or ISO-IGRP routing process in this router.	<code>show clns protocol [domain area-tag]</code>
Display all the destinations to which this router knows how to route packets.	<code>show clns route [nsap]</code>
Display information about the CLNS packets this router has seen.	<code>show clns traffic</code>
Display the IS-IS link state database.	<code>show isis database [level-1] [level-2] [I1] [I2] [detail] [lspid]</code>
Display the IS-IS Level 1 routing table.	<code>show isis routes</code>
Display a history of the shortest path first (SPF) calculations for IS-IS	<code>show isis spf-log</code>
Display all route maps configured or only the one specified.	<code>show route-map [map-name]</code>

Task	Command
Discover the paths taken by packets in the network (privileged).	trace
Discover the paths taken by packets in the network (user).	trace clns destination
Display the routing table in which the specified CLNS destination is found.	which-route { <i>nsap-address</i> <i>clns-name</i> }

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:

- Configuring NETs Examples
- Basic Static Routing Examples
- Static Intradomain Routing Example
- Static Interdomain Routing Example
- 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
- Configuring a Router in Two Areas Example
- Configuring ISO CLNS over X.25 Example
- Customizing Performance Parameters Example
- Configuring DECnet Cluster Aliases Example
- Route-Map Examples
- CLNS Filter Examples

Configuring NETs 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
!
```

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
```

Basic Static Routing Examples

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

Example 1


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

```
! configure the following network entity title for the routing process
clns net 47.0004.004d.0055.0000.0c00.BF3B.00
! enables clns packets to be routed
clns routing
! pass ISO CLNS traffic on ethernet 0 to end systems without routing
interface ethernet 0
clns enable
interface serial 0
! pass ISO CLNS traffic on serial 0 to end systems without routing
clns enable
! creates an interface static route
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 an 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 20-3 illustrates this network.

Figure 20-3 Static Routing Illustration



```

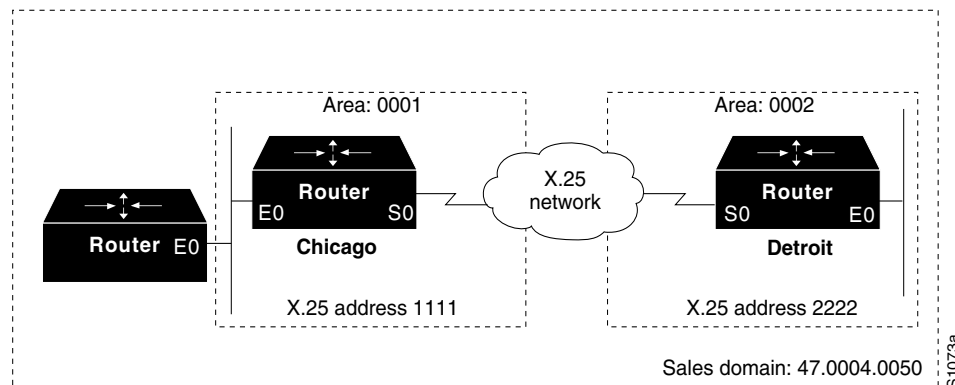
clns host foo 39.0001.1111.1111.1111.00
clns host bar 39.0002.2222.2222.2222.00
! assign a static address for the router
clns net foo
! enables CLNS packets to be routed
clns routing
!
interface ethernet 0
! pass ISO CLNS packet traffic to end systems without routing them
clns enable
!
interface ethernet 1
! pass ISO CLNS packet traffic to end systems without routing them
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 20-4 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 20-4 CLNS X.25 Intradomain Routing



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

```

! defines the name chicago to be used in place of the following NSAP
clns host chicago 47.0004.0050.0001.0000.0c00.243b.00
! defines the name detroit to be used in place of the following NSAP
clns host detroit 47.0004.0050.0002.0000.0c00.1e12.00
! enable routing of CLNS packets
clns routing
router iso-igrp sales
! configure net chicago, as defined above
net chicago
!
interface ethernet 0
! specify iso-igrp routing using the previously specified tag sales
clns router iso-igrp sales
!
interface serial 0
! set the interface up as a DTE with X.25 encapsulation
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 20-5 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 20-5 CLNS Interdomain Static Routing

Router A

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

```

! defining 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
!
interface ethernet 0
! specify iso-igrp routing using the previously specified tag orion
clns router iso-igrp orion
!
interface serial 1
! pass ISO CLNS traffic to end systems without routing
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 following network entity title for the routing process
net 47.0007.0200.0200.1112.1314.1516.00
! define the name foo to be used in place of Router A's NSAP
clns host foo 47.0006.0200.0100.0001.0102.0304.0506.00
!
interface ethernet 0
! specify iso-igrp routing using the previously specified tag pleiades
clns router iso-igrp pleiades
!
interface serial 0
! pass ISO CLNS traffic to end systems without routing
clns enable
! pass packets bound for foo in domain 47.0006.0200 through serial 0
clns route 47.0006.0200 foo

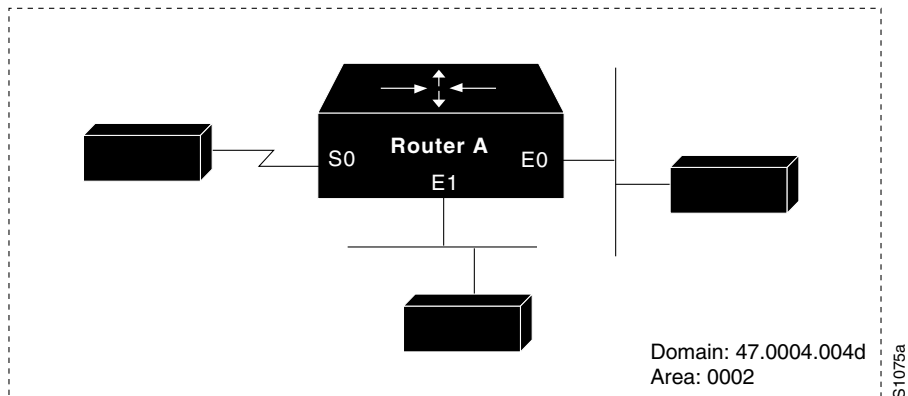
```

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

Dynamic Routing within the Same Area Example

Figure 20-6 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 20-6 CLNS Dynamic Routing within a Single Area



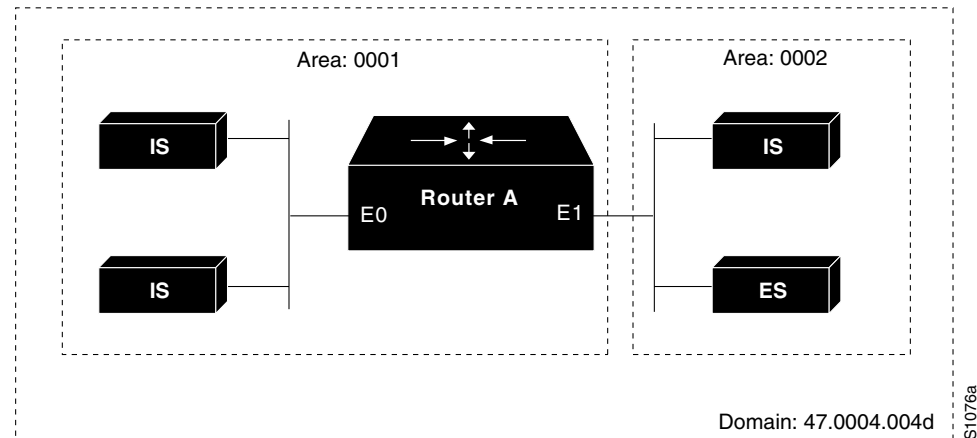
```

! enable clns packets to be routed
clns routing
! define a tag castor for the routing process
router iso-igrp castor
! configure the following net for the process in area 2, domain 47.0004.004d
net 47.0004.004d.0002.0000.0C00.0506.00
!
interface ethernet 0
! specify iso-igrp routing using the previously specified tag castor
clns router iso-igrp castor
!
interface ethernet 1
! specify iso-igrp routing using the previously specified tag castor
clns router iso-igrp castor
!
interface serial 0
! specify iso-igrp routing using the previously specified tag castor
clns router iso-igrp castor
    
```

Dynamic Routing in More Than One Area Example

Figure 20-7 and the example configuration that follows illustrate how to configure a router named Router A that exists in two areas.

Figure 20-7 CLNS Dynamic Routing within Two Areas



```

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

```

Dynamic Routing in Overlapping Areas Example

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

```


! enable routing of clns packets
clns routing
! define a tag capricorn for the routing process
router iso-igrp capricorn
! configure the following 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 following NET for the process in area 4, domain 47.0004.004d
net 47.0004.004d.0004.0000.0C00.0506.00
!
interface ethernet 0
! specify iso-igrp routing on interface ethernet 0 using the tag capricorn
clns router iso-igrp capricorn
!
interface ethernet 1
! specify iso-igrp routing on interface ethernet 1 using the tag capricorn
clns router iso-igrp capricorn
! specify iso-igrp routing on interface ethernet 1 using the tag cancer
clns router iso-igrp cancer
!
interface ethernet 2
! specify iso-igrp routing on interface ethernet 2 using the tag cancer
clns router iso-igrp cancer

```

Dynamic Interdomain Routing Example

Figure 20-8 and the configurations that follow illustrate how to configure three domains that are to be transparently connected.

Figure 20-8 CLNS Dynamic Interdomain Routing



Router Chicago

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

```
! enable routing of clns packets
clns routing
! define a tag A for the routing process
router iso-igrp A
! configure the following 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 following 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
!
interface ethernet 0
! specify iso-igrp routing with the tag A
clns router iso-igrp A
!
interface serial 0
! specify iso-igrp routing with the tag B
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.

```
clns routing
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.

```

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

```

clns routing
router isis
net 47.0004.004d.0001.0000.0c00.1111.00
interface ethernet 0
clns router isis
interface ethernet 1
clns router isis
interface serial 0
clns router isis
! 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.

```

clns routing
! 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
isis-type level-1-2
interface ethernet 0
clns router isis areal
! 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 areal
isis metric 2 level-2
isis circuit-type level-2-only

```

```
interface serial 0
clns router isis areal
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
```

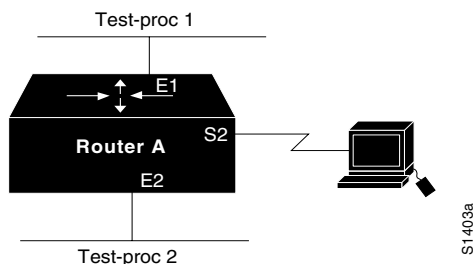
Configuring a Router in Two Areas Example

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 20-9 illustrates this configuration.

Figure 20-9 ISO-IGRP Configuration



```
clns routing
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 20-9 still applies. Ethernet interface 2 is configured for IS-IS routing and is assigned the tag of test-proc2.

```

clns routing
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

```

Configuring 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 20-10 illustrates this configuration.

Figure 20-10 Routers Acting as DTEs and DCEs



Router A

```

clns routing
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 31101 broadcast
```

Router B

```
clns routing
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
```

Customizing Performance Parameters Example

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

```
clns routing
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
interface serial 2
! specify an mtu of 978 bytes; generally, do not alter the default mtu value
clns mtu 978
```

Configuring 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
!
interface ethernet 0
! enable cluster aliasing on interface ethernet 0
clns cluster-alias
!
interface ethernet 1
! enable cluster aliasing on interface ethernet 1
clns cluster-alias
```

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...

```

CLNS Filter Examples

The following example returns a permit action if an address starts with either 47.0005 or 47.0023. It returns an implicit deny action on any other address.

```

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

```

The following example returns a deny action if an address starts with 39.840F, but returns a permit action for 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

```

