



Troubleshooting ISO CLNS

This chapter presents protocol-related troubleshooting information for International Organization for Standardization (ISO) Connectionless Network Service (CLNS) protocol connectivity and performance problems. ISO CLNS is a network layer standard that is part of the Open System Interconnection (OSI) protocol suite.

The Cisco IOS 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 routing 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 network service access point (NSAP)-to-X.121 mapping. The LAPB, X.25, Frame Relay, and SMDS encapsulations interoperate with other vendors.

Cisco's CLNS implementation is also compliant with the Government Open Systems Interconnection Profile (GOSIP) Version 2. As part of its CLNS support, Cisco routers fully support the following 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 NSAP addresses.
- ISO 10589—Documents Intermediate System-to-Intermediate System (IS-IS) Intra-domain Routing Exchange Protocol.

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

ISO CLNS Technology Basics

The world of OSI networking has a unique terminology:

- *End system* (ES) refers to any nonrouting network device.
- *Intermediate system* (IS) refers to a router.
- *Area* is a group of contiguous networks and attached hosts that are specified by a network administrator or manager to be an area.
- *Domain* is a collection of connected areas. Routing domains provide full connectivity to all end systems within them.

ISO CLNS Addressing

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 (see Figure 12-1). This byte is called the *n-selector*. Its function is similar to the port number in other protocol suites.

Cisco's implementation supports all NSAP address formats that are defined by ISO 8348/Ad2; however, Cisco provides 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 initial domain part (IDP) is made up of 1-byte 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 authority and format identifier (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 12-1 illustrates the ISO-IGRP NSAP addressing structure.

Figure 12-1 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 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.

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 12-2 illustrates the IS-IS NSAP addressing structure.

Figure 12-2 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 can match all but the n-selector (S) field in the DSP.
- ISO-IGRP requires at least 10 bytes of length; 1 for domain, 2 for area, 6 for system ID, and 1 for n-selector.
- Cisco's implementation of IS-IS requires at least 8 bytes; 1 for area, 6 for system ID, and 1 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 second example is the OSI network NSAP address format:

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

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 Table 12-1, which is an example of a static routing table, the next-hop NETs are listed for completeness but are not necessary to understand the routing algorithm. Table 12-2 offers examples of how the longest matching NSAP prefix can be matched with routing table entries in Table 12-1.

Table 12-1 Sample Routing Table Entries

Entry	NSAP Address Prefix	Next-Hop NET
1	47.0005.000c.0001	47.0005.000c.0001.0000.1234.00

Table 12-1 Sample Routing Table Entries

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 12-2 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.

Troubleshooting ISO CLNS

This section presents protocol-related troubleshooting information for ISO CLNS protocol connectivity and performance problems. It describes specific ISO CLNS symptoms, the problems that are likely to cause each symptom, and the solutions to those problems.



Note

Discussions of host configuration problems in this chapter assume that the host is a UNIX system. Equivalent actions might also be applicable to non-UNIX hosts, but the discussions do not specifically address non-UNIX end-station problems.

The following sections cover the most common network issues in ISO CLNS networks:

- ISO CLNS: Host Cannot Access Hosts on Local or Remote Network
- ISO CLNS: Host Cannot Access Hosts in Same Area
- ISO CLNS: Host Cannot Access Hosts in Different Area
- ISO CLNS: Connections Fail Using Certain Protocols
- ISO CLNS: Users Cannot Make Connections over Parallel Path
- ISO CLNS: Redistribution Causes Routing Problems
- ISO CLNS: Poor Performance

ISO CLNS: Host Cannot Access Hosts on Local or Remote Network

Symptom: Hosts cannot communicate with other hosts. Hosts might be located on the local or a remote network. Connections to some hosts on a network might be possible, whereas connections to other hosts on the same network fail.

Table 12-3 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-3 ISO CLNS: Host Cannot Access Hosts on Local or Remote Network

Possible Problem	Solution
Missing or misconfigured default gateway specification	<ol style="list-style-type: none"> Determine whether a default gateway is specified in the adjacency table of the host attempting to make a connection. Use the following UNIX command: <pre>host% netstat -rn</pre> Check the output of this command for a default gateway specification. Syntax Description: <ul style="list-style-type: none"> netstat—Displays protocol statistics and current TCP/IP¹ network connections r—Displays the contents of the routing table n—Displays addresses and port numbers in numeric form If the default gateway specification is incorrect, or if it is not present at all, you can change or add a default gateway using the following UNIX command at the local host: <pre>host% route add default <i>address</i> 1</pre> where <i>address</i> is the IP address of the default gateway (the router local to the host). The value 1 indicates that the specified gateway is one hop away. It is recommended that you specify a default gateway as part of the boot process. Specify the ISO CLNS address of the gateway in the following UNIX host file: <pre>/etc/defaultrouter</pre> This filename might be different on your UNIX system.
End system has no Level 1 router	<ol style="list-style-type: none"> Use the show clns neighbors detail privileged exec command to show all ESs² and ISs³ to which the router is directly connected. Make sure there is at least one Level 1 router on the same network as the end system.

Table 12-3 ISO CLNS: Host Cannot Access Hosts on Local or Remote Network (continued)

Possible Problem	Solution
Level 1 router or ES has bad address	<ol style="list-style-type: none"> 1. Verify that the Level 1 router has the same address as the ES. 2. Verify that all bytes of the NSAP⁴ address, up to but not including the system ID, are the same on both the router and the ES. The domain and area addresses must match, and the station IDs must be unique. (The value of the n-selector byte has no impact in this case.)
ES host is not running ES-IS ⁵ protocol	<ol style="list-style-type: none"> 1. Use the appropriate host commands to verify that an ES-IS process is running. If necessary, initiate the ES-IS process on the host. 2. Check the adjacency database on the host and verify that it has an entry for its directly connected router. 3. Use the debug clns packet privileged exec command on the Level 1 router to verify that it sees and forwards packets from the ES. <p>Caution: Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use debug commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. Moreover, it is best to use debug commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use.</p> <ol style="list-style-type: none"> 4. If necessary, statically configure the router to recognize the ES by using the clns es-neighbor interface configuration command. The following is the syntax for the clns es-neighbor command: <pre>clns es-neighbor nsap snpa</pre> <p>Syntax Description:</p> <ul style="list-style-type: none"> • <i>nsap</i>—Specific NSAP to map to a specific MAC⁶ address. • <i>snpa</i>—Data link (MAC) address. <p>Example:</p> <p>The following example defines an ES neighbor on Ethernet interface 0:</p> <pre>interface ethernet 0 clns es-neighbor 47.0004.004D.0055.0000.0C00.A45B.00 0000.0C00.A45B</pre> <p>In this case, the end systems with the following NSAP, or NET,⁷ are configured with an Ethernet MAC address of 0000.0C00.A45B:</p> <ul style="list-style-type: none"> • 47.0004.004D.0055.0000.0C00.A45B.00

continues

Table 12-3 ISO CLNS: Host Cannot Access Hosts on Local or Remote Network (continued)

Possible Problem	Solution
Router between hosts is down	<ol style="list-style-type: none"> Use the trace exec command to check connectivity between routers and the source ES. If the trace fails at a router, use the show clns neighbors exec command to see which neighboring routers and ESs are recognized. <p>Sample Display:</p> <p>The following is sample output from the show clns neighbors command. This display is a composite of the show clns es-neighbor and show clns is-neighbor commands:</p> <pre> router# show clns neighbors System Id SNPA Interface State Holdtime Type Protocol 0000.0000.0007 aa00.0400.6408 Ethernet0 Init 277 IS ES-IS 0000.0C00.0C35 0000.0c00.0c36 Ethernet1 Up 91 L1 IS-IS 0800.2B16.24EA aa00.0400.2d05 Ethernet0 Up 29 L1L2 IS-IS 0800.2B14.060E aa00.0400.9205 Ethernet0 Up 1698 ES ES-IS 0000.0C00.3E51 *HDLC* Serial1 Up 28 L2 IS-IS 0000.0C00.62E6 0000.0c00.62e7 Ethernet1 Up 22 L1 IS-IS 0A00.0400.2D05 aa00.0400.2d05 Ethernet0 Init 24 IS ES-IS </pre> <ol style="list-style-type: none"> If neighboring routers and end systems are up, perform one of the following procedures: <ul style="list-style-type: none"> For ISO-IGRP,⁸ check the routing table and see whether the routes are being learned. Use the show clns route exec command to display the routing tables. For IS-IS,⁹ check the LSP¹⁰ database to see whether the links are being reported in link state advertisements. Check the IS-IS routing table to see whether the routes are being installed in the routing table. Use the show isis database detail exec command to display the routing tables.
Route redistribution problem	Misconfigured route redistribution can cause connectivity problems. For specific troubleshooting information, see the section “ISO CLNS: Redistribution Causes Routing Problems” later in this chapter.

- TCP/IP = Transmission Control Protocol/Internet Protocol
- ES = end system
- IS = intermediate system
- NSAP = Network service access point
- ES-IS = End System-to-Intermediate System
- MAC = Media Access Control
- NET = network entity title
- IGRP = Interior Gateway Routing Protocol

9. IS-IS = Intermediate System-to-Intermediate System
10. LSP = Link State Protocol

ISO CLNS: Host Cannot Access Hosts in Same Area

Symptom: Hosts cannot access other hosts in the same area. The hosts might be on the same network or they might be in a different network in the same area.

Table 12-4 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-4 ISO CLNS: Host Cannot Access Hosts in Same Area

Possible Problem	Solution
Area address is configured incorrectly on the host	<ol style="list-style-type: none"> 1. Check all Level 1 routing tables and link-state databases. 2. Verify that the hosts are in the same area. 3. Check that the NSAP address is entered correctly on the hosts.
Different area addresses are merged into a single area, but the router is configured incorrectly	<ol style="list-style-type: none"> 1. Use the show running-config privileged exec command to see router configurations. Check whether multiple area addresses are configured. 2. If multiple network addresses are configured, verify that the router is configured to support a multihomed area (a single area that has more than one area address; see Figure 12-3). 3. To communicate, routers must establish a Level 1 adjacency. Therefore, area addresses in a multihomed area must overlap across routers. <p>For example, in the multihomed area shown in Figure 12-3, to configure Area 1 and Area 2 as a multihomed area, both Router A and Router B must be configured to be in both areas. IS-IS routing supports the assignment of multiple area addresses on the same router. This concept is referred to as <i>multihoming</i>. Multihoming provides a mechanism for smoothly migrating network addresses, as follows:</p> <ul style="list-style-type: none"> • Splitting up an area—Nodes within a given area can accumulate until 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 that have a common area address into a single area.


continues

Table 12-4 ISO CLNS: Host Cannot Access Hosts in Same Area (continued)

Possible Problem	Solution
Different area addresses are merged into a single area, but the router is configured incorrectly (continued)	<ul style="list-style-type: none"> 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. <p>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.</p> <p>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.</p> <ol style="list-style-type: none"> Alternatively, one router can be configured in both areas, while the other router remains configured for a single area. Provided that the area numbers on routers overlap, the routers will establish a Level 1 adjacency, allowing them to communicate.
ES host is not running ES-IS protocol	<ol style="list-style-type: none"> Use the appropriate host commands to verify that an ES-IS process is running. If necessary, initiate the ES-IS process on the host. Check the adjacency database on the host and verify that it has an entry for its directly connected router. Use the debug clns packet privileged exec command on the Level 1 router to verify that it sees and forwards packets from the ES. <p>Caution: Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use debug commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. Moreover, it is best to use debug commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use.</p> <ol style="list-style-type: none"> If necessary, statically configure the router to recognize the ES by using the clns es-neighbor interface configuration command. The following is the syntax for the clns es-neighbor command: <pre>clns es-neighbor nsap snpa</pre>

Table 12-4 ISO CLNS: Host Cannot Access Hosts in Same Area (continued)

Possible Problem	Solution
ES host is not running ES-IS protocol (continued)	<p>Syntax Description:</p> <ul style="list-style-type: none"> <i>nsap</i>—Specific NSAP to map to a specific MAC address. <i>snpa</i>—Data link (MAC) address. <p>Example:</p> <p>The following example defines an ES neighbor on Ethernet interface 0:</p> <pre>interface ethernet 0 clns es-neighbor 47.0004.004D.0055.0000.0C00.A45B.00 0000.0C00.A45B</pre>
Route redistribution problem	Misconfigured route redistribution can cause connectivity problems. For specific troubleshooting information, see the section “ISO CLNS: Redistribution Causes Routing Problems” later in this chapter.

Figure 12-3 Multihomed Area Sample Network

ISO CLNS: Host Cannot Access Hosts in Different Area

Symptom: Host cannot access hosts in a different area. Hosts in the same area are accessible.

Table 12-5 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-5 ISO CLNS: Host Cannot Access Hosts in Different Area

Possible Problem	Solution
Level 2 routers are not routing packets to the correct area	<ol style="list-style-type: none"> Use the trace command to verify that Level 1 routers are routing packets to the nearest Level 2 router. Sample Display: The following display shows an example of ISO CLNS trace output: <pre> router# trace Protocol [ip]: clns Target CLNS address: thoth Timeout in seconds [3]: Probe count [3]: Minimum Time to Live [1]: Maximum Time to Live [30]: Type escape sequence to abort. Tracing the route to THOTH (55.0006.0100.0000.0000.0001.8888.1112.1314.1516) HORUS(55.0006.0100.0000.0000.0001.6666.3132.3334.353 6) 32 msec ! 28 msec 28 msec ! 2 ISIS(55.0006.0100.0000.0000.0001.7777.2122.2324.2526) 56 msec ! 80 msec 56 msec ! 3 THOTH(55.0006.0100.0000.0000.0001.8888.1112.1314.151 6) 80 msec ! 80 msec ! 8 </pre> Use the trace exec command to verify that Level 2 routers are routing packets to the correct destination area. If packets are not being routed to the correct area, check the Level 2 routing tables (ISO-IGRP¹) or the Level 2 link state databases (IS-IS) to see whether the packets are being forwarded to another area. If necessary, reconfigure routers with the correct area addresses and Level 2 (IS-IS) routing information.
ES host is not running ES-IS protocol	<ol style="list-style-type: none"> Use the appropriate host commands to verify that an ES-IS process is running. If necessary, initiate the ES-IS process on the host. Check the adjacency database on the host and verify that it has an entry for its directly connected router. Use the debug clns packet privileged exec command on the Level 1 router to verify that it sees and forwards packets from the ES.

Table 12-5 ISO CLNS: Host Cannot Access Hosts in Different Area (continued)

Possible Problem	Solution
ES host is not running ES-IS protocol (continued)	<p>Caution: Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use debug commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. Moreover, it is best to use debug commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use.</p> <p>4. If necessary, statically configure the router to recognize the ES by using the clns es-neighbor interface configuration command.</p>
Route redistribution problem	Misconfigured route redistribution can cause connectivity problems. For specific troubleshooting information, see the section “ISO CLNS: Redistribution Causes Routing Problems” later in this chapter.

Table 12-5 ISO CLNS: Host Cannot Access Hosts in Different Area (continued)

Possible Problem	Solution
Router between hosts is down	<p>1. Use the trace exec command to check connectivity between routers and the source ES.</p> <p>Sample Display:</p> <p>The following display shows an example of ISO CLNS trace output.</p> <pre> router# trace Protocol [ip]: clns Target CLNS address: thoth Timeout in seconds [3]: Probe count [3]: Minimum Time to Live [1]: Maximum Time to Live [30]: Type escape sequence to abort. Tracing the route to THOTH (55.0006.0100.0000.0000.0001.8888.1112.1314.1516) HORUS(55.0006.0100.0000.0000.0001.6666.3132.3334.353 6) 32 msec ! 28 msec 28 msec ! 2 ISIS(55.0006.0100.0000.0000.0001.7777.2122.2324.2526) 56 msec ! 80 msec 56 msec ! 3 THOTH(55.0006.0100.0000.0000.0001.8888.1112.1314.151 6) 80 msec ! 80 msec ! 8 </pre> <p>2. If the trace fails at a router, use the show clns neighbors exec command to see which neighboring routers and ESs are recognized.</p>
Router between hosts is down	<p>3. If neighboring routers and end systems are up, perform one of the following procedures:</p> <ul style="list-style-type: none"> • For ISO-IGRP, check the routing table and see whether the routes are being learned. Use the show clns route exec command to display the routing tables. • For IS-IS, check the LSP² database to see whether the links are being reported in link state advertisements. Check the IS-IS routing table to see whether the routes are being installed in the routing table. Use the show isis database detail exec command to display the routing tables.

1. IGRP = Interior Gateway Routing Protocol

2. LSP = Link State Protocol

ISO CLNS: Connections Fail Using Certain Protocols

Symptom: Host connections fail using certain protocols. Hosts might be able to connect to other hosts using some protocols but are unable to connect using others.

Table 12-6 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-6 ISO CLNS: Connections Fail Using Certain Protocols

Possible Problem	Solution
Host is not configured to support the service	Verify that the needed protocols are correctly installed and configured on the host system. Consult your vendor's documentation for information on configuring hosts.
Misconfigured access list	<ol style="list-style-type: none"> 1. Use the trace exec command to determine the path taken to reach remote hosts. 2. If you discover a router that is stopping traffic, use the show access-lists privileged exec command to see whether any access lists are configured on the router. 3. Disable all access lists on the router using no access-group interface configuration commands on the appropriate interfaces. 4. Determine whether hosts can now use the protocol in question. If traffic can get through, it is likely that an access list is blocking protocol traffic. 5. Make sure the access list does not filter traffic from ports that are used by the protocol in question. Configure explicit permit statements for traffic that you want the router to forward normally. 6. Enable the access list and verify that the protocol still functions correctly. If problems persist, continue isolating and analyzing access lists on all routers in the path from source to destination.

ISO CLNS: Users Cannot Make Connections over Parallel Path

Symptom: In environments with multiple paths between networks, when one link goes down, connections across a parallel link are not possible.



Note

IS-IS has equal-cost load balancing for both Level 1 and Level 2 routes. If there are parallel paths in an IS-IS network and one goes down, the other should serve as a backup that is ready to be used immediately.

Table 12-7 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-7 ISO CLNS: Users Cannot Make Connections over Parallel Path

Possible Problem	Solution
Routing has not converged	<ol style="list-style-type: none"> 1. Use the show clns route privileged exec command to view the CLNS routing table. Examine the table for routes listed as “possibly down.” This indicates that the routing protocol has not converged. 2. Wait for the routing protocol to converge. Use the show clns route command again to see whether the routes are now up. <p>Note: ISO-IGRP does load balancing only for domain prefix routes. If you are doing Level 1 or Level 2 routing in ISO-IGRP, only a single path is maintained. If that path goes down, you must wait for the network to converge before the alternate path is available.</p>
Misconfigured access list	<ol style="list-style-type: none"> 1. Use the trace exec command to determine the path taken to reach remote hosts. 2. If you discover a router that is stopping traffic, use the show access-lists privileged exec command to see whether any access lists are configured on the router. 3. Disable all access lists on the router using no access-group interface configuration commands on the appropriate interfaces. 4. Determine whether hosts can now use the protocol in question. If traffic can get through, it is likely that an access list is blocking protocol traffic. 5. Make sure the access list does not filter traffic from ports that are used by the protocol in question. Configure explicit permit statements for traffic that you want the router to forward normally. 6. Enable the access list and verify that the protocol still functions correctly. If problems persist, continue isolating and analyzing access lists on all routers in the path from source to destination.
Hardware or media problem	<p>For information on troubleshooting hardware problems, see Chapter 3. “Troubleshooting Hardware and Booting Problems.” For information on troubleshooting media problems, refer to the media troubleshooting chapter that covers the media type used in your network.</p>

ISO CLNS: Redistribution Causes Routing Problems

Symptom: Route redistribution does not work properly and causes routing problems. Traffic does not get through a router that is redistributing routes between two different routing areas or domains—typically IS-IS and ISO-IGRP. Observed symptoms range from poor performance to no communication at all.

Table 12-8 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-8 ISO CLNS: Redistribution Causes Routing Problems

Possible Problem	Solution
Misordered sequence numbers	<p>The sequence numbers used in route-map router configuration commands determine the order in which conditions are tested. Misordered sequence numbers can cause redistribution problems.</p> <ol style="list-style-type: none"> 1. Use the show running-config privileged exec command to display the router configuration. Look for route-map router configuration command entries. 2. If route-map commands are configured, look at the sequence numbers that are assigned. Lower sequence numbers are tested before higher sequence numbers, regardless of the order in which they are listed in the configuration. 3. If conditions are not being tested in the order you want, you must modify the sequence numbers to change the testing order. The syntax for the route-map command to adjust the sequence number is as follows: <p style="text-align: center;">route-map <i>map-tag</i> {permit deny} <i>sequence-number</i></p> <p>Syntax Description:</p> <ul style="list-style-type: none"> • <i>map-tag</i>—Meaningful name for the route map. The redistribute command uses this name to reference this route map. Multiple route maps can share the same map tag name. Can either be an expression or a filter set.
Misordered sequence numbers (continued)	<ul style="list-style-type: none"> • permit—If the match criteria are met for this route map and permit is specified, the route is redistributed as controlled by the set actions. If the match criteria are not met, and permit is specified, the next route map with the same map tag is tested. If a route passes none of the match criteria for the set of route maps sharing the same name, it is not redistributed by that set. • deny—If the match criteria are met for the route map and deny is specified, the route is not redistributed and no further route maps sharing the same map tag name will be examined. • <i>sequence-number</i>—Number that indicates the position a new route map is to have in the list of route maps already configured with the same name. If given with the no form of this command, it specifies the position of the route map that should be deleted.

Table 12-8 ISO CLNS: Redistribution Causes Routing Problems (continued)

Possible Problem	Solution
Missing or misconfigured default-metric command	<p>1. Use the show running-config exec command to view the router configuration. Look for a default-metric router configuration command entry.</p> <p>2. If the default-metric router configuration command or the distance router configuration command is missing, add the appropriate version of the missing command.</p> <p>Syntax:</p> <p>The following is the syntax for the default-metric command:</p> <pre>default-metric <i>number</i></pre> <p>Syntax Description:</p> <ul style="list-style-type: none"> <i>number</i>—Default metric value appropriate for the specified routing protocol. <p>Syntax:</p> <p>The following is the syntax for the distance command:</p> <pre>distance <i>value</i> [clns]</pre> <p>Syntax Description:</p> <ul style="list-style-type: none"> <i>value</i>—Administrative distance, indicating the trustworthiness of a routing information source. This argument has a numeric value between 0 and 255. A higher relative value indicates a lower trustworthiness rating. Preference is given to routes with smaller values. The default, if unspecified, is 110. clns—(Optional) CLNS-derived routes for IS-IS. <p>Refer to the <i>Cisco IOS Network Protocols Configuration Guide, Part 2</i> and <i>Network Protocols Command Reference, Part 2</i>, for information about adjusting ISO CLNS default metrics.</p>

continues

Table 12-8 ISO CLNS: Redistribution Causes Routing Problems (continued)

Possible Problem	Solution
Missing or misconfigured distance command	<p data-bbox="570 308 1284 407">1. Use the show running-config exec command to view the router configuration. Look for a distance router configuration command entry.</p> <p data-bbox="570 413 1284 548">2. If the distance command is missing, configure a distance specification on the router. Use the distance router configuration command to configure the administrative distance for CLNS routes learned.</p> <p data-bbox="570 554 1284 588">Syntax:</p> <p data-bbox="570 594 1284 636">The following is the syntax for the distance command:</p> <pre data-bbox="613 659 857 693">distance <i>value</i> [clns]</pre> <p data-bbox="570 716 1284 749">Syntax Description:</p> <ul data-bbox="570 756 1284 1003" style="list-style-type: none"> <li data-bbox="570 756 1284 961">• <i>value</i>—Administrative distance, indicating the trustworthiness of a routing information source. This argument has a numeric value between 0 and 255. A higher relative value indicates a lower trustworthiness rating. Preference is given to routes with smaller values. The default, if unspecified, is 110. <li data-bbox="570 968 1284 1003">• clns—(Optional) CLNS-derived routes for IS-IS. <p data-bbox="570 1010 1284 1043">Example:</p> <p data-bbox="570 1050 1284 1155">In the following example, the distance value for CLNS routes learned is 90. Preference is given to these CLNS routes rather than routes with the default administrative distance value of 110:</p> <pre data-bbox="613 1178 813 1272">router isis distance 90 clns</pre>

Table 12-8 ISO CLNS: Redistribution Causes Routing Problems (continued)

Possible Problem	Solution
Redistribution feedback loop exists	<p>Redistribution between an IS-IS cloud and an ISO-IGRP cloud should be performed only at a single point. If it is not, routing information can be advertised back into one of the clouds, causing routing feedback loops.</p> <p>Examples:</p> <p>The following example illustrates redistribution of ISO-IGRP routes of Michigan and ISO-IGRP routes of Ohio into the IS-IS area tagged USA:</p> <pre>router isis USA redistribute iso-igrp Michigan redistribute iso-igrp Ohio</pre> <p>The following example illustrates redistribution of IS-IS routes of France and ISO-IGRP routes of Germany into the ISO-IGRP area tagged Backbone:</p> <pre>router iso-igrp Backbone redistribute isis France redistribute iso-igrp Germany</pre>
Redistribution feedback loop exists (<i>continued</i>)	<p>If you must redistribute at another point, use default metrics to perform the redistribution in one direction only.</p> <p>Refer to the Cisco IOS <i>Network Protocols Configuration Guide, Part 2</i>, and <i>Network Protocols Command Reference, Part 2</i>, for information about adjusting ISO CLNS default metrics.</p>

ISO CLNS: Poor Performance

Symptom: Users experience poor performance or sudden loss of connections. One or more routers might be receiving duplicate routing updates and might see routers and ESs on multiple interfaces.

Table 12-9 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 12-9 ISO CLNS: Poor Performance

Possible Problem	Solution
Multiple ISO-IGRP processes are configured on a single interface	<ol style="list-style-type: none">1. Use the show clns interface exec command to view the interface configuration. Look for multiple ISO-IGRP processes that are configured on a single interface.2. If multiple ISO-IGRP processes are configured on a single interface, different Level 2 updates are being sent out through the same interface. Multiple Level 2 updates on the same interface can cause congestion problems, especially if the network is large and links are flapping outside the damping intervals. Flapping is a routing problem where an advertised route between two nodes alternates (flaps) back and forth between two paths due to a network problem that causes intermittent interface failures.3. Remove one of the ISO-IGRP processes from the interface configuration using the appropriate no clns router iso-igrp interface configuration command.

Table 12-9 ISO CLNS: Poor Performance (continued)

Possible Problem	Solution
Bridge or repeater in parallel with router	<p>A bridge or repeater in parallel with a router can cause updates and traffic to be seen from both sides of an interface.</p> <ol style="list-style-type: none"> Use the show clns is-neighbors detail and the show clns neighbors detail exec commands to see through which routers and protocols the router's adjacencies were learned. <p>The following is sample output from the show clns neighbors detail command:</p> <pre> router# show clns neighbors detail System Id SNPA Interface State Holdtime Type Protocol 000.0000.0007 aa00.0400.6408 Ethernet0 Init 291 IS ES-IS Area Address(es): 47.0005.80FF.F500.0000.0003.0020 0000.0C00.0C35 0000.0c00.0c36 Ethernet1 Up 94 L1 IS-IS Area Address(es): 47.0004.004D.0001 39.0001 0800.2B16.24EA aa00.0400.2d05 Ethernet0 Up 9 L1L2 IS-IS Area Address(es): 47.0004.004D.0001 0800.2B14.060E aa00.0400.9205 Ethernet0 Up 1651 ES ES-IS Area Address(es): 49.0040 0000.0C00.3E51 *HDLC* Serial1 Up 27 L2 IS-IS Area Address(es): 39.0004 0000.0C00.62E6 0000.0c00.62e7 Ethernet1 Up 26 L1 IS-IS Area Address(es): 47.0004.004D.0001 0A00.0400.2D05 aa00.0400.2d05 Ethernet0 Init 29 IS ES-IS Area Address(es): 47.0004.004D.0001 </pre> <p>Look for routers that are known to be on a remote network. A router listed in the adjacency table but that is not on a directly connected network indicates a problem.</p> <p>You can also look for paths to networks (or areas) on multiple interfaces.</p> <ol style="list-style-type: none"> If you determine that there is a parallel bridge or repeater, remove the device or configure filters that block routing updates from being learned from the device.
Route redistribution problem	<p>Misconfigured route redistribution can cause performance problems. For specific troubleshooting information, see the section "ISO CLNS: Redistribution Causes Routing Problems" earlier in this chapter.</p>