

NOTE

This chapter is a supplement to the Cisco Press book *Building Scalable Cisco Networks* (ISBN 1-57870-228-3). It covers the IS-IS material in accordance with the latest version of *Building Scalable Cisco Internetworks*. This chapter was created to follow Chapter 4, “Interconnecting Multiple OSPF Areas.” This placement is reflected in this chapter’s starting assumption in the configuration exercise. This supplemental chapter is not meant to replace any material in the original book, but it should be read between Chapters 4 and 5 of the book.

In addition to this supplemental chapter, you will find online addenda to Appendixes D, G, H, and I. The addendum to Appendix D, “Summary of BSCN Router Commands,” provides you with a descriptive list of the commands you will find in this chapter. The addendum to Appendix G, “Answers to the Review Questions,” provides you with the answers to the review questions at the end of this chapter. The addendum to Appendix H, “Configuration Exercise Equipment Requirements and Backbone Configurations,” provides the configuration commands for the backbone router for the configuration exercise at the end of this chapter. The addendum to Appendix I, “Glossary,” provides a listing of some of the important terms and acronyms used in this chapter.

This chapter introduces the Intermediate System-to-Intermediate System (IS-IS) protocol. This chapter includes the following sections:

- Overview of OSI Protocols and IS-IS Routing
- Operation of IS-IS for CLNS/CLNP
- IP and OSI Routing with Integrated IS-IS
- Basic Integrated IS-IS Router Configuration
- Modeling WAN Networks in Integrated IS-IS
- Summary
- Configuration Exercise: Configuring a Multiarea IS-IS Network
- Answers to Configuration Exercise: Configuring a Multiarea IS-IS Network
- Review Questions

Configuring IS-IS Protocol

This chapter provides an overview of Intermediate System-to-Intermediate System (IS-IS) technology and its structures and protocols. It also gives basic configuration examples. The chapter begins with Open System Interconnection (OSI) routing and then focuses on Integrated IS-IS, which supports IP routing. Basic IS-IS and Integrated IS-IS router configuration commands, examples, and some troubleshooting guidelines are presented. The major part of this chapter is dedicated to an explanation of IS-IS concepts and capabilities, including the hierarchy and addressing of OSI-based networks.

When you finish this chapter, you will be able to explain the basic OSI terminology and network layer protocols used in OSI, the way in which networks and interfaces are represented in IS-IS, the basic principles of area routing, and the use of IS-IS in nonbroadcast multiaccess (NBMA) environments. You will be able to identify similarities and differences between Integrated IS-IS and OSPF, as well as characteristics of an effective addressing plan for IS-IS deployment. You will be able to list the types of IS-IS routers along with their roles in IS-IS area design, and describe the hierarchical structure of IS-IS areas, the concept of establishing adjacencies, and the concepts of routing information and database synchronization. Finally, given a set of network requirements, you will be able to configure Integrated IS-IS and verify proper operation (within described guidelines) of Cisco routers.

Overview of OSI Protocols and IS-IS Routing

This section provides an overview of the OSI protocols and IS-IS routing.

OSI Protocols

ISO and OSI: What's the Difference?

The International Organization for Standardization (ISO) was formed to develop standards for data networking. (As an interesting aside, note that the word *ISO* is Greek for “same.”)

The Open System Interconnection (OSI) protocols represent an international standardization program that facilitates multivendor equipment interoperability.

The OSI protocols are part of an international program to develop data-networking protocols and other standards that facilitate multivendor equipment interoperability. The OSI program grew out of a need for international networking standards and is designed to facilitate communication between hardware and software systems despite differences in underlying architectures.

The ISO has been charged with developing standards for data networking.

The OSI specifications were conceived and implemented by two international standards organizations: the ISO and the International Telecommunication Union Telecommunication Standardization Sector (ITU-T).

A Condensed History of IS-IS

IS-IS was ad hoc in its evolution, whereas OSPF was more formal:

1985: Originally called DECnet Phase V Routing

1988: Adopted by ISO and renamed IS-IS

1990: Publication of RFC 1142, “OSI IS-IS Intradomain Routing Protocol”

1990: Publication of RFC 1195, “Use of OSI IS-IS for Routing in TCP/IP and Dual Environments”

1991: Cisco IOS Software starts supporting IS-IS

1995: ISPs start adopting IS-IS

2000: Publication of IETF draft “IS-IS Extensions for Traffic Engineering”

2001: Publication of IETF draft “IS-IS Extensions in Support of Generalized MPLS”

The world of OSI internetworking includes various network services with these characteristics:

- Independence of underlying communications infrastructure
- End-to-end transfer
- Transparency
- Quality of service (QoS) selection
- Addressing

Transparency

A protocol is *transparent* when it does not place any constraints on transmitted data. It means that headers as well as data must be transported unmodified end to end.

Who Uses IS-IS?

IS-IS is popular among telcos and large ISPs. This popularity finds its roots with ISPs that were around at the beginning of the Internet and chose IS-IS over OSPF for their IGP. It is believed that at that time, IS-IS had fewer technical limitations than OSPF as an IGP. Those ISPs have since become today's tier 1 carriers, so any appliance targeting tier 1 carriers must offer IS-IS.

The OSI protocol suite supports numerous standard protocols at the physical, data link, network, transport, session, presentation, and application layers.

OSI network layer addressing is implemented by using two types of hierarchical addresses: network service access point (NSAP) addresses and a specific subset of NSAPs called network-entity titles (NETs). An NSAP is a conceptual point on the boundary between the network and transport layers. The NSAP is the location at which OSI network services are provided to the transport layer. Each transport layer entity is assigned a single NSAP, which is individually addressed in an OSI internetwork using NSAP addresses.

The OSI protocol suite specifies two routing protocols at the network layer: End System-to-Intermediate System (ES-IS) and IS-IS. In addition, the OSI suite implements two types of network services: connectionless service and connection-oriented service.

Differences Between Connection-Oriented and Connectionless Services

Connection-oriented services must first establish a connection with the desired service before passing any data. A connectionless service can send the data without establishing a connection first. In general, connection-oriented services provide some level of delivery guarantee, whereas connectionless services do not.

Details on these protocols and addresses are discussed later in this chapter, in the sections, "OSI Network Layer" and "OSI Addressing," respectively.

OSI Protocol Terminology

In an OSI network, four significant architectural entities exist: hosts, areas, a backbone, and a domain. The following describes these entities and tells how routers fit into an OSI network:

- A *domain* is any portion of an OSI network that is under a common administrative authority.

- Within any OSI domain, one or more *areas* can be defined. An area is a logical entity; it is formed by a set of contiguous routers and the data links that connect them. All routers in the same area exchange information about all the hosts that they can reach.
- The areas are connected to form a *backbone*. All routers on the backbone know how to reach all areas.
- An *end system* (ES) is any nonrouting *host* or node. An *intermediate system* (IS) is a *router*. These terms are the basis for the OSI ES-IS and IS-IS protocols.

NOTE Nowadays, domains tend to be called autonomous systems.

Mapping the OSI Protocol Suite to the OSI Reference Model

The OSI protocol suite supports numerous standard protocols at each of the seven OSI reference model layers.

Differences Between the OSI Protocol Suite and the OSI Reference Model

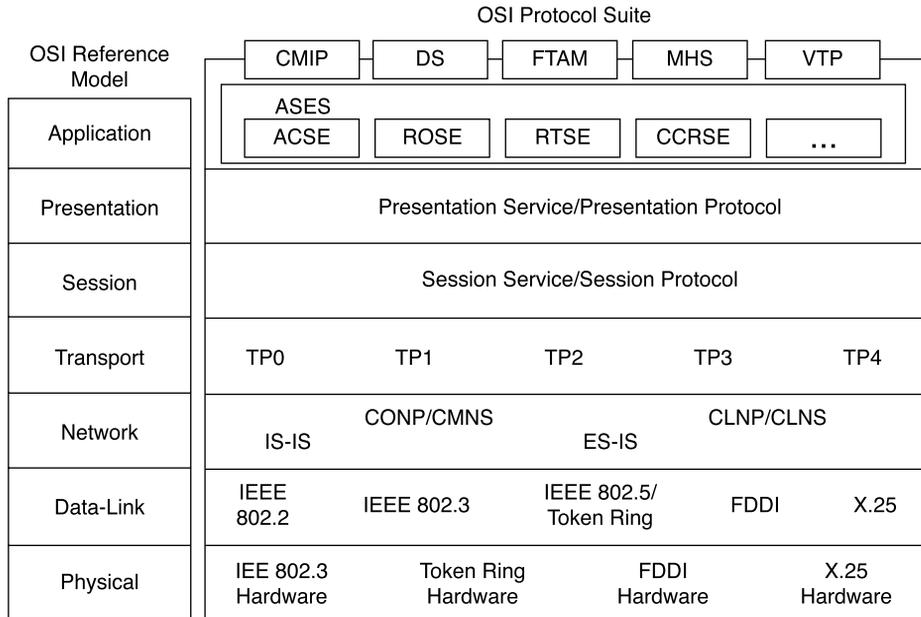
The ISO developed the Open Systems Interconnection networking suite in the 1980s. It has two major components:

- The abstract *model* of networking, known as the OSI Reference Model, or seven-layer model
- A set of concrete networking *protocols*, known as the OSI protocol suite, which include CLNP, ES-IS, and so on

The OSI Reference Model has enjoyed a far greater acceptance than the OSI protocol itself.

Figure S-1 illustrates the entire OSI protocol suite and its relation to the layers of the OSI reference model.

Figure S-1 How the OSI Protocol Suite Maps to the OSI Reference Model



OSI Network Layer

This section describes the services and protocols at the OSI Network layer.

OSI Services and Network Protocols

Two types of OSI network layer services are available to the OSI transport layer:

- **Connectionless Network Service (CLNS)**—CLNS performs datagram transport and does not require a circuit to be established before data is transmitted.
- **Connection-Mode Network Service (CMNS)**—CMNS requires explicit establishment of a path or circuit between communicating transport layer entities before transmitting data.

Whereas CLNS and CMNS define the actual services provided to the OSI transport layer entities that operate immediately above the network layer, Connectionless Network Protocol (CLNP) and Connection-Oriented Network Protocol (CONP) name the protocols that these services use to convey data at the network layer. CLNP is the OSI equivalent of IP.

The differences between CMNS and CONP are as follows:

- CONP is an OSI network layer protocol that carries upper-layer data and error indications over connection-oriented links. CONP is based on the X.25 Packet-Layer Protocol (PLP) and is described in the ISO 8208 standard “X.25 Packet-Layer Protocol for DTE.” CONP provides the interface between CMNS and upper layers. It is a network layer service that acts as the interface between the transport layer and CMNS; it is described in the ISO 8878 standard.
- CMNS performs functions related to the explicit establishment of paths through CONP. When support is provided for CMNS, routing uses the X.25 protocols as the relaying functions. CMNS functions include connection setup, maintenance, and termination; it also provides a mechanism for requesting a specific QoS.

The differences between CLNP and CLNS are as follows:

- CLNP is an OSI network layer protocol that carries upper-layer data and error indications over connectionless links. CLNP provides the interface between CLNS and upper layers.
- CLNS provides network layer services to the transport layer through CLNP. When support is provided for CLNS, routing uses routing protocols to exchange routing information. CLNS does not perform connection setup or termination because paths are determined independently for each packet that is transmitted through a network. In addition, CLNS provides best-effort delivery, which means that no guarantee exists that data will not be lost, corrupted, misordered, or duplicated. CLNS relies on transport layer protocols to perform error detection and correction.

Summary of OSI Protocols and Services

Type	Protocol	Service
Connection-oriented	CONP	CMNS
Connectionless	CLNP	CLNS

OSI Routing Protocols

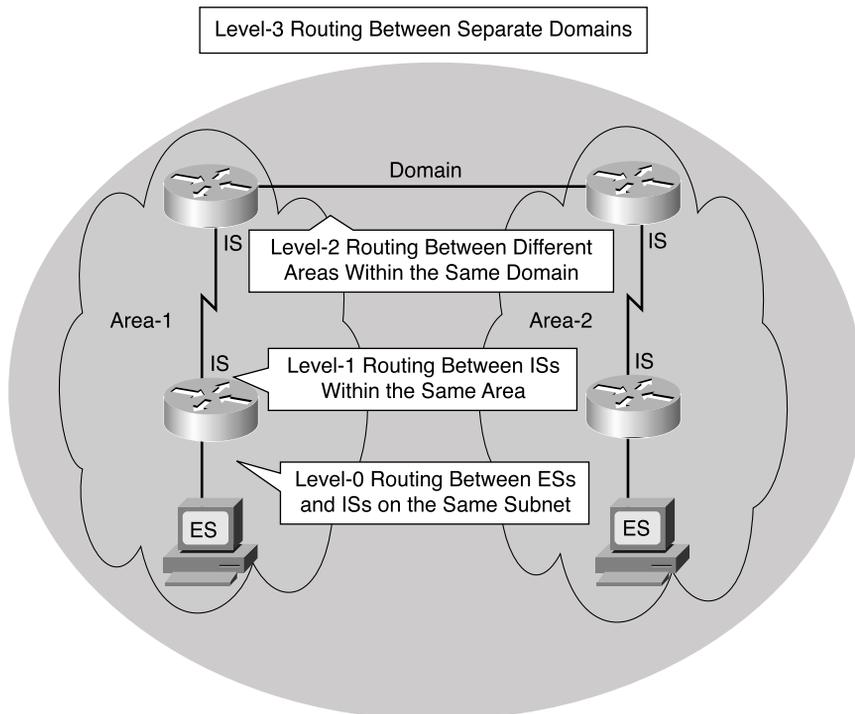
ISO has developed standards for two types of routing protocols:

- **ES-IS discovery protocol**—ES-IS performs “routing” between End Systems and Intermediate Systems referred as Level 0 “routing.” ES-IS is analogous to the Address Resolution Protocol (ARP) in IP. Although it is not explicitly a routing protocol, ES-IS is included here because it is commonly used with routing protocols to provide end-to-end data movement through an internetwork.

- IS-IS routing protocols**—IS-IS performs hierarchical (Level 1, Level 2, and Level 3) routing between intermediate systems. Level 3 routing is done between separate domains. However, note that the IS-IS routing protocol is not itself capable of Level 3 routing. As described in the section, “Interconnecting IS-IS Domains” later in this chapter, other protocols are required for interdomain routing.

The hierarchical routing levels used in OSI are illustrated in Figure S-2.

Figure S-2 *OSI Hierarchical Routing*



To simplify router design and operation, OSI distinguishes among Level 1, Level 2, and Level 3 routing. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Level 3 routing is done between separate domains. Hierarchical routing simplifies backbone design because Level 1 ISs need to know only how to get to the nearest Level 2 IS.

In OSI, each ES lives in a particular area. OSI routing begins when the ESs discover the nearest IS by listening to intermediate system hello (ISH) packets. When an ES wants to send a packet to another ES, it sends the packet to an IS on its directly attached network; this is Level 0 routing. The router then looks up the destination address and forwards the packet along the best route. If the destination ES is on the same subnetwork, the local IS knows this from listening to end system hello (ESH) packets and forwards the packet

appropriately. The IS also might provide a redirect message back to the source to tell it that a more direct route is available.

If the destination address is an ES on another subnetwork in the same area, the IS knows the correct route (Level 1 routing) and forwards the packet appropriately.

If the destination address is an ES in another area, the Level 1 IS sends the packet to the nearest Level 2 IS (Level 2 routing). Forwarding through Level 2 ISs continues until the packet reaches a Level 2 IS in the destination area. Within the destination area, ISs forward the packet along the best path until the destination ES is reached.

Routing between separate domains is referred to as Level 3 routing.

Routing in an OSI CLNS/CLNP Environment

For routing in the ISO CLNS/CLNP environment, Cisco routers support the following protocols:

- **IS-IS**—IS-IS is a dynamic link-state routing protocol used in an ISO CLNS environment for routing CLNP. Routers usually operate as ISs and can exchange reachability information with other ISs using the IS-IS protocol. As an IS, a Cisco router can be a Level 1 router, a Level 2 router, or a Level 1–2 router. In the latter case, the router can advertise itself at Level 1 as an exit point from the area. Integrated IS-IS allows the IS-IS protocol to propagate routing information for other protocols as well as—or instead of—CLNS. Specifically, IS-IS can route CLNS, IP, or both (this latter is called *dual* mode).
- **ISO-IGRP**—Cisco IOS Software offers a proprietary routing protocol for CLNS. As its name suggests, ISO-IGRP is based on Cisco’s Interior Gateway Routing Protocol (IGRP). It uses distance vector technology to propagate routing information. As such, it shares some of the limitations of its IP counterpart, including long convergence times (because of periodic updates and long invalid times and hold times).
- **Static CLNS routes**—As with IP, static CLNS routes can be created. (Although this is not really a protocol, static routes can be considered to be a type of routing protocol.)

Recommended Reading

Various aspects of IS-IS are described in these ISO documents:

- **ISO 8473**—Documents the ISO CLNP
- **ISO/IEC 8348, Appendix A**—Documents NSAP addresses
- **ISO 9542**—Documents the ES-IS routing exchange protocol
- **ISO/IEC 10589**—Documents the IS-IS intradomain routing exchange protocol

Additionally, the function of Integrated IS-IS—the use of OSI IS-IS for routing in TCP/IP and dual environments, as described in the next section—is described in RFC 1195. (RFC 1195 can be found at www.cis.ohio-state.edu/cgi-bin/rfc/rfc1195.html.)

Integrated IS-IS

As previously mentioned, IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, it distributes routing information for routing CLNP data for the ISO CLNS environment.

Integrated IS-IS is an implementation of the IS-IS protocol for routing multiple network protocols; it is an extended version of IS-IS for mixed ISO CLNS and IP environments. Integrated IS-IS tags CLNP routes with information regarding IP networks and subnets. It can be used purely for IP routing, purely for ISO routing, or for a combination of the two.

Integrated IS-IS provides an alternative to OSPF in an IP environment, mixing ISO CLNS and IP routing in one protocol.

NOTE

Like all modern IP routing protocols, Integrated IS-IS supports the following features:

- Variable-length subnet masks (VLSMs). The mask and the prefix are sent in the routing updates.
 - Redistribution of IP routes into and out of IS-IS.
 - Summarization of IP routes.
-

Integrated IS-IS Versus OSPF

Integrated IS-IS and OSPF are both link-state protocols with similarities as follows:

- Link-state representation, aging, and metrics
- Link-state databases and SPF algorithms
- Update, decision, and flooding processes

As discussed in Chapters 3 and 4 of *Building Scalable Cisco Networks*, OSPF is based on a central backbone (area 0), with all other areas ideally being physically attached to area 0. In OSPF, the border between areas is inside routers, the Area Border Routers (ABRs), and each link belongs to one area. This central backbone configuration means that certain design constraints will inevitably exist. When this type of hierarchical model is used, a good, consistent IP addressing structure is necessary to summarize addresses into the backbone and reduce the amount of information that is carried in the backbone and advertised across the network.

In comparison, IS-IS also has a hierarchy with Level 1 and Level 2 routers, but in IS-IS, the area borders lie on links rather than in routers. As shown in Figure S-2, each IS-IS router belongs to exactly one Level 2 area. Significantly fewer link-state packets (LSPs), also known as link-state protocol data units (PDUs), get used; thus, many more routers can reside in a single area. This capability makes IS-IS more scalable than OSPF. IS-IS allows a more flexible approach to extending the backbone by adding additional Level 2 routers. In IS-IS, this process is less complex than with OSPF.

Areas in OSPF and IS-IS

In OSPF, the area boundary is “inside” the router; interfaces belong to an area. In IS-IS, the “whole” router, not just an interface, belongs to an area.

With regard to CPU use and the processing of routing updates, IS-IS is more efficient than OSPF. In IS-IS, one LSP is sent per IS-IS router in each area (including redistributed prefixes [routes]), compared to the many OSPF LSAs that would be sent. Not only are there fewer LSPs to process, but the mechanism by which IS-IS installs and withdraws prefixes is less processor intensive.

Both OSPF and IS-IS are link-state protocols and thus provide fast convergence. The convergence time depends on a number of factors (timers, number of nodes, type of router, and so on).

Based on the default timers, IS-IS detects a failure quicker than OSPF and thus should converge more rapidly. Of course, if there are many neighbors and adjacencies to consider, the convergence time also might depend on the processing power of the router. IS-IS tends to be less CPU intensive than OSPF.

The timers in IS-IS allow for better fine-tuning than what is available for OSPF. There are other adjustable timers, so finer granularity can be achieved. By adjusting the timers, convergence time can be significantly decreased. However, this speed could come at the expense of stability. The network operator must understand the implications of changing the timers before making any adjustments.

Both IS-IS and OSPF are scalable, and the scalability of link-state protocols has been proven in the current live ISP backbones.

OSPF does have more features than IS-IS, including route tags, stub and not-so-stubby areas (NSSA), and OSPF on-demand circuits.

Level 1, Level 2, and Level 1–2 Routers

An IS-IS network is termed a domain; this is the equivalent of an autonomous system (AS) in OSPF. Within the domain, a two-level hierarchy exists:

- Level 1 ISs are responsible for routing to ESs inside an area. This is similar to OSPF internal nonbackbone routers in a totally stubby area.
- Level 2 ISs route between areas only. This is similar to internal backbone routers in OSPF.
- Level 1–2 ISs route between areas and the backbone. They participate in the Level 1 intra-area routing and the Level 2 interarea routing. This is the equivalent of Area Border Routers (ABRs) in OSPF.

Level 1 routers are also referred to as station routers because they enable stations (ESs) to communicate with each other and the rest of the network.

NOTE

End stations don't communicate by way of CLNP.

A contiguous group of Level 1 routers defines an area. The Level 1 routers maintain the Level 1 database, which defines the picture of the area itself and its exit points to neighboring areas.

Level 2 routers are also referred to as area routers because they interconnect the Level 1 areas. Level 2 routers store a separate database that contains only the interarea topology information.

Level 1–2 routers keep two separate link-state databases; this allows them to act as if they were two IS-IS routers, as follows:

- They support a Level 1 function to communicate with the other Level 1 routers in their area and maintain the Level 1 LSP information in a Level 1 topology database. They inform other Level 1 routers that they are an exit point from the area.
- They support a Level 2 function to communicate with the rest of the backbone and maintain a Level 2 topology database separately from their Level 1 database.

IS-IS does not share the concept of an area 0 with OSPF. Instead, an IS-IS backbone can appear as a set of distinct areas interconnected by a chain of Level 2 routers, weaving their way through and between the Level 1 areas. The IS-IS backbone consists of a set of Level 1–2 routers and Level 2 routers, and must be contiguous.

IS-IS uses a two-level hierarchy. The link-state information for these two levels is distributed separately, giving rise to Level 1 LSPs and Level 2 LSPs.

LSPs on point-to-point links are sent to a unicast address. LSPs on broadcast media (LANs) are sent to a multicast address.

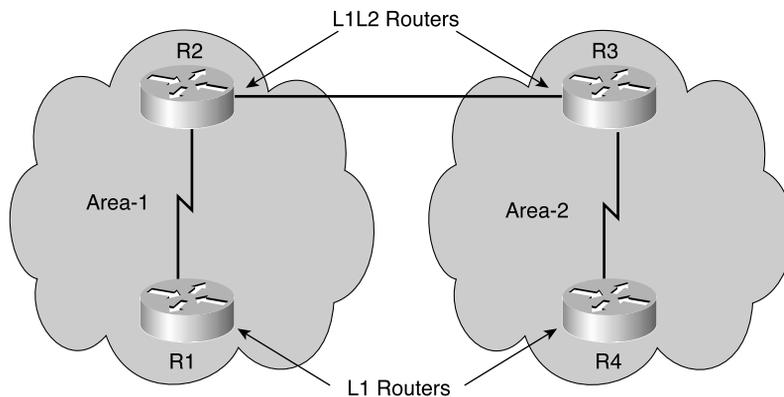
As with OSPF, one router on a LAN sends the LSP information on behalf of that LAN. In IS-IS, this router is called the designated intermediate system (DIS). It is a pseudonode, the representation of the LAN, and sends separate Level 1 and Level 2 LSPs on behalf of the network.

NOTE No backup DIS exists in IS-IS, in contrast to the backup DR for OSPF. If a DIS dies, a new election takes place.

Examples of IS-IS Hierarchical Routing

Figure S-3 shows the physical view of an example IS-IS area configuration. Physically, a Level 1–2 router connects to Level 1 routers inside its area and to Level 2 routers in the backbone. In the figure, R2 and R3 are Level 1–2 routers; R1 and R4 are Level 1 routers. R2 and R3 belong to their respective Level 1 areas and provide a physical connection between them.

Figure S-3 Example 1 IS-IS Area Configuration—Physical View



NOTE Recall that the boundary between areas in IS-IS exists on a link between routers and not on a router itself, as in OSPF.

Figure S-4 shows the logical view of the same example shown in Figure S-3. In Figure S-4, R2 and R3 are Level 1–2 routers; R1 and R4 are Level 1 routers. R2 and R3 are still Level 1 routers, but, in addition, they provide an entry point to the Level 2 backbone interconnecting both Level 1 areas.

Figure S-4 Example 1 IS-IS Area Configuration—Logical View

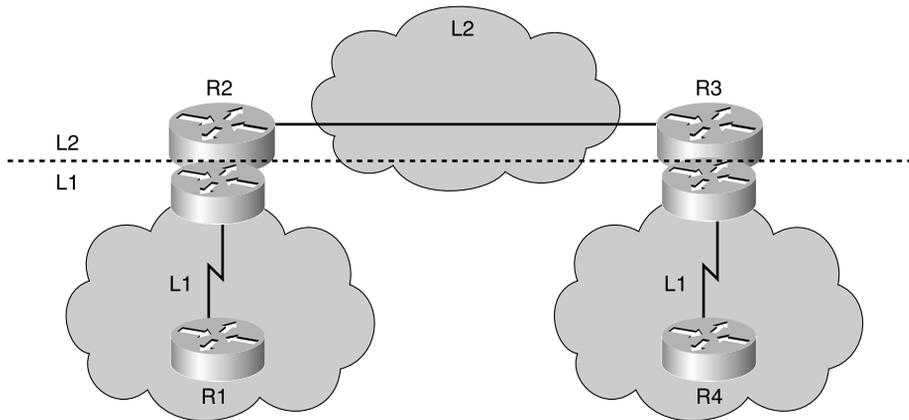
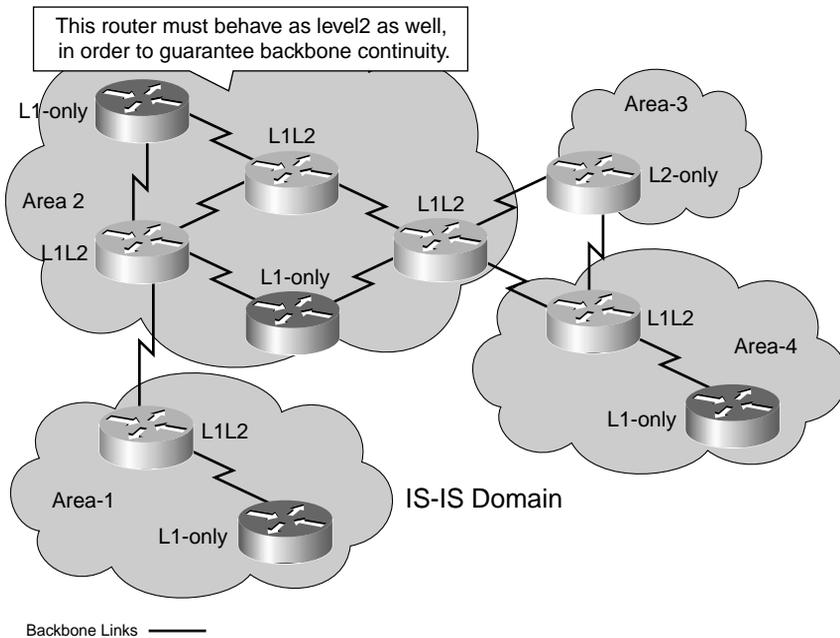


Figure S-5 shows another example.

Figure S-5 Example 2 IS-IS Level 2 and Level 1–2 Routers Forming a Level 2 Backbone



In Figure S-5, area 1 contains two routers:

- One router borders to area 2 and, therefore, is a Level 1–2 IS.
- The other router is contained totally within the area and, therefore, is Level 1 only.

Area 2 has many routers:

- Some routers are specified as Level 1 only and can route internally to that area only (and to the exit points).
- Level 1–2 routers form a chain across the area linking to the neighbor areas.
- Although the middle of these three Level 1–2 routers does not link directly to another area, it must support Level 2 routing so that the backbone is contiguous. If this middle router fails, the other Level 1–only routers (though providing a physical path across the area) could not perform the Level 2 function, and the backbone would be broken.

Area 3 contains one router that borders to area 3, but it has no intra-area neighbors and is, therefore, Level 2 only. If another router was added to area 3, the border router would revert to Level 1–2.

Figure S-5 also shows that the border between the areas in an IS-IS network exists on the *links between Level 2 routers* (in contrast to OSPF, where the border exists inside the Area Border Router [ABR] itself).

Operation of IS-IS for CLNS/CLNP

This section describes the operation of IS-IS. As mentioned earlier, IS-IS is used for routing CLNS/CLNP data, while Integrated IS-IS is used for routing IP. Integrated IS-IS is described in the next section.

OSI Addressing

OSI network layer addressing is done through NSAP addresses that identify any system in the OSI network. These OSI addresses are often simply referred as *NSAPs*.

NOTE

A variety of NSAP formats are used for different systems, and each OSI routing protocol uses a different representation of NSAP. NSAPs usually are represented in hexadecimal format with a variable length of up to 40 hex digits. NSAPs also are used in ATM. Later in this section, you will see the format interpreted by Cisco IOS Software.

The LSPs, hello PDUs, and other routing PDUs are OSI-format PDUs; therefore, every IS-IS router requires an OSI address even if it is routing only IP. IS-IS uses the OSI address in

the LSPs to identify the router, build the topology table, and build the underlying IS-IS routing tree.

The NSAPs contain the following:

- The OSI address of the device
- The link to the higher-layer process

NOTE

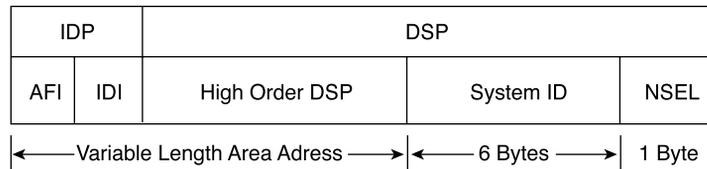
The NSAP address can be thought of as the equivalent of a combination of IP address and upper-layer protocol in an IP header.

Dijkstra's algorithm, which is used for both IS-IS and OSPF, requires a point-to-point connection between devices. OSPF uses the router ID to represent each router for this connection, while IS-IS uses the NET (a subset of the NSAP discussed in the following section).

NSAP Address Structure

Cisco routers can route CLNS data that uses addressing conforming to the ISO 10589 standard. The NSAP structure is illustrated in Figure S-6.

Figure S-6 *Structure of NSAP Addresses*



An OSI NSAP address can be up to 20 octets long and consists of the following parts, as shown in Figure S-6:

- The authority and format ID (AFI) specifies the format of the address and the authority that assigned that address. The AFI is 1 byte.
- The interdomain ID (IDI) identifies this domain. The IDI can be up to 10 bytes.
- The AFI and IDI together make up the interdomain part (IDP) of the NSAP address. This can loosely be equated to an IP classful major network.
- The high-order domain-specific part (DSP) (HODSP) is used for subdividing the domain into areas. This can be considered loosely as the OSI equivalent of a subnet in IP.

- The system ID identifies an individual OSI device (an ES or an IS). In OSI, a *device* has an address, just as it does in the DECnet protocol. This is different from IP, in which an *interface* has an address. OSI does not specify a fixed length for the system ID, but it does specify that it be consistent for all devices. Cisco IOS Software fixes the system ID as the 6 bytes preceding the 1-byte NSAP selector (NSEL).

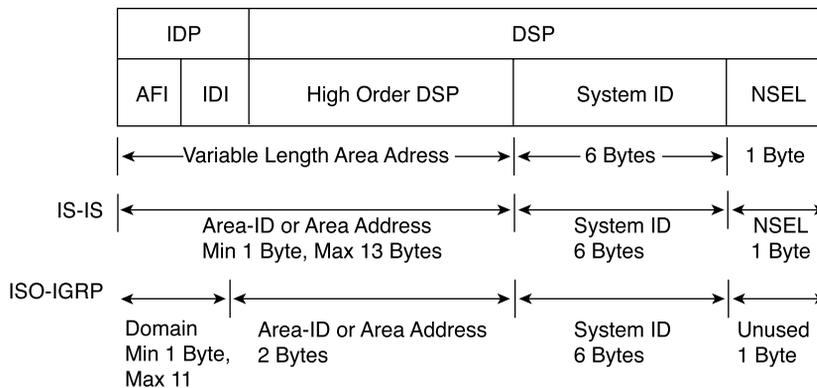
NOTE A Media Access Control (MAC) address is often used for the system ID.

- The NSEL (also known as the N-selector, the service identifier or the process ID) identifies a process on the device. It is a loose equivalent of a port or socket in IP. The NSEL is 1 byte long. It is not used in routing decisions. When the NSEL is set to 00, the address identifies the device itself—its network level address. In this case, the NSAP is known as a network entity title (NET).
- The HODSP, system ID, and NSEL together make up the domain-specific part (DSP) of the NSAP address.

IS-IS Versus ISO-IGRP NSAPs

IS-IS and ISO-IGRP interpret NSAPs differently, as shown in Figure S-7.

Figure S-7 How IS-IS and ISO-IGRP Interpret NSAP Addresses



IS-IS uses a two-layer architecture, joining the IDP and HODSP fields and treating them as its area address (Level 2), with the remaining system ID used for Level 1 routing. When used with IS-IS, therefore, the NSAP is divided into three parts, as shown in Figure S-7: 1 octet for NSEL, 6 octets for the system ID, and from 1 to 13 octets for the area address or area ID field. An NSAP has a variable length of 8 to 20 octets. It is usually longer than 8 bytes to permit some granularity in the allocation of areas.

NOTE Figure S-7 shows field length in bytes (or octets), where the NSAP address usually is represented in hexadecimal. NSAP addresses, like ATM addresses, are 160 bits long:

$$40 \text{ Hex Digits} \times \frac{4 \text{ bits}}{\text{Hex Digits}} = 20 \text{ octets} \times \frac{8 \text{ bits}}{\text{octets}} = 160 \text{ bits}$$

ISO-IGRP routes are based on a three-level architecture: domain (using the AFI and IDI fields for Level 3), area (using the HODSP field for Level 2), and system ID (for Level 1). What IS-IS treats simply as the area ID, ISO-IGRP splits into a *domain* and an *area*. ISO-IGRP sets the 2 bytes to the left of the system ID as the area ID or area address field, allowing for a theoretical 65,535 areas in an ISO-IGRP network. Everything else (a maximum of 11 bytes) is treated as a domain ID. Therefore, the minimum length for an ISO-IGRP NSAP is 10 bytes (1-byte NSEL, 6-byte system ID, 2-byte area, and minimum 1-byte domain).

ISO-IGRP sends routing information based on domain (variable length), area (length fixed by the protocol at 2 bytes), and finally system ID (fixed at 6 bytes). The NSEL is not used by ISO-IGRP.

Network Entity Title

As discussed earlier, if the NSEL field is 00, the NSAP refers to the device itself—that is, it is the equivalent of the Layer 3 OSI address of that device.

This address with the NSEL set to 00 is known as the NET. The NET is used by routers to identify themselves in the LSPs. Therefore, it forms the basis for the OSI routing calculation.

NET Is NSAP with NSEL = 00

A key point is that an NSAP address in which the NSEL is set to 00 is called the NET.

NETs and NSAPs are specified in all hexadecimal digits and must start and end on a byte boundary.

Official NSAP prefixes are required for CLNS routing. Addresses starting with the value of 49 (AFI = 49) are considered to be private addresses (analogous to RFC 1918 for IP addresses). These addresses are routed by IS-IS; however, this group of addresses should not be advertised to other CLNS networks.

Addresses starting with AFI values 39 and 47 represent the ISO data country code and ISO international code designator, respectively.

As shown in Figure S-7 and the examples in the next section, the Cisco IOS Software IS-IS routing process interprets the NSAP address as follows (from the right, or least-significant digit, end):

- The last byte is the NSEL field and must be specified as a single byte, with two hex digits, preceded by a period (.). In a NET, this N-selector field is set to (00).
- The preceding 6 bytes (this length is fixed by Cisco IOS Software) are the system ID. It is customary to use either a MAC address from the router or (for Integrated IS-IS) an IP address (for example, the IP address of a loopback interface) as part of the system ID.
- The IS-IS routing process of Cisco IOS Software treats the rest of the address as the area ID, or area address, as follows:
 - It is 1 to 13 bytes long.
 - Using a 1-byte field for area limits the scope for area definitions. Thus, the customary area ID consists of 3 bytes, with an AFI of 1 byte and 2 additional bytes for the area ID. For example, in the address 49.0001.0000.0c12.3456.00, the AFI is 49 and the additional 2 bytes are 0001, for an effective area ID of 49.0001.
 - Cisco IOS Software attempts to summarize the area ID as much as possible. For example, if an IS-IS network is organized with major areas subdivided into minor areas, and this is reflected in the area ID assignments, then:
 - Between the minor areas, Cisco IOS Software will route based on the whole area ID.
 - Between the major areas, Cisco IOS Software will summarize the area ID portion up to the major area boundary.

NSAP Examples

The following examples illustrate how an NSAP address is interpreted by IS-IS and ISO-IGRP.

The NSAP 49.0001.aaaa.bbbb.cccc.00 consists of the following:

- For IS-IS:
 - Area = 49.0001
 - System ID = aaaa.bbbb.cccc
 - N-selector = 00

- For ISO-IGRP:
 - Domain = 49
 - Area = 0001
 - System ID = aaaa.bbbb.cccc
 - N-selector = ignored by ISO-IGRP

The NSAP 39.0f01.0002.0000.0c00.1111.00 consists of the following:

- For IS-IS:
 - Area = 39.0f01.0002
 - System ID = 0000.0c00.1111
 - N-selector = 00
- For ISO-IGRP:
 - Domain = 39.0f01
 - Area = 0002
 - System ID = 0000.0c00.1111
 - N-selector = ignored by ISO-IGRP

Identifying Systems in IS-IS

In IS-IS, the area ID is associated with the IS-IS routing process; a router can be a member of only one Level 2 area. Recall from earlier in the chapter that a router can belong to only one area. The area ID or area address uniquely identifies the routing area, and the system ID identifies each node.

Restrictions on Areas and System IDs

Restrictions on areas and system IDs are as follows:

- All routers in an area must use the same area address. Indeed, the shared area address actually defines the area.
- An ES might be adjacent to a Level 1 router only if they both share a common area address. In other words, ESs recognize only ISs (and ESs on the same subnetwork) that share the same area address.
- Area routing (Level 1) is based on system IDs. Therefore, each device (ES and IS) must have a unique system ID within the area, and all system IDs must be the same length. Cisco mandates a 6-byte system ID.
- All Level 2 ISs come to know about all other ISs in the Level 2 backbone. Therefore, they, too, must have unique system IDs within the area.

System IDs

The system ID must be unique inside an area. As noted earlier, it is customary to use either a MAC address from the router or (particularly for Integrated IS-IS) to code the IP address (for example, of a loopback interface) into the system ID.

It is generally recommended that the system IDs remain unique across the domain; that way there can never be a conflict at Level 1 or Level 2 if a device is moved into a different area, for example.

All the system IDs in a domain must be of equal length. This is an OSI directive; Cisco enforces this by fixing the length of the system ID at 6 bytes in all cases.

Subnetwork Point of Attachment and Circuits

Two other terms used in IS-IS are *subnetwork point of attachment* (SNPA) and *circuit*.

An SNPA is the point at which subnetwork services are provided. This is similar to the Layer 2 address corresponding to the Layer 3 (NET or NSAP) address. The SNPA is usually taken from the following:

- The MAC address on a LAN interface.
- The virtual circuit ID for X.25 or ATM.
- The data-link connection identifier (DLCI) for Frame Relay.
- For High-Level Data Link Control (HDLC) interfaces, the SNPA is simply HDLC.

A link is the path between two neighbor ISs and is defined as being “up” when communication is possible between the two neighbors’ SNPAs.

A circuit is an interface; interfaces are uniquely identified by a Circuit ID. The router assigns a one octet Circuit ID to each interface on the router, as follows:

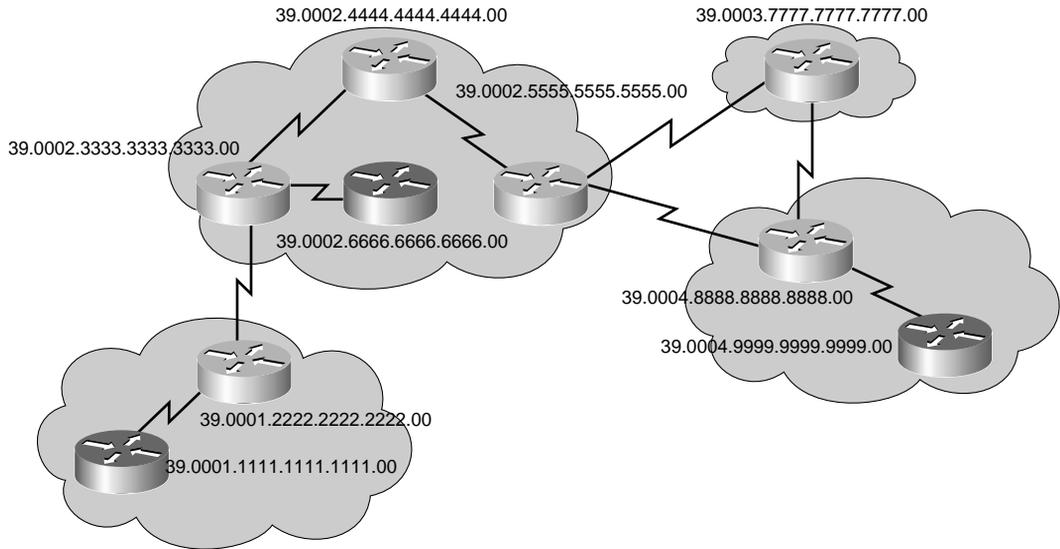
- In the case of point-to-point interfaces, this is the sole identifier for the circuit—for example, 03.
 - In the case of LAN interfaces (and other broadcast multi-access interfaces), this circuit ID is tagged to the end of the system ID of the DIS to form a 7-byte LAN ID. An example is 1921.6811.1001.03, where 03 is the circuit ID.
-

Example of NET Addresses in IS-IS Network

Figure S-8 shows examples of the NETs for routers in an IS-IS domain. Observe the following in Figure S-8:

- The 6-byte system IDs are unique across the network.
- The 3-byte area IDs are common to areas and distinct between areas.
- The 1-byte N-selectors are set to 00, indicating that these are NETs.

Figure S-8 NSAP Addresses in an IS-IS Network



IS-IS PDUs

The OSI stack defines a unit of data as a protocol data unit (PDU). A frame therefore is regarded by OSI as a data-link PDU, and a packet (or datagram, in the IP world) is regarded as a network PDU.

Three types of PDUs (with 802.2 Logical Link Control encapsulation) are shown in Figure S-9. As the figure shows, the IS-IS and ES-IS PDUs are encapsulated directly in a data-link PDU (there is no CLNP header and no IP header), while true CLNP (data) packets contain a full CLNP header between the data-link header and any higher-layer CLNS information.

Figure S-9 OSI Protocol Data Units

IS-IS:	Data-Link Header (OSI Family 0xFEFE)	IS-IS Header (First Byte Is 0x83)	IS-IS TLVs
ES-IS:	Data-Link Header (OSI Family 0xFEFE)	IS-IS Header (First Byte Is 0x82)	ES-IS TLVs
CLNP	Data-Link Header (OSI Family 0xFEFE)	IS-IS Header (First Byte Is 0x81)	CLNS

IS-IS and ES-IS PDUs contain multiple sets of variable-length fields, depending on the function of the PDU. Each field contains a type code, a length, and then the appropriate values—hence the abbreviation TLV, for Type, Length, Value fields—as shown in Figure S-9.

Four general types of packets exist, and each type can be Level 1 or Level 2:

- **LSP**—Used to distribute link-state information
- **Hello PDU (ESH, ISH, IS-IS Hello [IIH])**—Used to establish and maintain adjacencies
- **Partial sequence number PDU (PSNP)**—Used to acknowledge and request link-state information
- **Complete sequence number PDU (CSNP)**—Used to distribute a router’s complete link-state database

LSPs and hello PDUs are detailed in the following sections; the use of PSNP and CSNP is described in the section, “Link-State Database Synchronization.”

Link State Packets

This section describes the IS-IS LSPs.

Network Representation

In OSI, there are two main types of physical links:

- **Broadcast**—Multiaccess media types that support addresses referring to groups of attached systems and are typically LANs.
- **Nonbroadcast**—Media types that must address ESs individually and that are typically WAN links. These include point-to-point links, multipoint links, and dynamically established links.

Consequently, IS-IS supports only two media representations for its link states:

- *Broadcast* for LANs
- *Point-to-point* for all other media

NOTE

IS-IS has no concept of an NBMA network. It is recommended that point-to-point links (for example, subinterfaces) be used instead of NBMA networks such as native ATM, Frame Relay, or X.25.

LSP Contents

In IS-IS, a router describes itself with an LSP. The router's LSP contains the following:

- An LSP header, describing these items:
 - The PDU type and length
 - The LSP ID and sequence number
 - The remaining lifetime for this LSP (used to age out LSPs)
- Type Length Value (TLV) variable-length fields:
 - The router's neighbor ISs (used to build the map of the network)
 - The router's neighbor ESs
 - Authentication information (used to secure routing updates)
 - Attached IP subnets (optional for Integrated IS-IS)

The LSP sequence numbers enable receiving routers to ensure that they use only the latest LSPs in their route calculations, thus avoiding duplicate LSPs being entered in the topology tables.

When a router reloads, the sequence number is set initially to 1. The router might then receive its own old LSPs back from its neighbors (which has the last good sequence number before the router reloaded). It records this number and reissues its own LSPs with the next-highest sequence number.

The Remaining Lifetime field in the LSP is used by the LSP aging process to ensure that outdated and invalid LSPs are removed from the topology table after a suitable period. The LSP remaining lifetime counts down from 1200 seconds (20 minutes) to 0.

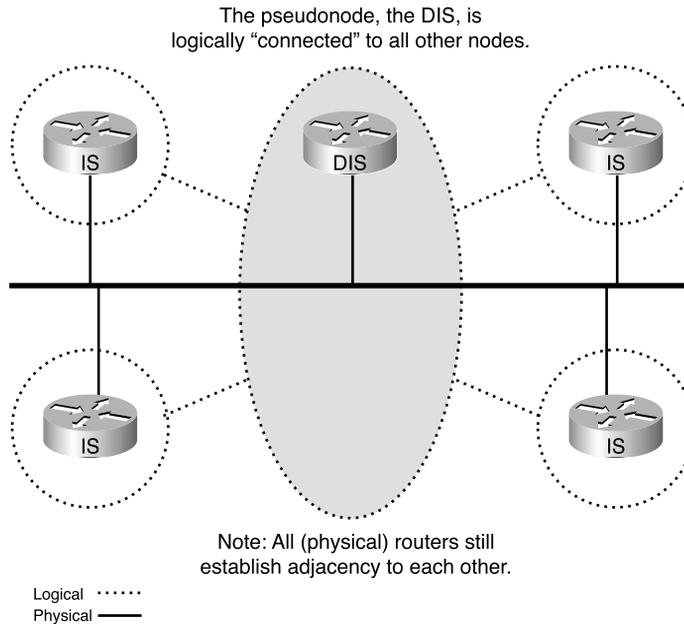
NOTE

IS-IS uses an LSP refresh interval of 15 minutes, as specified by ISO 10589. Each IS-IS router that originates an LSP is responsible for updating its entries using this timer. The Remaining Lifetime timer is how long an LSP is kept as valid in an IS-IS LSP database.

LAN Representation

Dijkstra's algorithm, used for IS-IS, requires a virtual router (pseudonode) for broadcast media to build a weighted directed graph of the shortest paths from a single source vertex to all other vertices.

For this reason, the DIS is elected to generate an LSP representing a virtual router connecting all attached routers to a star-shape topology. The DIS is shown in Figure S-10. The decision process for the election of the DIS is based first on the router with the highest configured priority and second on the router with the highest MAC address.

Figure S-10 IS-IS Designated Intermediate System Is Elected to Represent the LAN

In IS-IS, all routers on the LAN establish adjacencies with all other routers and with the DIS. Thus, if the DIS fails, another router can take over immediately with little or no impact on the topology of the network.

This is different from OSPF behavior. In OSPF, once the designated router (DR) and a backup DR (BDR) are elected, the other routers on the LAN establish adjacencies only with the DR and the BDR (the BDR is elected, and in the case of DR failure, is then promoted to be the DR).

LSP Variables

IS-IS LSPs include specific information about the router's attachments. This information is included in multiple TLV fields in the main body of the LSP:

- The links to neighbor routers (ISs), including the metrics of those interfaces
- The links to neighbor ESs

NOTE

If Integrated IS-IS is operational, the attached IP subnets are described as ESs, using a special TLV specified for IP information.

The metrics of IS-IS links are associated with the *outgoing* interface toward the neighbor IS (router). Up to four metrics can be specified, as follows:

- **Default metric (required): cost**—No automatic calculation of the metric for IS-IS takes place, compared to some routing protocols that calculate the link metric automatically based on bandwidth (OSPF) or bandwidth/delay (EIGRP). Using narrow metrics (the default), an interface cost is between 1 and 63 (a 6-bit metric value). All links use the metric of 10 by default. The total cost to a destination is the sum of the costs on all outgoing interfaces along a particular path from the source to the destination, and the least-cost paths are preferred.
- **Delay, expense, and error (optional)**—These metrics are intended for use in type of service (ToS) routing. These could be used to calculate alternative routes referring to the DTR (delay, throughput, and reliability) bits in the IP ToS field.

Extended Metric

In IS-IS, using the old-style narrow cost metric discussed previously, the total path metric is limited to 1023 (the sum of all link metrics along a path between the calculating router and any other node or prefix). This small metric value proved insufficient for large networks and provided too little granularity for new features such as traffic engineering and other applications, especially with high-bandwidth links.

Cisco IOS Software addresses this issue with the support of a 24-bit metric field, the so-called wide metric. Using the new metric style, link metrics now have a maximum value of 16,777,215 ($2^{24}-1$) with a total path metric of 4,294,967,295 ($2^{32}-1$).

Running different metric styles within one network poses a serious problem: Link-state protocols calculate loop-free routes because all routers (within one area) calculate their routing tables based on the same link-state database. This principle is violated if some routers look at old-style (narrow) and some at new-style (wider) TLVs. However, if the same interface cost is used for both the old- and new-style metrics, the SPF computes a loop-free topology.

Hello Messages

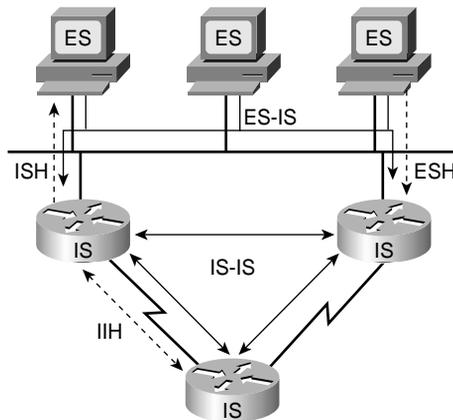
IS-IS uses hello PDUs to establish adjacencies with other routers (ISs) and ESs. Hello PDUs carry information about the system, its parameters and capabilities.

IS-IS has three types of hello PDUs, as follows:

- ESH, sent by an ES to an IS
- ISH, sent by an IS to an ES
- IIH, used between two ISs

The three types of hello PDUs are shown in Figure S-11.

Figure S-11 *Three Types of IS-IS Hello PDUs*



IS-IS Communication

ISs use IIHs to establish and maintain their neighbor relationships. When an adjacency is established, the ISs exchange link-state information using LSPs.

ISs also send out ISHs. ESs listen for these ISHs and randomly pick an IS (the one that sent the first ISH they hear) to forward all their packets to. Hence, OSI ESs require no configuration to forward packets to the rest of the network.

ISs (routers) listen to the ESHs and learn about all the ESs on a segment. ISs include this information in their LSPs.

For particular destinations, ISs might send redirect (RD) messages to ESs to provide them with an optimal route off the segment. This process is similar to IP Redirect.

Adjacencies

Separate adjacencies are established for Level 1 and Level 2. If two neighboring routers in the same area run both Level 1 and Level 2, they establish two adjacencies, one for each level. The Level 1 and Level 2 adjacencies are stored in separate Level 1 and Level 2 adjacency tables.

On LANs, the two adjacencies are established with specific Layer 1 and Layer 2 IIH PDUs. Routers on a LAN establish adjacencies with all other routers on the LAN and send LSPs to all routers on the LAN (unlike OSPF, in which routers establish adjacencies only with the designated router).

On point-to-point links, there is a common IIH format, part of which specifies whether the hello relates to Level 1, Level 2, or both.

By default, hello PDUs are sent every 10 seconds; the timeout to declare a neighbor down (that is, missing three hello packets) is 30 seconds. These timers are adjustable. (The commands to adjust the timers are not discussed in this book.)

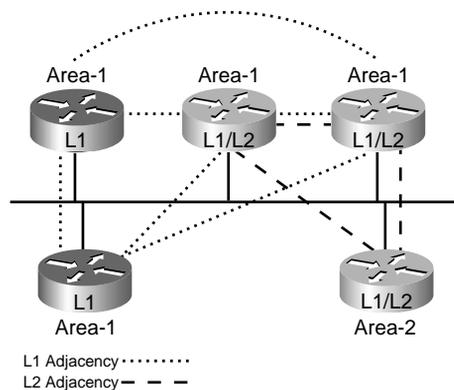
LAN Adjacencies

IIH PDUs announce the area ID. Separate IIH packets announce the Level 1 and Level 2 neighbors. Adjacencies are established based on the area address announced in the incoming IIHs and the type of the router.

For example, in Figure S-12, routers from two different areas are connected to the same LAN. On this LAN, the following is true:

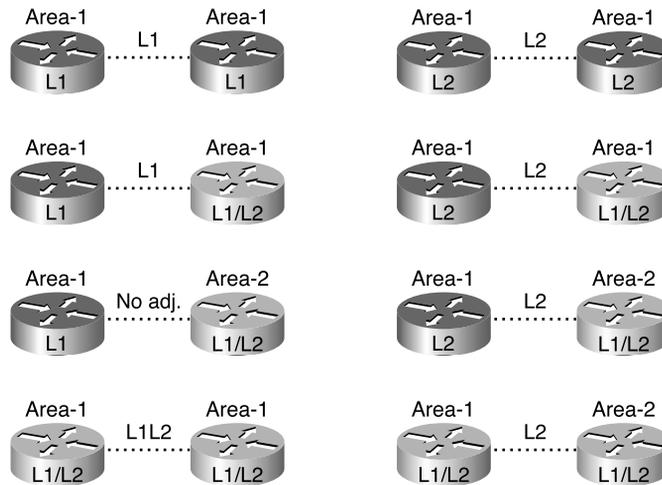
- The routers from one area accept Level 1 IIH PDUs only from their own area and, therefore, establish adjacencies only with their own area routers.
- The routers from a second area similarly accept Level 1 IIH PDUs only from their own area.
- The Level 2 routers (or the Level 2 process within any Level 1–2 router) accept only Level 2 IIH PDUs and establish only Level 2 adjacencies.

Figure S-12 IS-IS Adjacencies Are Based on Area Address and Type of Router



WAN Adjacencies

On point-to-point links (that is, on a WAN), the IIH PDUs are common to both Level 1 and Level 2 but announce the level type and the area ID in the hellos. This is illustrated in Figure S-13.

Figure S-13 *On a WAN, Area Address and Router Type Are Announced in a Common IIH PDU*

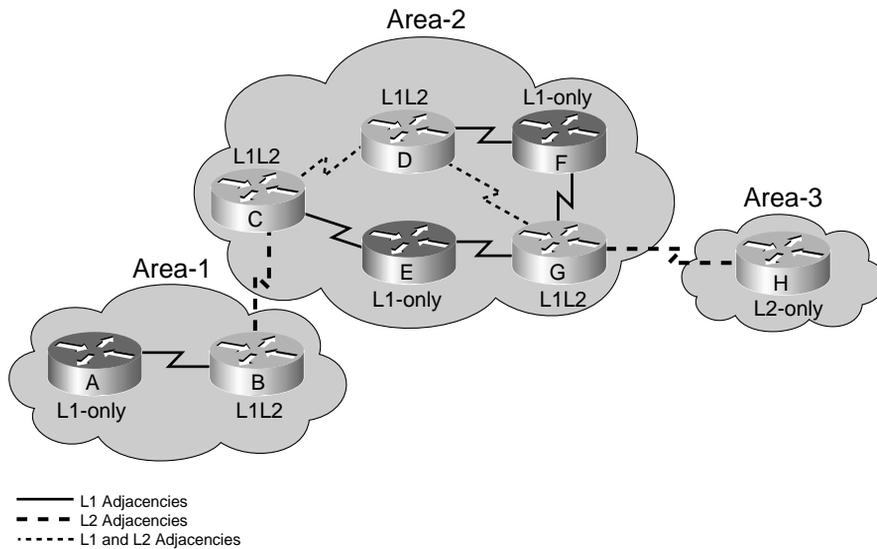
As shown in Figure S-13, the following is true:

- Level 1 routers in the same area (which includes links between Level 1–only and Level 1–2 routers) exchange IIH PDUs specifying Level 1 and establish a Level 1 adjacency.
- Level 2 routers (in the same area or between areas, and including links between Level 2–only and Level 1–2 routers) exchange IIH PDUs specifying Level 2 and establish a Level 2 adjacency.
- Two Level 1–2 routers in the same area establish both Level 1 and Level 2 adjacencies, and maintain these with a common IIH PDU specifying both the Level 1 and Level 2 information.
- Two Level 1–2 routers in different areas establish only a Level 2 adjacency.
- Two Level 1 routers that might be physically connected but are not in the same area (including a Level 1–only to a Level 1–2 router in a different Level 1 area) exchange Level 1 IIH PDUs but ignore these because the area IDs do not match. Therefore, they do not establish adjacency.

Level 2 Adjacencies

Figure S-14 shows examples of the following:

- Level 1–only routers establishing Level 1 adjacencies
- Level 2 routers establishing only Level 2 adjacencies (between areas)
- Level 1–2 routers establishing both Level 1 and Level 2 adjacencies with their Level 1–2 neighbors in the same area

Figure S-14 *Level 2 Adjacencies Must Be Continuous***NOTE**

In OSPF, there is a backbone *area*; here there is a backbone *path*. The path of connected Level 2 routers is called the backbone. All individual areas and the backbone must be contiguous. Level 2 adjacency exists independent of the area and must be contiguous. In the example in Figure S-14, the backbone is maintained by routers B, C, D, G, and H. A backbone is a set of contiguous Level 1–2 and Level 2 routers.

Link-State Database Synchronization

IS-IS link-state database synchronization is accomplished using special PDUs: PSNPs and CSNPs. These special PDUs bear the generic name of sequence number PDUs (SNPs),

SNPs (PSNPs and CSNPs) ensure that LSPs are sent reliably. SNPs contain LSP descriptors—not the actual, detailed LSP information, but headers describing the LSPs.

PSNPs usually contain only one LSP descriptor block. They are used as follows:

- To acknowledge receipt of an LSP
- To request a complete LSP for an entry missing in the originating router's topology database

CSNPs are a list of the LSPs held by a router.

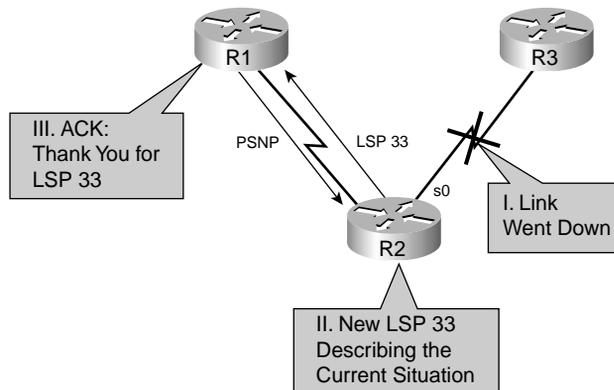
CSNPs are sent periodically on LANs. Receiving routers can compare the list of LSPs in the CSNP with their link-state database and request (with a PSNP) any missing LSPs.

CSNPs are sent on point-to-point links when the link comes active. In Cisco IOS Software, periodic CSNPs can be configured on point-to-point links.

Figure S-15 shows an example of link-state database synchronization on a point-to-point link. In this figure, the following is true:

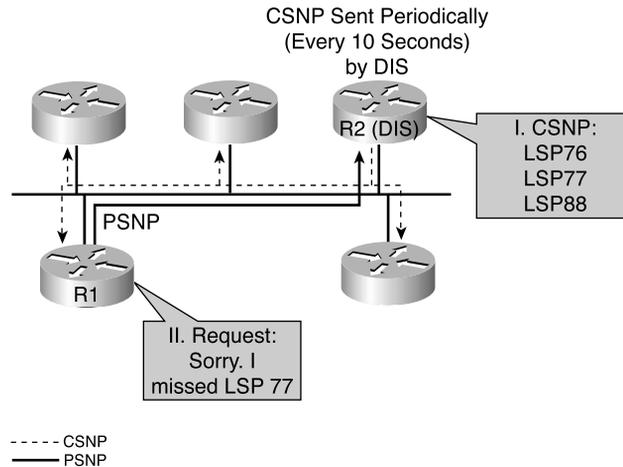
- A link fails.
- The R2 router notices this failure and issues a new LSP noting the change.
- The R1 router receives the LSP, stores it in its topology table, and sends a PSNP back to R2 to acknowledge receipt of the LSP.

Figure S-15 *Link-State Database Synchronization on a Point-to-Point Link*



On a LAN, the DIS periodically (every 10 seconds) sends CSNPs listing the LSPs that it holds in its link-state database. This is multicast to all IS-IS routers on the LAN.

Figure S-16 shows an example of link-state database synchronization on a LAN. In the example, the R2 router is the DIS. R2 sends a CSNP. The R1 router compares this list of LSPs with its topology table and realizes that it is missing one LSP. Therefore, it sends a PSNP to the DIS (R2) to request the missing LSP. The DIS reissues that LSP, and the R1 router acknowledges it with a PSNP (these last two steps are not shown in this figure but are similar to those in Figure S-15).

Figure S-16 *Link-State Database Synchronization on a LAN*

IP and OSI Routing with Integrated IS-IS

Integrated IS-IS supports three types of networks:

- OSI only (CLNS)
- IP only
- Dual (that is, both OSI and IP simultaneously)

The Integrated IS-IS LSPs can contain multiple variable-length TLV fields, with some containing OSI-specific state information and some containing IP-specific state information.

This section explains how IP routing is accomplished in an Integrated IS-IS network and tells why CLNS addresses must be configured, even if the network is running only IP.

IP Networks in Integrated IS-IS

Integrated IS-IS LSPs describe IP information in a similar manner to the way that IS-IS describes ESs. There are specific TLV types for IP information.

However, even if Integrated IS-IS is being used only for IP routing, OSI protocols are used to form the neighbor relationship between routers (routers still establish ES/IS adjacencies and use IS-IS hello packets). Therefore, a NET address is required for SPF calculations (using Dijkstra's algorithm) and Layer 2 forwarding. The NET address is used to identify the router in Integrated IS-IS.

This section describes how the OSI routing process works in IS-IS or Integrated IS-IS.

Building the OSI Forwarding Database

The following process is used to build the OSI forwarding database (the CLNS routing table) in a router running IS-IS or Integrated IS-IS:

- The link-state database is used to calculate the Shortest Path First (SPF) tree to OSI destinations (the NETs, or OSI addresses, of the routers). The link metrics are totaled along each path to decide which is the shortest path to any given destination.
- Separate link-state databases exist for Level 1 and Level 2 routes. Therefore, in Level 1–2 routers, SPF is run twice (once for each level) and separate SPF trees are created for each level.
- ES reachability is calculated with a partial route calculation (PRC) based on the previously mentioned Level 1 and Level 2 SPF trees. (There are no OSI ESs if the network is a pure IP Integrated IS-IS environment.)
- The best paths are inserted in the CLNS routing table (OSI forwarding database).

OSI Routing Process

Routing inside a Level 1 area is based on the system ID of the destination OSI (NSAP) address.

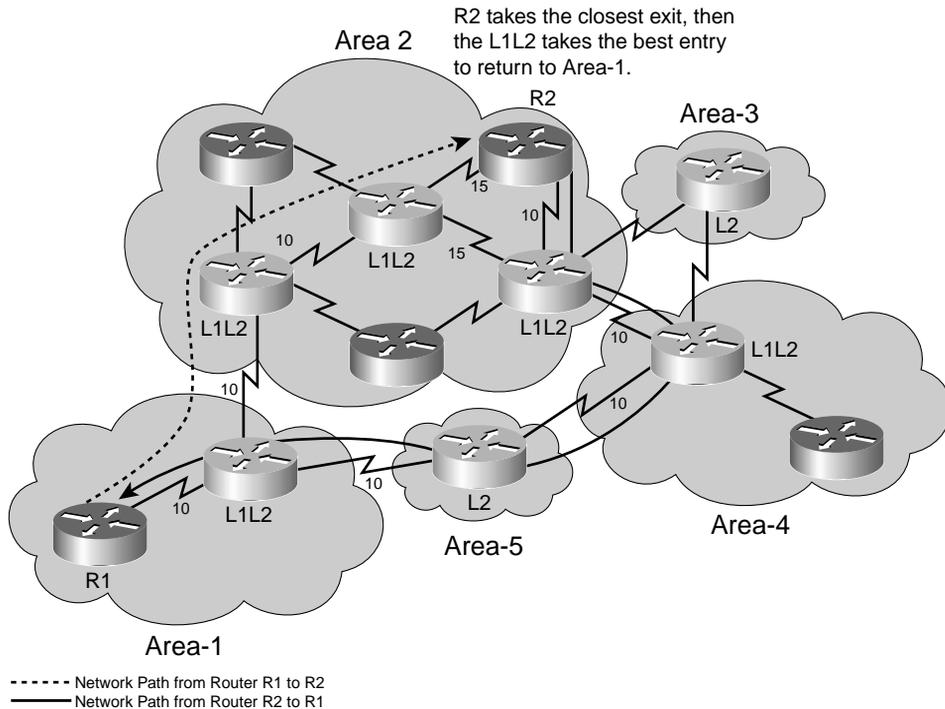
Level 1–2 routers send default routes to the Level 1 routers in their area. When a Level 1 router has a packet that is destined for another area, the following rules apply:

- The Level 1 router routes the packet to the nearest Level 1–2 router. The Level 1 router finds the closest exit point from the area, based on the best default route to the Level 1–2 routers in its area.
- The Level 1–2 router routes the packets into the Level 2 backbone, based on the destination area ID. The packet travels across the Level 2 backbone to the destination area.
- When it arrives in the destination area, Level 1 routing is again used to route the packet to its final destination inside that area.

The interface between the Level 1 world and the Level 2 world takes place on a Level 1–2 router. The Level 1–2 router behaves as if it were both a Level 1 router (routing to Level 1 destinations) and a Level 2 router (routing between areas).

Level 2 routing is based on the area ID. If a Level 1–2 router receives a packet (from a Level 2 neighbor) destined for its own area, it routes it as for Level 1, based on the system ID.

The fact that Level 1 routers see a default route only to the nearest Level 1–2 routers can lead to suboptimal routing, as shown in Figure S-17.

Figure S-17 Example of IS-IS Suboptimal Interarea Routing

In Figure S-17, router R1 routes packets to router R2 through area 1's Level 1–2 router. This router looks at the destination area of the packet and routes it directly into area 2. In area 2, the packets are routed as Level 1 to R2 (even though the initial next hop is another Level 1–2 router, the routing is Level 1.)

Packets returning from R2 to R1 are routed by R2 to its nearest Level 1–2 router. This router happens to see the best route to area 1 as being through area 4 and routes the return packets by a different route than the incoming packets traveled. As can be seen in Figure S-17, the path taken is not actually the least cost from R2 to R1.

Asymmetric routing (packets in different directions taking different paths) is not necessarily detrimental to the network.

Interconnecting IS-IS Domains

An IS-IS domain is a collection of IS-IS areas and is the equivalent of an IP AS.

IS-IS can support the interconnection of multiple domains. In a pure OSI (CLNS) environment, ISO-IGRP, a Cisco proprietary protocol, can be used. ISO-IGRP interprets

the initial domain identifier (IDI) portion of CLNS routes and allows routing between domains. Static CLNS routes can also be used. The standard OSI Interdomain Routing Protocol (IDRP) provides the same function; however, IDRP is not supported by Cisco IOS Software.

In an IP environment, an IP interdomain protocol is required. The most common of these is the Border Gateway Protocol (BGP).

Example of OSI (CLNS) Intra-Area and Interarea Routing

This section presents an example of OSI (CLNS) intra-area and interarea routing, using IS-IS or Integrated IS-IS.

CLNS Troubleshooting Commands Used in This Example

For this example, outputs of some troubleshooting commands on a Cisco router are used to illustrate how OSI routing works. The commands used in this example are detailed in this subsection.

NOTE

Some of the CLNS troubleshooting commands enable you to type a CLNS name instead of an NSAP address, and some commands display a name as well as the NSAP address. This mapping of name to address can be accomplished using the **clns host** global configuration command; this command works similar to the **ip host** global configuration command and is described later, in the section, “Other Integrated IS-IS Configuration Commands.”

Use the **show isis topology** [*nsap*] [**level-1**] [**level-2**] [**I1**] [**I2**] EXEC command to display a list of the paths to all connected routers. This command is described in Table S-1.

Table S-1 **show isis topology** Command Description

show isis topology Command	Description
<i>nsap</i>	Host name or NSAP of a router.
level-1	(Optional) The IS-IS link-state database for Level 1.
level-2	(Optional) The IS-IS link-state database for Level 2.
I1	(Optional) Abbreviation for the option level-1 .
I2	(Optional) Abbreviation for the option level-2 .

NOTE The **show isis topology** command is not documented in the Cisco IOS Software Command Reference manuals; however it is shown in the Cisco IOS Software Configuration Guides, and it does work on the routers.

Use the **show clns route** [*nsap*] EXEC command to display all the CLNS destinations to which this router knows how to route packets. The *nsap* optional parameter is the CLNS NSAP address or host name. The output of this command displays the IS-IS Level 2 routing table, as well as static and ISO-IGRP learned prefix routes. This routing table stores IS-IS area addresses and prefix routes; destinations are sorted by category.

Use the **show isis route** EXEC command to display the IS-IS Level 1 forwarding table for IS-IS-learned routes.

NOTE The **show isis route** command is documented as **show isis routes** in the Cisco IOS Software Command Reference manuals; however **show isis route** is the command that works on the routers.

This command is useful only if IS-IS is running in OSI mode. It does not display any routes if IS-IS is running only in IP mode; in this case, the command output indicates that you should use the **show isis topology** command instead.

Use the **which-route** {*nsap-address* | *clns-name*} EXEC command if you want to know which next-hop router will be used or if you have multiple processes running and want to troubleshoot your configuration. This command displays the routing table in which the specified CLNS destination is found. This command is described in Table S-2.

Table S-2 **which-route** Command Description

which-route Command	Description
<i>nsap-address</i>	CLNS destination network address
<i>clns-name</i>	Destination host name

Use the **show ip route** [*address* [*mask*] [**longer-prefixes**]] | [*protocol* [*process-id*]] command to display all or specific portions of the IP routing table. This command is described in Table S-3.

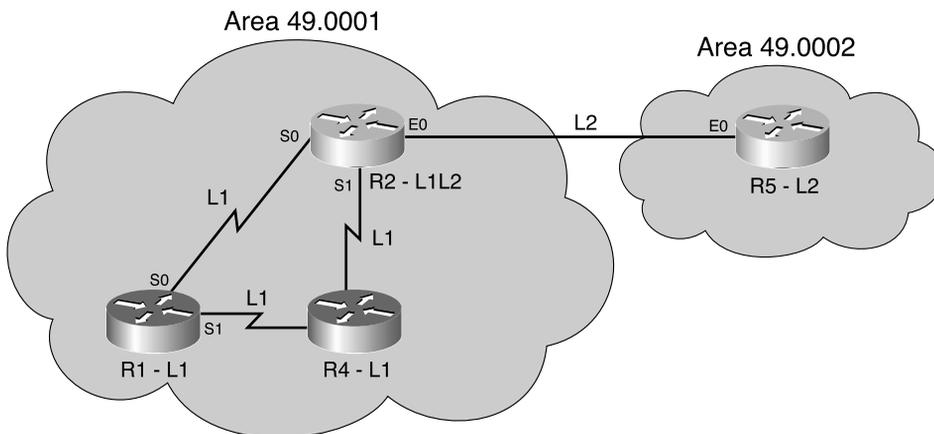
Table S-3 show ip route Command Description

show ip route Command	Description
<i>address</i>	(Optional) Address about which routing information should be displayed.
<i>mask</i>	(Optional) Argument for a subnet mask.
longer-prefixes	(Optional) The <i>address</i> and <i>mask</i> pair becomes a prefix, and any routes that match that prefix are displayed.
<i>protocol</i>	(Optional) Name of a routing protocol, or the keywords connected , static , or summary . If you specify a routing protocol, use one of the following keywords: bgp , egp , eigrp , hello , igrp , isis , ospf , or rip .
<i>process-id</i>	(Optional) Number used to identify a process of the specified protocol.

Example Network

Figure S-18 shows the network used in this example.

Figure S-18 Example Network for Intra-area and Inter-area Routing Discussion



Example Troubleshooting Command Output

The example output of the **show isis topology** command shown in Example S-1 is taken from router R1, a Level 1 router, in the example in Figure S-18.

Example S-1 **show isis topology** Command Output from Router R1 in Figure S-18

```
R1#show isis topology
IS-IS paths to level-1 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
R1	--			
R2	10	R2	Se0	*HDLC*
R4	10	R4	Se1	*HDLC*

The example output of the **show isis topology** command shown in Example S-2 is taken from router R2, a Level 1–2 router, in the example in Figure S-18.

Example S-2 **show isis topology** Command Output from Router R2 in Figure S-18

```
R2#show isis topology
IS-IS paths to level-1 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
R1	10	R1	Se0	*HDLC*
R2	--			
R4	10	R4	Se1	*HDLC*

```
IS-IS paths to level-2 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
R2	--			
R5	10	R5	Et0	0010.7bb5.9e20

The **show isis topology** command displays the least-cost paths to the destination NETs. The fields in the outputs in Example S-1 and Example S-2 are as follows:

- The system ID shows the NET of the destination. Cisco IOS Software uses dynamic host name mapping (as defined in RFC 2763) to map this to a host name, where that host name is available to the router. A router's host name is included in its outgoing LSP.
- The metric shows the sum of the metrics on the least-cost path to the destination.
- The next-hop router (an IS) is shown, along with the interface through which that next hop is reached and the SNPA of that next hop. (Recall that for a serial line, the SNPA is simply HDLC.)

Example S-2 shows that separate topology databases exist for Level 1 and Level 2 in router R2, a Level 1–2 router.

The **show clns route** command displays the CLNS destinations to which this router can route packets. The example output of the **show clns route** command shown in Example S-3 is taken from router R1, a Level 1 router, in the example in Figure S-18.

Example S-3 **show clns route** Command Output from Router R1 in Figure S-18

```
R1#show clns route

CLNS Prefix Routing Table
49.0001.0000.0000.0001.00, Local NET Entry
```

As Example S-3 shows, router R1 in Figure S-18 shows only its local NET entry because it is a Level 1-only router and, therefore, has no Level 2 area routes to display.

The **show isis route** command shows the Level 1 routes to IS-IS neighbors. The example output of the **show isis route** command shown in Example S-4 is taken from router R1 in the example in Figure S-18.

Example S-4 **show isis route** Command Output from Router R1 in Figure S-18

```
R1#show isis route

IS-IS Level-1 Routing Table - version 312
System Id Next-Hop Interface SNPA Metric State
R2 R2 Se0 *HDLC* 10 Up L2-IS
R4 R4 Se1 *HDLC* 10 Up
R1 --
Default route out of area - (via 2 L2-attached ISs)
System Id Next-Hop Interface SNPA Metric State
R2 R2 Se0 *HDLC* 10 Up
```

As Example S-4 shows, router R1 in Figure S-18 can see the other Level 1 routers in its area, routers R2 and R4. Level 1–2 routers appear in the Level 1 routing table (by virtue of their Level 1 connection) with a note at the end of their entry to show that they also act at Level 2. In this case, the Level 1–2 router is R2, as shown in Example S-4. The closest Level 1–2 router also appears as the default route out of the area. Again, this is router R2 in Example S-4.

In the **show isis route** command, as in the **show isis topology** command, the next-hop IS, the interface over which that next hop is reached, its SNPA, and the summed metric to that destination is shown for all IS routes. In Example S-4, the state of all the neighbors is up—this indicates that the hello process has established an adjacency.

For comparison, the example output of the **show clns route** command shown in Example S-5 is taken from router R2 in the example in Figure S-18. This output shows the local NET entry and also shows the Level 2 routes to its own areas and the neighbor areas.

Example S-5 **show clns route** Command Output from Router R2 in Figure S-18

```
R2#show clns route
CLNS Prefix Routing Table
49.0001.0000.0000.0002.00, Local NET Entry
49.0002 [110/10]
  via R5, IS-IS, Up, Ethernet0
49.0001 [110/0]
  via R2, IS-IS, Up
```

NOTE

As displayed in Example S-5, router R2, a Level 1–2 router, regards the route to its own area (area 49.0001) as being through itself, further emphasizing that the Level 1 and Level 2 processes in a Level 1–2 router operate separately.

The example output of the **show isis route** command shown in Example S-6 is taken from router R2 in Figure S-18. This output shows R2's routes to its IS-IS neighbors, routers R4 and R1.

Example S-6 **show isis route** Command Output from Router R2 in Figure S-18

```
R2#show isis route
IS-IS Level-1 Routing Table - version 47
System Id  Next-Hop  Interface  SNPA  Metric  State
R4         R4         Se1       *HDLC*  10     Up
R1         R1         Se0       *HDLC*  10     Up
```

An alternative method of finding the route to a destination NET or NSAP is to use the **which-route** command. The example output of the **which-route** command shown in Example S-7 is taken Router R1 in the example in Figure S-18.

Example S-7 **which-route** Command Output from Router R1 in Figure S-18

```
R1#which-route 49.0001.0000.0000.0002.00
Route look-up for destination 49.0001.0000.0000.0002.00
  Found route in IS-IS level-1 routing table

Adjacency entry used:
System Id  Interface  SNPA  State  Holdtime  Type  Protocol
0000.0000.0002 Se0       *HDLC*  Up     26      L1    IS-IS
Area Address(es): 49.0001
Uptime: 00:09:50
```

continues

Example S-7 *which-route* Command Output from Router R1 in Figure S-18 (Continued)

```

R1#which-route 49.0002.0000.0000.0005.00
Route look-up for destination 49.0002.0000.0000.0005.00
  Using route to closest IS-IS level-2 router

Adjacency entry used:
System Id  Interface  SNPA  State  Holdtime  Type  Protocol
0000.0000.0002  Se0    *HDLC*  Up     27        L1    IS-IS
Area Address(es): 49.0001
Uptime: 00:09:57

```

In Example S-7, the **which-route** command is entered on the Level 1–only router, R1. The first command requests the route to router R2 (the NSAP shown in the command is R2’s NSAP). The second command requests the route to router R5. The **which-route** command output states whether the destination is reachable by Level 1 or by the default exit point to Level 2, and it returns information on the next hop to the destination.

For comparison, the example output of the **which-route** commands shown in Example S-8 is taken from router R5 in the example in Figure S-18; R5 is a Level 2 router.

Example S-8 *which-route* Command Output from Router R5 in Figure S-18

```

R5#which-route 49.0001.0000.0000.0002.00
  Found route in CLNS L2 prefix routing table

Route entry used:
i 49.0001 [110/10] via R2, Ethernet0/0
Adjacency entry used:
System Id Interface  SNPA          State Hold. Type  Prot
R2        Et0/0    0000.0c92.e515  Up    24   L2   IS-IS
Area Address(es): 49.0001

R5#which-route 49.0001.0000.0000.0001.00
  Found route in CLNS L2 prefix routing table

Route entry used:
i 49.0001 [110/10] via R2, Ethernet0/0
Adjacency entry used:
System Id Interface  SNPA          State Hold. Type  Prot.
R2        Et0/0    0000.0c92.e515  Up    21   L2   IS-IS
Area Address(es): 49.0001

```

The first command in Example S-8 requests the route to R2. The second command requests the route to R1. As shown in this example output, the **which-route** command output on a Level 2 router states that the route was matched by an entry from the CLNS Level 2 routing table, and it returns information on the next hop to the destination.

Building an IP Forwarding Table

So far in this section, the process and outputs have referred to the OSI part of the Integrated IS-IS process. These are the same as for pure OSI IS-IS routing.

However, in the IP world, when running Integrated IS-IS, IP information is included in the LSPs. IP reachability behaves in IS-IS as if it were ES information.

It is important to note that IP information takes no part in the calculation of the SPF tree—it is simply information about leaf connections to the tree. Therefore, updating the IP reachability is only a PRC (similar to ES reachability).

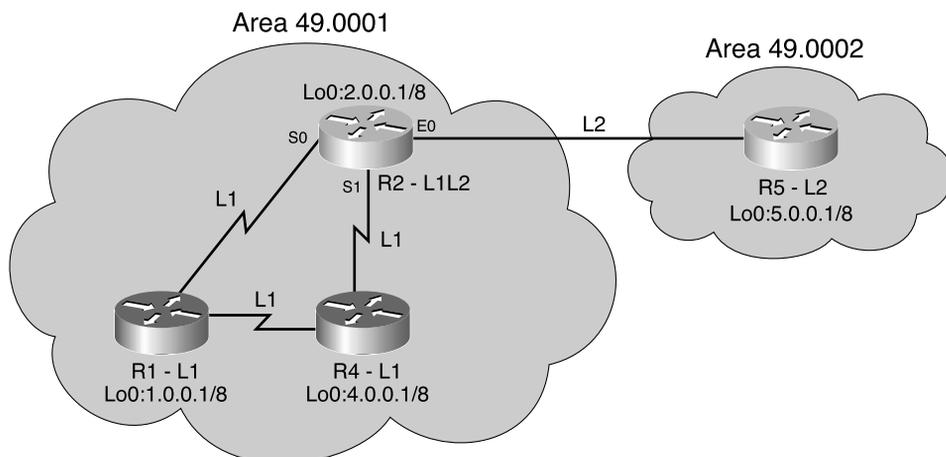
IP routes are generated by the PRC and offered to the IP routing table, where they are accepted based on routing table rules comparing, for example, administrative distance. When entered in the routing table, IP IS-IS routes are shown as being through Level 1 or Level 2, as appropriate.

The separation of IP reachability from the core IS-IS network architecture gives Integrated IS-IS better scalability than OSPF. OSPF sends LSAs for individual IP subnets. If an IP subnet fails, the LSA is flooded through the network and, in all circumstances, all routers must run a full SPF calculation.

In comparison, in an Integrated IS-IS network, the SPF tree is built from CLNS information. If an IP subnet fails in Integrated IS-IS, the LSP is flooded as for OSPF. However, if the failed subnet is a leaf IP subnet (that is, the loss of the subnet has not affected the underlying CLNS architecture), the SPF tree is unaffected and, therefore, only a PRC happens.

Figure S-19 illustrates an IP network running Integrated IS-IS; the IP addresses shown are the addresses of the loopback 0 interface on each router. Example S-9 shows the routing table on router R2 in this example network; the routes are IS-IS routes.

Figure S-19 *An IP Network Running Integrated IS-IS*



Example S-9 `show ip route` Command Output from Router R2 in Figure S-19

```
R2#show ip route
<output omitted>
i L1 1.0.0.0/8 [115/10] via 10.12.0.1, Ser0
i L1 4.0.0.0/8 [115/10] via 10.24.0.4, Ser1
i L2 5.0.0.0/8 [115/10] via 11.0.0.10, Eth0
```

In Example S-9, the output is interpreted as follows:

- The **i** indicates that the route was sourced from IS-IS.
- **L1** and **L2** show whether the IS-IS path to these destination IP networks is by way of IS-IS Level 1 or Level 2 routing, respectively.
- The next-hop IP addresses are the IP addresses of the corresponding next-hop IS-IS neighbor routers.

Basic Integrated IS-IS Router Configuration

This section covers the commands used to configure and troubleshoot Integrated IS-IS on a Cisco router.

Integrated IS-IS Configuration

This section identifies the steps and the basic commands used to configure Integrated IS-IS on a Cisco router.

Integrated IS-IS Configuration Steps

The following steps should be taken when configuring Integrated IS-IS:

- Step 1** Define areas, prepare an addressing plan for the routers (including defining the NETs), and determine interfaces that will run Integrated IS-IS.
- Step 2** Enable IS-IS as an IP routing protocol on the routers, and assign a tag to the process (if required).
- Step 3** Configure the NETs on the routers. This identifies the routers for IS-IS.
- Step 4** Enable Integrated IS-IS on the proper interfaces on the routers. Do not forget interfaces to stub IP networks, such as loopback interfaces (although there will not be any CLNS neighbors on these interfaces).

Basic Integrated IS-IS Configuration Commands

To enable Integrated IS-IS on a router for IP routing, you need only three commands, as described in this section. Note that there are many more commands to tune the IS-IS processes, but only three are required to start Integrated IS-IS.

The **router is-is** [*tag*] global configuration command enables Integrated IS-IS on the router. The optional tag can be used to identify multiple IS-IS processes by giving a meaningful name for a routing process. If it is not specified, a null tag (0) is assumed and the process is referenced with a null tag. This name must be unique among all IP router processes for a given router.

NOTE

When routing of CLNS packets is also required, use the **clns routing** global configuration command. (This command was on by default on the routers used for testing, but the Cisco IOS Software documentation says that it is not on by default.)

After the Integrated IS-IS process is enabled, the router must be identified for IS-IS by assigning a NET to the router with the **net** *network-entity-title* router configuration command. In this command, *network-entity-title* is the NET that specifies the area address and the system ID for the IS-IS routing process. This argument can be either an address or a name.

Finally, interfaces that are to use IS-IS to distribute their IP information (and additionally might be used to establish IS-IS adjacencies) must be configured using the **ip router isis** [*tag*] interface configuration command. If there is more than one IS-IS process on the router (as specified using the **router isis** command), interfaces must state which IS-IS process they belong to by specifying the appropriate tag.

NOTE

Configuring Integrated IS-IS to run on an interface is slightly different than configuring interfaces for most other IP routing protocols. In most other protocols, the interfaces are defined by **network** commands in the router configuration mode. There is no **network** command under the **router isis** command.

NOTE

When routing of CLNS packets is also required, use the **clns router isis** [*tag*] interface configuration command.

Other Integrated IS-IS Configuration Commands

By default, Cisco IOS Software enables both Level 1 and Level 2 operations on IS-IS routers. If a router is to operate only as an area router or only as a backbone router, this can be specified by entering the **isis-type** {**level-1** | **level-1-2** | **level-2-only**} router configuration command. This command is described in Table S-4. To specify that the router will act only as an area (or Level 1) router, use **level-1**. To specify that the router will act only as a backbone (or Level 2) router, use **level-2-only**.

Table S-4 **isis-type** Command Description

isis-type Command	Description
level-1	Router acts as a station router. This router will learn about destinations only inside its area. For interarea routing, it depends on the closest Level 1–2 router.
level-1-2	Router acts as both a station router and an area router. This router will run two instances of the routing algorithm. This is the default.
level-2-only	Router acts as an area router only. This router is part of the backbone and does not talk to Level 1–only routers in its own area.

Similarly, although the router might be a Level 1–2 router, it might be required to establish Level 1 adjacencies only over certain interfaces and Level 2 adjacencies over other interfaces. The **isis circuit-type** {**level-1** | **level-1-2** | **level-2-only**} interface configuration command can be used to specify either Level 1– or Level 2–only interfaces. This command is described in Table S-5. Because the default is Level 1–2, Cisco IOS Software attempts to establish both types of adjacency over the interface if this command is not specified.

Table S-5 **isis circuit-type** Command Description

isis circuit-type Command	Description
level-1	A Level 1 adjacency might be established if there is at least one area address in common between this system and its neighbors. Level 2 adjacencies will never be established over this interface.
level-1-2	A Level 1 and Level 2 adjacency is established if the neighbor is also configured as level-1-2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default.
level-2-only	Level 2 adjacencies are established if the other routers are Level 2 or Level 1–2 routers and their interfaces are configured for Level 1–2 or Level 2. Level 1 adjacencies will never be established over this interface.

Unlike some other IP protocols, IS-IS takes no account of line speed or bandwidth when setting its link metrics. All interfaces are assigned a metric of 10 by default. To change this value, you need to use the **isis metric default-metric {level-1 | level-2}** interface configuration command. The metric can have different values for Level 1 and Level 2 over the same interface. This command is described in Table S-6.

Table S-6 **isis metric** Command Description

isis metric Command	Description
<i>default-metric</i>	Specifies the metric assigned to the link and used to calculate the cost from each other router through the links in the network to other destinations. You can configure this metric for Level 1 or Level 2 routing. The range is from 0 to 63. The default value is 10.
level-1	Specifies that this metric should be used only in the SPF calculation for Level 1 (intra-area) routing.
level-2	Specifies that this metric should be used only in the SPF calculation for Level 2 (interarea) routing.

To define a name-to-NSAP mapping that can then be used with commands requiring NSAPs, use the **clns host name nsap** global configuration command. The assigned NSAP name is displayed, where applicable, in **show** and **debug EXEC** commands. This command is described in Table S-7.

Table S-7 **clns host** Command Description

clns host Command	Description
<i>name</i>	Desired name for the NSAP. The first character can be either a letter or a number, but if you use a number, the operations that you can perform are limited.
<i>nsap</i>	NSAP to which that the name maps.

Use the **summary-address address mask {level-1 | level-1-2 | level-2} prefix mask** router configuration command to create aggregate addresses for IS-IS or OSPF. The **no summary-address** command restores the default. This command is described in Table S-8.

Table S-8 **summary-address** Command Description

summary-address Command	Description
<i>address</i>	Summary address designated for a range of addresses.
<i>mask</i>	IP subnet mask used for the summary route.
level-1	Only routes redistributed into Level 1 are summarized with the configured address/mask value.
level-1-2	The summary route is applied both when redistributing routes into Level 1 and Level 2 IS-IS, and when Level 2 IS-IS-advertised Level 1 routes reachable in its area.

continues

Table S-8 *summary-address Command Description (Continued)*

summary-address Command	Description
level-2	Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address/mask value, and redistributed routes into Level 2 IS-IS are summarized also.
<i>prefix</i>	IP route prefix for the destination.
<i>mask</i>	IP subnet mask used for the summary route.

To configure the priority of designated routers, use the **isis priority value {level-1 | level-2}** interface configuration command. To reset the default priority, use the **no** form of this command. This command is described in Table S-9.

Table S-9 *isis priority Command Description*

Isis priority Command	Description
<i>value</i>	Sets the priority of a router and is a number from 0 to 127. The default value is 64.
level-1	Sets the priority for Level 1 independently.
level-2	Sets the priority for Level 2 independently.

Integrated IS-IS Configuration Examples

This section includes some example configurations of Integrated IS-IS. The first example shows the minimum commands required to run Integrated IS-IS, while the second example shows a two-area configuration.

Basic Integrated IS-IS Configuration Example

Example S-10 shows a simple Integrated IS-IS configuration, specifying only the IS-IS process and the NET, and enabling IS-IS on the interfaces. The router with this configuration acts as an IP-only Level 1–2 router.

Example S-10 *Basic Integrated IS-IS Configuration*

```

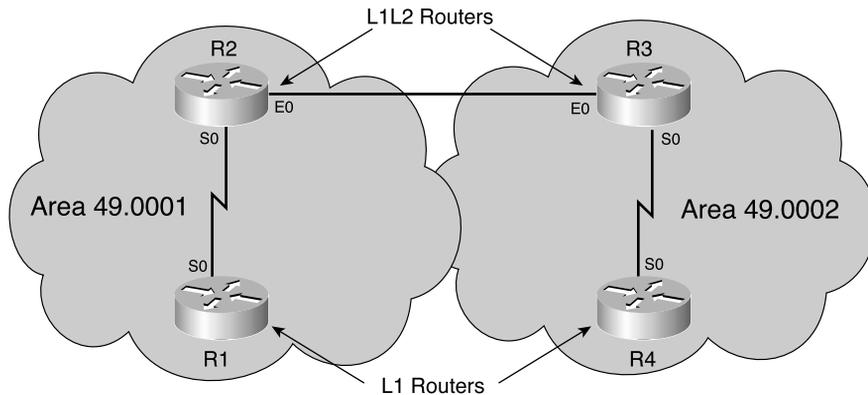
interface ethernet 0
 ip address 10.1.1.1 255.255.255.0
 ip router isis
!
interface serial 0
 ip address 10.1.2.1 255.255.255.0
 ip router isis
!
router isis
 net 01.0001.0000.0000.0002.00

```

Two-Area Integrated IS-IS Configuration Example

This example shows how to configure a simple two-area IS-IS network, optimizing the Level 1 and Level 2 operations of the links and routers. Figure S-20 shows the network used in this example.

Figure S-20 *Two-Area Integrated IS-IS Network*



In Figure S-20, router R1 is in area 49.0001 with no links outside that area and, therefore, needs to operate only as a Level 1 router. The configuration for R1 is shown in Example S-11.

Example S-11 *Configuration of Router R1 in Figure S-20*

```
hostname R1

!
interface Serial0
 ip address 192.168.120.1 255.255.255.0
 ip router isis
!
router isis
 net 49.0001.1921.6800.1005.00
 is-type level-1
```

The **is-type level-1** command under **router isis** in Example S-11 ensures that the router creates only a Level 1 database and takes part only in Level 1 adjacencies over its interfaces. Note that it is not necessary to also specify the **isis circuit-type** command on the interfaces in this case—Level 1 on the interfaces is implied by setting the IS-IS process as Level 1 only.

Router R2 in Figure S-20 is a member of area 49.0001 but also connects that area with the neighboring area 49.0002. Thus, R2 is required to act as both a Level 1 and a Level 2 router.

This is the default operation of the **is-type** command, so this command is not required in the router definition. The configuration of R2 is shown in Example S-12.

Example S-12 *Configuration of Router R2 in Figure S-20*

```
hostname R2!  
interface Ethernet0  
  ip address 192.168.220.2 255.255.255.0  
  ip router isis  
  isis circuit-type level-2-only  
!  
interface Serial0  
  ip address 192.168.120.2 255.255.255.0  
  ip router isis  
  isis circuit-type level-1  
!  
router isis  
  net 49.0001.1921.6800.1006.00
```

To optimize the operation of the interfaces to its two neighbors, router R2 in this example includes configuration to specify the type of adjacency to be established. Serial 0, the interface toward router R1 (in the same area), specifies **isis circuit-type level-1**, while the Ethernet 0 interface, the interface toward router R3 (in a different area, and therefore Level 2), has **isis circuit-type level-2-only** specified.

NOTE

The interfaces on both routers R1 and R2 could also have had the **isis metric** command configured, to reflect the different metrics on the serial and Ethernet interfaces.

Troubleshooting Integrated IS-IS

This section identifies commands that can be used to troubleshoot Integrated IS-IS on a Cisco router.

CLNS Troubleshooting Commands

Troubleshooting Integrated IS-IS, even in an IP-only network, requires some investigation of CLNS data. For example, the IS-IS neighbor relationships are established over OSI, not over IP, so showing IS-IS neighbors requires using the **show clns neighbors** command (as described later in this section). Indeed, two ends of a CLNS adjacency can actually have IP addresses on different subnets, with no impact to the operation of IS-IS (although IP next-hop resolution could be an issue).

Some of the CLNS troubleshooting commands are shown in the context of an earlier example, in the section, “CLNS Troubleshooting Commands Used in This Example.” In this section, more CLNS **show** commands are described.

The **show clns** EXEC command displays general information about the CLNS network.

The **show clns protocol** [*domain* | *area-tag*] EXEC command displays information for the specific IS-IS processes in the router. This command is described in Table S-10.

Table S-10 **show clns protocol** *Command Description*

show clns protocol Command	Description
<i>domain</i>	(Optional) Particular ISO IGRP routing domain
<i>area-tag</i>	(Optional) Particular IS-IS area

The **show clns interface** [*type number*] EXEC command displays CLNS-specific information about the interfaces running IS-IS. This command is described in Table S-11.

Table S-11 **show clns interface** *Command Description*

show clns interface Command	Description
<i>type</i>	(Optional) Interface type
<i>number</i>	(Optional) Interface number

The **show clns neighbors** [*type number*] [**detail**] EXEC command is very useful because it displays the neighbor ISs—that is, the routers with which this router has IS-IS adjacencies. (ES neighbors, if there are any, are also displayed.) This command is described in Table S-12.

Table S-12 **show clns neighbors** *Command Description*

show clns neighbors Command	Description
<i>type</i>	(Optional) Interface type.
<i>number</i>	(Optional) Interface number.
detail	(Optional) When specified, the area addresses advertised by the neighbor in the hello messages are displayed. Otherwise, a summary display is provided.

The optional keyword **detail** in the **show clns neighbors** command displays comprehensive information about the neighbors rather than listing a summary of the neighbors, as is the case without that keyword specified. The list of neighbors can be narrowed to those neighbors across a particular interface by specifying the interface *type* and *number* parameters in the command.

To display IS-IS–related information for IS-IS router adjacencies, use the **show clns is-neighbors** [*type number*] [**detail**] EXEC command. Neighbor entries are sorted according to the area in which they are located. This command is described in Table S-13.

Table S-13 **show clns is-neighbors** *Command Description*

show clns is-neighbors Command	Description
<i>type</i>	(Optional) Interface type.
<i>number</i>	(Optional) Interface number.
detail	(Optional) When specified, the areas associated with the intermediate systems are displayed. Otherwise, a summary display is provided.

CLNS and IS-IS Troubleshooting Commands

Further commands to troubleshoot the Integrated IS-IS network are described in this section.

The **show isis route** EXEC command displays the IS-IS Level 1 routing table (that is, routes to all other system IDs in the area). An example output from this command is shown earlier, in the section “Example of OSI (CLNS) Intra-Area and Interarea Routing.”

The **show clns route** [*nsap*] EXEC command displays the IS-IS Level 2 routing table (as well as static and ISO-IGRP–learned prefix routes). More details on this command and an example output are shown earlier, in the section, “Example of OSI (CLNS) Intra-Area and Interarea Routing.”

The **show isis database** [**level-1**] [**level-2**] [**l1**] [**l2**] [**detail**] [**lspid**] EXEC command displays the contents of the IS-IS link-state database. This command is described in Table S-14.

Table S-14 **show isis database** *Command Description*

show isis database Command	Description
level-1	(Optional) Displays the IS-IS link-state database for Level 1.
level-2	(Optional) Displays the IS-IS link-state database for Level 2.
l1	(Optional) Abbreviation for the option level-1 .
l2	(Optional) Abbreviation for the option level-2 .

Table S-14 show isis database Command Description (Continued)

show isis database Command	Description
detail	(Optional) When specified, the contents of each LSP are displayed. Otherwise, a summary display is provided.
lspid	(Optional) The link-state PDU identifier. When specified, the contents of a single LSP are displayed by its ID number.

To force IS-IS to refresh its link-state database and recalculate all routes, issue the **clear isis** *[tag | *]* command, specifying the IS-IS process tag or * to clear all IS-IS entries.

NOTE The **clear isis** command is not documented in the Cisco IOS Software documentation, but it does work on the routers.

To display how often and why the router has run a full SPF calculation, use the **show isis spf-log EXEC** command.

IP Troubleshooting Commands

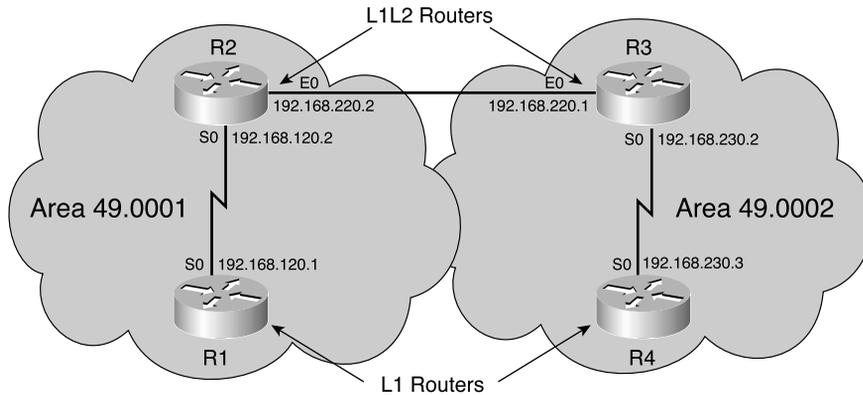
To troubleshoot the IP functionality of the Integrated IS-IS network, you can use standard IP **show** commands.

The **show ip protocols EXEC** command displays the active routing protocols, what interfaces they are active on, what networks they are routing for, and other parameters related to the routing protocols.

The **show ip route** *[address [mask] [longer-prefixes]] | [protocol [process-id]]* command displays the IP routing table. The detail for a particular route or a list of all routes in the routing table from a particular process can be specified. This command is described in detail earlier in the section, “Example of OSI (CLNS) Intra-Area and Interarea Routing.”

Integrated IS-IS Troubleshooting Command Output Examples

The network used to obtain these example outputs is the same as the one used for the earlier configuration examples shown in Figure S-20. The network is shown again in Figure S-21, with the IP addresses added.

Figure S-21 Two-Area Integrated IS-IS Network with IP Addressing

The system IDs of the routers in Figure S-21 are as follows:

- Router R1: 1921.6800.1005
- Router R2: 1921.6800.1006
- Router R3: 1921.6800.1007
- Router R4: 1921.6800.1008

Example S-13 shows the **show cns protocol** command output from router R2 in Figure S-21.

Example S-13 show cns protocol Command Output from Router R2 in Figure S-21

```
R2#show cns protocol
IS-IS Router: <Null Tag>
  System Id: 1921.6800.1006.00  IS-Type: level-1-2
  Manual area address(es):
    49.0001
  Routing for area address(es):
    49.0001
  Interfaces supported by IS-IS:
    Serial0 - IP
    Ethernet0 - IP
  Redistributing:
    static
  Distance: 110
```

Example S-13 shows the following:

- The Integrated IS-IS process tag (if present). Router R2 has a null tag.
- The system ID, level type(s), and area ID for this router. Router R2's system ID is 1921.6800.1006.00, it is a Level 1–2 router, and it is in area 49.0001. (Note that, in this command, the System ID includes both the System ID field and the NSEL field; in this case, 00.)

- The interfaces using Integrated IS-IS for routing (including whether that is for IP or CLNS or both). On router R2, interfaces Serial 0 and Ethernet 0 are using Integrated IS-IS for IP routing only.
- Any redistribution of other route sources. Router R2 is redistributing static routes, by default.
- The administrative distance for CLNS routes. This is similar to the administrative distances used for IP routing. For IS-IS in a CLNS environment, the default administrative distance is 110, as shown for router R2 in this example. (Use the **distance value [clns]** router configuration command to change this default.)

Example S-14 shows the **show clns neighbors** command output from routers R1 and R2.

Example S-14 show clns neighbors Command Output from Routers R1 and R2 in Figure S-21

R1#show clns neighbors						
System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6800.1006	Se0	*HDLC*	Up	28	L1	IS-IS
R2#show clns neighbors						
System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6800.1007	Et0	0010.7b81.d6ec	Up	24	L2	IS-IS
1921.6800.1005	Se0	*HDLC*	Up	21	L1	IS-IS

Example S-14 indicates the following:

- The IS-IS neighbors. Router R1 has one neighbor, router R2; Router R2 has two neighbors, routers R1 and R3.
- The neighbors SNPA and state.
- The hold time, the number of seconds before this adjacency entry times out. It indicates the remaining time to wait for receipt of a hello before the neighbor is declared *down*.
- The neighbors' level and type. Router R1 sees R2 as a Level 1 router. Router R2 sees R1 as a Level 1 router and R3 as a Level 2 router.

Example S-15 shows the **show clns interface** command output from router R2.

Example S-15 show clns interface Command Output from Router R2 in Figure S-21

```
R2#show clns interface s0
Serial0 is up, line protocol is up
Checksums enabled, MTU 1500, Encapsulation HDLC
ERPDU enabled, min. interval 10 msec.
RDPDU enabled, min. interval 100 msec., Addr Mask enabled
Congestion Experienced bit set at 4 packets
CLNS fast switching enabled
CLNS SSE switching disabled
DEC compatibility mode OFF for this interface
Next ESH/ISH in 21 seconds
```

continues

Example S-15 show cns interface *Command Output from Router R2 in Figure S-21 (Continued)*

```
Routing Protocol: IS-IS
Circuit Type: level-1
Interface number 0x1, local circuit ID 0x100
Level-1 Metric: 10, Priority: 64, Circuit ID: 1921.6800.1006.00
Number of active level-1 adjacencies: 1
Next IS-IS Hello in 7 seconds
```

Example S-15 indicates the following:

- That the Serial 0 interface is running IS-IS and is attempting to establish only Level 1 adjacencies
- The interface numbers and circuit ID for IS-IS purposes
- The metric(s) for the interface and a priority for DIS negotiation (not relevant in this case because it is a serial interface)
- Information about hello timers and the number of adjacencies that have been established

Example S-16 shows the **show ip protocols** command output from router R2.

Example S-16 show ip protocols *Command Output from Router R2 in Figure S-21*

```
R2#show ip protocols
Routing Protocol is "isis"
  Sending updates every 0 seconds
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: isis
  Address Summarization:
    None
  Routing for Networks:
    Ethernet0
    Serial0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.120.1   115          00:04:53
    192.168.220.1   115          00:04:58
  Distance: (default is 115)
```

Example S-16 shows that Integrated IS-IS is running on router R2. It also shows the interfaces taking part in Integrated IS-IS and the sources of routing information (the neighbor routers). The default IP administrative distance for Integrated IS-IS is 115.

Example S-17 shows the **show ip route isis** command output from routers R1 and R2. This command displays only the IS-IS routes from the IP routing table.

Example S-17 **show ip route isis** Command Output from Routers R1 and R2 in Figure S-21

```
R1#show ip route isis
i*L1 0.0.0.0/0 [115/10] via 192.168.120.2, Serial0
R2#show ip route isis
i L2 192.168.230.0/24 [115/20] via 192.168.220.1, Ethernet0
```

The route in Example S-17 on router R1 is from Level 1, as indicated by the **i L1** tag; this is a default route to router R2. The route in Example S-17 on router R2 is from Level 2, as indicated by the **i L2** tag. As for all IP routes, the administrative distance and the metric are shown in square brackets, as in **[115/20]**. Integrated IS-IS has an IP administrative distance of 115, by default. The metric shown for each route is the IS-IS cost to the destination.

NOTE

There are two different administrative distances. The CLNS administrative distance, shown in the **show clns protocol** command, defaults to 110; the IP administrative distance for IS-IS routes, shown in the **show ip route** command, defaults to 115.

Example S-18 shows the **show clns** command output from router R2.

Example S-18 **show clns** Command Output from Router R2 in Figure S-21

```
R2#show clns
Global CLNS Information:
 2 Interfaces Enabled for CLNS
NET: 49.0001.1921.6800.1006.00
Configuration Timer: 60, Default Holding Timer: 300, Packet Lifetime 64
ERPDU's requested on locally generated packets
Intermediate system operation enabled (forwarding allowed)
IS-IS level-1-2 Router:
  Routing for Area: 49.0001
```

Example S-18 shows that router R2 has two interfaces enabled for CLNS and that it is a Level 1–2 router in area 49.0001.

Example S-19 shows the **show clns is-neighbors** command output from routers R1 and R2.

Example S-19 **show clns is-neighbors** Command Output from Routers R1 and R2 in Figure S-21

```
R1#show clns is-neighbors
System Id      Interface  State  Type Priority  Circuit Id      Format
1921.6800.1006 Se0      Up     L1    0         00              Phase V
R2#show clns is-neighbors
System Id      Interface  State  Type Priority  Circuit Id      Format
1921.6800.1007 Et0      Up     L2    64        1921.6800.1006.0Phase V
1921.6800.1005 Se0      Up     L1    0         00              Phase V
```

As shown in Example S-19, router R1 has one IS neighbor, router R2. Router R2 has two IS neighbors, routers R1 and R3.

Example S-20 shows the **show isis database** command output from router R2.

Example S-20 show isis database Command Output from Router R2 in Figure S-21

```
R2#show isis database
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1921.6800.1005.00-00 0x00000004   0x485B        936            0/0/0
1921.6800.1006.00-00* 0x00000005   0x2E18        1155           1/0/0
1921.6800.1006.01-00* 0x00000001   0xFC74        462            0/0/0

IS-IS Level-2 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1921.6800.1006.00-00* 0x00000003   0x28FA        1180           0/0/0
1921.6800.1006.01-00* 0x00000002   0x7C36        1196           0/0/0
1921.6800.1007.00-00 0x00000003   0xF3BF        462            0/0/0
```

Example S-20 shows that router R2, a Level 1–2 router, has two separate databases, one for Level 1 and another for Level 2.

Example S-21 shows the **show isis spf-log** command output from router R2. This command displays how often and why the router has run a full SPF calculation.

Example S-21 show isis spf-log Command Output from Router R2 in Figure S-21

```
R2#show isis spf-log

Level 1 SPF log
When   Duration  Nodes  Count  Last trigger LSP  Triggers
00:17:52  0         1      4  1921.6800.1006.00-00  NEWAREA NEWADJ NEWLSP TLVCONTENT
00:17:47  4         2      1  1921.6800.1005.00-00  TLVCONTENT
00:12:24  4         3      2  1921.6800.1006.01-00  NEWLSP TLVCONTENT
00:12:13  4         3      2  1921.6800.1006.00-00  ATTACHFLAG LSPHEADER
00:04:32  4         3      1                               PERIODIC

Level 2 SPF log
When   Duration  Nodes  Count  Last trigger LSP  Triggers
00:17:53  0         1      1  1921.6800.1006.00-00  NEWLSP
00:12:24  4         2      3  1921.6800.1006.01-00  NEWADJ NEWLSP TLVCODE
00:12:19  4         3      1  1921.6800.1007.00-00  NEWLSP
00:04:33  8         3      1                               PERIODIC
```

As Example S-21 shows, router R2 keeps separate logs for the Level 1 and Level 2 SPF algorithm.

Modeling WAN Networks in Integrated IS-IS

This section discusses how to configure Integrated IS-IS over a WAN.

Categories of WANs for Integrated IS-IS

WANs typically are implemented as either point-to-point or point-to-multipoint, and most support multiple connections. These WANs typically do not support broadcasting, though, and thus are classified as NBMA.

Integrated IS-IS considers three categories of WANs, as follows:

- **Point-to-point leased circuits**—There are few issues for configuring Integrated IS-IS on these circuits.
- **Dialup connections**—Configuring IS-IS over dialup should be avoided if possible, except to use as a backup.
- **Switched WANs**—Various design options exist for NBMA networks.

These three categories are discussed in further detail in the following sections.

Point-to-Point

Point-to-point WANs are typically leased circuits between two routers. A point-to-point WAN has two devices attached, one at each end of the circuit. Usually such links will run Cisco HDLC or Point-to-Point Protocol (PPP). This corresponds exactly to the Integrated IS-IS classification of a point-to-point network.

NOTE

A point-to-point circuit still is regarded as an NBMA network, just as a back-to-back Ethernet connection is still a LAN. Both are examples of multiple-access networks that have only two devices attached.

On a point-to-point link a single IIH PDU is sent. This specifies whether the adjacency is at Level 1, Level 2 or both.

When the adjacency is established, each neighbor sends a CSNP describing the contents of its link-state database. Each router then requests any missing LSPs from the neighbor using PSNPs and acknowledges the receipt of the LSPs with PSNPs.

This activity reduces the amount of routing traffic across the point-to-point link—each router exchanges only the information missing from its link-state database rather than the entire link-state database of its neighbor router.

Dialup

Dialup networks using dial-on-demand routing (DDR) can be configured as either point-to-point or point-to-multipoint WAN implementations, as follows:

- Legacy DDR dialup connections (that is, using **dialer map** commands) are NBMA (even though they might use PPP as their line protocol) because a single interface can support multiple destinations.
- Dialer profiles and dialer virtual profiles are point-to-point connections (because one dialer profile represents one remote profile), but these can suffer from loss-of-neighbor delays similar to those for NBMA networks.
- Dialer virtual profiles are point-to-point connections in which the interface drops immediately if the remote end disconnects. This leads to faster neighbor-loss detection and faster convergence.

NOTE

Dialup interfaces are not dealt with further in this book. As a general rule, you should avoid using IS-IS over dialup, except to provide dial-backup functionality.

For further information on dialup networks and configuring them, see the Cisco Press Book *Building Cisco Remote Access Networks*.

Switched WANs

IS-IS can work over an NBMA multipoint network only if the network is configured with a full mesh. Anything less than a full mesh could cause serious connectivity and routing issues. However, even if a full mesh is configured, this is no guarantee that a full mesh will exist at all times. A failure in the underlying switched WAN network or a misconfiguration on one or more routers could break the full mesh either temporarily or permanently. Therefore, you should avoid NBMA multipoint configurations for IS-IS networks. Use point-to-point subinterfaces instead.

Point-to-point subinterfaces usually should be configured with their own IP subnet (typically with a 30-bit, or /30, subnet mask). In modern IP networks using private addressing or variable-length subnetting, usually plenty of spare IP addresses can be applied to point-to-point subinterfaces.

Alternatively, because Integrated IS-IS uses CLNS packets for its route propagation, the **ip unnumbered** interface configuration command can be used on point-to-point interfaces. However, this works only on more recent Cisco IOS Software releases (Version 12.0 and later); earlier releases fail to establish an IS-IS adjacency because the IP subnets do not match at either end of the link.

Configuring Integrated IS-IS in an NBMA WAN

This section describes how to configure Integrated IS-IS in an NBMA environment and shows some example configurations.

Configuration Steps

To enable Integrated IS-IS over switched WAN media, do the following:

Step 1 Start the Integrated IS-IS process and assign NETs as usual.

Step 2 On each NBMA, interface do the following:

- Design a mesh between the NBMA peers (full or partial).
- Configure point-to-point subinterfaces for each NBMA virtual circuit (VC), and assign IP addresses.
- Define the mapping of Level 3 protocols/addresses to the VC. If manual mappings are used (for example, **x25 map** or **frame-relay map** commands), the CLNS mapping must specify the keyword **broadcast** to support routing packets. However, the IP mapping does not require this (it is used only for next-hop resolution, not for routing packets).
- Start IS-IS processing on the subinterface using the **ip router isis** command. This command must *not* be used on the main interface of the router, or that (multipoint) interface will generate a pseudonode LSP for itself.

Further tuning can be done to control the flooding of link-state information using timer and blocking commands.

NOTE

You can reduce LSP flooding in two ways:

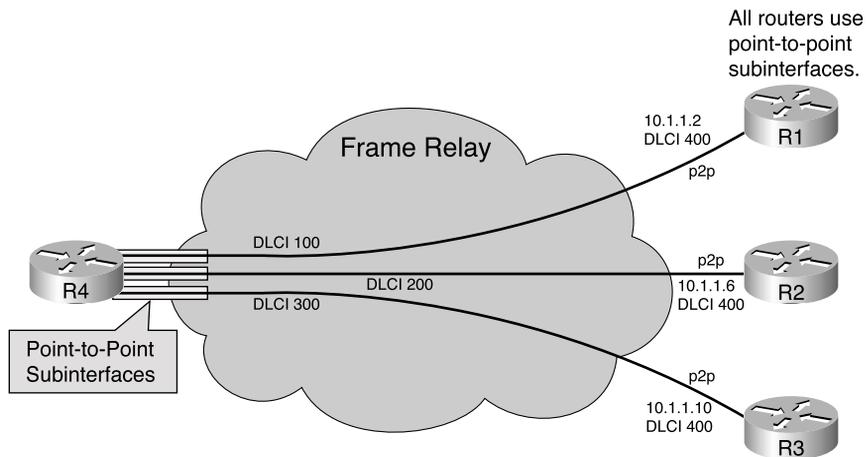
Blocking flooding on specific interfaces—The advantage of full blocking over mesh groups is that it is easier to configure and understand, and fewer LSPs are flooded. Blocking flooding on all links permits the best scaling performance but results in a less robust network structure. Permitting flooding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

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

Running IS-IS in NBMA using Point-to-Point Subinterfaces

Figure S-22 is an example of a router network connected over Frame Relay and using point-to-point subinterfaces. In this network, Frame Relay is modeled as a collection of subnets, with each Frame Relay permanent virtual circuit (PVC) treated as its own point-to-point network with its own subnet.

Figure S-22 *Frame Relay Network Using Point-to-Point Subinterfaces*



The network in Figure S-22 is an example of a star network topology. It is important to note that the routers at the points of the star (R1, R2, and R3) are also configured with point-to-point subinterfaces, even though (unlike the central router, R4) they connect to only one VC. This is the best practice for all routing protocols (because it allows further VCs to be added without affecting the existing VC), but it is imperative to IS-IS. A main interface (for example, interface Serial 0) is a multipoint interface, even if it happens to have only one VC configured. If the single VC were configured under a main interface, IS-IS would treat this as a broadcast network and attempt to elect a DIS. Also, the adjacency would not establish because the multipoint end would send broadcast, network-style hellos, but the central router would send point-to-point hello PDUs.

Example S-22 shows the configuration of the central router, router R4, in Figure S-22.

Example S-22 *Configuration of Router R4 in Figure S-22*

```

interface Serial0/0
  encapsulation frame-relay
!
interface Serial0/0.1 point-to-point
  ip address 10.1.1.1 255.255.255.252
  ip router isis
  frame-relay interface-dlci 100
!
interface Serial0/0.2 point-to-point
  ip address 10.1.1.5 255.255.255.252
  ip router isis
  frame-relay interface-dlci 200
!
interface Serial0/0.3 point-to-point
  ip address 10.1.1.9 255.255.255.252
  ip router isis
  frame-relay interface-dlci 300

```

Example S-22 shows the following:

- The encapsulation type (frame-relay) is set under the main interface (Serial0/0). No IP or IS-IS configuration is included under the main interface.
- Three subinterfaces are defined, one for each VC. Each subinterface specifies the following:
 - The IP address for that point-to-point link. There is a different subnet for each subinterface.
 - Integrated IS-IS as the routing protocol over that subinterface (this is done using the **ip router isis** command).
 - The VC to use for that point-to-point subinterface, using the **frame-relay interface-dlci** command. This is the only command needed to enable both IP and CLNS across this VC. The router automatically enables across this VC all the protocols that are enabled on the point-to-point subinterface, and it specifies broadcast for all those protocols.

Example S-23 shows the output of the **show frame-relay map** EXEC command for router R4 in the example network shown in Figure S-22.

Example S-23 *show frame-relay map Command Output from Router R4 in Figure S-22*

```

R4#show frame-relay map
Serial0/0.1 (up): point-to-point dlci, dlci 100(0x64,0x1840), broadcast
  status defined, active
Serial0/0.2 (up): point-to-point dlci, dlci 200(0xC8,0x3080), broadcast
  status defined, active
Serial0/0.3 (up): point-to-point dlci, dlci 300(0xA4,0x4580), broadcast
  status defined, active

```

As shown in Example S-23, the **show frame-relay map** command displays the status of each Frame Relay VC, providing information as follows:

- Its assigned subinterface—for example, Serial0/0.1.
- Its type—in this case, point-to-point, meaning that it has been assigned to a point-to-point subinterface.
- The VC identification—for example, dlci 100.
- Whether it supports broadcast (for example, routing) packets.
- Its status. A status of defined means it has been configured on the Frame-Relay switch, and a status of active indicates that this VC is operational. A status of inactive means that the router on the other end is not available. A status of deleted means that the DLCI was not advertised in the last LMI full status update.

Example S-24 shows the output of the **debug isis adj-packets EXEC** command on router R4 in the example network shown in Figure S-22.

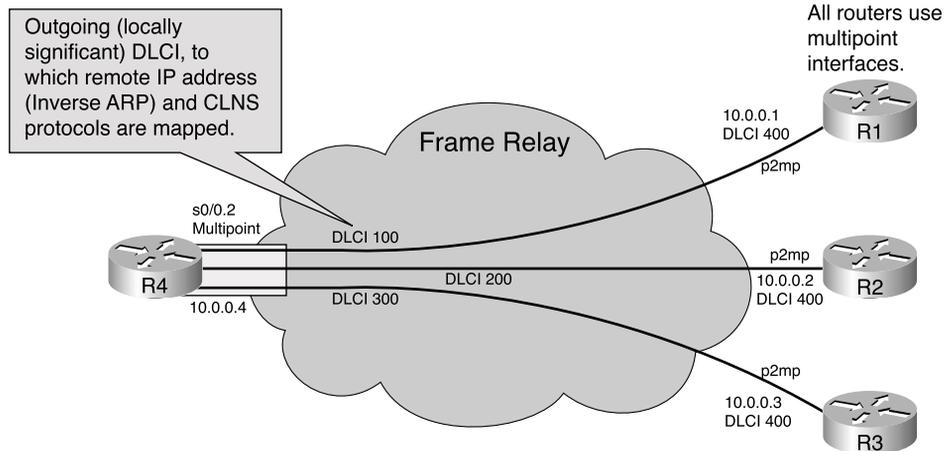
Example S-24 debug isis adj-packets *Command Output from Router R4 in Figure S-22*

```
R4#debug isis adj-packets
ISIS-Adj: Sending serial IIH on Serial0/0.1, length 1499
ISIS-Adj: Rec serial IIH from DLCI 100 (Serial0/0.1), cir type L1L2,
        cir id 00, length 1499
ISIS-Adj: rcvd state UP, old state UP, new state UP
ISIS-Adj: Action = ACCEPT
<output omitted>
ISIS-Adj: Sending serial IIH on Serial0/0.2, length 1499
ISIS-Adj: Rec serial IIH from DLCI 200 (Serial0/0.2), cir type L1L2, cir id 01,
        length 1499
ISIS-Adj: Sending serial IIH on Serial0/0.3, length 1499
ISIS-Adj: Rec serial IIH from DLCI 300 (Serial0/0.3), cir type L1L2, cir id 02,
        length 1499
```

Example S-24 shows the neighborhood establishment across the Serial0/0.1 subinterface, sending and receiving serial (that is, point-to-point) IIH PDUs and declaring the adjacency to be up. Ongoing hello conversations for the other subinterfaces also are shown.

Running IS-IS in NBMA Using Multipoint Interfaces

Figure S-23 shows an alternative NBMA version of the network shown in Figure S-22. In this case, all the Frame Relay ports are configured as multipoint interfaces, either as a multipoint subinterface (on the central router, router R4) or possibly as the main interfaces on the other routers. All interfaces share the same IP subnet.

Figure S-23 Frame Relay Network Using Multipoint Interfaces

In Figure S-23, the network is still a star topology, as in the previous example. In a multipoint environment it is important that a full mesh be implemented; therefore, all routers in this network must have VCs interconnecting them, although these are not shown in the figure.

If this were a true hub-and-spoke environment and the spoke sites had no need to communicate with each other, this topology could work with only the indicated DLCIs rather than having a full mesh. In this case, the central router (R4) must become the DIS for the NBMA network (because it is the only router visible to all others), so a suitable IS-IS priority should be set on the Frame Relay interface. Routes would be installed in each spoke router toward the other spoke routers through their local IP addresses. However, packets to these destinations would be dropped because there are no direct VCs between the spokes.

Example S-25 shows the configuration of the multipoint interface on the central router, R4.

Example S-25 Configuration of Router R4 in Figure S-23

```
interface Serial0/0
 encapsulation frame-relay
 !
interface Serial0/0.2 multipoint
 ip address 10.0.0.4 255.0.0.0
 ip router isis
 frame-relay map clns 100 broadcast
 frame-relay map clns 200 broadcast
 frame-relay map clns 300 broadcast
 frame-relay interface-dlci 100
 frame-relay interface-dlci 200
 frame-relay interface-dlci 300
 !
router isis
 net 00.0001.0000.0000.0004.00
```

As can be seen in Example S-25, IP and CLNS maps must be configured separately in a multipoint environment:

- The **frame-relay interface-dlci** command is used to enable IP across the Frame Relay PVCs. Inverse ARP resolves the remote end IP addresses. On a point-to-point subinterface, this command enables all traffic, but in a multipoint environment, this enables only IP.
- Alternatively, the IP maps could be entered explicitly using the **frame-relay map ip ip address dlci** command. In this case, the **broadcast** keyword might not be necessary for IP because only directed IP packets will use this VC. (Remember that the Integrated IS-IS routing protocol is not sending IP packets.)
- To enable CLNS, which must be done separately from IP in a multipoint environment, the **frame-relay map clns** command is used. CLNS is used for the IS-IS routing packets, so the **broadcast** keyword must be specified.

Examples S-26 and S-27 show how the same monitoring commands as used in the previous example produce a slightly different output for the multipoint environment.

Example S-26 is the **show frame-relay map** command output from router R4 in Figure S-23.

Example S-26 show frame-relay map Command Output from Router R4 in Figure S-23

```
R4#show frame-relay map
Serial0/0.4 (up): CLNS dlci 400(0x190,0x6400), static,
                  broadcast,
                  CISCO, status defined, active
Serial0/0.4 (up): ip 10.1.4.3 dlci 400(0x190,0x6400), dynamic,
                  broadcast,, status defined, active
```

In Example S-26, the **show frame-relay map** command again displays the status of each Frame Relay VC. This time however, separate entries are created for the IP and CLNS mappings (even though they use the same VC). The CLNS map shows that it is created as a *static* map and that *broadcast* was specified. The IP map is *dynamic* because the IP address was resolved by Inverse ARP.

Example S-27 is the **debug isis adj-packets** command output from router R4 in Figure S-23.

Example S-27 debug isis adj-packets Command Output from Router R4 in Figure S-23

```
R4#debug isis adj-packets
ISIS-Adj: Sending L2 LAN IIH on Serial0/0.2, length 1500
ISIS-Adj: Rec L2 IIH from DLCI 400 (Serial0/0.2), cir type L1L2,
          cir id 0000.0000.0004.03, length 1500
ISIS-Adj: Sending L1 LAN IIH on Serial0/0.2, length 1500
ISIS-Adj: Rec L1 IIH from DLCI 400 (Serial0/0.2), cir type L1L2,
          cir id 0000.0000.0004.03, length 1500
```

In Example S-27, the **debug isis adj-packets** command again shows the neighborhood establishment. This time, the adjacency uses LAN IIH PDUs because this is a multipoint environment.

Detecting Mismatched Interfaces

When interfaces are misconfigured in an NBMA environment, the network will not function as desired.

An example of a misconfiguration is when one end of a link is specified as a point-to-point subinterface but the other end is using a point-to-multipoint interface. Issuing the **show clns neighbors** command on each router indicates the mismatch. An example is shown in Example S-28. Here router R2 is specified as a point-to-point subinterface, while router R4 is using a point-to-multipoint interface configuration.

Example S-28 `show clns neighbors` Command Output from Misconfigured Routers

R2#show clns neighbors						
System Id	Interface	SNPA	State	Holdtime	Type	Protocol
0000.0000.0004	Se0/0.2	DLCI 300	Up	8	L1	IS-IS
R5	Et0/0	0050.3ef1.5960	Up	8	L2	IS-IS
R1	Se0/0.1	DLCI 100	Up	23	L1	IS-IS
R4#show clns neighbors						
System Id	Interface	SNPA	State	Holdtime	Type	Protocol
R6	Et0/0	0010.117e.74a8	Up	26	L2	IS-IS
R3	Se0/0.3	DLCI 400	Up	28	L2	IS-IS
0000.0000.0002	Se0/0.2	DLCI 300	Init	29	L1	IS-S
0000.0000.0001	Se0/0.1	DLCI 200	Up	290	IS	ES-IS

The shaded lines in Example S-28 indicate the routers' view of each other. Router R2 (the point-to-point end) shows the adjacency as up. However, the adjacency is stuck in the init state on router R4 (the multipoint end).

The misconfiguration results from the fact that the two ends of the VC are set to different types, which is an illegal configuration. The point-to-point end sends serial IIH PDUs. The multipoint end sends LAN IIH PDUs.

The ISO standard (ISO 10589) defines a three-way handshake for initiating LAN adjacencies, as follows:

- The adjacency starts in the down state. The IS sends out LAN IIH PDUs, identifying itself.

- If a LAN IIH PDU is received, the adjacency is installed in the init state. This router then sends out an IIH PDU to the neighbor, including the neighbor's SNPA in the hello packet. The neighbor does the same thing with this router's SNPA.
- The IS receives a second IIH from the neighbor router with its own SNPA identified in the packet. On receipt of this, the IS understands that the new neighbor knows of its presence and, therefore, declares the adjacency up.

According to the ISO standard, this process is omitted for a point-to-point adjacency. However, Cisco IOS Software implements the same three-way handshake by adding a Point-to-Point adjacency state TLV (TLV 240) in the serial hello PDUs. In a similar manner to the LAN adjacency, the router checks for its own SNPA in the neighbor's hello PDU before declaring the adjacency up.

The result of the mismatch configuration described in this example depends on the Cisco IOS Software release level. If using a release earlier than Release 12.1(1)T, the following happens (this is the case shown in Example S-28):

- Router R4 (configured for multipoint) receives the point-to-point hello from router R2 but treats it as a LAN hello and puts the adjacency in the init state. It looks for its own SNPA in the received hello PDUs (in a LAN hello PDU, this would be identified in TLV 6, IS Neighbors, but this TLV is not present in a serial hello) but never finds them. Therefore, the adjacency remains in the init state.
- Router R2 (configured for point-to-point) receives a LAN hello PDU and treats it as a point-to-point hello. It checks the hello for a TLV 240 (point-to-point adjacency state) and fails to find one. For backward compatibility, or perhaps to allow the link to be made to a non-Cisco IS-IS device, the router assumes that this is an ISO-specified point-to-point link, ignores the Cisco three-way handshake, and allows the adjacency to establish, setting it to up.

If using a release of Cisco IOS Software since Release 12.1(1)T, the following happens:

- Router R4 (configured for multipoint) receives the point-to-point hello, realizes that it is the wrong hello type, and installs the neighbor as an ES. Router R2 would show up in the **show clns neighbors** command output on router R4 with a protocol of ES-IS.
- Router R2 (configured for point-to-point) receives the LAN hello, recognizes the mismatch, and ignores the neighbor. Router R4 would not appear at all in router R2's **show clns neighbors** command output. Using the **debug isis adj-packets** command would show the incoming LAN IIH PDU and router R2 declaring the mismatch.

Summary

In this chapter, you learned the basics of the OSI protocols and the IS-IS and Integrated IS-IS routing protocols. You also learned how to configure and troubleshoot Integrated IS-IS for IP on a Cisco router.

An IS is a router. A domain is any portion of an OSI network that is under a common administrative authority. Within any OSI domain, one or more areas can be defined. An area is a logical entity; it is formed by a set of contiguous routers and the data links that connect them. All routers in the same area exchange information about all the hosts that they can reach. The areas are connected to form a backbone. All routers on the backbone know how to reach all areas.

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, it distributes routing information for routing CLNP data for the ISO CLNS environment.

Integrated IS-IS is an implementation of the IS-IS protocol for routing multiple network protocols; it is an extended version of IS-IS for mixed ISO CLNS and IP environments, or for IP only.

OSI network layer addressing is implemented with NSAP addresses that identify any system in the OSI network. If the NSEL field of the NSAP is 00, the NSAP refers to the device itself—that is, it is the equivalent of the Layer 3 OSI address of that device. This address with the NSEL set to 00 is known as the NET. The NET is used by routers to identify themselves in the LSPs and, therefore, forms the basis for the OSI routing calculation. (The NET is a similar concept to the router identifier used by OSPF.)

Every IS-IS router requires an OSI address even if it is routing only IP. IS-IS uses the OSI address in the LSPs to identify the router, build the topology table, and build the underlying IS-IS routing tree.

It is important to note that IP information takes no part in the calculation of the SPF tree—it is simply information about leaf connections to the tree.

Troubleshooting Integrated IS-IS, even in an IP-only world, requires some investigation of CLNS data. For example, the IS-IS neighbor relationships are established over OSI, not over IP.

Configuration Exercise: Configuring a Multiarea IS-IS Network

Configuration Exercises

In this book, configuration exercises are used to provide practice in configuring routers with the commands presented. If you have access to real hardware, you can try these exercises on your routers. Refer to Appendix H in the Cisco Press book *Building Scalable Cisco Networks* (and the online addendum to Appendix H) for a list of recommended equipment and configuration commands for the backbone routers. However, even if you don't have access to any routers, you can go through the exercises and keep a log of your own running configurations on separate sheets of paper.

Commands used and answers to the configuration exercises are provided at the end of the exercise.

In these exercises, you are in control of a pod of 3 routers; there are assumed to be 12 pods in the network. The pods are interconnected to a backbone. In most of the exercises, there is only one router in the backbone; in some cases, another router is added to the backbone. Each of the configuration exercises in this book assumes that you have completed the previous exercises on your pod.

NOTE

This configuration exercise assumes that you have completed the OSPF exercises in Chapter 4 of the Cisco Press book *Building Scalable Cisco Networks*.

The configuration for the backbone_r1 router for this exercise is provided on the Cisco Press website in the addendum to Appendix H.

In this exercise, you will configure IS-IS within your pod to the backbone_r1 router.

Objectives

In the following configuration exercise, you will do the following:

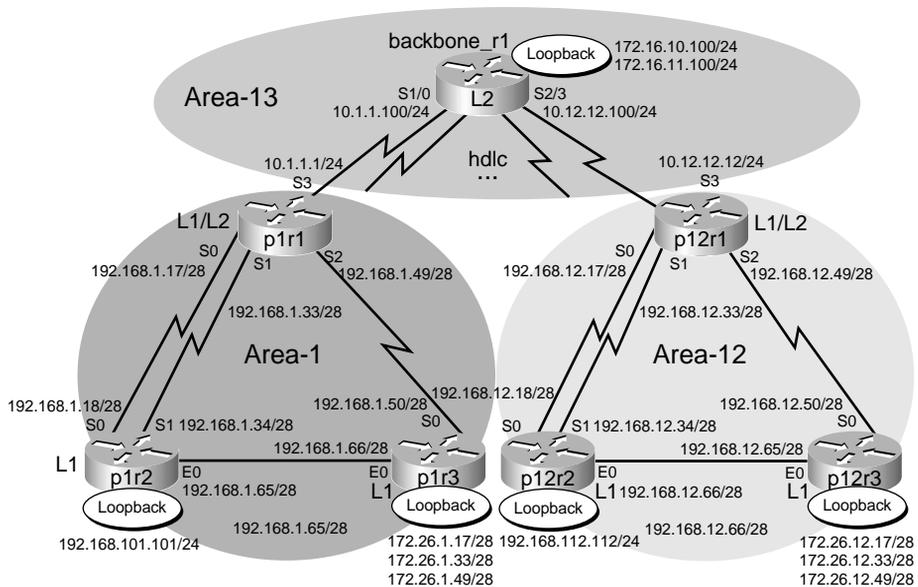
- Configure the pxr1, pxr2, and pxr3 routers in your pod to be in a specific IS-IS area.
- Enable IS-IS on the proper interfaces on the pxr1, pxr2, and pxr3 routers.
- Configure the pxr2 and pxr3 routers to be Level 1 routers.

- Configure the pxr1 router to perform route summarization.
- Verify proper Level 1 adjacencies among the pxr1, pxr2, and pxr3 routers within your pod.
- Verify proper Level 2 adjacencies between your pxr1 router and the backbone_r1 router.
- Verify that the pxr1 router is advertising a default route to the pxr2 and pxr3 routers.
- Verify which router is the designated IS router on the Ethernet link between pxr2 and pxr3.
- Verify and change the Ethernet interface IS-IS router priority.
- Verify proper network connectivity.
- Use **show** and **debug** commands to verify IS-IS operations.

Visual Objective

Figure S-24 illustrates the topology used in the network.

Figure S-24 Configuration Exercise Topology



Command List

In this exercise, you will use commands in Table S-15. Refer to this list if you need configuration command assistance during the exercise.

Table S-15 *Configuration Exercise Command List*

Command	Description
router isis	Enables IS-IS
net 00.000x.xyxy.xyxy.xyxy.00	Configures the NET
ip router isis	Enables IS-IS on the interface
is-type level-1	Configures the router as a Level 1 IS-IS router
summary-address 192.168.x.0 255.255.255.0	Configures route summarization
isis priority 63	Changes the IS-IS router priority on the Ethernet interface to influence the DR election
show clns interface s0	Displays CLNS interface status and configuration
show clns is-neighbors	Displays the IS-IS neighbor table
show isis database	Displays the IS-IS link-state database
show isis spf-log	Displays how often and why the router has run a full SPF calculation
debug isis adj-packets	Displays IS-IS adjacency-related packets
debug isis update-packets	Displays IS-IS update-related packets

Setup

To set up, do the following:

Step 1 Disable OSPF on all the routers within your pod.

Step 2 Shut the pxr1 S3 interface.

Task 1: Enabling IS-IS Within Your Pod

Complete the following steps:

Step 1 Enable IS-IS on the pxr1, pxr2, and pxr3 routers within your pod.

Step 2 Configure the pxr1, pxr2, and pxr3 routers within your pod with the IS-IS NETs shown in the following table.

Pod	pxr1 NET	pxr2 NET	pxr3 NET
1	00.0001.1111.1111.1111.00	00.0001.1212.1212.1212.00	00.0001.1313.1313.1313.00
2	00.0002.2121.2121.2121.00	00.0002.2222.2222.2222.00	00.0002.2323.2323.2323.00
3	00.0003.3131.3131.3131.00	00.0003.3232.3232.3232.00	00.0003.3333.3333.3333.00
4	00.0004.4141.4141.4141.00	00.0004.4242.4242.4242.00	00.0004.4343.4343.4343.00
5	00.0005.5151.5151.5151.00	00.0005.5252.5252.5252.00	00.0005.5353.5353.5353.00
6	00.0006.6161.6161.6161.00	00.0006.6262.6262.6262.00	00.0006.6363.6363.6363.00
7	00.0007.7171.7171.7171.00	00.0007.7272.7272.7272.00	00.0007.7373.7373.7373.00
8	00.0008.8181.8181.8181.00	00.0008.8282.8282.8282.00	00.0008.8383.8383.8383.00
9	00.0009.9191.9191.9191.00	00.0009.9292.9292.9292.00	00.0009.9393.9393.9393.00
10	00.0010.A1A1.A1A1.A1A1.00	00.0010.A2A2.A2A2.A2A2.00	00.0010.A3A3.A3A3.A3A3.00
11	00.0011.B1B1.B1B1.B1B1.00	00.0011.B2B2.B2B2.B2B2.00	00.0011.B3B3.B3B3.B3B3.00
12	00.0012.C1C1.C1C1.C1C1.00	00.0012.C2C2.C2C2.C2C2.00	00.0012.C3C3.C3C3.C3C3.00

NOTE The backbone_r1 router is configured with a NET of 00.0013.1111.2222.3333.00.

Step 3 Enable IS-IS on the interfaces specified in the following table on the pxr1, pxr2, and pxr3 routers within your pod.

Router	Interfaces on Which to Enable IS-IS				
pxr1	S0	S1	S2		
pxr2	S0	S1	E0	Loopback10	
pxr3	S0	E0	Loopback11	Loopback12	Loopback13

NOTE For this task, the pxr1 S3 interface is shut down.

Step 4 Display the routing table of the pxr1, pxr2, and pxr3 routers within your pod, and verify that you have full connectivity within your pod.

Step 5 Examine the pxr1 routing table and answer the following questions:

Does IS-IS load-balance by default?

What is the IS-IS routing metric based on by default?

Use the **show cns int s0** command to determine the Level 1 and Level 2 default metric.

What is the default administrative distance of IS-IS routes?

Step 6 Save the current configurations of all the routers within your pod to NVRAM.

Task 2: Enabling Connectivity to the backbone_r1 Router

Complete the following steps:

Step 1 No shut the S3 interface on the pxr1 router.

Step 2 Enable IS-IS on the S3 interface on the pxr1.

Step 3 Display the routing table of the pxr1, pxr2, and pxr3 routers within your pod, and verify that you have full connectivity to the backbone_r1 router.

Do you see the backbone_r1 router's loopback interfaces in your routing table? (Note that you also might see routes to the other pods.)

Does IS-IS perform autosummarization across the network boundary by default?

Does IS-IS use the same default administrative distance for Level 1 and Level 2 routes?

Make sure that you can ping the loopback interfaces on the backbone_r1 router from the pxr1, pxr2, and pxr3 routers within your pod.

Step 4 Save the current configurations of all the routers within your pod to NVRAM.

Task 3: Changing the IS-IS Router Type

Complete the following steps:

Step 1 By default, the pxr1, pxr2, and pxr3 routers within your pod are what type of IS-IS router (Level 1, Level 2, or Level 1–2)?

What command can be used to determine the IS-IS router type?

Step 2 Examine the IS-IS link state database of the pxr1, pxr2, and pxr3 routers within your pod. Do you see both the Level 1 and Level 2 link-state database on all the routers?

In the link-state database, which link-state packet (LSP) is identified with an asterisk?

Examine the routing table of the pxr2 and pxr3 routers within your pod. Do you see both the Level 1 and Level 2 routes?

Step 3 Configure the pxr2 and pxr3 routers within your pod to be Level 1 IS-IS routers.

Step 4 What are the advantages of configuring the pxr2 and pxr3 routers within your pod to be Level 1 IS-IS routers?

Step 5 Re-examine the link-state database of the pxr2 and pxr3 routers within your pod. Do the pxr2 and pxr3 routers contain only the Level 1 link-state database now?

Step 6 Re-examine the routing table of the pxr2 and pxr3 routers now. What are the differences after they became Level 1-only IS-IS routers?

Step 7 In the pxr2 and pxr3 routing tables, do you see a default route through the pxr1 router?

Step 8 From the pxr2 or pxr3 routers, can you still ping the backbone_r1 router's loopback interfaces?

Step 9 Save the current configurations of all the routers within your pod to NVRAM.

Task 4: Configure Route Summarization

Complete the following steps:

Step 1 From pxr3, telnet to the backbone_r2 router using the password cisco.

Step 2 Display the routing table of the backbone_r1 router. Do you see your 192.168.x.y/28 and 172.26.x.y/28 subnets?

Step 3 Exit the telnet to the backbone_r1 router.

Step 4 Configure the Level 1–2 router within your pod to perform route summarization. Summarize the 192.168.x.y/28 and 172.26.x.y/28 subnets within your pod to 192.168.x.0/24 and 172.26.x.0/24 (where *x* is your pod number).

Step 5 Telnet to the backbone_r1 router.

- Step 6** Re-examine the routing table of the backbone_r1 router. Do you see your summarized routes?
- Step 7** From the backbone_r1 router, can you ping all the interfaces within your pod?
- Step 8** Exit the telnet to the backbone_r1 router.
- Step 9** Save the current configurations of all the routers within your pod to NVRAM.

Task 5: Using IS-IS show and debug Commands

Complete the following steps:

- Step 1** At the pxr2 router, enable the **debug isis adj-packets** command and observe the debug output.

What type of hello is pxr2 sending to pxr3 over the Ethernet interface?
- Step 2** Shut the pxr2 E0 interface and observe the debug output.
- Step 3** No shut the pxr2 E0 interface and observe the debug output. From the debug output, determine which router (pxr2 or pxr3) is elected as the DR.

NOTE

The correct term in IS-IS is DIS, not DR. However, the debug output on the router uses the term DR, so we will also use the term DR in this task.

In the debug output, you should see something similar to the output shown in Example S-29 if pxr3 is the DR.

Example S-29 debug isis adj-packets Command Output from Router pxr2 If Router pxr3 Is the DR

```
02:23:12: ISIS-Adj: Adjacency state goes to Up
02:23:12: ISIS-Adj: Run level 1 DR election for Ethernet0
02:23:12: ISIS-Adj: New level 1 DR 3333.3333.3333 on Ethernet0
pxr2#
```

Or, in the debug output, you should see something similar to the output shown in Example S-30 if pxr2 is the DR.

Example S-30 debug isis adj-packets Command Output from Router pxr2 If Router pxr2 Is the DR

```
05:19:43: ISIS-Adj: Run level 1 DR election for Ethernet0
05:19:43: ISIS-Adj: No change (it's us)
pxr2#
```

What determines who will be elected as the DR on the Ethernet?

- Step 4** Disable **debug isis adj-packets** at the pxr2 router.
- Step 5** At the DR (as determined from Step 3), enter the **show clns int e0** command. What is the default IS-IS router priority set to?
- Step 6** At the DR (as determined from Step 3), change the E0 interface IS-IS router priority to 63.
- Step 7** At the DR (as determined from Step 3), enter the **show clns int e0** command again and verify that the DR is now changed to the other router.

The Circuit ID field in the **show clns int e0** output identifies which router is the DR. Example S-31 shows the output of this command on the pxr3 router. In this example, the DR is the pxr2 router and the pxr3 router has an IS-IS priority of 63.

Example S-31 *show clns Command Output from Router pxr3 If Router pxr2 Is the DR*

```
pxr3#show clns int e0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
  Congestion Experienced bit set at 4 packets
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 21 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x1, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 63, Circuit ID: 2222.2222.2222.01
    Number of active level-1 adjacencies: 1
  Next IS-IS LAN Level-1 HELLO in 8 seconds
```

- Step 8** At the pxr2 router, enable the **debug isis update-packets** command. Shut and no shut the E0 interface on the pxr2 router and observe the debug output.
- Step 9** Disable **debug isis update-packets** at the pxr2 router.
- Step 10** At the pxr2 router, use the proper **show** command to verify how often and why the router has run a full SPF calculation.
- Step 11** At the pxr2 router, use the proper **show** command to examine the IS-IS neighbor table. How many entries do you see? Are they all Level 1 adjacency?
- Step 12** At the pxr1 router, use the proper **show** command to examine the IS-IS neighbor table. How many entries do you see? Do you see both Level 2 and Level 1 adjacencies?

Step 13 Display the running configuration of the p1r1 router. Is **cls routing** enabled by default?

Step 14 Save the current configurations of all the routers within your pod to NVRAM.

Completion Criteria

You have successfully completed this configuration exercise if you correctly supplied the commands required to configure IS-IS within your pod and to the backbone_r1 router, and if you were able to correctly answer the questions in the exercise. At the end of this exercise, all the routers should have full connectivity to each other. All routers will be running the IS-IS routing protocol.

Answers to Configuration Exercise: Configuring a Multiarea IS-IS Network

This section provides the answers to the questions in the configuration exercise. The answers are in **bold**.

Answers to Setup

To set up, do the following:

Step 1 Disable OSPF on all the routers within your pod.

The following example shows how to perform the required setup on the pod 1 routers.

```
p1r1(config)#no router ospf 200
```

```
p1r2(config)#no router ospf 200
```

```
p1r3(config)#no router ospf 200
```

Step 2 Shut the p1r1 S3 interface.

The following example shows how to perform the required setup on the p1r1 router.

```
p1r1(config)#int s3  
p1r1(config-if)#shutdown
```

Answers to Task 1: Enabling IS-IS Within Your Pod

Complete the following steps:

Step 1 Enable IS-IS on the pxr1, pxr2, and pxr3 routers within your pod.

Step 2 Configure the pxr1, pxr2, and pxr3 routers within your pod with the IS-IS NETs shown in the following table.

Pod	pxr1 NET	pxr2 NET	pxr3 NET
1	00.0001.1111.1111.1111.00	00.0001.1212.1212.1212.00	00.0001.1313.1313.1313.00
2	00.0002.2121.2121.2121.00	00.0002.2222.2222.2222.00	00.0002.2323.2323.2323.00
3	00.0003.3131.3131.3131.00	00.0003.3232.3232.3232.00	00.0003.3333.3333.3333.00
4	00.0004.4141.4141.4141.00	00.0004.4242.4242.4242.00	00.0004.4343.4343.4343.00
5	00.0005.5151.5151.5151.00	00.0005.5252.5252.5252.00	00.0005.5353.5353.5353.00
6	00.0006.6161.6161.6161.00	00.0006.6262.6262.6262.00	00.0006.6363.6363.6363.00
7	00.0007.7171.7171.7171.00	00.0007.7272.7272.7272.00	00.0007.7373.7373.7373.00
8	00.0008.8181.8181.8181.00	00.0008.8282.8282.8282.00	00.0008.8383.8383.8383.00
9	00.0009.9191.9191.9191.00	00.0009.9292.9292.9292.00	00.0009.9393.9393.9393.00
10	00.0010.A1A1.A1A1.A1A1.00	00.0010.A2A2.A2A2.A2A2.00	00.0010.A3A3.A3A3.A3A3.00
11	00.0011.B1B1.B1B1.B1B1.00	00.0011.B2B2.B2B2.B2B2.00	00.0011.B3B3.B3B3.B3B3.00
12	00.0012.C1C1.C1C1.C1C1.00	00.0012.C2C2.C2C2.C2C2.00	00.0012.C3C3.C3C3.C3C3.00

The following example shows how to enable IS-IS and configure the NETs on the pod 1 routers.

```
p1r1(config)#router isis
p1r1(config-router)#net 00.0001.1111.1111.1111.00

p1r2(config)#router isis
p1r2(config-router)#net 00.0001.1212.1212.1212.00

p1r3(config)#router isis
p1r3(config-router)#net 00.0001.1313.1313.1313.00
```

NOTE

The backbone_r1 router is configured with a NET of 00.0013.1111.2222.3333.00.

Step 3 Enable IS-IS on the interfaces specified in the following table on the pxr1, pxr2, and pxr3 routers within your pod.

Router	Interfaces on Which to Enable IS-IS				
pxr1	S0	S1	S2		
pxr2	S0	S1	E0	Loopback10	
pxr3	S0	E0	Loopback11	Loopback12	Loopback13

NOTE For this task, the pxr1 S3 interface is shut down.

The following example shows how to enable IS-IS on the appropriate interfaces of the pod 1 routers.

```
p1r1(config)#int s0
p1r1(config-if)#ip router isis
p1r1(config-if)#exit
p1r1(config)#int s1
p1r1(config-if)#ip router isis
p1r1(config-if)#exit
p1r1(config)#int s2
p1r1(config-if)#ip router isis
p1r1(config-if)#exit
p1r1(config)#
```

```
p1r2(config)#int s0
p1r2(config-if)#ip router isis
p1r2(config-if)#exit
p1r2(config)#int s1
p1r2(config-if)#ip router isis
p1r2(config-if)#exit
p1r2(config)#int e0
p1r2(config-if)#ip router isis
p1r2(config-if)#exit
p1r2(config)#int loopback 10
p1r2(config-if)#ip router isis
p1r2(config-if)#exit
p1r2(config)#
```

```
p1r3(config)#int s0
p1r3(config-if)#ip router isis
p1r3(config-if)#exit
p1r3(config)#int e0
p1r3(config-if)#ip router isis
p1r3(config-if)#exit
p1r3(config)#int loopback 11
p1r3(config-if)#ip router isis
p1r3(config-if)#exit
p1r3(config)#int loopback 12
p1r3(config-if)#ip router isis
p1r3(config-if)#exit
p1r3(config)#int loopback 13
p1r3(config-if)#ip router isis
p1r3(config-if)#exit
p1r3(config)#
```

Step 4 Display the routing table of the p1r1, p1r2, and p1r3 routers within your pod, and verify that you have full connectivity within your pod.

The following example output is from the p1r1 router:

```
p1r1#show ip route
<output omitted>
Gateway of last resort is not set

    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.32 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.16 [115/20] via 192.168.1.50, Serial2
    192.168.1.0/28 is subnetted, 4 subnets
i L1   192.168.1.64 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
        [115/20] via 192.168.1.50, Serial2
C       192.168.1.32 is directly connected, Serial1
C       192.168.1.48 is directly connected, Serial2
C       192.168.1.16 is directly connected, Serial0
i L1   192.168.101.0/24 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
p1r1#
```

The following example output is from the p1r2 router:

```
p1r2#show ip route
<output omitted>
Gateway of last resort is not set

    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.32 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.16 [115/20] via 192.168.1.66, Ethernet0
    192.168.1.0/28 is subnetted, 4 subnets
C       192.168.1.64 is directly connected, Ethernet0
C       192.168.1.32 is directly connected, Serial1
i L1   192.168.1.48 [115/20] via 192.168.1.33, Serial1
        [115/20] via 192.168.1.17, Serial0
        [115/20] via 192.168.1.66, Ethernet0
C       192.168.1.16 is directly connected, Serial0
C       192.168.101.0/24 is directly connected, Loopback10
p1r2#
```

The following example output is from the p1r3 router:

```
p1r3#show ip route
<output omitted>
Gateway of last resort is not set

    172.26.0.0/28 is subnetted, 3 subnets
C       172.26.1.48 is directly connected, Loopback13
C       172.26.1.32 is directly connected, Loopback12
C       172.26.1.16 is directly connected, Loopback11
    192.168.1.0/28 is subnetted, 4 subnets
C       192.168.1.64 is directly connected, Ethernet0
i L1   192.168.1.32 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
C       192.168.1.48 is directly connected, Serial0
i L1   192.168.1.16 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
i L1   192.168.101.0/24 [115/20] via 192.168.1.65, Ethernet0
p1r3#
```

These examples show that there is full connectivity within pod 1.

Step 5 Examine the p1r1 routing table and answer the following questions:

The following example output is from the p1r1 router:

```
p1r1#show ip route
<output omitted>
Gateway of last resort is not set

    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.32 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.16 [115/20] via 192.168.1.50, Serial2
    192.168.1.0/28 is subnetted, 4 subnets
i L1   192.168.1.64 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
        [115/20] via 192.168.1.50, Serial2
C       192.168.1.32 is directly connected, Serial1
C       192.168.1.48 is directly connected, Serial2
C       192.168.1.16 is directly connected, Serial0
i L1   192.168.101.0/24 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
p1r1#
```

Does IS-IS load-balance by default?

Yes, it does load-balance by default. In the previous example routing table, there are three equal-cost routes to 192.168.1.64 and two equal-cost routes to 192.168.101.0.

What is the IS-IS routing metric based on by default?

The default metric is a scalar cost where 1 hop = 10. Note that you count one hop for yourself or for the link (that is, the metric is always at least 20 for nonconnected networks).

Use the **show clns int s0** command to determine the Level 1 and Level 2 default metric.

The following example output is from the p1r1 router:

```
p1r1#show clns int s0
Serial0 is up, line protocol is up
  Checksums enabled, MTU 1500, Encapsulation HDLC
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
  Congestion Experienced bit set at 4 packets
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 18 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x100
    Level-1 Metric: 10, Priority: 64, Circuit ID: 1212.1212.1212.00
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: 1111.1111.1111.00
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 5 seconds
p1r1#
```

The Level 1 and Level 2 default metric is 10.

What is the default administrative distance of IS-IS routes?

The following example output is from the p1r1 router:

```
p1r1#show ip protocols
Routing Protocol is "isis"
  Sending updates every 0 seconds
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: isis
  Address Summarization:
    None
  Routing for Networks:
    Serial0
    Serial1
    Serial2
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.1.34     115          00:04:09
    192.168.1.50     115          00:03:21
    192.168.101.101 115          00:03:06
    172.26.1.17     115          00:03:11
  Distance: (default is 115)
```

This output shows that the default IP administrative distance is 115.

Step 6 Save the current configurations of all the routers within your pod to NVRAM.

The following example shows how to save the configuration of the p1r1 router using the copy run start command (this is an abbreviated form of the copy running-config startup-config command):

```
p1r1#copy run start
Destination filename [startup-config]?
Building configuration...
p1r1#
```

Answers to Task 2: Enabling Connectivity to the backbone_r1 Router

Complete the following steps:

Step 1 No shut the S3 interface on the p1r1 router.

Step 2 Enable IS-IS on the S3 interface on the p1r1.

The following example shows the configuration of the p1r1 router. The output from the router is also displayed:

```
p1r1(config)#int s3
p1r1(config-if)#no shutdown
p1r1(config-if)#
08:36:18: %LINK-3-UPDOWN: Interface Serial3, changed state to up
```

```

p1r1(config-if)#
08:36:19: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial3, changed state to
up
p1r1(config-if)#ip router isis
p1r1(config-if)#^Z
p1r1#

```

Step 3 Display the routing table of the p1r1, p1r2, and p1r3 routers within your pod, and verify that you have full connectivity to the backbone_r1 router.

The following example output is from the p1r1 router:

```

p1r1#show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
i L2   172.16.10.0 [115/20] via 10.1.1.100, Serial3
i L2   172.16.11.0 [115/20] via 10.1.1.100, Serial3
    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.32 [115/20] via 192.168.1.50, Serial2
i L1   172.26.1.16 [115/20] via 192.168.1.50, Serial2
    10.0.0.0/24 is subnetted, 1 subnets
C      10.1.1.0 is directly connected, Serial3
    192.168.1.0/28 is subnetted, 4 subnets
i L1   192.168.1.64 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
        [115/20] via 192.168.1.50, Serial2
C      192.168.1.32 is directly connected, Serial1
C      192.168.1.48 is directly connected, Serial2
C      192.168.1.16 is directly connected, Serial0
i L1   192.168.101.0/24 [115/20] via 192.168.1.34, Serial1
        [115/20] via 192.168.1.18, Serial0
p1r1#

```

The following example output is from the p1r2 router:

```

p1r2#show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
i L2   172.16.10.0 [115/30] via 192.168.1.33, Serial1
        [115/30] via 192.168.1.17, Serial0
i L2   172.16.11.0 [115/30] via 192.168.1.33, Serial1
        [115/30] via 192.168.1.17, Serial0
    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.32 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.16 [115/20] via 192.168.1.66, Ethernet0
    10.0.0.0/24 is subnetted, 1 subnets
i L1   10.1.1.0 [115/20] via 192.168.1.33, Serial1
        [115/20] via 192.168.1.17, Serial0
    192.168.1.0/28 is subnetted, 4 subnets
C      192.168.1.64 is directly connected, Ethernet0
C      192.168.1.32 is directly connected, Serial1
i L1   192.168.1.48 [115/20] via 192.168.1.33, Serial1
        [115/20] via 192.168.1.17, Serial0
        [115/20] via 192.168.1.66, Ethernet0
C      192.168.1.16 is directly connected, Serial0
C      192.168.101.0/24 is directly connected, Loopback10
p1r2#

```

The following example output is from the p1r3 router:

```

p1r3#show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
i L2   172.16.10.0 [115/30] via 192.168.1.49, Serial0
i L2   172.16.11.0 [115/30] via 192.168.1.49, Serial0
    172.26.0.0/28 is subnetted, 3 subnets
C     172.26.1.48 is directly connected, Loopback13
C     172.26.1.32 is directly connected, Loopback12
C     172.26.1.16 is directly connected, Loopback11
    10.0.0.0/24 is subnetted, 1 subnets
i L1   10.1.1.0 [115/20] via 192.168.1.49, Serial0
    192.168.1.0/28 is subnetted, 4 subnets
C     192.168.1.64 is directly connected, Ethernet0
i L1   192.168.1.32 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
C     192.168.1.48 is directly connected, Serial0
i L1   192.168.1.16 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
i L1 192.168.101.0/24 [115/20] via 192.168.1.65, Ethernet0
p1r3#

```

Do you see the backbone_r1 router's loopback interfaces in your routing table? (Note that you also might see routes to the other pods.)

Yes, the backbone_r1 loopback interfaces on networks 172.16.10.0 and 172.16.11.0 are visible in all the routing tables.

Does IS-IS perform autosummarization across the network boundary by default?

No, IS-IS does not perform autosummarization across the network boundary by default. The routing tables show all configured subnets of all networks.

Does IS-IS use the same default administrative distance for Level 1 and Level 2 routes?

Yes, IS-IS uses the same default admin distance for Level 1 and Level 2: 115.

Make sure you can ping the loopback interfaces on the backbone_r1 router from the p1r1, p1r2, and p1r3 routers within your pod.

The following examples are the results of the ping commands on the pod 1 routers to the backbone_r1 router's loopback interfaces. All pings are successful:

```

p1r1#ping 172.16.10.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms
p1r1#ping 172.16.11.100

```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 172.16.11.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms
p1r1#

p1r2#ping 172.16.10.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/60/64 ms
p1r2#ping 172.16.11.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/60/64 ms
p1r2#

p1r3#ping 172.16.10.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/59/60 ms
p1r3#ping 172.16.11.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/62/72 ms
p1r3#
```

- Step 4** Save the current configurations of all the routers within your pod to NVRAM.

The following example shows how to save the configuration of the p1r1 router using the copy run start command (this is an abbreviated form of the copy running-config startup-config command):

```
p1r1#copy run start
Destination filename [startup-config]?
Building configuration...
p1r1#
```

Answers to Task 3: Changing the IS-IS Router Type

Complete the following steps:

- Step 1** By default, the p1r1, p1r2, and p1r3 routers within your pod are what type of IS-IS router (Level 1, Level 2, or Level 1-2)?

What command can be used to determine the IS-IS router type?

The show clns is-neighbors command can be used to determine the IS-IS router type. The output from this command on the p1r1 router is shown here:

```
p1r1#show clns is-neighbors
```

System Id	Interface	State	Type	Priority	Circuit Id	Format
1212.1212.1212	Se1	Up	L1L2	0 /0	01	Phase V
1212.1212.1212	Se0	Up	L1L2	0 /0	00	Phase V
1313.1313.1313	Se2	Up	L1L2	0 /0	00	Phase V
1111.2222.3333	Se3	Up	L2	0	00	Phase V

```
p1r1#
```

By default, the routers are Level 1–2 (L1L2) routers. Note that the backbone_r1 router has a command to make it a Level 2–only router.

Step 2 Examine the IS-IS link state database of the p1r1, p1r2, and p1r3 routers within your pod. Do you see both the Level 1 and Level 2 link-state database on all the routers?

The following example output is from the p1r1 router:

```
p1r1#show isis database
IS-IS Level-1 Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00* 0x00000008  0x0B2C        838           1/0/0
1212.1212.1212.00-00 0x00000007  0x1C1C        839           1/0/0
1212.1212.1212.01-00 0x00000002  0xD237        434           0/0/0
1212.1212.1212.02-00 0x00000002  0xA265        1125          0/0/0
1313.1313.1313.00-00 0x00000009  0xBCB1        839           1/0/0
1313.1313.1313.01-00 0x00000001  0x4840        0 (441)       0/0/0
1313.1313.1313.02-00 0x00000002  0xA853        1117          0/0/0
1313.1313.1313.03-00 0x00000002  0xA159        1140          0/0/0
1313.1313.1313.04-00 0x00000001  0x9C5E        449           0/0/0

IS-IS Level-2 Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00* 0x0000000C  0xDAD4        827           0/0/0
1111.2222.3333.00-00 0x00000006  0x352F        826           0/0/0
1212.1212.1212.00-00 0x0000000A  0x6588        828           0/0/0
1212.1212.1212.01-00 0x00000003  0x6031        1115          0/0/0
1313.1313.1313.00-00 0x0000000A  0xEA43        828           0/0/0
1313.1313.1313.01-00 0x00000002  0x4641        0 (441)       0/0/0
p1r1#
```

The following example output is from the p1r2 router:

```
p1r2#show isis database
IS-IS Level-1 Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x00000008  0x0B2C        827           1/0/0
1212.1212.1212.00-00* 0x00000007  0x1C1C        832           1/0/0
1212.1212.1212.01-00* 0x00000002  0xD237        427           0/0/0
1212.1212.1212.02-00* 0x00000002  0xA265        1118          0/0/0
1313.1313.1313.00-00 0x00000009  0xBCB1        830           1/0/0
1313.1313.1313.01-00 0x00000001  0x4840        0 (429)       0/0/0
1313.1313.1313.02-00 0x00000002  0xA853        1107          0/0/0
1313.1313.1313.03-00 0x00000002  0xA159        1131          0/0/0
1313.1313.1313.04-00 0x00000001  0x9C5E        440           0/0/0

IS-IS Level-2 Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x0000000C  0xDAD4        816           0/0/0
1111.2222.3333.00-00 0x00000006  0x352F        815           0/0/0
1212.1212.1212.00-00* 0x0000000A  0x6588        821           0/0/0
1212.1212.1212.01-00* 0x00000003  0x6031        1108          0/0/0
1313.1313.1313.00-00 0x0000000A  0xEA43        819           0/0/0
1313.1313.1313.01-00 0x00000002  0x4641        0 (432)       0/0/0
p1r2#
```

The following example output is from the p1r3 router:

```
p1r3#show isis database
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x00000008  0x0B2C       819           1/0/0
1212.1212.1212.00-00 0x00000007  0x1C1C       822           1/0/0
1212.1212.1212.01-00 0x00000002  0xD237       415           0/0/0
1212.1212.1212.02-00 0x00000002  0xA265       1108          0/0/0
1313.1313.1313.00-00* 0x00000009  0xBCB1       824           1/0/0
1313.1313.1313.01-00* 0x00000001  0x4840       0 (421)       0/0/0
1313.1313.1313.02-00* 0x00000002  0xA853       1102          0/0/0
1313.1313.1313.03-00* 0x00000002  0xA159       1125          0/0/0
1313.1313.1313.04-00* 0x00000001  0x9C5E       434           0/0/0

IS-IS Level-2 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x0000000C  0xDAD4       808           0/0/0
1111.2222.3333.00-00 0x00000006  0x352F       807           0/0/0
1212.1212.1212.00-00 0x0000000A  0x6588       811           0/0/0
1212.1212.1212.01-00 0x00000003  0x6031       1098          0/0/0
1313.1313.1313.00-00* 0x0000000A  0xEA43       813           0/0/0
1313.1313.1313.01-00* 0x00000002  0x4641       0 (420)       0/0/0
p1r3#
```

Yes, the Level 1 and Level 2 link-state database on all the routers in the pod is visible.

In the link-state database, which link-state packet (LSP) is identified with an asterisk?

The LSPs with an asterisk are the local LSPs—in other words, LSPs originated from this router.

Examine the routing table of the p1r2 and p1r3 routers within your pod. Do you see both the Level 1 and Level 2 routes?

The following example output is from the p1r2 router:

```
p1r2#show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
i L2   172.16.10.0 [115/30] via 192.168.1.33, Serial1
       [115/30] via 192.168.1.17, Serial0
i L2   172.16.11.0 [115/30] via 192.168.1.33, Serial1
       [115/30] via 192.168.1.17, Serial0
    172.26.0.0/28 is subnetted, 3 subnets
i L1   172.26.1.48 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.32 [115/20] via 192.168.1.66, Ethernet0
i L1   172.26.1.16 [115/20] via 192.168.1.66, Ethernet0
    10.0.0.0/24 is subnetted, 1 subnets
i L1   10.1.1.0 [115/20] via 192.168.1.33, Serial1
       [115/20] via 192.168.1.17, Serial0
    192.168.1.0/28 is subnetted, 4 subnets
C      192.168.1.64 is directly connected, Ethernet0
C      192.168.1.32 is directly connected, Serial1
i L1   192.168.1.48 [115/20] via 192.168.1.33, Serial1
       [115/20] via 192.168.1.17, Serial0
       [115/20] via 192.168.1.66, Ethernet0
C      192.168.1.16 is directly connected, Serial0
C      192.168.101.0/24 is directly connected, Loopback10
p1r2#
```

The following example output is from the p1r3 router:

```

p1r3#show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
i L2   172.16.10.0 [115/30] via 192.168.1.49, Serial0
i L2   172.16.11.0 [115/30] via 192.168.1.49, Serial0
    172.26.0.0/28 is subnetted, 3 subnets
C      172.26.1.48 is directly connected, Loopback13
C      172.26.1.32 is directly connected, Loopback12
C      172.26.1.16 is directly connected, Loopback11
    10.0.0.0/24 is subnetted, 1 subnets
i L1   10.1.1.0 [115/20] via 192.168.1.49, Serial0
    192.168.1.0/28 is subnetted, 4 subnets
C      192.168.1.64 is directly connected, Ethernet0
i L1   192.168.1.32 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
C      192.168.1.48 is directly connected, Serial0
i L1   192.168.1.16 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
i L1  192.168.101.0/24 [115/20] via 192.168.1.65, Ethernet0
p1r3#

```

Yes, both the Level 1 and Level 2 routes are in the p1r2 and p1r3 routing tables.

Step 3 Configure the p1r2 and p1r3 routers within your pod to be Level 1 IS-IS routers.

The following example configuration shows how to configure p1r2 and p1r3 to be Level 1 IS-IS routers:

```

p1r2(config)#router isis
p1r2(config-router)#is-type level-1

p1r3(config)#router isis
p1r3(config-router)#is-type level-1

```

NOTE

You will have to wait a minute or so after changing the IS type for the routing tables to converge.

Step 4 What are the advantages of configuring the p1r2 and p1r3 routers within your pod to be Level 1 IS-IS routers?

The advantages include having shorter routing tables in the Level 1 routers (containing only area routes and default route) and being less processor intensive.

Step 5 Re-examine the link-state database of the p1r2 and p1r3 routers within your pod. Do the p1r2 and p1r3 routers contain only the Level 1 link-state database now?

The following example output is from the p1r2 router:

```
p1r2#show isis database
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x0000000D  0x0131        1117          1/0/0
1212.1212.1212.00-00* 0x0000000A  0x0C33        1111          0/0/0
1212.1212.1212.01-00* 0x00000005  0xCA3E        1125          0/0/0
1212.1212.1212.02-00* 0x00000002  0xA265        820           0/0/0
1313.1313.1313.00-00 0x0000000D  0xAAC9        1115          0/0/0
1313.1313.1313.01-00 0x00000001  0x4840        0 (131)       0/0/0
1313.1313.1313.02-00 0x00000002  0xA853        810           0/0/0
1313.1313.1313.03-00 0x00000002  0xA159        833           0/0/0
1313.1313.1313.04-00 0x00000002  0x9A5F        999           0/0/0
p1r2#
```

The following example output is from the p1r2 router:

```
p1r3#show isis database
IS-IS Level-1 Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
1111.1111.1111.00-00 0x0000000D  0x0131        1109          1/0/0
1212.1212.1212.00-00 0x0000000A  0x0C33        1099          0/0/0
1212.1212.1212.01-00 0x00000005  0xCA3E        1114          0/0/0
1212.1212.1212.02-00 0x00000002  0xA265        809           0/0/0
1313.1313.1313.00-00* 0x0000000D  0xAAC9        1110          0/0/0
1313.1313.1313.01-00* 0x00000001  0x4840        0 (122)       0/0/0
1313.1313.1313.02-00* 0x00000002  0xA853        802           0/0/0
1313.1313.1313.03-00* 0x00000002  0xA159        826           0/0/0
1313.1313.1313.04-00* 0x00000002  0x9A5F        992           0/0/0
p1r3#
```

Yes, the p1r2 and p1r3 routers contain only the Level 1 link-state database now.

- Step 6** Re-examine the routing table of the p1r2 and p1r3 routers now. What are the differences after they became Level 1–only IS-IS routers?

The following example output is from the p1r2 router:

```
p1r2#show ip route
<output omitted>
Gateway of last resort is 192.168.1.17 to network 0.0.0.0

 172.26.0.0/28 is subnetted, 3 subnets
 i L1   172.26.1.48 [115/20] via 192.168.1.66, Ethernet0
 i L1   172.26.1.32 [115/20] via 192.168.1.66, Ethernet0
 i L1   172.26.1.16 [115/20] via 192.168.1.66, Ethernet0
 10.0.0.0/24 is subnetted, 1 subnets
 i L1   10.1.1.0 [115/20] via 192.168.1.17, Serial0
        [115/20] via 192.168.1.33, Serial1
 192.168.1.0/28 is subnetted, 4 subnets
 C      192.168.1.64 is directly connected, Ethernet0
 C      192.168.1.32 is directly connected, Serial1
 i L1   192.168.1.48 [115/20] via 192.168.1.17, Serial0
        [115/20] via 192.168.1.33, Serial1
        [115/20] via 192.168.1.66, Ethernet0
 C      192.168.1.16 is directly connected, Serial0
 C      192.168.101.0/24 is directly connected, Loopback10
 i*L1  0.0.0.0/0 [115/10] via 192.168.1.17, Serial0
        [115/10] via 192.168.1.33, Serial1
p1r2#
```

The following example output is from the p1r3 router:

```

p1r3#show ip route
<output omitted>
Gateway of last resort is 192.168.1.49 to network 0.0.0.0

    172.26.0.0/28 is subnetted, 3 subnets
C       172.26.1.48 is directly connected, Loopback13
C       172.26.1.32 is directly connected, Loopback12
C       172.26.1.16 is directly connected, Loopback11
    10.0.0.0/24 is subnetted, 1 subnets
i L1   10.1.1.0 [115/20] via 192.168.1.49, Serial0
    192.168.1.0/28 is subnetted, 4 subnets
C       192.168.1.64 is directly connected, Ethernet0
i L1   192.168.1.32 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
C       192.168.1.48 is directly connected, Serial0
i L1   192.168.1.16 [115/20] via 192.168.1.65, Ethernet0
        [115/20] via 192.168.1.49, Serial0
i L1 192.168.101.0/24 [115/20] via 192.168.1.65, Ethernet0
i*L1 0.0.0.0/0 [115/10] via 192.168.1.49, Serial0
p1r3#

```

The difference is that the routing tables are shorter, containing only area routes and a default route.

Step 7 In the p1r2 and p1r3 routing tables, do you see a default route through the p1r1 router?

Yes, a default route to the p1r1 router is visible.

Step 8 From the p1r2 or p1r3 routers, can you still ping the backbone_r1 router's loopback interfaces?

Yes, as shown in the following example output, the p1r2 and p1r3 routers can still ping the backbone_r1 router's loopback interfaces:

```

p1r2#ping 172.16.10.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/74/136 ms
p1r2#ping 172.16.11.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
p1r2#

p1r3#ping 172.16.10.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/89/216 ms
p1r3#ping 172.16.11.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/60/60 ms
p1r3#

```

Step 9 Save the current configurations of all the routers within your pod to NVRAM.

The following example shows how to save the configuration of the p1r1 router using the copy run start command (this is an abbreviated form of the copy running-config startup-config command):

```
p1r1#copy run start
Destination filename [startup-config]?
Building configuration...
p1r1#
```

Answers to Task 4: Configure Route Summarization

Complete the following steps:

Step 1 From p1r3, telnet to the backbone_r1 router using the password cisco.

The following example output is from the p1r3 router:

```
p1r3#bbr1
Trying bbr1 (10.1.1.100)... Open
```

User Access Verification

```
Password:
backbone_r1>
```

Step 2 Display the routing table of the backbone_r1 router. Do you see your 192.168.x.y/28 and 172.26.x.y/28 subnets?

The following example output is from the backbone_r1 router:

```
backbone_r1>show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
C       172.16.10.0 is directly connected, Loopback100
C       172.16.11.0 is directly connected, Loopback101
    172.26.0.0/28 is subnetted, 3 subnets
 i L2   172.26.1.48 [115/30] via 10.1.1.1, Serial0/0
 i L2   172.26.1.32 [115/30] via 10.1.1.1, Serial0/0
 i L2   172.26.1.16 [115/30] via 10.1.1.1, Serial0/0
    10.0.0.0/24 is subnetted, 1 subnets
C       10.1.1.0 is directly connected, Serial0/0
    192.168.1.0/28 is subnetted, 4 subnets
 i L2   192.168.1.64 [115/30] via 10.1.1.1, Serial0/0
 i L2   192.168.1.32 [115/20] via 10.1.1.1, Serial0/0
 i L2   192.168.1.48 [115/20] via 10.1.1.1, Serial0/0
 i L2   192.168.1.16 [115/20] via 10.1.1.1, Serial0/0
 i L2  192.168.101.0/24 [115/30] via 10.1.1.1, Serial0/0
backbone_r1>
```

Yes, the 192.168.1.y/28 and 172.26.1.y/28 subnets appear in the backbone_r1 router's routing table.

Step 3 Exit the telnet to the backbone_r1 router.

The following example output is from the backbone_r1 router:

```
backbone_r1>exit
[Connection to bbr1 closed by foreign host]
p1r3#
```

Step 4 Configure the Level 1–2 router within your pod to perform route summarization. Summarize the 192.168.x.y/28 and 172.26.x.y/28 subnets within your pod to 192.168.x.0/24 and 172.26.x.0/24 (where x is your pod number).

The following example shows how to configure the p1r1 router, the Level 1–2 router in pod 1, to perform route summarization:

```
p1r1(config)#router isis
p1r1(config-router)#summary-address 192.168.1.0 255.255.255.0
p1r1(config-router)#summary-address 172.26.1.0 255.255.255.0
```

Step 5 Telnet to the backbone_r1 router.

The following example output is from the p1r1 router:

```
p1r1#bbr1
Trying bbr1 (10.1.1.100)... Open

User Access Verification

Password:
backbone_r1>
```

Step 6 Re-examine the routing table of the backbone_r1 router. Do you see your summarized routes?

The following example output is from the backbone_r1 router:

```
backbone_r1>show ip route
<output omitted>
Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
C       172.16.10.0 is directly connected, Loopback100
C       172.16.11.0 is directly connected, Loopback101
    172.26.0.0/24 is subnetted, 1 subnets
i L2   172.26.1.0 [115/30] via 10.1.1.1, Serial0/0
    10.0.0.0/24 is subnetted, 1 subnets
C       10.1.1.0 is directly connected, Serial0/0
i L2  192.168.1.0/24 [115/20] via 10.1.1.1, Serial0/0
i L2  192.168.101.0/24 [115/30] via 10.1.1.1, Serial0/0
backbone_r1>
```

Yes, the summarized routes to 172.26.1.0 and 192.168.1.0 are in the backbone_r1 router's routing table.

Step 7 From the backbone_r1 router, can you ping all the interfaces within your pod?

Yes, as the example output shows, the backbone_r1 router can ping all the interfaces on the pod 1 routers:

```
backbone_r1>ping 192.168.1.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.17, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/64/172 ms
backbone_r1>ping 192.168.1.18

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.18, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/69/84 ms
backbone_r1>ping 192.168.1.33

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.33, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/40/56 ms
backbone_r1>ping 192.168.1.34

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.34, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/68/84 ms
backbone_r1>ping 192.168.1.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.49, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/40/56 ms
backbone_r1>ping 192.168.1.50

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.50, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/68/84 ms
backbone_r1>ping 192.168.1.65

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.65, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/69/84 ms
backbone_r1>ping 192.168.1.66

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.66, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/67/80 ms
backbone_r1>ping 192.168.101.101

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.101.101, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/69/88 ms
backbone_r1>ping 172.26.1.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.26.1.17, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/106/268 ms
```

```
backbone_r1>ping 172.26.1.33

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.26.1.33, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/69/84 ms
backbone_r1>ping 172.26.1.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.26.1.49, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/68/84 ms
backbone_r1>
```

Step 8 Exit the telnet to the backbone_r1 router.

The following example output is from the backbone_r1 router:

```
backbone_r1>
backbone_r1>exit

[Connection to bbr1 closed by foreign host]
p1r1#
```

Step 9 Save the current configurations of all the routers within your pod to NVRAM.

The following example shows how to save the configuration of the p1r1 router using the copy run start command (this is an abbreviated form of the copy running-config startup-config command):

```
p1r1#copy run start
Destination filename [startup-config]?
Building configuration...
p1r1#
```

Answers to Task 5: Using IS-IS show and debug Commands

Complete the following steps:

Step 1 At the p1r2 router, enable the **debug isis adj-packets** command and observe the debug output.

The following example shows enabling the debug on the p1r2 router and the output that results:

```
p1r2#debug isis adj-packets
IS-IS Adjacency related packets debugging is on
p1r2#
01:45:23: ISIS-Adj: Sending serial IIH on Serial1
01:45:23: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2, cir id
01
01:45:23: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:23: ISIS-Adj: Action = ACCEPT
01:45:24: ISIS-Adj: Sending L1 IIH on Loopback10
p1r2#
01:45:24: ISIS-Adj: Sending L1 IIH on Ethernet0
```

```

p1r2#
01:45:26: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2, cir id
00
01:45:26: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:26: ISIS-Adj: Action = ACCEPT
01:45:27: ISIS-Adj: Sending L1 IIH on Loopback10
p1r2#
01:45:27: ISIS-Adj: Sending L1 IIH on Ethernet0
01:45:28: ISIS-Adj: Sending serial IIH on Serial0
p1r2#
01:45:30: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:30: ISIS-Adj: Sending L1 IIH on Ethernet0
01:45:31: ISIS-Adj: Rec L1 IIH from 0010.7b81.d6ec (Ethernet0), cir type L1, cir
id 1212.1212.1212.01
p1r2#

```

What type of hello is p1r2 sending to p1r3 over the Ethernet interface?

The p1r2 router is sending an L1 hello to the p1r3 router.

Step 2 Shut the p1r2 E0 interface and observe the debug output.

The following example shows shutting down the Ethernet 0 interface on the p1r2 router and the output that results:

```

p1r2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
01:45:38: ISIS-Adj: Rec L1 IIH from 0010.7b81.d6ec (Ethernet0), cir type L1,
cir
id 1212.1212.1212.01
01:45:39: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:39: ISIS-Adj: Sending L1 IIH on Ethernet0
01:45:39: ISIS-Adj: Sending serial IIH on Serial1
p1r2(config)#int e0
01:45:41: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2,
cir id
01
01:45:41: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:41: ISIS-Adj: Action = ACCEPT
01:45:42: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:42: ISIS-Adj: Sending L1 IIH on Ethernet0
p1r2(config-if)#shutdown
p1r2(config-if)#
01:45:45: ISIS-Adj: Sending serial IIH on Serial0
01:45:45: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:45: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2,
cir id
00
01:45:45: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:45: ISIS-Adj: Action = ACCEPT
01:45:46: %LINK-5-CHANGED: Interface Ethernet0, changed state to
administratively down
01:45:47: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0,
changed state to down
01:45:48: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:48: ISIS-Adj: Sending serial IIH on Serial1
01:45:49: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2,
cir id01
01:45:49: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:49: ISIS-Adj: Action = ACCEPT
01:45:50: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:53: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:54: ISIS-Adj: Sending serial IIH on Serial0

```

```

01:45:54: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2,
    cir id00
01:45:54: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:54: ISIS-Adj: Action = ACCEPT
01:45:56: ISIS-Adj: Sending serial IIH on Serial1
01:45:56: ISIS-Adj: Sending L1 IIH on Loopback10
01:45:58: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2,
    cir id01
01:45:58: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:45:58: ISIS-Adj: Action = ACCEPT
p1r2(config-if)#

```

Step 3 No shut the pxr2 E0 interface and observe the debug output. From the debug output, determine which router (pxr2 or pxr3) is elected as the DR.

NOTE

The correct term in IS-IS is DIS, not DR. However, the debug output on the router uses the term DR, so we will also use the term DR in this task.

In the debug output, you should see something similar to the output shown in Example S-32 if pxr3 is the DR.

Example S-32 debug isis adj-packets *Command Output from Router pxr2 If Router pxr3 Is the DR*

```

02:23:12: ISIS-Adj: Adjacency state goes to Up
02:23:12: ISIS-Adj: Run level 1 DR election for Ethernet0
02:23:12: ISIS-Adj: New level 1 DR 3333.3333.3333 on Ethernet0
pxr2#

```

Or, in the debug output, you should see something similar to the output shown in Example S-33 if pxr2 is the DR.

Example S-33 debug isis adj-packets *Command Output from Router pxr2 If Router pxr2 Is the DR*

```

05:19:43: ISIS-Adj: Run level 1 DR election for Ethernet0
05:19:43: ISIS-Adj: No change (it's us)
pxr2#

```

What determines who will be elected as the DR on the Ethernet?

The following example shows no shutting the Ethernet 0 interface on the p1r2 router and the output that results:

```

p1r2(config-if)#int e0
01:46:17: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:17: ISIS-Adj: Sending serial IIH on Serial0
p1r2(config-if)#no shutdown
p1r2(config-if)#
01:46:20: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:21: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:21: ISIS-Adj: Rec L1 IIH from 0010.7b81.d6ec (Ethernet0), cir type L1,
    cir id 1212.1212.1212.01
01:46:21: ISIS-Adj: New adjacency, level 1 for 0010.7b81.d6ec
01:46:22: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:22: %LINK-3-UPDOWN: Interface Ethernet0, changed state to up
01:46:22: ISIS-Adj: Sending serial IIH on Serial1
01:46:23: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2,
    cir id

```

```

00
01:46:23: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:23: ISIS-Adj: Action = ACCEPT
01:46:23: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0,
    changed state to up
01:46:23: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:24: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:24: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2,
    cir id
01
01:46:24: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:24: ISIS-Adj: Action = ACCEPT
01:46:26: ISIS-Adj: Sending serial IIH on Serial0
01:46:26: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:27: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:29: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:29: ISIS-Adj: Rec L1 IIH from 0010.7b81.d6ec (Ethernet0), cir type L1,
    cir
    id 1212.1212.1212.01
01:46:29: ISIS-Adj: Adjacency state goes to Up
01:46:29: ISIS-Adj: Run level 1 DR election for Ethernet0
01:46:29: ISIS-Adj: No change (it's us)
01:46:30: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:30: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2, cir id
00
01:46:30: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:30: ISIS-Adj: Action = ACCEPT
01:46:32: ISIS-Adj: Sending serial IIH on Serial1
01:46:32: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:32: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2, cir id
01
01:46:32: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:32: ISIS-Adj: Action = ACCEPT
01:46:33: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:35: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:35: ISIS-Adj: Sending serial IIH on Serial0
01:46:36: ISIS-Adj: Sending L1 IIH on Ethernet0
p1r2(config-if)#^Z
p1r2#
01:46:38: ISIS-Adj: Sending L1 IIH on Loopback10
01:46:38: ISIS-Adj: Rec L1 IIH from 0010.7b81.d6ec (Ethernet0), cir type L1,
    cir
    id 1212.1212.1212.01
01:46:39: %SYS-5-CONFIG_I: Configured from console by console
01:46:39: ISIS-Adj: Sending serial IIH on Serial1
01:46:39: ISIS-Adj: Sending L1 IIH on Ethernet0
01:46:40: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0), cir type L1L2,
    cir id00
01:46:40: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:40: ISIS-Adj: Action = ACCEPT
01:46:40: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1), cir type L1L2,
    cir id01
01:46:40: ISIS-Adj: rcvd state UP, old state UP, new state UP
01:46:40: ISIS-Adj: Action = ACCEPT
01:46:41: ISIS-Adj: Sending L1 IIH on Loopback10

```

As the example output shows, p1r2 remains the DR. Priority determines who is the DR. The default is 64; a priority of 0 means that the router cannot be elected. If two routers have the same priority, the one with the highest MAC address is the DR.

Step 4 Disable debug isis adj-packets at the p1r2 router.

The following example shows how to disable debugging on the p1r2 router:

```
p1r2#no debug all
All possible debugging has been turned off
p1r2#
01:46:42: ISIS-Adj: Sending L1 IIH on Ethernet0
p1r2#
```

Step 5 At the DR (as determined from Step 3), enter the **show clns int e0** command. What is the default IS-IS router priority set to?

The following example shows the output of the show clns int 0 command on the DR, p1r2:

```
p1r2#show clns int e0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
  Congestion Experienced bit set at 4 packets
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 30 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x2, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: 1212.1212.1212.01
    Number of active level-1 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 420 milliseconds
p1r2#
```

As shown, the default router priority is set to 64.

Step 6 At the DR (as determined from Step 3), change the E0 interface IS-IS router priority to 63.

The following example shows how to change the IS-IS router priority on the E0 interface of p1r2, the DR:

```
p1r2(config)#int e0
p1r2(config-if)#isis priority 63
```

Step 7 At the DR (as determined from Step 3), enter the **show clns int e0** command again and verify that the DR is now changed to the other router.

The Circuit ID field in the **show clns int e0** output identifies which router is the DR. Example S-34 shows the output of this command on the pxr3 router. In this example, the DR is the pxr2 router and the pxr3 router has an IS-IS priority of 63.

Example S-34 show clns Command Output from Router pxr3 If Router pxr2 Is the DR

```
pxr3#show clns int e0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
  Congestion Experienced bit set at 4 packets
  CLNS fast switching enabled
  CLNS SSE switching disabled
```

continues

Example S-34 show clns Command Output from Router pxr3 If Router pxr2 Is the DR (Continued)

```

DEC compatibility mode OFF for this interface
Next ESH/ISH in 21 seconds
Routing Protocol: IS-IS
  Circuit Type: level-1-2
  Interface number 0x1, local circuit ID 0x1
  Level-1 Metric: 10, Priority: 63, Circuit ID: 2222.2222.2222.01
  Number of active level-1 adjacencies: 1
Next IS-IS LAN Level-1 HELLO in 8 seconds

```

The following example is the output of the show clns int e0 command on the p1r2 router, the DR:

```

p1r2#show clns int e0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
  Congestion Experienced bit set at 4 packets
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 24 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x2, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 63, Circuit ID: 1313.1313.1313.01
    Number of active level-1 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 4 seconds
p1r2#

```

In this example, the priority is now 63 on interface E0. The circuit ID is that of router p1r3, indicating that p1r3 is now the DR.

- Step 8** At the pxr2 router, enable the **debug isis update-packets** command. Shut and no shut the E0 interface on the pxr2 router and observe the debug output.

The following example shows how to enable the debug on the p1r2 router and the output on the router:

```

p1r2#debug isis update-packets
IS-IS Update related packet debugging is on
p1r2#
01:58:08: ISIS-SNP: Rec L1 CSNP from 1313.1313.1313 (Ethernet0)
p1r2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
p1r2(config)#int e0
p1r2(config-if)#shutdown
p1r2(config-if)#
01:58:16: ISIS-SNP: Rec L1 CSNP from 1313.1313.1313 (Ethernet0)
01:58:16: ISIS-Update: Building L1 LSP
01:58:16: ISIS-Update: TLV contents different, code 80
01:58:16: ISIS-Update: TLV contents different, code 2
01:58:16: ISIS-Update: Full SPF required
01:58:16: ISIS-Update: Building L1 pseudonode LSP for Ethernet0
01:58:16: ISIS-Update: TLV contents different, code 2
01:58:16: ISIS-Update: Full SPF required
01:58:16: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 10,

```

```
ht 1199 on Serial1
01:58:16: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 10,
ht 1199 on Serial0
01:58:16: ISIS-Update: Sending L1 LSP 1212.1212.1212.01-00, seq 9,
ht 1199 on Serial1
01:58:16: ISIS-Update: Sending L1 LSP 1212.1212.1212.01-00, seq 9,
ht 1199 on Serial0
01:58:17: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial1)
01:58:17: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 10, ht 1196
01:58:17: ISIS-SNP: PSNP entry 1212.1212.1212.01-00, seq 9, ht 1197
01:58:17: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial0)
01:58:17: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 10, ht 1196
01:58:17: ISIS-SNP: PSNP entry 1212.1212.1212.01-00, seq 9, ht 1197
01:58:18: %LINK-5-CHANGED: Interface Ethernet0, changed state to
administratively down
01:58:18: ISIS-Update: Building L1 LSP
01:58:18: ISIS-Update: No change, suppress L1 LSP 1212.1212.1212.00-00, seq 11
01:58:19: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0,
changed state to down
p1r2(config-if)#no shutdown
p1r2(config-if)#
01:58:30: ISIS-Update: Building L1 LSP
01:58:30: ISIS-Update: TLV contents different, code 80
01:58:30: ISIS-Update: Leaf routes changed
01:58:30: ISIS-Update: Building L1 pseudonode LSP for Ethernet0
01:58:30: ISIS-Update: No change, suppress L1 LSP 1212.1212.1212.01-00, seq A
01:58:30: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 11,
ht 1199 on Serial1
01:58:30: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 11,
ht 1199 on Serial0
01:58:30: ISIS-Update: Rec L1 LSP 1313.1313.1313.01-00, seq 2, ht 1197,
01:58:30: ISIS-Update: from SNPA *HDLC* (Serial1)
01:58:30: ISIS-Update: LSP newer than database copy
01:58:30: ISIS-Update: TLV contents different, code 2
01:58:30: ISIS-Update: Full SPF required
01:58:31: ISIS-Update: Building L1 LSP
01:58:31: ISIS-Update: No change, suppress L1 LSP 1212.1212.1212.00-00, seq 12
01:58:31: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial1)
01:58:31: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 11, ht 1197
01:58:31: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial0)
01:58:31: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 11, ht 1196
01:58:31: ISIS-Update: Build L1 PSNP entry for 1313.1313.1313.01-00, seq 2
01:58:31: ISIS-Update: Sending L1 PSNP on Serial1
01:58:31: ISIS-Update: Build L1 PSNP entry for 1313.1313.1313.01-00, seq 2
01:58:31: ISIS-Update: Sending L1 PSNP on Serial0
01:58:32: %LINK-3-UPDOWN: Interface Ethernet0, changed state to up
01:58:33: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0,
changed state to up
01:58:35: ISIS-Update: Purging LSP 1212.1212.1212.01-00, seq 9
01:58:35: ISIS-Update: Building L1 LSP
01:58:35: ISIS-Update: Rate limiting L1 LSP 1212.1212.1212.00-00, seq 12
01:58:35: ISIS-Update: TLV contents different, code 2
01:58:35: ISIS-Update: Full SPF required
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.01-00, seq 9,
ht 0 on Ethernet0
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.01-00, seq 9,
ht 0 on Serial1
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.01-00, seq 9,
ht 0 on Serial0
01:58:35: ISIS-Update: Rec L1 LSP 1212.1212.1212.01-00, seq 9, ht 0,
01:58:35: ISIS-Update: from SNPA 0010.7b81.d6ec (Ethernet0)
01:58:35: ISIS-Update: LSP same as database copy
01:58:35: ISIS-Update: Rec L1 LSP 1212.1212.1212.01-00, seq 9, ht 0,
01:58:35: ISIS-Update: from SNPA *HDLC* (Serial1)
```

```

01:58:35: ISIS-Update: LSP same as database copy
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 12,
    ht 1199 on Ethernet0
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 12,
    ht 1199 on Serial1
01:58:35: ISIS-Update: Sending L1 LSP 1212.1212.1212.00-00, seq 12,
    ht 1199 on Serial0
01:58:35: ISIS-Update: Rec L1 LSP 1313.1313.1313.01-00, seq 3, ht 1196,
01:58:35: ISIS-Update: from SNPA 0010.7b81.d6ec (Ethernet0)
01:58:35: ISIS-Update: LSP newer than database copy
01:58:35: ISIS-Update: TLV contents different, code 2
01:58:35: ISIS-Update: Full SPF required
01:58:35: ISIS-Update: Sending L1 LSP 1313.1313.1313.01-00, seq 3,
    ht 1194 on Serial1
01:58:35: ISIS-Update: Sending L1 LSP 1313.1313.1313.01-00, seq 3,
    ht 1194 on Serial0
01:58:36: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial0)
01:58:36: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 12, ht 1197
01:58:36: ISIS-SNP: PSNP entry 1212.1212.1212.01-00, seq 9, ht 0
01:58:36: ISIS-SNP: PSNP entry 1313.1313.1313.01-00, seq 3, ht 1192
01:58:36: ISIS-Update: Build L1 PSNP entry for 1212.1212.1212.01-00, seq 9
01:58:36: ISIS-Update: Sending L1 PSNP on Serial1
01:58:36: ISIS-SNP: Rec L1 PSNP from 1111.1111.1111 (Serial1)
01:58:36: ISIS-SNP: PSNP entry 1212.1212.1212.00-00, seq 12, ht 1197
01:58:36: ISIS-SNP: PSNP entry 1313.1313.1313.01-00, seq 3, ht 1192
01:58:41: ISIS-SNP: Rec L1 CSNP from 1313.1313.1313 (Ethernet0)
p1r2(config-if)#exit
p1r2(config)#
01:58:50: ISIS-SNP: Rec L1 CSNP from 1313.1313.1313 (Ethernet0)

```

Step 9 Disable the **debug isis update-packets** at the p1r2 router.

The following example shows how to disable the debug on the p1r2 router:

```

p1r2#no debug all
All possible debugging has been turned off

```

Step 10 At the p1r2 router, use the proper **show** command to verify how often and why the router has run a full SPF calculation.

The following example shows the output of the show isis spf-log command on the p1r2 router:

```

p1r2#show isis spf-log

Level 1 SPF log
When      Duration  Nodes  Count  Last trigger LSP  Triggers
00:46:50    12       5      2  1313.1313.1313.02-00  NEWLSP
00:46:45    16       7      3  1313.1313.1313.03-00  NEWLSP TLVCONTENT
00:46:40    16       8      2  1313.1313.1313.04-00  NEWLSP TLVCONTENT
00:40:10    16       8      4  1313.1313.1313.00-00  ATTACHFLAG LSPHEADER
00:36:15    16       8      1  PERIODIC
00:31:28     8       3      5  1212.1212.1212.01-00  NEWLEVEL DELADJ LSPHEADER
00:30:37    40       8      2  1212.1212.1212.00-00  NEWADJ TLVCONTENT
00:30:32    36       4      3  1313.1313.1313.00-00  NEWADJ TLVCONTENT
00:30:22    20       8      1  1111.1111.1111.00-00  TLVCONTENT
00:30:17    24       8      3  1313.1313.1313.00-00  NEWADJ LSPHEADER TLVCONTEN
T
00:21:16    16       8      1  PERIODIC
00:14:29    20       8      3  1212.1212.1212.01-00  IPBACKUP IPQUERY TLVCONTEN
T
00:14:24    16       8      1  1313.1313.1313.00-00  TLVCONTENT
00:13:51    16       8      1  1313.1313.1313.00-00  TLVCONTENT

```

```

00:13:47      24      8      2 1212.1212.1212.01-00 NEWADJ TLVCONTENT
00:06:16      16      8      1                PERIODIC
00:04:27      16      8      4 1313.1313.1313.01-00 NEWLSP TLVCONTENT
00:01:58      20      9      3 1212.1212.1212.01-00 IPQUERY TLVCONTENT
00:01:43      24      8      4 1313.1313.1313.01-00 NEWADJ TLVCONTENT
p1r2#
    
```

Step 11 At the p1r2 router, use the proper **show** command to examine the IS-IS neighbor table. How many entries do you see? Are they all Level 1 adjacency?

The following example shows the output of the show clns is-neighbors command on the p1r2 router:

```

p1r2#show clns is-neighbors

System Id      Interface  State  Type Priority  Circuit Id      Format
1111.1111.1111 Se1        Up     L1    0         01              Phase V
1111.1111.1111 Se0        Up     L1    0         00              Phase V
1313.1313.1313 Et0        Up     L1    64        1313.1313.1313.0Phase V
p1r2#
    
```

There are three entries; all are Level 1.

Step 12 At the p1r1 router, use the proper **show** command to examine the IS-IS neighbor table. How many entries do you see? Do you see both Level 2 and Level 1 adjacency?

The following example shows the output of the show clns is-neighbors command on the p1r1 router:

```

p1r1#show clns is-neighbors

System Id      Interface  State  Type Priority  Circuit Id      Format
1212.1212.1212 Se1        Up     L1    0         01              Phase V
1212.1212.1212 Se0        Up     L1    0         00              Phase V
1313.1313.1313 Se2        Up     L1    0         00              Phase V
1111.2222.3333 Se3        Up     L2    0         00              Phase V
p1r1#
    
```

There are four entries. The adjacencies to the other routers in pod 1 are Level 1, and the adjacency to the backbone_r1 router is Level 2.

Step 13 Display the running configuration of the p1r1 router. Is **clns routing** enabled by default?

The following example output is from the p1r1 router. The clns routing command is on by default:

```

p1r1#show run
Building configuration...

Current configuration:
!
version 12.0
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
    
```

```
hostname p1r1
!
enable secret 5 $1$d1Qh$WrD8xeBBJLywzPs6AMZU0/
enable password sanfran
!
ip subnet-zero
no ip domain-lookup
ip host p1r1 192.168.1.17 192.168.1.33 192.168.1.49 10.1.1.1
ip host p1r2 192.168.1.65 192.168.1.18 192.168.1.34
ip host p1r3 192.168.1.66 192.168.1.50
ip host bbr1 10.1.1.100
c!ns routing
!
!
!
interface Ethernet0
no ip address
no ip directed-broadcast
shutdown
!
interface Serial0
bandwidth 64
ip address 192.168.1.17 255.255.255.240
no ip directed-broadcast
ip router isis
no ip mroute-cache
no fair-queue
clockrate 64000
!
interface Serial1
bandwidth 64
ip address 192.168.1.33 255.255.255.240
no ip directed-broadcast
ip router isis
clockrate 64000
!
interface Serial2
bandwidth 64
ip address 192.168.1.49 255.255.255.240
no ip directed-broadcast
ip router isis
clockrate 64000
!
interface Serial3
bandwidth 64
ip address 10.1.1.1 255.255.255.0
no ip directed-broadcast
ip router isis
!
interface BRI0
no ip address
no ip directed-broadcast
shutdown
!
router isis
summary-address 192.168.1.0 255.255.255.0
summary-address 172.26.1.0 255.255.255.0
net 00.0001.1111.1111.1111.00
!
ip classless
!
!
!
line con 0
```

```
exec-timeout 0 0
logging synchronous
transport input none
line aux 0
line vty 0 4
  password cisco
  login
!
end

p1r1#
```

Step 14 Save the current configurations of all the routers within your pod to NVRAM.

The following example shows how to save the configuration of the p1r1 router using the copy run start command (this is an abbreviated form of the copy running-config startup-config command):

```
p1r1#copy run start
Destination filename [startup-config]?
Building configuration...
p1r1#
```

NOTE

If you go on to do the Configuration Exercises in Chapter 5, “Configuring EIGRP,” of the Cisco Press book *Building Scalable Cisco Networks*, you will need to change setup Step 2 to

Step 2 Disable IS-IS on all the routers within your pod.

Review Questions

Answer the following questions, and refer to the online addendum to Appendix G, “Answers to the Review Questions,” for the answers.

- 1 What is common to OSPF and Integrated IS-IS?
- 2 How is the router identified in an IS-IS environment?
- 3 What is the difference between NSAP and NET?
- 4 What does a unique system ID define?
- 5 Which network representations are supported by IS-IS?
- 6 What is a pseudonode?
- 7 How do two Level 1 areas communicate?
- 8 How do systems find each other in IS-IS?
- 9 List the types of adjacencies between IS-IS systems.
- 10 How is IS-IS routing enabled on Cisco routers?