

This chapter covers the following topics, which you need to master to pass the CCNP/CCDP Routing exam. This list of exam objectives stated for the Building Scalable Cisco Internetworks (BSCI) exam (640-901) was obtained from the following web page:

http://www.cisco.com/warp/public/10/wwtraining/certprog/testing/current_exams/640-901.html

Cisco often updates these pages, and this URL may change. If you have problems, go to Cisco.com and follow the directions to the exam certifications.

- Explain basic OSI terminology and network layer protocols used in OSI
- Identify similarities and differences between Integrated IS-IS and OSPF
- Identify characteristics of an effective addressing plan for IS-IS deployment
- Explain how networks and interfaces are represented in IS-IS
- List the types of IS-IS routers and their role in IS-IS area design
- Describe the hierarchical structure of IS-IS areas
- Describe the concept of establishing adjacencies
- Describe the concepts of routing traffic transport and database synchronization
- Explain the basic principles of area routing
- Explain IS-IS NBMA (nonbroadcast multiaccess network) modeling solutions in switched WAN networks
- Given an addressing scheme and other laboratory parameters, identify the steps to configure Cisco routers for proper Integrated IS-IS operation
- Identify verification methods that ensure proper operation of Integrated IS-IS on Cisco routers

Configuring the Integrated IS-IS Protocol

The new BSCI exam covers the same material as the previous Building Scalable Cisco Networks (BSCN) exam with the addition of the Intermediate System-to-Intermediate System (IS-IS) protocol. The topics in this chapter detail the routing protocol Integrated IS-IS. This chapter assumes knowledge of routing protocols and, in particular, link-state routing protocols. This chapter introduces Integrated IS-IS by explaining the protocol's terminology and fundamental concepts. You learn both the basic operations of the protocol and its configuration.

This chapter is divided into two sections. The first part of the chapter deals theoretically with how the protocol works. The second part of the chapter covers how to implement and manage an Integrated IS-IS network.

This chapter addresses the exam topics listed on the chapter opening page, giving background information as required to explain an aspect of the protocol. Of course, Cisco maintains the right to change this list without notification and may add questions outside the areas stated, but using this list gives a basic foundation to the technology of Integrated IS-IS, which will prepare you for the questions posed.

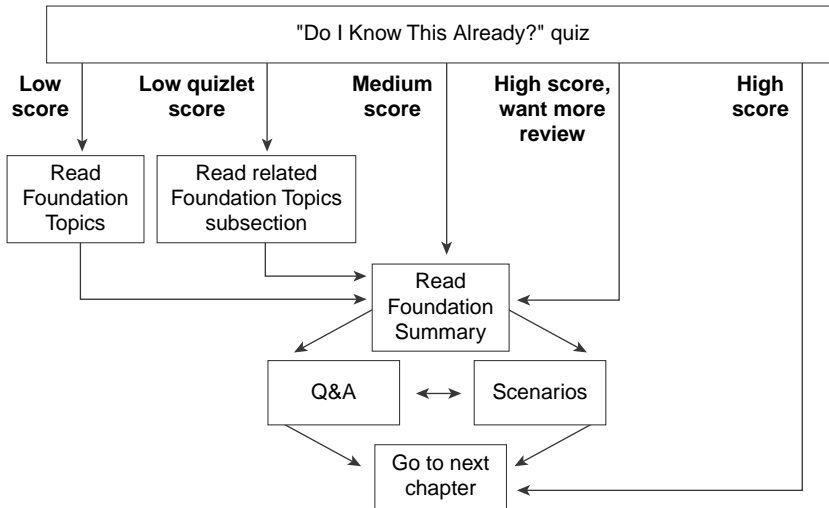
Because Integrated IS-IS is similar to the Open Shortest Path First (OSPF) protocol, you should read the chapters on OSPF and reinforce the fundamentals of link-state protocols. Differences between Integrated IS-IS and OSPF are clearly outlined within this chapter.

How Best to Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place, for easy reference.
- When you take a quiz, write down your answers. Studies show that retention significantly increases by writing down facts and concepts, even if you never look at the information again.
- Use the diagram in Figure S-1 to guide you to the next step.

Figure S-1 *How to Use This Chapter*



If you skip to the Foundation Summary, Q&A, and Scenarios sections and have trouble with the material there, you should go back to the Foundation Topics section.

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

This 16-question quiz helps you determine how to spend your limited study time. The quiz is sectioned into four smaller four-question “quizlets.” Figure S-1 outlines suggestions on how to spend your time in this chapter. Use Table S-1 to record your scores.

Table S-1 *Score Sheet for Quiz and Quizlets*

Quizlet Number	Topic	Questions	Score
1	Integrated IS-IS Addressing	1 to 4	
2	How Integrated IS-IS Discovers, Chooses, and Maintains Routes	5 to 8	
3	How to Configure Integrated IS-IS for Basic Operation	9 to 12	
4	How to Verify Integrated IS-IS Operation	13 to 16	
All questions	All	1 to 16	

1 What portion of the address does the initialism IDP identify?

2 If the NSEL is set to 0x00, what does this indicate?

3 How long is the AFI in a NSAP address and what function does it perform?

4 How long is the system ID defined in the NET on Cisco equipment?

5 How often, by default, does Integrated IS-IS send out Hello packets?

6 What is a neighbor in Integrated IS-IS?

7 What is an adjacency in Integrated IS-IS?

8 If the network is stable and sees no changes, how often will it send LSPs? Why are these updates sent out periodically?

9 What command is used to configure the Integrated IS-IS router process?

10 What is the default routing level on a Cisco router?

11 What command is used to configure Integrated IS-IS routing on the interface?

12 How is the NET address configured on the router?

13 What command is used to show the state of adjacencies on the router?

14 What command identifies the designated router for your LAN?

15 Explain briefly what **show isis database** reveals.

16 What command reveals the trigger for the last SPF calculation on the router?

The answers to this quiz are found at the end of this chapter in the “Answers to the Quiz Questions” section. The suggested choices for your next step are as follows:

- **2 or less on any quizlet**—Review the appropriate sections of the “Foundation Topics” portion of this chapter, based on Table S-1. Then move on to the “Foundation Summary” section, the “Q&A” section, and the “Scenarios” at the end of the chapter.
- **8 or less overall score**—Read the entire chapter. This includes the “Foundation Topics,” “Foundation Summary,” and “Q&A” sections and the “Scenarios” at the end of the chapter.
- **9 to 12 overall score**—Begin with the “Foundation Summary” section and then go to the “Q&A” section and the “Scenarios” at the end of the chapter. If you have trouble with these exercises, read the appropriate sections in “Foundation Topics.”
- **13 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and go to the “Q&A” section and the “Scenarios” at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Introduction to Integrated IS-IS

IS-IS is an Interior Gateway Protocol (IGP) developed in the 1980s by Digital Equipment and submitted to the International Organization for Standardization (ISO) as the routing protocol for Open System Interconnection (OSI). The creation of IS-IS was an attempt to produce a standards protocol suite that could allow internetworks to scale.

The development of IS-IS was motivated by the need for the following:

- A nonproprietary protocol
- A large addressing scheme
- A hierarchical addressing scheme
- A protocol that was efficient, allowing for fast, accurate convergence and low network overhead

The intention was to create one protocol suite that every system in the world could use. Unfortunately, to ensure a truly open standard, the ISO tried to accommodate every conceivable feature. The effort to include everything in IS-IS slowed the definition process and resulted in a relatively complex protocol.

The U.S. government tried to force OSI to be adopted as the international standard that it was conceived (as an initiative called the Government Open Systems Interconnections Profile [GOSIP]). The U.S. mandated that every system operated by the government had to be capable of running the OSI architecture. By forcing every government system to understand OSI, it was hoped that the protocols would become the practical standard as well as the academic solution to diverse proprietary implementations for computer networking. In the end, however, the initiative failed. The Internet, built on TCP/IP, prevailed as the practical alternative to an international standard. After all, it was there, and it worked.

However, IS-IS has always been used, although not as extensively as first hoped. Large Internet service providers (ISPs) have been using IS-IS since its inception in the 1980s, and recently it has begun to emerge in other markets. This new interest is for a variety of reasons, including the fact that IS-IS is a standard providing protocol independence, it has the capability to scale, and it has the capacity to define type of service (ToS) routing. ToS routing allows traffic engineering, which requires very complex routing decisions to be programmed into the protocol. Therefore, as of late, IS-IS has been taken from the shelf, dusted off, and put into use.

IS-IS Terminology

The terminology used by IS-IS may appear cumbersome and unfriendly. However, although the jargon may be unfamiliar, most of the concepts are no different from other routing protocols. Table S-2 explains briefly some of the commonly used IS-IS terms.

Table S-2 *IS-IS Terms*

Term	Definition
Adjacency	Local routing information that shows the reachability of a directly connected ES or IS. A separate adjacency is created for each neighbor on a circuit, and for each level of routing (that is, Level 1 and Level 2) on a broadcast circuit.
Administrative Domain	A group of routers that share the same routing protocol within one organization.
Area	A subdomain within an Administrative Domain. Routers in an area maintain detailed routing information about the area's internal composition. The routers also maintain routing information that allows them to reach other areas. The area address is contained in the NET and NSAP address.
Circuit	The local routing information for a single subnet point of attachment (SNPA).
Code/Length/Value (CLV)	These are the variable-length fields in a PDU. The Code field specifies the information in the Content field as a number. The Length field states the size of the Value field. The Value field contains the information itself.
Complete sequence number packet (CSNP)	CSNPs describe every link in the link-state database. CSNPs are sent on point-to-point links when the link comes up to synchronize the link-state databases. The designated router (DR), or designated intermediate system (DIS), on a multicast network sends out CSNPs every 10 seconds.
Connectionless Network Protocol (CLNP)	This is the ISO protocol used to carry data and error indications at the network layer. CLNP is similar to IP and has no facilities to detect errors in data transmission. It relies on the transport layer to provide guaranteed data delivery.

continues

Table S-2 *IS-IS Terms (Continued)*

Term	Definition
Connectionless Network Service (CLNS)	<p>CLNS uses a datagram transfer service and does not require a circuit to be established before data is transmitted.</p> <p>Whereas CLNP defines the actual protocol, CLNS describes a service provided up to the transport layer. Being a connectionless service, CLNP provides a “best effort” delivery of data; therefore, there is no guarantee that data will not be lost, corrupted, misordered, or duplicated. If you require guaranteed delivery, the transport layer or application layer needs to provide the service that will correct the problems when they arise.</p>
Designated intermediate system (DIS)	<p>The router (IS) on a LAN that is designated to perform additional duties. In particular, the DIS generates link-state PDUs on behalf of the LAN by treating the LAN as a pseudonode.</p>
Dual IS-IS	<p>IS-IS that supports both OSI and IP routing information. Areas within the autonomous system may either run OSI or IP or both. However, the configuration chosen must be consistent within the entire area.</p>
End system (ES)	<p>The end node or host, which has limited routing capabilities. The ES has the OSI or IP Layer 3 protocol running and can receive and send data.</p>
End System-to-Intermediate System (ES-IS)	<p>The protocol by which the OSI ES and the IS communicate to dynamically learn Layer 2 adjacencies.</p>
Hello	<p>Hello packets are used to discover and maintain adjacencies.</p>
Host address	<p>This is a subset of the NET address, which includes both the domain, area, and system ID.</p>
Integrated IS-IS	<p>Another term for Dual IS-IS. Indicates IS-IS can be used to support routing for two Layer 3 protocols (IP and CLNP) in the same network simultaneously.</p>

Table S-2 *IS-IS Terms (Continued)*

Term	Definition
Intermediate system (IS)	A router. The IS is a device capable of directing traffic to remote destinations.
Intermediate System-to-Intermediate System (IS-IS)	The OSI routing protocol that learns the location of the networks within the autonomous system so that data may be forwarded to the remote hosts.
IS-IS domain	A group of routers running the IS-IS protocols for exchanging routing information.
Level 1 (L1)	These routers are internal to the area, which means that they receive routing information for their area only and have no knowledge of the other areas' networks. To reach other areas, Level 1 routers maintain a default route to the nearest Level 2 router.
Level 1-2 (L1-2)	A router that connects areas. This router connects a Level 1 area to the Level 2 backbone. It will have a Level 1 routing table to route to ES and IS in its own area by system ID. It will maintain a Level 2 prefix table to route to other areas.
Level 2 (L2)	These routers are connected only to the backbone and provide transit traffic between areas.
Link	A physical connection to a neighbor. This link is then transmitted to all the other routers in the area via the LSP.
Link-state packet (LSP)	A packet that describes a router's links. There are separate LSPs for Level 1 and Level 2 updates.
Neighbor	A router on the same link with which an adjacency is formed and routing information is then exchanged.
Network entity title (NET)	Part of the OSI address. The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which defines the NSAP address of the system.
Network protocol data unit (NPDU)	See PDU.

continues

Table S-2 *IS-IS Terms (Continued)*

Term	Definition
Network selector (NSEL)	Sometimes referred to as the SEL field. This field describes the service at the network layer that the packet is to be sent. NSEL is similar to the Protocol field in IP.
Network service access point (NSAP)	Describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the SEL field set to a value other than 0x00.
Overload (OL) bit	The OL is set on an LSP if the router cannot store the entire link-state database. When other routers receive LSPs with this bit set, they will not send the router any transit traffic for fear that its routing table is incomplete. If the router is making decisions using incomplete data, its decisions may result in suboptimal paths or even routing loops. Traffic destined for the router can still be sent to the directly connected interfaces of a router transmitting the OL bit in its LSPs.
Partial sequence number packet (PSNP)	PSNPs are sent on point-to-point links to explicitly acknowledge each LSP the router receives. A router on a broadcast subnetwork sends a PSNP requesting the LSPs it needs to synchronize its link-state database.
Protocol data unit (PDU)	A unit of data passed from one layer of the OSI model to the same level of the OSI model on another node. Each layer prefixes the PDU to indicate the sending OSI layer so that the network layer sends NPDU and the data-link layer sends DLPDU.

Table S-2 *IS-IS Terms (Continued)*

Term	Definition
Pseudonode	The LAN identifier for a broadcast subnetwork. The pseudonode makes the broadcast medium appear as a virtual router and the routers appear as connected interfaces. The adjacency to a pseudonode is represented by the system ID of the DIS plus the circuit ID of the link. The pseudonode reduces the amount of resources needed to manage the adjacencies as the DIS generates link-state PDUs on behalf of the pseudonode. These LSPs are sent to all the connected ISs. The routers connected to the broadcast medium no longer have to maintain the adjacencies to all other routers on the medium (thus reducing memory and CPU and bandwidth resources).
Routeing Domain	Routeing Domain is the same as the Administrative Domain. It defines the boundaries of a network of interconnected routers operated and managed by the same administrative group. The spelling of Routeing Domain is not a typographic error but the British spelling adopted by the ISO committee.
Sequence number PDU (SNP)	SNPs are used to acknowledge the receipt of LSPs and to synchronize link-state databases.
Subnetwork	The data-link layer.
Subnetwork dependent layer	Interfaces with the data-link layer and hides the different kinds of data-link layers from the network layer. This sublayer transmits and receives PDUs from the subnetwork, translates DLPDUs into NPDUs, and hands them to the appropriate OSI process. The subnetwork dependent layer is also responsible for creating and maintaining adjacencies through the exchange of IS-IS Hello PDUs.

continues

Table S-2 *IS-IS Terms (Continued)*

Term	Definition
Subnetwork independent layer	Interfaces with the transport layer and provides it with network services. It describes how CLNS creates and maintains knowledge of the network by exchanging and processing routing information so that data may be transmitted efficiently to remote destination hosts and handed to the transport layer.
Subnetwork point of attachment (SNPA)	The data-link layer offers two services: the physical connection to the medium and the services offered to the physical layer and network layer. The SNPA refers to these services. The SNPA address is the physical address (for example, the MAC address on a LAN).
Type/Length/Value (TLV)	TLV is the same as a CLV, but some literature refers to the variable-length fields as Type/Length/Value in accordance with the IP terminology.

Similarities Between Integrated IS-IS and OSPF

Integrated IS-IS and OSPF share a common heritage. They are both link-state protocols and are based on the Dijkstra algorithm of Shortest Path First (SPF). In addition, they both have a two-level hierarchy. OSPF tends to be deployed mostly as an enterprise solution; whereas, Integrated IS-IS is used for IP routing in some ISP networks.

Table S-3 outlines other similarities between Integrated IS-IS and OSPF.

Table S-3 *Integrated IS-IS and OSPF Similarities and Terminology Comparison*

IS-IS Terminology	OSPF Terminology
Area	Stub area
Area ID	Area ID
Backbone area	Backbone area
DIS (designated intermediate system)	Designated router
Domain	Network
ES (end system)	Host
ES-IS (the address resolution feature of ES-IS)	ARP (Address Resolution Protocol)
IS (intermediate system)	Router

Table S-3 *Integrated IS-IS and OSPF Similarities and Terminology Comparison (Continued)*

IS-IS Terminology	OSPF Terminology
ISO Routing Domain	Autonomous system
Level 1	Internal nonbackbone stub area
Level 1-2	Area border router (ABR)
Level 2	Backbone router
LSP (link-state packet)	LSA (link-state advertisement)
CSNP and PSNP (complete and partial sequence number PDUs)	Link-state acknowledgement packet
PDU (protocol data unit)	Packet
NET (network entity title)	IP destination address (subnet and host), used in a similar way to router ID
NSAP (network service access point)	IP destination address + IP protocol number
Routing technology = link state	Routing technology = link state
<ul style="list-style-type: none"> • Classless routing protocol • Address summarization between areas • Uses a link-state database • Acknowledges LSPs • Shortest path is computed using Dijkstra's SPF algorithm • Hellos create and maintain adjacencies • Hellos and holdtime may be configured 	<ul style="list-style-type: none"> • Classless routing protocol • Address summarization between areas • Uses a link-state database • Acknowledges LSAs • Shortest path is computed using Dijkstra's SPF algorithm • Hellos create and maintain adjacencies • Hellos and holdtime may be configured
Subnet = data link	Subnet = IP network
SNPA (subnetwork point of attachment)	Layer 2 address; for example, the MAC address (Media Access Control) or DLCI (data-link connection identifier) in Frame Relay
System ID	Address of the host within the network, sometimes used as the router ID determining priorities
Virtual link (defined but not supported)	Virtual link

Differences Between OSPF and Integrated IS-IS

Although OSPF and Integrated IS-IS share the same common goals and use the same link-state technology to achieve them, the methods they use differ slightly. For example, the protocols differ in how the area address is assigned. In IS-IS the area and host address are assigned to the

entire router; whereas, in OSPF the address is assigned at the interface level. An IS-IS router is therefore in one area, while in OSPF a router may inhabit many areas.

This means that all Level 1 routers have to be within the same area, with a Level 1-2 router connecting them to another area. However, the Level 1-2 router needs to be in the same area as the Level 1 router with which it communicates. The Level 1-2 router can see the rest of the autonomous system and offers itself as the default route to the Level 1 area. This is very similar to the OSPF stub areas. The Level 2 router sends Level 2 updates in the other area, or prefix routes, just like the ABR in OSPF.

The role of the DR is subtly different. The DIS in IS-IS will exist for both Level 1 and Level 2 on multiaccess media, but there is no backup designated router (BDR). Also, in OSPF the DR is elected for life; in IS-IS, however, if another router comes on line with a higher priority, the existing DIS is deposed. Fewer adjacencies are formed in OSPF because the routers form adjacencies only with the DR and the BDR; whereas, in IS-IS every router makes an adjacency with every other router on the medium. However, IS-IS LSPs are sent out only by the DIS on behalf of the pseudonode.

A major difference is the encapsulation of the two protocols. IS-IS has protocol independence because it runs directly on top of the data-link layer. Fragmentation is the responsibility of the IS-IS process, and this allows for a streamlined protocol. More importantly, it makes it very simple to evolve the protocol as needed because it is not dependent and therefore is not limited by any third protocol. OSPF is encapsulated into IP and therefore is limited by the capabilities of that protocol.

The way that the LSPs are handled is also slightly different and influence, to an extent, the design of a network running either protocol. Unrecognized LSPs are ignored and flooded in IS-IS, for example; whereas, OSPF ignores and drops unrecognized LSAs.

Table S-4 lists the key differences.

Table S-4 *Integrated IS-IS and OSPF Technical Differences*

Technology	Integrated IS-IS	OSPF
Areas	Boundaries defined on the link.	Area boundaries are defined on the router.
	A router (IS) can be in one area.	Interfaces may be in different areas.
		A router may be in many areas.
		The Integrated IS-IS Level 1 area is similar to an OSPF stub area.

Table S-4 *Integrated IS-IS and OSPF Technical Differences (Continued)*

Technology	Integrated IS-IS	OSPF
Designated router (DR)	<p>If a new IS becomes active with the same or higher priority it becomes the new DIS, which results in a flood of LSPs.</p> <p>Adjacencies are created with all ISs on the broadcast media.</p> <p>Each IS sends a multicast LSP to all ISs on the media. The LSP is unacknowledged.</p>	<p>A same or higher priority does not dislodge the existing DR.</p> <p>Adjacencies on broadcast media are formed with the DR and BDR only.</p> <p>All LSAs are acknowledged.</p>
Encapsulation	<p>Integrated IS-IS runs on top of the data link layer (Layer 2).</p> <p>Integrated IS-IS is a network layer protocol with its own Layer 3 packet.</p> <p>Fragmentation is the responsibility of Integrated IS-IS.</p>	<p>OSPF is an IP application.</p> <p>Has an OSPF header and travels inside an IP packet.</p> <p>Fragmentation is the responsibility of IP.</p>
LAN flooding	<p>All ISs maintain adjacencies with all other ISs on a broadcast network.</p> <p>DIS sends CSNP to all ISs.</p> <p>Periodic CSNPs ensure the databases are synchronized.</p>	<p>Multicast updates and Hellos sent to DRs.</p> <p>Unicast acknowledgment sent from the DRs.</p>
LSAs	<p>Two types of LSP.</p> <p>LSPs are CLV encoded.</p> <p>Unrecognized LSPs are ignored and flooded.</p> <p>LSPs are always flooded across all media by the originating IS.</p>	<p>Seven types of LSA.</p> <p>Unrecognized LSAs are not flooded.</p> <p>Many small LSAs for summary and external updates.</p> <p>LSA updates generated by each router.</p>

ISO Addressing for Integrated IS-IS

Because IS-IS is the product of a committee, it has the feel of an academic solution that is intended to resolve every eventuality. Its addressing scheme thinks not just globally, for instance, but universally. The large address space is one characteristic that is responsible for its new popularity.

When Integrated IS-IS is routing IP traffic, the routing information is carried in the IS-IS updates; therefore, the participating routers need an ISO address. The ISO address comes in two forms: the NSAP and the NET address (depending on the type of device being addressed). The addressing scheme of IS-IS is long and complex, but clear rules define the address space, allowing the IS-IS protocol to find the routers to forward the data traffic to the end destination.

The IS-IS address is a variable-length address from 8 to 20 octets in length, much larger than the 4 octets in the TCP/IP address. ISO 10589 defines three parts to the address—Area, ID, and SEL—as shown in Figure S-2.

Figure S-2 *The Three Sections of an ISO Address*

AREA	ID	SEL
------	----	-----

IS-IS addressing can become complicated because these three fields are subdivided to allow greater granularity in routing. If you're confused, however, just remember that it all condenses to the three elements of the address, which perform the following functions:

- **Area**—The Area field is used to route between areas using Level 2 routing.
- **ID**—The ID field is used to route to a host or router within the area using Level 1 routing.
- **SEL**—The SEL field is used to route to an entity within the host or ES.

Although the three parts to the address describe how to get to the area and how to find the host and then the application within the host, finding the destination host only uses the first two parts; the last part of the address is used after the end host has received the packet. Therefore, Integrated IS-IS has two levels of hierarchy, basically how to get to the area and then the host. However, the first part of the address, routing to the area, may require many decisions to be made. These decisions may involve determining the country and then the organization and many suborganizations. So although the address has only two levels of hierarchy, many levels of hierarchy may be hidden within the first part of the address, which states how to get to the autonomous system and forms the external portion of the area.

Figure S-3 shows the division between external and internal routing—the Initial Domain Part (IDP) and the Domain Specific Part (DSP)—and how this division is broken down to reveal greater granularity for routing. In the figure, the three original elements have been placed underneath as reference.

Figure S-3 *An ISO Address*

IDP		DSP		
AFI (1 octet)	IDI	High Order DSP	System ID (1-8 octets)	NSEL (1 octet)
AREA			ID	SEL

The following list explains the parts of an ISO address and their relationship:

- **IDP**—External routing. The IDP is used to route to the domain, or autonomous system. The IDP is given by the ISO and identifies the organization responsible for assigning the format for the rest of the address by defining the DSP structure. The IDP comprises the following two parts:
 - **AFI (authority and format identifier)**—AFI is the first octet of the address. It is defined as 1 octet and it states the authority in charge of issuing addresses and the format of the rest of the address, within the constraints of IDI and DSP, to describe Area, ID, and SEL.
 - **IDI (initial domain identifier)**—IDI is the suborganization to the AFI; for example, the U.S. government is the AFI, and the General Services Agency is the IDI, otherwise known as GOSIP.
- **DSP**—DSP is used to route within the autonomous system. The authority specified by the IDP may further subdivide the DSP and may assign subauthorities responsible for parts of the DSP. The DSP can be broken down into the following fields:
 - **High Order DSP**—High Order DSP is typically the area within the autonomous system
 - **System ID**—System ID may have a value between 1 to 8 octets. It must have the same value throughout all systems in the autonomous system. Cisco uses six octets, which is a common solution because it allows the use of the MAC address to autoconfigure the system.
 - **NSEL**—NSEL identifies the particular service at the network layer to which to hand the packet.

The confusion is that the address can take so many forms; remember, however, that there are only two layers of hierarchy. By providing such flexibility in the address space, the ISO has ensured a decentralized address allocation and management as well as the ability to scale the network.

As with TCP/IP, the addressing scheme within an autonomous system may be the creative genius of the network administrator or the addressing scheme can be obtained from the AFI, an authorized ISO body such as ANSI or GOSIP.

NETs and NSAP

NETs and NSAPs are ISO addresses. The differences between the NET and NSAP addresses are subtle. The NET address is the address of the host, where the value in the NSEL field is set to 0x00. Therefore, there is no upper-layer protocol identified within the host. With no application identified within the end host, the packet can be routed to the destination, but it cannot be handed off to a process after it has been delivered. However, routers do not have upper-layer protocols to identify because they are transitory ISs. Therefore, the NSAP of the router is referred to as a NET as the NSEL field is set to 00.

The NSAP is the full ISO address. It defines not only the area and destination host within the area, but it also specifies where to send the incoming packet after it has reached the host. The NSEL field at the end of the ISO address specifies the upper-layer protocol and is similar to the Protocol field of the IP header.

Rules of ISO Addressing

The following list indicates a few rules that clarify ISO addressing:

- The ISO address is assigned to the system, not to the interface.
- Typically, the router has one NET address. The limit of three NETs is in a conventional IS-IS implementation; the limit is three NETs per area in a multiarea Integrated IS-IS implementation. Multiple addresses are used during transitions.
- If multiple NETs are configured on the same router, they must all have the same system ID.
- The area address must be the same for all routers in the same area.
- All Level 2 routers must have a system ID that is unique for the entire domain.
- All Level 1 routers must have a system ID that is unique for the entire area.
- The system ID must be the same length for all ISs and ESs within a routing domain.

Example of a NET Address

The following are examples of NET addresses. The first two addresses are used for routing within the autonomous system and therefore the IDI portion of the address has not been defined:

- A simple OSI NET address, using the host MAC address as the system ID:

47.0005.aa00.0301.16cd.00

← **To the Domain** →← **Within the Domain** →

AFI	IDI	HO-DSP	System ID	SEL
47.		0005.	aa00.0301.16cd.	00
Area			System ID	SEL

- A simple OSI NET address, using the host's loopback IP address of 144.132.16.19 as the system ID:

47.0001.1441.3201.6019.00

← **To the Domain** →← **Within the Domain** →

AFI	IDI	HO-DSP	System ID	SEL
47.		0001.	0001.1441.3201.6019.	00
Area			System ID	SEL

- A GOSIP version 2 address, showing the external routing information:

47. 0005.80ff.f800.0000. 0001.0000.0c00.1234.00

← **To the Domain** →← **Within the Domain** →

AFI	IDI	HO-DSP	System ID	SEL
47.	0005.80ff.f800.0000.	0001.	0000.0c00.1234.	00
Area			System ID	SEL

Integrated IS-IS Hierarchical Structure

As the addressing structure shows, there are potentially many levels of hierarchy within the addressing scheme, even though the protocol has only two layers of hierarchy. To accommodate the two levels of hierarchy, two types of router are defined: a Level 1 router that deals with the first level of routing, finding the end destination within the area, and a Level 2 router that finds the area within which the end destination resides. Both of these routers are combined in the Level 1-2 router, which runs both the Level 1 and the Level 2 processes and could be viewed as a third type of router.

The Level 1 Router

The Level 1 router locates the destination host within the area and as such is known as the *intra-area router*. It is similar to a stub router in OSPF because its knowledge of the network is limited to the area, using a default route to the nearest Level 2 router as the means of routing traffic external to the area.

Every Level 1 router has a link-state database containing all the routing information for the area. The nature of the link-state protocol means that each database is identical. Because the database is limited to routes within the area, there is a requirement that neighbors must be in the same area so that the routers can communicate.

The Level 2 Router

To route traffic between areas, a Level 2 router is needed. Routing between areas is referred to as *interarea routing*. This router is similar to a backbone internal router in OSPF, and as in OSPF the backbone must be contiguous. Level 2 routers communicate via Hellos that are understood only by other Level 2 routers. To fracture the area means the loss of routing information. As in Level 1 routing, the link-state database is identical on all Level 2 routers, although the database contains prefixes of addresses in other areas as opposed to internal area addresses.

The Level 1-2 Router

The router that has everything is of course the Level 1-2 router. This is both an intra- and interarea router. Its characteristics are similar to those of an ABR in OSPF. This router may have neighbors in different areas because it sends both Level 1 and Level 2 Hellos and can thus communicate with everyone. It holds both a Level 1 database for each area to which it is connected and a Level 2 database with all the information for interarea routing.

The Level 1-2 configuration is very convenient because the router informs Level 1 routers that it is a Level 2 router and can forward traffic to other areas as well as inform other Level 2 routers of the areas to which it is connected. Although convenient, it consumes more resources on the router in terms of memory and CPU and also utilizes more bandwidth in maintaining the link-state databases. This configuration is the default configuration on Cisco routers.

Basic Principles of Area Routing

Having reviewed the complexities of ISO addressing and the corresponding hierarchical nature of Integrated IS-IS routing, it is necessary to pull it all together. This section explains the basics of Integrated IS-IS routing. The following list shows what happens when a router receives a packet to forward and how decisions are made in area routing based on the ISO destination address of the incoming packet:

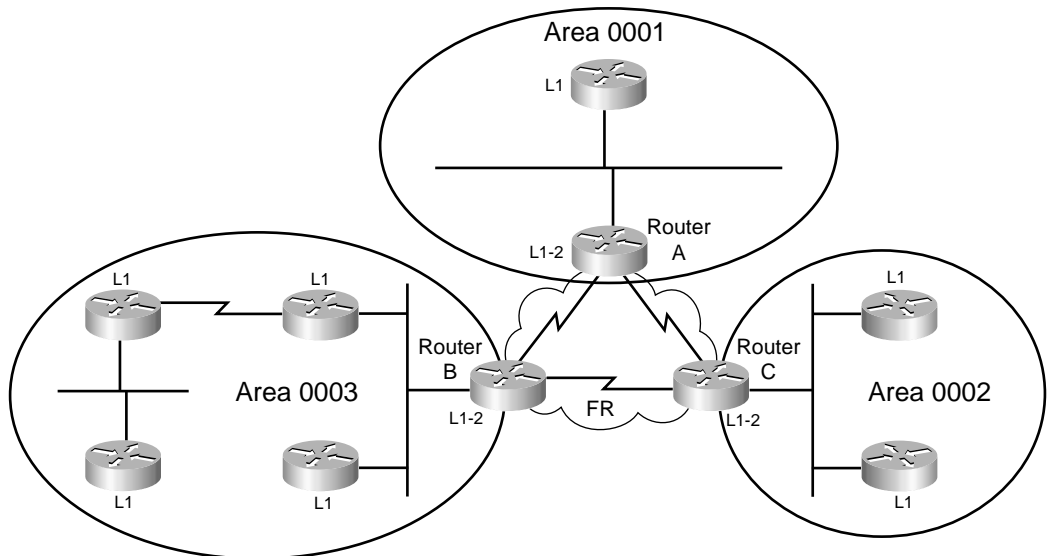
- 1 When a router receives traffic to route to a remote destination, it performs the customary routing table lookup.
- 2 The router strips off the system ID and the SEL to reveal the Area portion of the address. If the Area address is the same as that of the router, it routes the packet toward the host system ID using the Level 1 database.
- 3 If the Area is different, the router does one of the following:
 - (a) Sends the packets to the nearest Level 2 router (if the router is a Level 1 router).
 - (b) Looks up the route in the forwarding database to find a match, if the router is a Level 2 router.

- (c) Resolves the address to the longest match, ensuring the greatest level of accuracy. Summarization, as described in Chapter 3, “IP Address,” is used by Integrated IS-IS, and therefore prefix routing can reduce the routing tables and speed up the forwarding process.

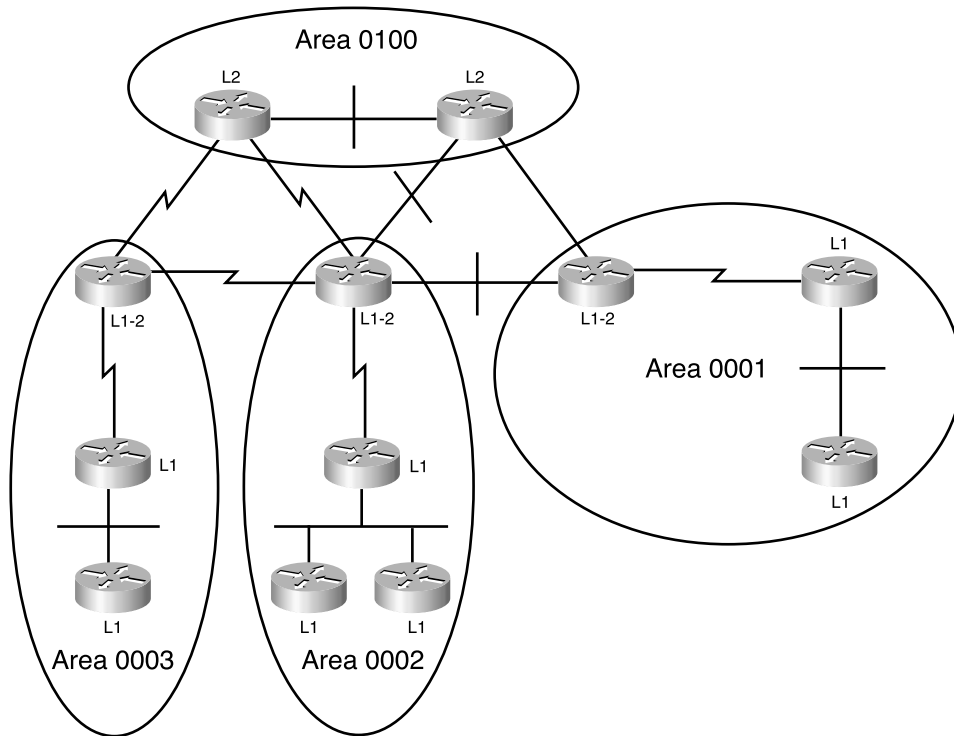
The decisions made in area routing are explained in more depth in the section titled “Integrated IS-IS Operation.”

Areas in Integrated IS-IS are defined on the link, requiring the router to define the area and not the interface as in OSPF. In Figure S-4 shows how backbone routers are used to connect different areas.

Figure S-4 *Areas and Backbone Routers*



For the Level 2 routing updates to be exchanged, all the routers capable of sending Level 2 updates must be contiguous; Figure S-5 shows this.


Figure S-5 *A Contiguous IS-IS Backbone*

Because Level 1 routers communicate only to other Level 1 routers and likewise Level 2 routers communicate only between themselves, it can get confusing. Figure S-6 shows the logical communications or adjacencies set up between the routers shown in Figure S-5.

Integrated IS-IS Networks and Interfaces

Knowledge of the network is formed by the exchange of routing information between the routers. Essentially, every router tells the routers connected to it who it is, its interfaces, and to whom it is connected. After every router shares this information, all routers and their attached networks are known and a detailed topology map is created.

Routers sharing a common data-link layer become Integrated IS-IS neighbors if the Hello packets that they exchange meet the criteria for forming an adjacency. Although the process of finding a neighbor differs slightly depending on the media, the information sent in the Hellos is essentially the same. Each Hello states the originator of the Hello and the capabilities of its interface. If each receiving router's interface has the same capabilities, an adjacency is created. After an adjacency has been formed, the Integrated IS-IS neighbors exchange routing information in the form of LSPs. This is explained in more detail in the section on PDU format called "Network Layer Protocols Used in Integrated IS-IS."

Figure S-6 *Level 1 and Level 2 Adjacencies*

For an adjacency to be formed and maintained, both interfaces must agree on the following:

- The maximum packet size (MTU) that can be transmitted by the interface must be the same.
- Each router needs to be configured as the same level of routing—that is, either Level 1 or Level 2—so that they can decode the Hellos sent by the other router.
- If the routers are both Level 1, they must be in the same area.
- Level 1 routers form adjacencies with each other, and Level 2 routers form adjacencies with other Level 2 routers. For a Level 1 router to communicate with a Level 2 router, one of the routers needs to be configured as a Level 1-2 router. Therefore, to connect to another area at least one of the routers must be configured as a Level 1-2, allowing the Level 2 router to receive the packets destined for another area from the Level 1 router.
- The system ID must be unique to each router.
- If authentication is configured, it must be configured identically on both routers.
- The Hello timers (including the holddown timer) must match. If one router has a Hello timer of 40 seconds, the defaults on the other router would time out the holddown timer and purge the LSP, resulting in a flapping link and endless SPF calculations. Although this misconfiguration does not affect the neighbors being established, the link dies as soon as the timers expire.

Integrated IS-IS defines two network types—broadcast subnetworks and point-to-point networks—compared to the five types used in OSPF. A broadcast network, as in OSPF, is a multiaccess data link that supports broadcasts and multicasts. The point-to-point links are deemed to be nonbroadcast and may be permanent virtual circuits (PVCs) or dynamically set up switched virtual circuits (SVCs).

The following sections describe adjacencies on point-to-point links, broadcast links, and nonbroadcast multiaccess (NBMA) links.

Establishing Adjacencies on a Point-to-Point Link

A point-to-point link connects two routers. After a Hello packet has been received, each side declares the other side reachable. The routers are neighbors. At this point, each side sends a CSNP. The CSNP is a list of all the links held in the links-state database, which triggers a synchronization of the link-state database on each machine. This process is described in detail in the section “The Update Process.”

Periodic Hellos maintain the adjacency. If a router does not hear a Hello within the Hello holdtime, the router is declared dead and the database purged of any entries associated with the router. Cisco sets the default Hello multiplier to three. The holdtime is defined as the Hello time multiplied by the Hello multiplier, making the holdtimer expire every 30 seconds.

Establishing Adjacencies on a Broadcast Link

On broadcast links, all the routers running Integrated IS-IS receive packets sent by one router—the DIS—to control the amount of traffic that needs to be generated to maintain the adjacencies and thus the databases. The DIS has the responsibility of flooding the LSPs to all connected systems running Integrated IS-IS. More accurately, the DIS floods the LSPs for the pseudonode.

The pseudonode represents the LAN, with each router simulating an interface on the imaginary router. This imaginary router is called the pseudonode. As if it were a real router, the pseudonode floods a new pseudonode LSP when there is a change in the status of its connections (for example, when a new neighbor comes online).

The adjacencies with the other routers are maintained by the DIS, which sends out Hellos every 3.3 seconds, three times the speed of other routers. This is to ensure the integrity of the adjacencies. If there is a problem with the DIS, or a router with a higher priority appears, it is quickly identified and a new router is elected in the place of the old Integrated IS-IS, which is forced into retirement. The election is based on priority and stated on the interface in the event that all the interfaces are set at the default of 64. Then, the highest numeric SNPA (data-link address) determines the DIS.

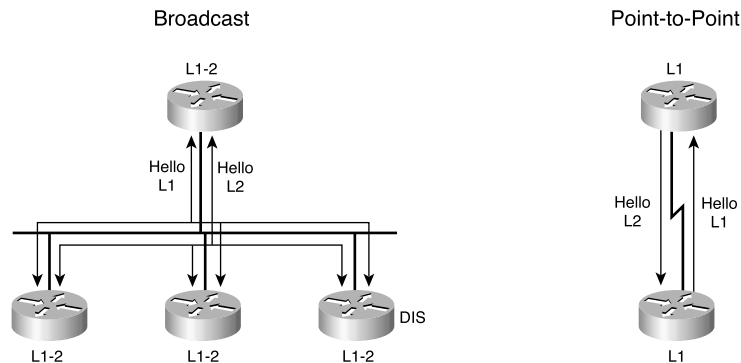
Establishing Adjacencies on a Nonbroadcast Multiaccess (NBMA) Link

The creation and maintenance of adjacencies on both point-to-point and multiaccess is straightforward. It becomes more complicated when the format of the technology cloud allows either link technology to be used.

A nonbroadcast multiaccess (NBMA) link is neither a broadcast medium nor is it a point-to-point link; it is a bit of both. Frame Relay, ATM, and X.25 are examples of NBMA. Using PVCs, NBMA provides multiple connections, which could be viewed as a LAN. The confusion occurs when Integrated IS-IS sees the link is multiaccess. Having no knowledge of multiaccess WAN clouds, Integrated IS-IS believes that the medium is some form of LAN and therefore has broadcast capabilities. However, the medium is a WAN, not a LAN, and although this can be simulated, the WAN cloud has no inherent broadcast capabilities.

To avoid complications and possible errors, Cisco recommends that the links be configured as a series of point-to-point links. Figure S-7 shows a comparison of the broadcast and point-to-point topologies supported by IS-IS.

Figure S-7 *Network Topologies Supported by IS-IS*



Network Layer Protocols Used in Integrated IS-IS

The PDU is created by the network layer and encapsulated directly into the data-link frame. All the Integrated IS-IS packets share the same 8-octet header. After the fixed header, there are a number of optional variable-length fields that contain specific routing related information. These variable-length fields are called Type/Length/Value (TLV) or Code/Length/Value (CLV).

The fields found in the fixed header of every IS-IS PDU are each 1 octet in length. Table S-5 shows the fixed header.

Table S-5 *The Fixed Header Common to All IS-IS PDUs*

Field	Length of Field in Octets	Description
Intradomain Routing Protocol	1	All Integrated IS-IS PDUs have a value of 0x83. This identifies the packet.
Length Indication	1	States the length of the fixed header.
Version/Protocol ID	1	Set to 1.
ID Length	1	The size of the system ID in the NSAP. This can be an integer between 1 and 8. The Cisco default setting is 6, which is represented by 0, to show that the default has not been changed.
Reserved/Packet Type	1	The first 3 bits are reserved, set to 0, and ignored. The Packet Type indicates whether this is a Hello, LSP, or SNP.
Version	1	Set to 1.
Reserved	1	Set to 0 and ignored.
Maximum Area Addresses	1	States the number of area addresses permitted for this area. Multiple area addresses are configured during transitions. Cisco set the maximum address size to 3, which is represented by 0.

After this common header are the PDU-specific fields and the variable-length fields, creating the different packet types seen in Integrated IS-IS. The PDU-specific fields are discussed in the sections relevant to the different packet types—Hellos, LSPs, and SNPs. The variable-length fields are discussed in the section on TLVs.

There are three Integrated IS-IS packets, as the following list describes:

- **Hello**—These packets create and maintain neighbor relationships and adjacencies. There are three types of Integrated IS-IS Hello packet. The type of packet is defined in the fixed header under the Type field and allows the packet to be handed off to the appropriate process. The different types are as follows:
 - **LAN Level 1**—Type 15. Generated by Level 1 and Level 1-2 routers
 - **LAN Level 2**—Type 16. Generated by Level 2 and Level 1-2 routers
 - **Point-to-point**—Type 17. Generated by Level 1, Level 2, and Level 1-2 routers

- **Link-state packet (LSP)**—LSPs hold information on the neighbors connected to the router. There are two types of LSP, as follows:
 - **Level 1**—Type 18. Generated by Level 1 and Level 1-2 routers.
 - **Level 2**—Type 20. Generated by Level 2 and Level 1-2 routers
- **Sequence number packet (SNP)**—SNPs describe the LSPs in the transmitting router's link-state database. The information is condensed and is never flooded but only sent between neighbors. SNPs ensure link-state databases synchronization by:
 - Distributing groups of LSPs on a LAN without explicit individual acknowledgements
 - Acknowledging individual LSPs
 - Requesting LSPs at startup

There are two types of SNP for each level of routing, as follows:

- **Complete SNP(CSNP)**—Includes every LSP in the database:
 - Level 1**—Type 24
 - Level 2**—Type 25
- **Partial SNP (PSNP)**—Includes a subset of LSPs, used to request individual LSPs and to acknowledge receipt of these LSPs:
 - Level 1**—Type 26
 - Level 2**—Type 27

The following sections describe Hello packets, LSPs, and SNPs in more detail.

The Format of the Hello Packet

There are three different Hellos—point-to-point Hello, LAN Level 1 Hello, and LAN Level 2 Hello. Because the point-to-point and broadcast media work differently, the adjacencies are formed in a different manner. There needs to be separate Hellos for the point-to-point networks and the broadcast networks. A point-to-point network has only one other router to communicate with and, therefore, after the level of routing has been established, the appropriate updates can be sent. However, a broadcast network is a multiaccess network and can have a mixture of both Level 1 and Level 2 routers. For this reason, the broadcast or LAN network has two Hello formats, the Level 1 format and the Level 2 format. The Hellos for the broadcast media are referred to as LAN Hellos.

Point-to-point Hello packets are used over point-to-point links. Table S-6 shows the point-to-point Hello.

Table S-6 *The Point-to-Point Hello*

Field	Length of Field in Octets	Description
Fixed Integrated IS-IS header	8	Common to all Integrated IS-IS PDUs
Circuit Type	1	States whether the transmitting router is a Level 1, Level 2, or Level 1-2
Source ID	ID length	The system ID from the NSAP of the transmitting router
Holding Time	2	How long the neighbors must wait for a Hello before they can declare the transmitting router to be dead
Packet Length	2	The length of the Hello packet in octets
Local Circuit ID	1	Identifier for the transmitting interface, which is unique to the transmitting router

LAN Hello packets are used over broadcast links. The PDU is the same for all router types, although the values within the fields differ.

Table S-7 shows the LAN Hello.

Table S-7 *LAN Hello Packet Format*

Field	Length of Field in Octets	Description
Fixed Integrated IS-IS header	8	Common to all Integrated IS-IS PDUs.
Circuit Type	1	States whether the transmitting router is a Level 1, Level 2, or Level 1-2.
Source ID	ID length	The system ID from the NSAP of the transmitting router.
Holding Time	2	How long the neighbors must wait for a Hello before they can declare the transmitting router to be dead.
Packet Length	2	The length of the Hello packet in octets.
Priority	2	Used in the election of the DIS (highest priority wins).
LAN ID	ID length + 1	The DIS uses its system ID plus an additional octet to name the LAN. The extra octet is used to identify the LAN from other LAN connections on the DIS.

The Format of the LSP

The LSP from a Level 1 router is flooded to all routers in the area. The LSP contains a list of all the adjacencies.

Likewise, a Level 2 router floods the LSP to every other Level 2 router in the domain. However, this LSP contains the list of adjacencies to other Level 2 routers and the areas that the transmitting router can reach. The TLVs hold the Level 1 and Level 2 information, allowing the LSP format to be the same for both Level 1 and Level 2 routers.

Table S-8 shows the format of the LSP.

Table S-8 *LSP Packet Format*

Field	Length of Field in Octets	Description
Fixed Integrated IS-IS header	8	Common to all Integrated IS-IS PDUs.
Packet Length	2	The length of the entire LSP.
Remaining Lifetime	2	How long in seconds before the LSP is purged from the database.
LSP ID	ID length + 2	A three-part ID: <ul style="list-style-type: none"> • 6 octets for the transmitting system ID or pseudonode ID of the transmitting DIS • 1 octet = zero for router 1 octet = nonzero for pseudonode • 1 octet = fragment bit, if the LSP information is too great to be contained in one PDU
Sequence Number	4	Used to determine the latest LSP version.
Checksum	2	The checksum on the contents of the LSP.

continues

Table S-8 *LSP Packet Format (Continued)*

Field				Length of Field in Octets	Description
P	ATT	OL	IS Type	1	<p>Together the Partition (P) bit, Attached (ATT) bit, Overload (OL) bit, and IS Type take up 1 octet.</p> <p>P—The Partition bit is used by Level 2 routers to identify whether automatic partition repair is supported. Cisco does not support this feature and the value is always 0.</p> <p>ATT—The Attached bit is used in Level 1 LSPs generated by Level 1-2 routers. It shows Level 1 routers a potential exit from the area. Level 1 routers use this bit to determine the closest Level 2 router.</p> <p>Overload bit (OL)—If the generating router has run out of memory for the link-state database, this bit is set to 1. Other routers treat this router as a host, forwarding packets destined to networks directly connected to routers setting the OL bit. If the router has run out of memory for the link-state database, the database may be incomplete and unreliable as a forwarding router.</p> <p>IS Type is used to indicate whether the router is a Level 1 or Level 2 router.</p>

The Format of the SNP

As with the LSP, the SNP format is the same for both Level 1 and Level 2 routers. The CSNP has two additional fields to the PSNP. The extra fields define the address range, which states all the LSPs that are included in the CSNP. This is to ensure that if the database is so large that several PDUs need to be sent, the information can be reassembled at the receiving router.

Table S-9 shows the format of the SNP.

Table S-9 *SNP Packet Format*

Field	Length of Field in Octets	Description
Fixed Integrated IS-IS header	8	Common to all Integrated IS-IS PDUs.
Packet Length	2	The length of the entire LSP.
Source ID	ID length + 1	The system ID from the NSAP of the transmitting router.
Start LSP ID	ID length + 2	Refer to the definition of LSP ID in Table S-8.
End LSP ID	ID length + 2	Refer to the definition of LSP ID in Table S-8.

TLVs

TLVs, sometimes called CLVs, are one of the main strengths of Integrated IS-IS. TLVs provide flexibility and extend the functionality of the protocol. The Type/Length/Value fields are variable in length and are appended to the various packet formats. This means that the protocol can adapt to the changing needs and advances in technology by defining a new TLV.

The structure of the TLV is as follows:

- **Type or Code**—State the Type of field. This code identifies the TLV and all the characteristics that pertain to it. (For example, TLV 128 defines the capability to carry IP routes in IS-IS packets.) In essence the TLV 128 is Integrated IS-IS.
- **Length**—The length of the following field in octets. This is important because the next field may be any length. By identifying the length of the field, error detection is possible.
- **Value**—The information, be it IP routes, IS-IS neighbors, or authentication.

One of the most radical TLVs was the introduction of TLV 128, as defined in RFC 1195 in 1992. This extended Integrated IS-IS to support IS-IS and IP routing. In the future, it is expected that a TLV will be created to support IPv6.

To read the definition of these TLVs and in particular the various codes, you must read the standards. ISO/IEC 10589 defines the TLV codes 1 to 10; whereas, RFC 1195 defines the TLV codes 128 to 133.

It is important to note which TLVs are supported by equipment because this determines the design and configuration of the network. The receiving router ignores TLVs that are not supported.

Table S-10 describes the main TLVs, states the document that defines them, and names the PDU(s) to which they are appended.

Table S-10 *TLVs Implemented by Cisco*

TLV	Name	Source Document	Used by PDU	Description
1	Area Addresses	ISO 10589	Hello and LSP Levels 1 and 2	The area addresses configured on the router. Cisco supports a maximum of three area addresses.
2	IS Neighbors	ISO 10589	LSP Level 1 and 2	Lists the neighbors of the transmitting router. The neighbor ID is the system ID plus an extra octet. If the neighbor is a pseudonode, the last octet has a positive value. If the neighbor is a router, the last octet has a value of 0x00. Of the many options within this TLV, the default metric is the only one supported and contains a value between 0 and 63.
3	ES Neighbors	ISO 10589	LSP Level 1	Similar to IS neighbors, assumes the same cost to many ES neighbors as connected via broadcast media. ES neighbors with different costs appear as separate entries.
5	Prefix Neighbors	ISO 10589	LSP Level 2	The same as an ES neighbor but stating an address prefix or domain and area rather than an ES.
6	IS Neighbors	ISO 10589	Hello	Lists all the system IDs of the routers from whom a valid Hello has been received. Level 1 routers list Level 1 neighbors; whereas, Level 2 routers list Level 2 neighbors.
8	Padding	ISO 10589	Hello	Used to pad a Hello PDU so that it can be a minimum length. The padding is ignored. Cisco uses all 0s to pad packets.
9	LSP Entries	ISO 10589	SNP	The state of an LSP is given via the Remaining Lifetime, LSP ID, Sequence Number, and Checksum. This identifies each LSP and ensures no duplication or corruption.

Table S-10 *TLVs Implemented by Cisco (Continued)*

TLV	Name	Source Document	Used by PDU	Description
10	Authentication	ISO 10589	Hello, LSP Level 1 and 2, and SNP	The first octet states the type of authentication used. Of the 255 possible values, only one type of authentication—clear text—is defined in ISO 10589. An interface may be configured with a transmit password. If the receiving interface does not receive the expected password, the Hello is rejected.
128	IP Internal Reachability Information	RFC 1195	LSP Level 1 and 2	Lists all IP addresses that the transmitting router knows about from interfaces within the area.
129	Protocols supported	RFC 1195	Hello	States the protocols the transmitting router supports (whether CLNS only, IP only, or both CLNS and IP).
131	Inter-Domain Routing Protocol Information	RFC 1195	LSP Level 2	Allows information from external routing protocols to be carried in Level 2 LSPs. Similar to route tags in RIPv2, EIGRP, and OSPF.
132	IP Interface Address	RFC 1195	Hello and LSP Level 1 and 2	The IP address or addresses of the transmitting interface.

Integrated IS-IS Operation

Before you can fully grasp the details of IS-IS and how it routes traffic, you must understand the routing protocol operation. A high-level overview of the Integrated IS-IS operation follows:

- 1 Routers send Hellos out of all Integrated IS-IS interfaces to discover neighbors and form adjacencies.
- 2 Routers sharing a common data link become neighbors.
- 3 Routers build LSPs based on local Integrated IS-IS interfaces and prefixes learned from other adjacent routers. These are sent to all neighbors.
- 4 Routers flood received LSPs to all adjacent routers except to the neighbor from whom the LSP was received.
- 5 When new or different LSPs are received, the router builds a link-state database incorporating those LSPs.

- 6 The router calculates the SPF for each destination and constructs the Shortest Path Table (SPT) and the forwarding database.

The routing process is divided into four stages: the update, decision, forwarding, and receive processes.

The Update Process

The router can forward data packets to the remote destination only if it has a sound understanding of the network topology. It can then determine the best path to the remote network and send the data packet on the next stage of its journey.

To ensure that every router has the information required, each router generates its own LSP. The LSP lists the router's neighbors, which are propagated throughout the network. The flooding of LSPs ensures every router has the same set of LSPs and that the link-state databases are identical.

LSPs are generated whenever there is a change in the network, often due to a configuration change. However, any of the following instances trigger a new LSP to be flooded throughout the network:

- An adjacency either comes up or down (for example, a new router comes online).
- An interface on the router changes state or is assigned a new metric.
- An IP route changes (for example, due to redistribution).

During the flooding process, a router both receives and sends LSPs. The action of each instance has clear rules of engagement.

Sending and Receiving an LSP

On receipt of the new LSP, the router stores it in the link-state database and marks it for flooding. If the LSP is already present in the database, the router just acknowledges and ignores it. If it is a new LSP, the router creates an LSP describing its own adjacencies. The router sends the new LSP and the duplicated LSP to its neighbors, who in turn flood to their neighbors and onward to the very edges of the network. Because Level 1 and Level 2 routers have their own link-state databases, Level 1 LSPs are flooded throughout the area; whereas, Level 2 LSPs are sent across all Level 2 adjacencies.

The process of propagating LSPs differs slightly, depending on which medium the LSP was received. The following sections describe the propagation process for point-to-point and broadcast links.

Propagating LSPs on a Point-to-Point Interface

The point-to-point link does not need to ensure that multiple systems have synchronized databases. With only one other router to work with, some reliance is given to the router's capability to determine the need to update and thus bandwidth may be optimized.

The following list describes the flooding process on a point-to-point link:

- When an adjacency is established, both sides send a CSNP packet with a compressed version of their link-state database.
- If there are any LSPs in the receiving router's database that were not in the CSNP it received, it sends a copy of the missing LSPs to the other router.
- Likewise, if the database is missing any LSPs received in the CSNP, the receiving router requests the detailed LSP to be sent.
- The individual LSPs are requested, sent, and acknowledged via PSNPs.
- When an LSP is sent, the router sets a timer and if no explicit acknowledgement has been received before the timer expires, the LSP is re-sent. This timer is the **minimumLSPTransmission-interval** and may be configured; the default on a Cisco router is five seconds.

Propagating LSPs on a Broadcast Link

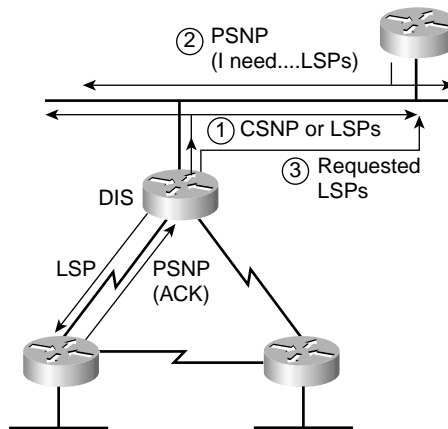
A broadcast link may need to send out both Level 1 and Level 2 updates. The link sends these updates using multicast MAC addresses to all Level 1 routers and all Level 2 routers. The DIS takes on much of the responsibility for synchronizing the databases on behalf of the pseudonode. The DIS has three tasks, as follows:

- Creating and maintaining adjacencies
- Creating and updating the pseudonode LSP
- Flooding the LSPs over the LAN

The following lists the main steps in the flooding process:

- On receipt of a CSNP, the router compares each compressed LSP with the link-state database.
- If the database has a newer version of the LSP sent in the CSNP or if there is no instance of the LSP in the CSNP, the router multicasts the LSP onto the LAN.
- If the database is missing an LSP that was in the CSNP, it sends a PSNP requesting the full LSP. Although the router multicasts, it is only the DIS that takes action.

Figure S-8 summarizes the flow of CSNPs and PSNPs on broadcast and point-to-point links.

Figure S-8 *The Propagation of CSNPs and PSNPs*

Determining Whether the LSP in the Database Is Valid

The LSP contains three fields that help determine whether the LSP that has been received is more recent than that held in the database and whether it is intact or has been corrupted. These three fields are as follows:

- Remaining Lifetime**—This is used to age-out old LSPs. If an LSP has sat in the database for 20 minutes, it is assumed that the originating router has died. The refresh timer is set to 15 minutes.

If the lifetime expires, the LSP has the content removed, leaving only the header. The lifetime is set to show that it is a new LSP, and then it is flooded through the network. All receiving routers accept the mutilated LSP, recognize that this means the route is bad and purge the existing LSP from their database.
- Sequence Number**—This is an unsigned 32-bit linear number. The first LSP is allocated the sequence number 1, and the following LSPs are incremented by 1.
- Checksum**—If a router receives an LSP and the checksum does not compute correctly, the LSP is flushed and the lifetime is set to 0. The router floods the LSP, all routers purge the LSP, and the originating router retransmits a new LSP. This recovery action is the same as that of expiring lifetime.

The Decision Process

After the link-state databases have been synchronized, it is necessary to decide which path to take to reach the destination. Because the routers and hosts may have multiple connections to each other, there may be many paths from which to choose.

To make this decision as to the best path, link-state protocols employ the algorithm defined by Dijkstra. This algorithm is extremely complex, but its core function is to create a tree that mirrors the shortest paths to all destinations. Each router builds a shortest path tree with itself as root. This is achieved by taking all the LSPs from the link-state database and using the Dijkstra algorithm to create the shortest path tree (SPT). The SPT is used in turn to create the forwarding table, also known as the routing table.

Several tables are created during this process. The PATH table is the shortest path tree during construction, and the TENT is a tentative database (a scratchpad).

The steps by which the router's SPT and the forwarding database are built are as follows:

- Step 1** Place self at root (0, 0) in PATH.
- Step 2** The SPF takes each LSP from the link-state database and selects the best or shortest path by considering the metric for each path in turn. The lowest metric is selected as the shortest path.
- Step 3** For the node just placed in the PATH table, look at LSP. Add the metric cost of getting to the node to the cost shown in LSP (the cost from the node to the destination advertised). If this path is not already in PATH or TENT with a better cost, put in TENT.
- Step 4** If the TENT table is empty, stop.
- Step 5** If the TENT table is not empty, find the lowest cost entry and move the PATH to the TENT table.
- Step 6** Go to Step 2.

If there is more than one path to a remote destination, the criteria by which the lowest cost paths are selected and placed in the forwarding database are as follows:

- 1 If there is more than one path with the lowest value metric, Cisco equipment places up to six equal-cost paths into the table. The default number of equal-cost paths is four.
- 2 Optional metrics are chosen before the default metric, but because Cisco supports only the default metric, this is a moot point.
- 3 Internal paths are chosen before external paths because going outside the autonomous system is likely to be a suboptimal route and may be the result of a routing loop.
- 4 Level 1 paths within the area are more attractive. If the path is within the area, it is not only more efficient to route directly to it, but going outside the area and returning may be the cause of a routing loop, demanding greater resources and time.

- 5 The address with the longest match or most specific address in IP is the address with the longest IP subnet mask. This ensures that the closest router is chosen because prefix routing is configured by summarization that can occur only on area boundaries.
- 6 If type of service (ToS) has been configured, the path with ToS is chosen over those paths with no ToS.
- 7 If multiple paths have ToS configured, the path with the shortest route is selected.
- 8 If the ToS is the same, up to six entries are placed into the forwarding database and the routers load balance between them.
- 9 If there is no path, the forwarding database sends the packet to the nearest Level 2 router, which is the default router.

Metrics or Cost

The metric defines the cost of the path. Integrated IS-IS has four metrics, only one of which is required, which is good because only one is supported. The metrics defined in ISO 10589 are as follows:

- **Default**—Sometimes referred to as cost. Every Integrated IS-IS router must support this metric. Cisco set the default for all interfaces to be 10.
- **Delay**—This optional metric reflects the transit delay.
- **Expense**—This optional metric reflects the monetary expense of the network.
- **Error**—The reliability of the path is determined as the metric.

The ISO metrics are configured on the outgoing interface using an integer value between 0 and 63. A 10-bit field describes the total path to a destination, allowing a value between 0 and 1023. Considering this to be inadequate, Cisco has increased the metric size to be 24 bits. The default configuration is using the 6-bit or narrow metric.

In determining the shortest path, the lowest metric is chosen, internal paths are chosen over external paths, and Level 1 routes have precedence over Level 2 routes.

The default or cost metric is the only metric supported by Cisco because each metric used in Integrated IS-IS requires a different link-state database calculation, for both the Level 1 and Level 2 routes.

Other rules have little to do with the decision process but are responsible for ensuring the LSPs are valid. Essentially, Integrated IS-IS ignores incomplete fragments. Any fragmented packet entering the router is ignored if there is no starting fragment with the packet header (because this indicates packet loss or corruption). The O fragment also defines whether the OL bit is set, the Integrated IS-IS type, and whether there are entries in the Area Address options field. Although this does not eliminate the reassembly of corrupted packets, it certainly reduces the load.

The Forwarding Process

The IP subnets are carried in the LSPs in the IP Reachability TLV (code 128) and are treated in the same manner as the OSI addresses. The decision process for building the SPT is the routing decision.

After the SPT has been built, the forwarding database can be created. The forwarding table is essentially a lookup table for the longest match, where load balancing occurs on multiple equal-cost paths. The forwarding table for Integrated IS-IS is more relevant to CLNS than to IP because the IP routing information is entered directly into the IP routing table, where IP routes are leaves to the tree of IS-IS.

Although the forwarding database is concerned solely with forwarding the data, it selects the longest match and load balances across equal-cost links.

The Receive Process

If the frame is valid, it passes user data and error reports to the forwarding process, while routing information Hellos, LSPs, SNPs are sent to the update process.

The receive process is not described in depth here because it is primarily concerned with CLNS routing and does not concern IP. A detailed description is given in the ISO 10589 standards document.

Integrated IS-IS Design Considerations

Optimizing the network depends on careful planning and design. Although each network is constrained by physical and technical limitations, you must strive to design your network to meet the needs of its users and accommodate the demands of various applications.

In Integrated IS-IS, the fundamental design considerations are areas and addressing. The next sections describe the Integrated IS-IS design considerations of area design, route summarization, and NBMA modeling solutions.

Area Design of Integrated IS-IS Routers

When designing a network, you have to compromise. Typically, the trade-off is between reliability and speed. What is most efficient for the network is determined by the requirements of the network and the resources available.

In designing Integrated IS-IS networks and the hierarchical design, you need to consider the data flow in addition to the resources required by the routing protocol. Tuning the update process may be sufficient, although this results in the compromise of resources and reliability. If you reduce the update timers, the databases converge more quickly, but the network could be depleted of necessary resources to route data.

In area design, two SPF databases need to be maintained, which requires the use of additional resources, not only in terms of CPU and memory, but also in bandwidth for the propagation of LSPs.

Some typical designs include the following:

- A totally flat network that uses only Level 1 routing. This design will not scale because any change in the network requires a flood of LSPs to all routers, which consequently run the SPF algorithm.

However, this simplified design has some advantages: there will be only one SPF database and no problem of suboptimal routing between areas.

- A totally flat network using only Level 2 routing. As the network expands, Level 1 areas may be added. The Level 2 area has complete internal knowledge with the streamlined advantages of one SPF instance. The Level 1 area may well be a company connecting to the ISP, in which case the areas would be added as new customers came online.
- A totally flat network using the Cisco default of every router running Level 1-2 routing. This allows for easy migration to a hierarchical design and overcomes the problem of suboptimal routing. However, this design requires resources to maintain two SPF databases.
- A hierarchical network where the core is running Level 2 routing, with Level 1 areas connecting into the core. Level 1-2 routers are used to connect the areas. This design results in additional resources and the possibility of suboptimal routing. Configuring the metric for the outbound interface determines the routing decision. This requires a careful knowledge of the topology of the network to ensure that the problem is not compounded. Recent developments have allowed for route leaking to be created, which allows the Level 2 routers to pass some specific routing information into Level 1. This facilitates the optimal routing decision to be selected.

By default Cisco routers running Integrated IS-IS operate as Level 1-2 routers. You can configure the router to reflect a true Level 1 and Level 2 design, but the default operation overcomes many potential pitfalls, including the following:

- **Area partitions**—In hierarchical routing between areas, if there is a loss of connection between the Level 1 router and the Level 2 router, the subarea could be completely annexed unless there is an alternative route.
- **Loss of data**—If the area becomes partitioned, a section of the network becomes unavailable.
- **Suboptimal routing decisions**—Level 1 areas have knowledge only of networks within their own areas. To reach another area, packets are sent to the nearest Level 2 router. Without additional configuration, the Level 1 router determines the nearest Level 2 router to be the one with the lowest hop count. The metrics used are the default metric of 10 on each outbound interface; therefore, the best route translates to that with the lowest hop count. As you know, the router two hops away may include 16-Mbps Token Ring and a 56-kbps link as opposed to the 3 hops of Fast Ethernet and ATM.

Route Summarization

As discussed in Chapter 3, summarization has many advantages. Summarization reduces the resource needs of the network and hides network problems within an area and thus the SPF calculations. The more details the router knows of the network, the more it must do to maintain an accurate knowledge of that network; that is, anytime it fails to hear an LSP, however remote and small its detail of the network, the entire network must flood LSPs and recalculate the SPT. Summarization allows the areas to manage the internal knowledge of the network and to summarize that knowledge across area boundaries.

The summarization command is given in the section “Configuring Integrated IS-IS.” The rules for summarization using OSPF in multiple areas apply to Integrated IS-IS.

The rules for summarizing IP routes in Integrated IS-IS are as follows:

- Level 1-2 routers can summarize the routes held within their area. The summarized routes are propagated to a Level 2 router. This is an efficient method of establishing prefix routing into other areas. The summarization is configured at the Level 1-2 router at the edge of the area. This is similar to the OSPF summarization by the ABR.
- If one Level 1-2 router has route summarization configured, you must configure route summarization on every Level 1-2 router injecting updates into Level 2. If one router is advertising without summarization, all traffic destined for the area is sent to that router because Integrated IS-IS uses longest-match routing. Therefore, the router advertising the more specific route is the chosen.
- Level 1 routes cannot be summarized within the area because this is not permitted by the protocol.

Integrated IS-IS NBMA Modeling Solutions in a Switched WAN Network

In the good old days, say eight or more years ago, most organizations had a broadcast network, such as Ethernet, for the LAN and point-to-point links, such as a serial link, for the WAN. The broadcast link was multiaccess and the point-to-point link connected two systems, typically over a very low-bandwidth link. Accommodating these different technologies when building a routing protocol made clear sense, but technologies have developed in complexity since then, particularly in the WAN environment.

Frame Relay and ATM are examples of nonbroadcast multiaccess networks, which are not accommodated in Integrated IS-IS. OSPF has a point-to-multipoint configuration option, but Integrated IS-IS does not. The solution in Integrated IS-IS is to configure the link as multipoint, allowing the election of a DIS. This is discussed in the section “Configuring Integrated IS-IS.”

The alternative is to configure the interfaces with subinterfaces that are point-to-point. The following section provides a sample configuration.

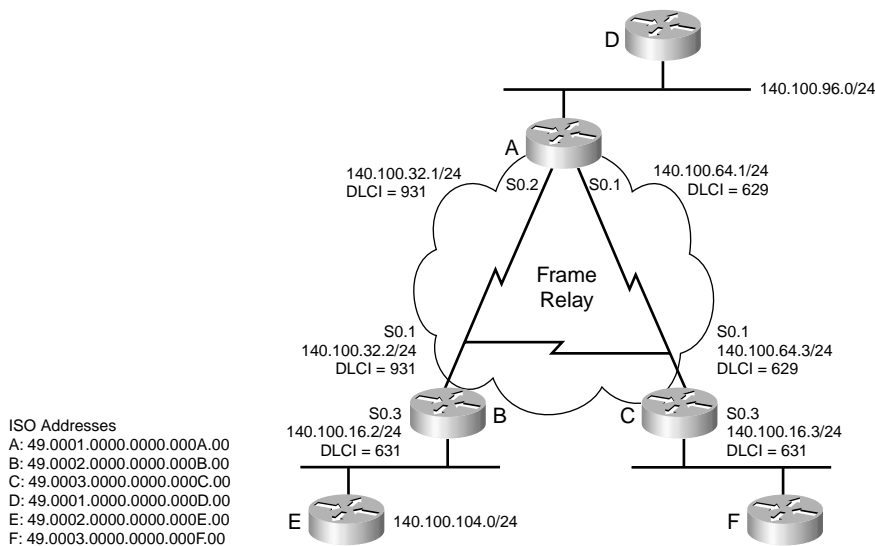
Configuring Integrated IS-IS

The preparation for configuring any routing protocol requires a sound understanding of the network topology and a coherent addressing scheme. When you have these, the basic configuration of Integrated IS-IS is as follows:

- Step 1** Enable the router process Integrated IS-IS with the **router isis** command.
- Step 2** Configure the NET address, thus assigning the area with the **net network-address** router subcommand.
- Step 3** Enable Integrated IS-IS for IP on the relevant interfaces with the **ip router isis** interface subcommand.

Figure S-9 illustrates a simple network to support the working configuration example. The figure shows the topology of the network and the addressing scheme that was used.

Figure S-9 Basic Integrated IS-IS Configuration



Example S-1 shows the basic configuration required to run Integrated IS-IS. The relevant commands are highlighted for easy identification, and a brief explanation of the code is inserted after an exclamation mark to make the code easier to read. Note the router process is started and the NET address defined immediately beneath this router process because the IS-IS address is

assigned to the router and as such is a part of the routing process. IS-IS is started on the interface after the IS-IS router process is defined with a NET address.

Example S-1 *Basic Integrated IS-IS Configuration*

```

Working Example for Router A
interface Ethernet0
 ip address 140.100.96.1 255.255.255.0
!assign the IP address and mask
 ip router isis
!start IS-IS on the interface
!
interface Serial0
 no ip address
 encapsulation frame-relay
 no fair-queue
 clockrate 56000
!
interface Serial0.1 point-to-point
 ip address 140.100.64.1 255.255.255.0
!assign the IP address and mask
 ip router isis
!start IS-IS on the interface
 frame-relay interface-dlci 629
!
interface Serial0.2 point-to-point
 ip address 140.100.32.1 255.255.255.0
!assign the IP address and mask
 ip router isis
!start IS-IS on the interface
 frame-relay interface-dlci 931
!
!
router isis
 net 49.0001.0000.0000.000a.00
!Start the IS-IS routing process and define the ISO address for the entire router.
!
 ip classless

```

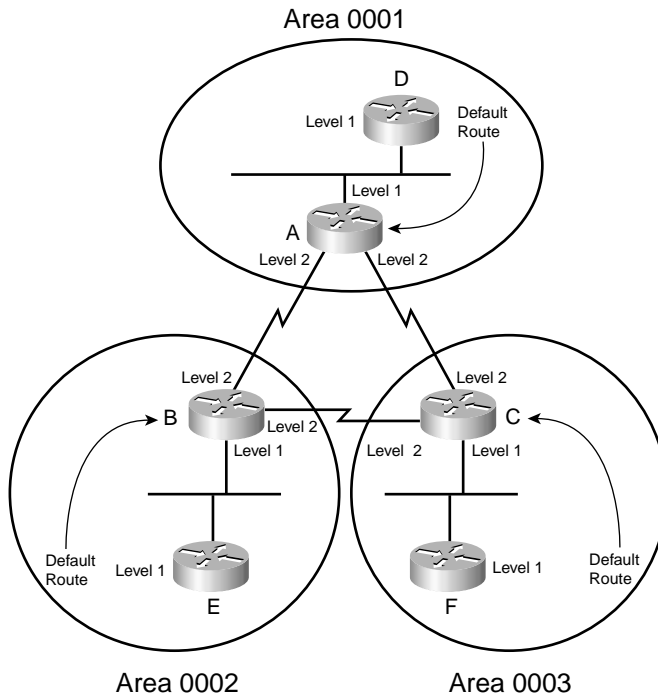
Remember that by default Cisco routers are configured as Level 1-2 routers to avoid suboptimal routing. Also the **cls routing** command is not required for IP-only IS-IS. In some versions of the Cisco IOS software, **cls routing** is written into the configuration file automatically by Cisco, when the routing process for IS-IS is configured.

Changing the Router from Level 1-2

Changing the routing level in Integrated IS-IS is done to overcome Integrated IS-IS consuming resources on both the routers and the media. Using the topology in Figure S-10, it is possible to configure routers D, E, and F as Level 1 routers because they are internal to their areas. The

command is just **is-type level-1**. This is a subrouter command and as such is placed under the **router isis** command.

Figure S-10 *Changing the Level of Routing Both for the Entire Router and at the Interface Level*



It is also possible to change the level of routing to Level 1 at the interface level by issuing the **isis circuit-type level-1** command under the interface. In Figure S-10, this is done on the Routers A, B, and C, the Ethernet interface pointing is configured to be running Level 1 routing, while the serial interface runs only Level 2 routing.

Level 1-2 routers send LSPs with an ATT bit in their Level 1 LSPs, indicating that they are attached to another area. This creates a default route in the Level 1 router, pointing to the nearest Level 2 router. This allows the nearest Level 2 router to be the transit router for all data destined for another area.

As illustrated in Figure S-10, Example S-2 shows Router A configured as Level 1 on the Ethernet interface pointing to Router D; whereas, the other interfaces are configured as Level 2-only. The relevant commands are highlighted for easy identification.

Example S-2 *Changing Router A Routing Level on an Interface Basis*

```
Working Example for Router A
interface Ethernet0
 ip address 140.100.96.1 255.255.255.0
 ip router isis
 isis circuit-type level-1
 !Configure Level 1 routing on the interface
 !
interface Serial0
 no ip address
 encapsulation frame-relay
 no fair-queue
 clockrate 56000
 !
interface Serial0.1 point-to-point
 ip address 140.100.64.1 255.255.255.0
 ip router isis
 frame-relay interface-dlci 629
 isis circuit-type level-2-only
 ! Configure Level 2 routing on the interface
 !
interface Serial0.2 point-to-point
 ip address 140.100.32.1 255.255.255.0
 ip router isis
 frame-relay interface-dlci 931
 isis circuit-type level-2-only
 !Configure Level 2 routing on the interface
 !
 !
router isis
 net 49.0001.0000.0000.000a.00
 !
 ip classless
```

Example S-3 shows the configuration command to change the level of routing on Router D to be Level 1. Note that this command appears under the **router isis** command, changing the behavior of the entire process rather than at the interface level. This example is illustrated in Figure S-10. The relevant commands are highlighted for easy identification.


Example S-3 *Changing the Level of Routing for the Entire Router*

```
Working Example for Router D
interface Ethernet0
 ip address 140.100.96.2 255.255.255.0
 no ip directed-broadcast
 ip router isis
 !
interface Serial0
 no ip address
 no ip directed-broadcast
 no ip mroute-cache
 shutdown
 no fair-queue
 !
interface Serial0.1
 no ip directed-broadcast
 !
interface Serial1
 no ip address
 no ip directed-broadcast
 shutdown
 !
interface BRI0
 no ip address
 no ip directed-broadcast
 shutdown
 !
router isis
 net 49.0001.0000.0000.000d.00
 is-type level-1
 !Define the router as a Level 1 router
 !
 ip classless
```

Configuring Summarization

Configuring summarization is straightforward. The complexity is in the addressing scheme design. To configure summarization, enter the **summary-address** command followed by the summary address and mask under the router process. Figure S-11 shows the topology used to support the working example shown in Example S-4.

In Figure S-11, the IP routes from Router B are summarized into areas 0001 and 0003. Router B is capable of summarizing routes because the router is both a Level 1 and Level 2 router, acting as an ABR in OSPF. Example S-4 shows Router B summarizing routes 140.100.104.0 through to 140.100.107.0 with a mask summarized from /24 to /22. The relevant commands are highlighted for easy identification.

Figure S-11 Summarizing IP Addresses Between IS-IS Areas**Example S-4** Summarization of IP Routes from Router B

```

Working Example for Router B
interface Ethernet0
 ip address 140.100.104.1 255.255.255.0
 no ip directed-broadcast
 ip router isis
 isis circuit-type level-1
 !
interface Serial0
 no ip address
 no ip directed-broadcast
 encapsulation frame-relay
 no ip mroute-cache
 no fair-queue
 clockrate 56000
 !
interface Serial0.2 point-to-point
 ip address 140.100.32.2 255.255.255.0
 no ip directed-broadcast
 ip router isis
 isis circuit-type level-2-only
 frame-relay interface-dlci 931
 !
interface Serial0.3 point-to-point
 ip address 140.100.16.2 255.255.255.0

```

continues

Example S-4 *Summarization of IP Routes from Router B (Continued)*

```


no ip directed-broadcast
ip router isis
isis circuit-type level-2-only
frame-relay interface-dlci 631
!
!
router isis
summary-address 140.100.104.0 255.255.252.0
! Routes 140.100.104.0-140.100.107.0 summarized with the /22 bit mask.
net 49.0002.0000.0000.000b.00
!
ip classless

```

Configuring NBMA

Frame Relay is a WAN technology that has evolved beyond the point-to-point capabilities. When configured in a full mesh, the cloud is multiaccess, although each circuit is discrete and therefore not a true broadcast medium. The configuration for IS-IS over the fully meshed Frame Relay cloud is illustrated in Figure S-12, and there is a working example of the configuration in Example S-5.

In Figure S-12, the Frame Relay cloud has three fully meshed routers, which can operate as a LAN as far as IS-IS PDUs are concerned.

Figure S-12 *NBMA Frame Relay Cloud Running Broadcast Integrated IS-IS*

In Example S-5, the **frame-relay map ip** command maps the IP destination address to the outgoing DLCI and defines the interface as a broadcast interface. Integrated IS-IS uses the links as if they were truly a broadcast link and elects a DIS.

The **frame-relay map clns** command maps to the CLNS process on the destination router. Without the second command, no routes appear in the IP routing table because CLNS does not receive the frames to populate the IP routing table. Remember that these are IP routes carried in the IS-IS routing protocol. It is IS-IS that updates the IP routing table.

Example S-5 *NBMA Frame Relay Cloud, Running Broadcast Integrated IS-IS*

```

Working Example for Router A
!
interface Ethernet0
 ip address 140.100.96.1 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
interface Serial0
 ip address 140.100.64.1 255.255.255.0
 ip router isis
 encapsulation frame-relay
 clockrate 56000
 frame-relay map clns 629 broadcast
!Map the layer 2 DLCI address in Frame Relay to the clns process of Router C
 frame-relay map clns 931 broadcast
!Map the layer 2 DLCI address in Frame Relay to the clns process of Router B
 frame-relay map ip 140.100.64.2 931 broadcast
!Map the layer 2 DLCI address in Frame Relay to the Destination IP address of Router
B
 frame-relay map ip 140.100.64.3 629 broadcast
!Map the layer 2 DLCI address in Frame Relay to the Destination IP address of Router
C

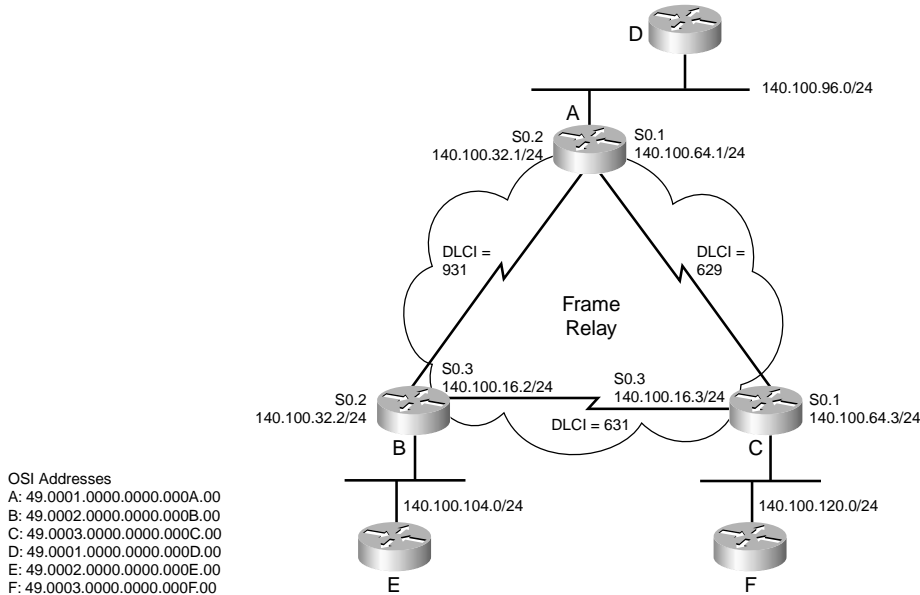
 isis circuit-type level-2-only
!
interface Serial1
 no ip address
 shutdown
!
!
router isis
 net 49.0001.0000.0000.000a.00
!
 ip classless

```

The alternative solution is to define subinterfaces and to configure each subinterface as point-to-point. This requires an IP subnet per link. Because the link is point-to-point, there is no need to configure **frame-relay map** commands because there is no choice of destination to define. The point-to-point link is just a pipe that goes to one destination. As shown in the configuration, it is only necessary to configure the interface as point-to-point, start Frame Relay, and define the DLCI. As well as configuring Frame Relay, you must also start the IS-IS process for the interface.

Figure S-13 supports the working example for this configuration. It shows the DLCI addresses and the IP addresses for the point-to-point links as well as the ISO addresses for the Routers A, B, and C.

Figure S-13 NBMA Frame Relay Cloud Running Point-to-Point Integrated IS-IS



Example S-6 shows the configuration for Router A to run Integrated IS-IS across the Frame Relay cloud as if it were a broadcast network.

Example S-6 NBMA Frame Relay Cloud, Running Broadcast Integrated IS-IS

```
Working Example for Router A
!
interface Ethernet0
 ip address 140.100.96.1 255.255.255.0
 ip router isis
!
interface Serial0
 no ip address
 encapsulation frame-relay
!Configure Frame Relay for the interface
 no fair-queue
 clockrate 56000
!
interface Serial0.1 point-to-point
!Configure the subinterface to be point-to-point
```

Example S-6 *NBMA Frame Relay Cloud, Running Broadcast Integrated IS-IS (Continued)*

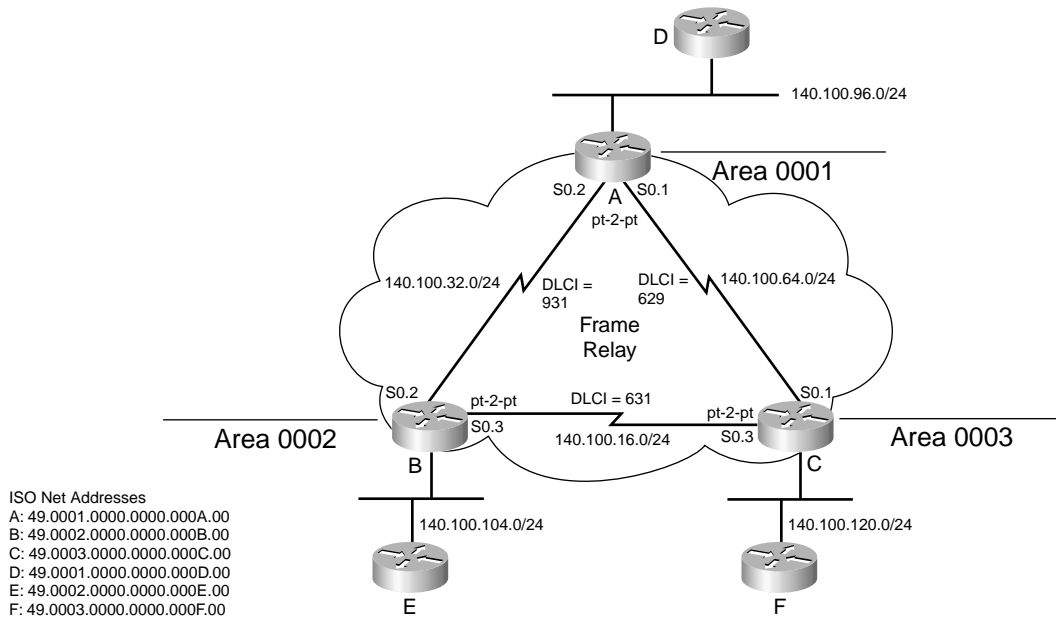
```
ip address 140.100.64.1 255.255.255.0
ip router isis
frame-relay interface-dlci 629
!Configure Frame Relay and define the DLCI to the destination
interface Serial0.2 point-to-point
! Configure the subinterface to be point-to-point
ip address 140.100.32.1 255.255.255.0
ip router isis
frame-relay interface-dlci 931
! Configure Frame Relay and define the DLCI to the destination
!
router isis
net 49.0001.0000.0000.000a.00
!
ip classless
```

Verifying the Integrated IS-IS Operation

The ability to monitor the network enables you to optimize the network and detect problems early. Useful commands to verify the operation of Integrated IS-IS include the following **show** commands:

- **show clns neighbor**
- **show clns interface E0**
- **show isis database**
- **show isis database detail *LSPID***

The following sections explain each of these **show** commands in more detail. The commands explained in this section correspond to the topology illustrated in Figure S-14 and the configuration shown in Example S-7 for Router A.

Figure S-14 *The Network Topology for the show Commands***Example S-7** *The Configuration for Router A in Figure S-14*

```

Working Example for Router A
interface Ethernet0
 ip address 140.100.96.1 255.255.255.0
 ip router isis
!
interface Serial0
 no ip address
 encapsulation frame-relay
 no fair-queue
 clockrate 56000
!
interface Serial0.1 point-to-point
 ip address 140.100.64.1 255.255.255.0
 ip router isis
 frame-relay interface-dlci 629
!
interface Serial0.2 point-to-point
 ip address 140.100.32.1 255.255.255.0
 ip router isis
 frame-relay interface-dlci 931
!
!
router isis
 net 49.0001.0000.0000.000a.00
!
ip classless

```

The show clns neighbors Command

The **show clns neighbors** command has some of the contents of the neighbors table and the state of the link. Note the SNPA is the MAC address of the interface. The type of routing that is used is Level 1-2.

The EXEC command has the following syntax:

```
show clns area-tag neighbors [type number] [area] [detail]
```

Table S-11 explains the syntax of this command.

Table S-11 Explanation of the **show clns neighbors** Command

Field	Description
<i>area-tag</i>	Used primarily in multiarea Integrated IS-IS configurations. This is a name for an Integrated IS-IS routing process on a router. Each name must be unique to the process on the router. If an area tag is specified, output is limited to the specified area.
<i>type number</i>	(Optional) Type of interface (for example, Ethernet). (Optional) Interface number (for example, Ethernet 1).
area	(Optional) Shows the CLNS multiarea adjacencies.
detail	(Optional) Shows details of each adjacency rather than a summary display.

Example S-8 shows output for the **show clns neighbors** command.

Example S-8 Output for the **show clns neighbors** Command

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
0000.0000.000B	Se0.2	DLCI 931	Up	22	L2	IS-IS
0000.0000.000C	Se0.1	DLCI 629	Up	23	L2	IS-IS
0000.0000.000D	Et0	00e0.1e3d.d56f	Up	8	L2	IS-IS

The output of the **show clns neighbors** command shows that Router A has three neighbors. The system ID shows that the serial subinterface S0.1 has heard an LSP from 0000.0000.000C Se0.1, which has the data-link DLCI Frame Relay address of 629. The protocol is IS-IS and it is running Level 2 routing. The link is up and has 23 seconds before another Hello needs to be received. Because the Hello timer is set by default to send Hellos every 10 seconds, it should receive another Hello in 7 seconds, which will reset the holdtime.

This command is good to quickly check connectivity. This output shows that Frame Relay is correctly configured and working, as is IS-IS.

Adding the parameter **detail** to the **show clns neighbors** command gives information about each neighbor and the connection to that neighbor. Example S-9 shows output for the **show clns**

neighbors detail command. Here the area address for the neighbor, the IP address of the transmitting interface, and the length of time that the interface has been up are shown. This command gives information that enables you to verify the addressing scheme of the network.

Example S-9 *Output for the show clns neighbors detail Command*

```

Working Example for Router A
RouterA#show clns neighbor detail

System Id      Interface      SNPA              State  Holdtime  Type Protocol
0000.0000.000B Se0.2         DLCI 931          Up     27        L2   IS-IS
  Area Address(es): 49.0002
  IP Address(es):  140.100.32.2*
  Uptime: 00:05:17
0000.0000.000C Se0.1         DLCI 629          Up     28        L2   IS-IS
  Area Address(es): 49.0003
  IP Address(es):  140.100.64.3*
  Uptime: 00:05:22
0000.0000.000D Et0           00e0.1e3d.d56f   Up     7         L2   IS-IS
  Area Address(es): 49.0003
  IP Address(es):  140.100.96.2*
  Uptime: 00:15:01

```

Table S-12 explains the meaning of the fields in the output screen. The command explained is with the **detail** parameter to cover all the fields shown in both output examples.

Table S-12 *Explanation of the show clns neighbors detail Command Output Screen*

Field	Description
System Id	The system address. It identifies a system in an area. Cisco defines this to be a 6-octet identifier.
SNPA	Subnetwork point of attachment. This is the data-link address.
Interface	Interface through which the neighbor was learned.
State	State of the ES or IS. The states are as follows: <ul style="list-style-type: none"> • Init—The status of the adjacency is initializing. The router is waiting for an Integrated IS-IS Hello message. Until the Hello is received, the neighbor is not adjacent. • Up—Considers the other system reachable. • Down—Considers the other system unreachable.
Holdtime	Number of seconds before this adjacency entry times out.
Type	The types of adjacency are as follows: <ul style="list-style-type: none"> • ES—End-system adjacency. This could be discovered via the ES-IS protocol or statically configured. • Router adjacency—This could be discovered via the ES-IS protocol or statically configured.

Table S-12 *Explanation of the show clns neighbors detail Command Output Screen (Continued)*

Field	Description
L1	(Optional) Displays the router adjacencies for Level 1.
L1L2	(Optional) Displays the router adjacencies for Level 1 and 2.
L2	(Optional) Displays the router adjacencies for Level 2.
Protocol	The routing protocol used to learn the adjacency. The protocol might be ES-IS, IS-IS, ISO IGRP, Static, or DECnet.

The show clns interface Command

The misconfiguration of the interface for Integrated IS-IS results in the inability to create adjacencies. Typically the error is a simple mismatch of parameters, and this can be seen in the **show clns interface** command.

The neighbor database told you of one neighbor and the interface indicates one adjacency. Because this is a LAN interface, it is possible to identify the DIS. The circuit ID shows the pseudonode ID, which has a value greater than 0x00 in the octet after the system ID. Note a DIS is elected for both Level 1 and Level 2 routing. Because no priority has been manually configured, the tiebreaker used to elect the DIS was the highest SNPA on the segment. In this case, the SNPA is the MAC address. Note the default metric is 10 and the priority is 64.

The EXEC command has the following syntax:

```
show clns interface [type number]
```

Table S-13 explains the syntax of this command

Table S-13 *Explanation of the show clns interface Command*

Field	Description
<i>type</i>	(Optional) Type of interface (for example, Ethernet)
<i>number</i>	(Optional) Interface number (for example, Ethernet 1)

In this example, Frame Relay is configured with point-to-point links. Because there is only one other router on this link, there is no need for a DIS to be elected. Therefore, the circuit ID shows the system ID of a router, rather than a pseudonode. Remember it is the octet following the system ID that indicates whether this ID is a pseudonode representing the multiaccess link. If the ID is that of a pseudonode, the system ID is that of the DIS with the next octet showing a nonzero value, such as 0x01.

This makes more sense when you look at the Ethernet interface. The Ethernet 0 interface has the Level 1 circuit ID as A.01. This indicates that the DIS for Level 1 is router A. The octet following the circuit ID of 01 has a nonzero value indicating the ID of a pseudonode.

Because this is a Level 1-2 router, there is also a circuit ID for the Level 2 adjacency. This is relevant only on the Ethernet interface because it is the only multiaccess link. Note that the value for the Level 2 DIS is that of Router D. Router D has been configured as a Level 1 router and as such cannot communicate Level 2 updates. It has effectively set itself up as the DIS to maintain consistency in the circuit ID on a multiaccess link.

Example S-10 shows output for the **show clns interface** command.

Example S-10 *Output Screen for the show clns interface Command*

```
Working Example for Router A
RouterA#show clns interface
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 47 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: A.01
    Number of active level-1 adjacencies: 0
    Level-2 Metric: 10, Priority: 64, Circuit ID D.01
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 1 seconds
    Next IS-IS LAN Level-2 Hello in 1 seconds
Serial0 is up, line protocol is up
  CLNS protocol processing disabled
Serial0.1 is up, line protocol is up
  Checksums enabled, MTU 1500, Encapsulation FRAME-RELAY
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching disabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 58 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x1, local circuit ID 0x100
    Level-1 Metric: 10, Priority: 64, Circuit ID: A.00
    Number of active level-1 adjacencies: 0
    Level-2 Metric: 10, Priority: 64, Circuit ID: A.00
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 2 seconds
Serial0.2 is up, line protocol is up
  Checksums enabled, MTU 1500, Encapsulation FRAME-RELAY
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching disabled
  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 0x101
    Level-1 Metric: 10, Priority: 64, Circuit ID: A.00
```

Example S-10 *Output Screen for the show clns interface Command (Continued)*

```

Number of active level-1 adjacencies: 0
Level-2 Metric: 10, Priority: 64, Circuit ID: A.00
Number of active level-2 adjacencies: 1
Next IS-IS Hello in 886 milliseconds

```

Table S-14 explains the meaning of the fields in the output screen.

Table S-14 *Explanation of the Output Screen from the show clns interface Command*

Field	Description
Checksums enabled	This may be either enabled or disabled.
MTU	The number following MTU is the maximum transmission size for a packet on this interface. Note the MTU is not 1500 because 3 bytes are taken by the OSI header.
Encapsulation	The encapsulation is always SAP (ISO1).
Routing Protocol	This indicates whether ES-IS or IS-IS is running. This determines the type of Hellos sent.
Circuit Type	This indicates whether the link is enabled for L1, L2, or L1-2 routing. This determines the Hellos that are sent.
Level-1 Metric	The metric value for the outbound interface for Level 1 routing. Note that the default setting is 10.
Priority	The priority setting for DIS election. Note the default of 64 is used.
Circuit ID	This identifies the DIS for L1 if there is one on the link.
Number of active Level-1 adjacencies	The number of L1 adjacencies formed on this link.
Level-2 Metric	The metric value for the outbound interface for Level 2 routing. Note that the default setting is 10.
Priority	The priority setting for DIS election. Note the default of 64 is used.
Circuit ID	This identifies the DIS for L2 if there is one on the link.
Number of active Level-2 adjacencies	The number of L2 adjacencies formed on this link.
Next IS-IS LAN Level-1 Hello	The number of seconds before the next Hello is expected. This is a good troubleshooting tool because you can repeat the command to see whether the Hellos are received.
Next IS-IS LAN Level-2 Hello	The number of seconds before the next Hello is expected. This is a good troubleshooting tool because you can repeat the command to see whether the Hellos are received.

The show isis database Command

The **show isis database** command shows the LSPs held in the local database. The LSP ID shows the system ID of the generating router and whether this LSP is from a router or a pseudonode. The last octet shows whether the LSP was too large to fit into one PDU; a nonzero value in this field indicates that this is a fragment of an LSP. The number states the fragment number.

Because IS-IS is a link-state protocol, the database should be identical on every router of the same level within the area.

The EXEC command has the following syntax:

```
show isis area-tag database [Level-1] [Level-2] [l1] [l2] [detail] [lspid]
```

Table S-15 explains the syntax of this command.

Table S-15 Explanation of the **show isis database** Command

Field	Description
<i>area-tag</i>	Used primarily in multiarea Integrated IS-IS configurations. This is a name for an Integrated IS-IS routing process on a router. Each name must be unique to the process on the router. If an area tag is specified, output is limited to the specified area. Required for multiarea Integrated IS-IS configuration. Optional for conventional Integrated IS-IS configuration.
Level-1	(Optional) Displays the Integrated IS-IS link-state database for Level 1.
Level-2	(Optional) Displays the Integrated IS-IS link-state database for Level 2.
l1	(Optional) Abbreviation for Level-1 .
l2	(Optional) Abbreviation for Level-2 .
detail	(Optional) Shows the contents of each LSP rather than a summary display.
lspid	(Optional) Identifier for the link-state PDU. Shows the contents of the specified individual LSP.

Example S-11 shows output for the **show isis database** command.

Example S-11 Output Screen for the **show isis database** Command

```
RouterA#show isis database
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
A.00-00        * 0x00000017  0x76D5        876           1/0/0
IS-IS Level-2 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
A.00-00        * 0x00000018  0xB74F        881           0/0/0
0000.0000.000B.00-00  0x0000001A  0xB561        872           0/0/0
0000.0000.000B.01-00  0x00000016  0x6045        1095          0/0/0
C.00-00        0x0000001E  0x6267        869           0/0/0
```

Example S-11 Output Screen for the **show isis database** Command (Continued)

C.01-00	0x00000002	0xF25F	958	0/0/0
0000.0000.000E.00-00	0x00000018	0x010A	858	0/0/0
0000.0000.000D.00-00	0x0000001A	0x413C	985	0/0/0
0000.0000.000D.04-00	0x00000017	0xFCA0	1006	0/0/0

Table S-16 explains the meaning of the fields in the output screen. Note that both Level 1 and Level 2 databases are shown because the router is running the default configuration of Level 1-2 routing. All the LSP information is contained in these databases, including the LSPs generated by the router itself. An asterisk marks these entries in the output.

Table S-16 Explanation of the **show isis database** Command

Field	Description
LSPID	The LSPID indicates the transmitting router. The system ID is followed by two octets. If the first octet has a value greater than 0x00, this indicates that the ID is that of the pseudonode and the system ID is that of the DIS. The last octet is the Fragment bit. If the value is 0x00, the entire LSP was carried in one LSP. If the value is greater than 0x00, it indicates that this PDU is the <i>n</i> th fragment of the LSP.
LSP Seq Num	Sequence number for the LSP. This allows routers to determine whether they have received the latest information from the source and ensures the integrity of the link-state database.
LSP Checksum	Checksum of the entire LSP packet.
LSP Holdtime	Amount of time the LSP remains valid (in seconds). An LSP holdtime of 0 indicates that this LSP was purged and is being removed from the link-state database.
ATT	The LSP indicates that this router is a Level 2 router with a path out of the area. Level 1 routers use the presence of this bit to identify the closest Level 1-2 router to send their out-of-area data.
P	Partition repair capability. Not supported by Cisco.
OL	The Overload bit indicates that the router has an incomplete database due to memory overload and is therefore not used for transit data.

The show isis database detail Command

The **show isis database detail** command shows the complete LSP and the values for the individual fields. The EXEC command has the same syntax structure as shown in the **show isis database** command, and an explanation of syntax is given in Table S-14. Table S-15 explains the command output fields.

Example S-12 shows output for the **show isis database detail** command. The highlighted lines show that there are two databases: one for the Level 1 routing and the other for the Level 2

routing. It shows for each LSP the area and the IP address of the transmitting interface and the metric cost to the IP routes it knows. The default metric is a cost of 10; therefore, a metric of 20 indicates a route that is 2 hops away.

Example S-12 Output Screen for the **show isis database detail** Command

```
RouterA#show isis database detail

IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num    LSP Checksum   LSP Holdtime   ATT/P/OL
A.00-00        * 0x00000017   0x76D5         873            1/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname:A
  IP Address:   140.100.32.1
  Metric: 10    IP 140.100.96.0 255.255.255.0
  Metric: 10    IP 140.100.64.0 255.255.255.0
  Metric: 10    IP 140.100.32.0 255.255.255.0
  Metric: 10    IS A.01
IS-IS Level-2 Link State Database:
LSPID          LSP Seq Num    LSP Checksum   LSP Holdtime   ATT/P/OL
A.00-00        * 0x00000018   0xB74F         877            0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: A
  IP Address:   140.100.32.1
  Metric: 10    IS 0000.0000.000B.00
  Metric: 10    IS C.00
  Metric: 10    IS 0000.0000.000D.04
  Metric: 10    IP 140.100.96.0 255.255.255.0
  Metric: 10    IP 140.100.64.0 255.255.255.0
  Metric: 10    IP 140.100.32.0 255.255.255.0
0000.0000.000B.00-00  0x0000001A    0xB561         868            0/0/0
  Area Address: 49.0002
  NLPID:        0xCC
  IP Address:   140.100.16.2
  Metric: 10    IS 0000.0000.000B.01
  Metric: 10    IS A.00
  Metric: 10    IS C.00
  Metric: 10    IP 140.100.104.0 255.255.255.0
  Metric: 20    IP 140.100.105.0 255.255.255.0
  Metric: 20    IP 140.100.106.0 255.255.255.0
  Metric: 20    IP 140.100.107.0 255.255.255.0
  Metric: 10    IP 140.100.32.0 255.255.255.0
  Metric: 10    IP 140.100.16.0 255.255.255.0
0000.0000.000B.01-00  0x00000016    0x6045         1089           0/0/0
  Metric: 0     IS 0000.0000.000B.00
  Metric: 0     IS 0000.0000.000E.00
C.00-00        0x0000001E    0x6267         863            0/0/0
  Area Address: 49.0003
  NLPID:        0xCC
  Hostname: C
  IP Address:   140.100.100.1
  Metric: 10    IS C.02
  Metric: 10    IS A.00
```

Example S-12 Output Screen for the **show isis database detail** Command (Continued)

```

Metric: 10      IS 0000.0000.000B.00
Metric: 10      IP 140.100.100.0 255.255.255.0
Metric: 10      IP 140.100.64.0 255.255.255.0
Metric: 10      IP 140.100.16.0 255.255.255.0
C.01-00        0x00000002      0xF25F      951      0/0/0
Metric: 0      IS C.00
0x00000018    0x010A      850      0/0/0
Area Address: 49.0002
0000.0000.000E.00-00 NLPID:      0xCC
IP Address: 140.100.105.1
Metric: 10      IS 0000.0000.000B.01
Metric: 10      IP 140.100.104.0 255.255.255.0
Metric: 10      IP 140.100.105.0 255.255.255.0
Metric: 10      IP 140.100.106.0 255.255.255.0
Metric: 10      IP 140.100.107.0 255.255.255.0
Metric: 20      IP 140.100.32.0 255.255.255.0
Metric: 20      IP 140.100.16.0 255.255.255.0
0000.0000.000D.00-00 0x0000001A      0x413C      976      0/0/0
Area Address: 49.0003
NLPID:      0xCC
IP Address: 140.100.97.1
Metric: 10      IS 0000.0000.000D.04
Metric: 10      IP 140.100.96.0 255.255.255.0
Metric: 10      IP 140.100.97.0 255.255.255.0
Metric: 10      IP 140.100.98.0 255.255.255.0
Metric: 10      IP 140.100.99.0 255.255.255.0
0000.0000.000D.04-00 0x00000017      0xFCA0      996      0/0/0
Metric: 0      IS 0000.0000.000D.00
Metric: 0      IS A.00

```

Table S-17 explains the meaning of the fields in the output screen.

Table S-17 Explanation of the **show isis database detail** Command

Field	Description
Area Address	Area addresses that may be reached from this router. L1 LSPs describe the area addresses configured manually on the originating router. L2 LSPs describe the area addresses for the area to which this route belongs.
Metric	The cost of the outbound interface between the originating router and the advertised neighbor, or the metric of the compound cost between advertising router to the advertised destination.

Troubleshooting the Integrated IS-IS Operation

Unfortunately, even after the most careful planning, configurations may fail to work and the most scrutinized networks can break. The **show** commands are essential for hunting down problems, and careful documentation of the steps taken enable the administrator to build a solid understanding of the problem and systematically eliminate possible problem sources. Having this log is invaluable when, if all else fails, you have to elicit the help of the Technical Assistance

Center (TAC) engineers at Cisco. The following commands are useful for troubleshooting Integrated IS-IS:

- **show isis spf-log**
- **debug** commands

The following sections describe these commands in greater detail.

The show isis spf-log Command

The **show isis spf-log** command explains a great deal about the SPF calculations on the router. It gives the reasons that triggered SPF for the last 20 occurrences.

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

```
show isis spf-log
```

The highlighted lines show the normal periodic SPF calculations. It also shows that Router A sent an LSP because it had set the ATT bit, and later it generated a new LSP when new adjacencies came online. This output shows the healthy working of an IS-IS network, but the command is very useful if the network appears unstable with routes appearing and disappearing. Example S-13 shows output for the **show isis spf-log** command.

Example S-13 Output Screen for the **show isis spf-log** Command

```
RouterA#show isis spf-log

  Level 1 SPF log
  When   Duration  Nodes  Count  First trigger LSP  Triggers
04:23:24    0        1      1      A.00-00          TLVCODE
04:08:46    0        1      1                      PERIODIC
03:53:46    0        1      1                      PERIODIC
03:38:46    0        1      1                      PERIODIC
03:23:46    0        1      1                      PERIODIC
03:08:46    0        1      1                      PERIODIC
02:53:46    0        1      1                      PERIODIC
02:38:46    0        1      1                      PERIODIC
02:23:46    0        1      1                      PERIODIC
02:08:46    0        1      1                      PERIODIC
01:53:46    0        1      1                      PERIODIC
01:38:46    0        1      1                      PERIODIC
01:23:46    0        1      1                      PERIODIC
01:08:46    0        1      1                      PERIODIC
00:53:46    0        1      1                      PERIODIC
00:38:47    0        1      1                      PERIODIC
00:23:47    0        1      1                      PERIODIC
00:15:14    0        1      2      A.00-00          ATTACHFLAG LSPHEADER
00:08:46    0        1      1                      PERIODIC

  Level 2 SPF log
  When   Duration  Nodes  Count  First trigger LSP  Triggers
03:53:48    0        1      1                      PERIODIC
```

Example S-13 Output Screen for the **show isis spf-log** Command (Continued)

03:38:48	0	1	1		PERIODIC
03:23:48	0	1	1		PERIODIC
03:08:48	0	1	1		PERIODIC
02:53:48	0	1	1		PERIODIC
02:38:48	0	1	1		PERIODIC
02:23:48	0	1	1		PERIODIC
02:08:48	0	1	1		PERIODIC
01:53:48	0	1	1		PERIODIC
01:38:48	0	1	1		PERIODIC
01:23:48	0	1	1		PERIODIC
01:08:48	0	1	1		PERIODIC
00:53:48	0	1	1		PERIODIC
00:38:48	0	1	1		PERIODIC
00:23:48	0	1	1		PERIODIC
00:15:22	0	3	5	A.00-00	NEWADJ LSPHEADER TLVCONTENT
00:08:48	0	3	1		PERIODIC
00:05:44	4	4	4	A.00-00	NEWADJ TLVCONTENT
00:05:38	4	7	5	0000.0000.000B.00-00	LSPHEADER TLVCONTENT

Table S-18 explains the meaning of the fields in the output screen.

Table S-18 Explanation of the **show isis spf-log** Command

Field	Description
When	The amount of time since a full SPF calculation occurred (hh:mm:ss). The last 20 occurrences are shown.
Duration	Number of milliseconds it took to complete this SPF run. This is elapsed time as opposed to CPU time.
Nodes	Number of routers and pseudonodes (LANs) calculated in the SPF run.
Count	This count shows the number of events (such as receiving new LSPs) that occurred while the router was waiting before running full SPF. The router waits five seconds after the first triggered LSP to ensure that all the information is received.
Last trigger LSP	Whenever a full SPF calculation is triggered by the arrival of a new LSP, the ID of the LSP is stored. The LSPID can indicate where to start troubleshooting for any routing instability in an area. If multiple LSPs are causing an SPF run, only the LSPID of the last received LSP is remembered.

The debug Commands

The **debug** command is a helpful troubleshooting tool, but it does have certain disadvantages that you must be aware of. The **debug** command has the highest process priority and forces the router into process switching. It is therefore capable of consuming all the resources on the router, thus becoming the problem as opposed to helping to solve the problem. It is important to just turn on **debug** for the specific task to be monitored and to turn it off as soon as the data

has been gathered. The **no** form of this command disables debugging output. It is wise to direct the output to a log file because each character sent to the console forces processor interrupt; in this way, the data can be perused with care.

The **debug** commands are EXEC commands. Table S-19 lists the options available for monitoring Integrated IS-IS.

Table S-19 *The debug Command Options for Integrated IS-IS*

Command Option	Description
debug isis adj-packets	Displays information on all adjacency-related activity. This includes the following: <ul style="list-style-type: none">• Hello packets (sent and received)• Any changes in the state of an adjacency in Integrated IS-IS.
debug isis spf-statistics	Displays statistical information about building routes between routers. Using the statistical information provided, one can determine how long it takes to place a Level 1 IS or Level 2 IS on the shortest path tree (SPT) using the IS-IS protocol.
debug isis update-packets	Displays SNPs and (CSNPs and PSNPs) LSPs that are detected by the router.

Conclusion

The complexity of Integrated IS-IS is its very strength. The flexibility that enables you to add extensions as technology requires by using the TLV codes is the main reason for its recent increase in popularity. The fact that Integrated IS-IS is a standards-based protocol is another strength. The enormity of its address range is a resource that may never be completely tapped.

Foundation Summary

The “Foundation Summary” is a collection of quick reference information that provides a convenient review of many key concepts in this chapter. For those of you who already feel comfortable with the topics in this chapter, this summary will help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparations before the exam, these tables and figures are a convenient way to review the day before the exam.

The ISO Address Format (NSAP)

Figure S-15 shows the format of a CLNS address.

Figure S-15 *An ISO Address*

IDP		DSP		
AFI (1 octet)	IDI	High Order DSP	System ID (1-8 octets)	NSEL (1 octet)
AREA			ID	SEL

For routing purposes, ISO 10589 specifies three parts to the address, namely how to get to the area how to find the host and then the application within the host. However, there are only two elements in the address used to find a destination host; the address that locates the area and the host within that area. Therefore, Integrated IS-IS has two levels of hierarchy, as follows:

- **IDP**—The IDP is used to route to the domain, or autonomous system. The IDP is given by the ISO and identifies the body responsible for assigning the format for the rest of the address by defining the DSP structure. The following two parts comprise IDP:
 - **AFI (authority and format identifier)**—AFI is the first octet of the address. It is defined as 1 octet and it states the authority in charge of issuing addresses and the format of the rest of the address, within the constraints of IDI and DSP, to describe Area, ID, and SEL.
 - **IDI (initial domain identifier)**—IDI is the suborganization to the AFI. (For example, the U.S. government is the AFI and the General Services Agency is the IDI, otherwise known as GOSIP.)
- **DSP**—DSP is used to route within the autonomous system. The authority specified by the IDP may further subdivide the DSP and may assign subauthorities responsible for parts of the DSP. The DSP can be broken down into the following fields.

— **High Order DSP**—High Order DSP is typically the area within the autonomous system.

— **System ID**—System ID may have a value between 1 to 8 octets. It must have the same value throughout all systems in the autonomous system. Cisco uses six octets, which is a common solution because it allows the use of the MAC address to autoconfigure the system.

— **NSEL**—NSEL identifies the particular service at the network layer to which to hand the packet.

Characteristics of a Level 1 router include the following:

- An intra-area router.
- Similar to an OSPF stub router.
- Knowledge of the network limited to the area.
- Has a link-state database with all the routing information for the area.
- The address of the closest Level 2 router to which to send traffic destined for another area.
- Neighbors must be in the same area.
- A DIS is elected on LANs.

Characteristics of a Level 2 router include the following:

- An interarea router.
- Similar to a backbone internal router in OSPF.
- Level 2 routers must be contiguous (that is, the area cannot be fractured).
- Neighbors may be in different areas.
- Has a Level 2 database with all the information for interarea routing.

Characteristics of a Level 1-2 router include the following:

- An intra- and interarea router.
- Similar to an area boundary router (ABR) in OSPF.
- Neighbors may be in different areas.
- Has a Level 2 database with all the information for interarea routing.
- Has a Level 1 database for each area to which it is connected.
- It informs Level 1 routers that it is a Level 2 router and can forward traffic to other areas.
- It informs other Level 2 routers of the areas to which it is connected.

Table S-20 summarizes the commands covered in this chapter.

Table S-20 *Summary of Commands*

Command	Function	Location of Command
Router# router isis	Starts the Integrated IS-IS process on the router	An EXEC command in configuration mode
Router(config-router)# net network-address	Defines the NET address of the router	A subrouter command, placed under router isis
Router(config-router)# is-type level-1	Defines whether the router is acting as an L1, L2, or L1-2 router	A subrouter command, placed under router isis
Router(config-router)# summary-address address mask	Enables you to manually summarize the networks	A subrouter command, placed under router isis
Router(config-if)# default-information originate	Advertises the default route	A subrouter command, placed under router isis
Router(config-if)# ip router isis	Turns on Integrated IS-IS routing for the interface	An interface command, placed under interface [<i>type number</i>]
Router(config-if)# clns router isis	Turns CLNS routing at the interface to allow the L2 LSPs to be unraveled so that the ATT bit may be seen	An interface command, placed under interface [<i>type number</i>]
Router(config-if)# frame-relay map ip destination address DLCI broadcast	Maps the Frame Relay address to the destination and defines the link as a broadcast link to be treated as a LAN	An interface command, placed under interface [<i>type number</i>]
Router# show clns neighbor detail	Displays information drawn from the neighbor table	An EXEC command
Router# show clns interface	Displays information about Integrated IS-IS Hello and timers seen on the interface	An EXEC command
Router# show isis database	Shows the link-state database	An EXEC command
Router# show isis spf-log	Show why and how often the SPF algorithm was calculated (listing the last 20 occurrences)	An EXEC command

Q&A

The following questions test your understanding of the topics covered in this chapter. After you have answered the questions, check your answers in the later “Answers to the Quiz Questions” section. If you get an answer wrong, review the answer and ensure that you understand the reason for your mistake. If you are confused by the answer, refer to the appropriate text in the chapter to review the concepts.

- 1 Which system generates the pseudonode?

- 2 Which timers are tracked in the neighbor table?

- 3 What is a CSNP? When is it used?

- 4 What is a PSNP? When is it used?

- 5 Every Integrated IS-IS PDU has the same fixed 8-octet header. What does it contain?

- 6 When does Integrated IS-IS recalculate the SPT?

7 What happens in the event of the DIS dying?

8 Integrated IS-IS may be used to send information about which routed protocols?

9 How long is the holddown timer, by default?

10 What link-state algorithm is run to create entries for the routing table?

11 When is an adjacency said to be INIT?

12 For Frame Relay, when would you configure the **map** command with the broadcast parameter?

13 Which command is used to display all update packets that are both received and sent by a router?

14 In the system ID 49.0000.0000.000B.00.03, what does the 03 of the last octet indicate?

15 What is the relationship between the Hello timer and when the path is considered to have died?

16 What is the purpose of the ATT bit in the Level 2 LSP?

17 Integrated IS-IS areas are similar to which type of areas in OSPF?

18 What is the purpose of the OL bit and how is it sent to other routers?

19 Describe one design restriction in configuring Level 2 routing.

20 Given the following address:

49.0001.2222.2222.2222.00

Is this a NET or NSAP address? Give reasons for your choice.

21 On which router may you configure route summarization?

22 What is a pseudonode and what is its purpose?

23 State two reasons that a router may not be able to find a neighbor

24 Which command shows the LSPs in detail?

25 Explain briefly why two routers cannot have the same system ID within the area.

26 What does TLV stand for? Briefly explain its purpose.

27 How many link-state databases does a Level 1-2 router have?

28 Why does the design of Level 1 routers storing the nearest Level 2 router as the exit point from the area sometimes result in suboptimal routing?

29 Integrated IS-IS packets run directly on top of which layer?

30 Where is the NET address configured?

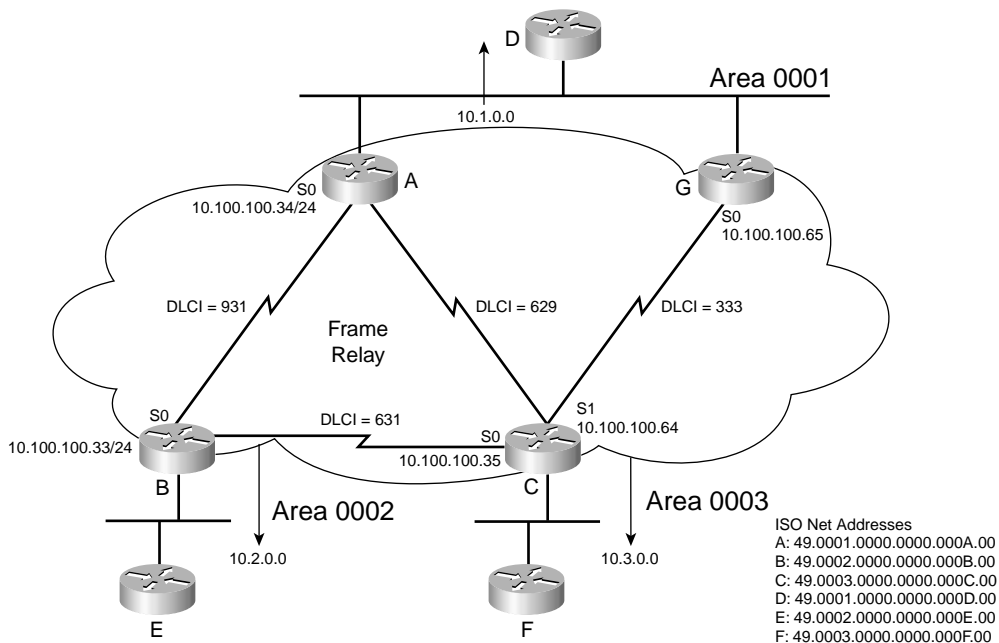
Scenarios

The following scenarios and questions are designed to draw together the content of the chapter and to exercise your understanding of the concepts. There is not necessarily a right answer. The thought process and practice in manipulating the concepts are the goals of this section. The answers to the scenario questions are found at the end of this chapter.

Scenario S-1

The ISP Flying Data has recently converted from OSPF to Integrated IS-IS. The migration was relatively painless. The company has used the private network 10.0.0.0 for IP and the private OSI addressing with AFI 49. They have created a hierarchical addressing structure. See Figure S-16 to see this addressing scheme.

Figure S-16 Diagram for Scenario S-1




The addressing of the network was a major project, with all the pitfalls that accompany such a major exercise. The network is now stable, and it is time to configure the WAN connections using multipoint Frame Relay:

- 1 Issue the commands that will allow Router A to use Integrated IS-IS routing across the NBMA cloud as if the cloud were a broadcast medium. Refer to Figure S-16 for the addressing scheme.
- 2 The WAN is a Frame Relay cloud, and Router G has a point-to-point link with Router C. Issue the commands for Router C that configure the link for Integrated IS-IS as a point-to-point link.
- 3 To reduce bandwidth consumption and to hide some network detail, summarization has been suggested as a solution over the WAN links. Issue the commands for Router A that will summarize the networks behind this router with a prefix of /16 across the WAN.

Scenario S-2

Given the configuration of Integrated IS-IS in Example S-14 and the output screen in Example S-15, perform the tasks and answer the questions listed. The WAN has light user traffic and has a fully meshed configuration, as shown in Figure S-17.

Figure S-17 Diagram for Scenario S-2



Example S-14 Router B's Configuration File

```
B#show running config
Building configuration...

Current configuration:
!
version 12.0
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname B
!
!
ip subnet-zero
ip tcp synwait-time 5
no ip domain-lookup
clns routing
!
!
interface Ethernet0
 ip address 140.100.104.1 255.255.255.0
 no ip directed-broadcast
 ip router isis
!
interface Serial0
 no ip address
 no ip directed-broadcast
 encapsulation frame-relay
 no ip mroute-cache
 no fair-queue
 clockrate 56000
!
interface Serial0.2 point-to-point
 ip address 140.100.32.2 255.255.255.0
 no ip directed-broadcast
 ip router isis
 frame-relay interface-dlci 931
!
interface Serial0.3 point-to-point
 ip address 140.100.16.2 255.255.255.0
 no ip directed-broadcast
 ip router isis
 frame-relay interface-dlci 631
!
interface Serial1
 no ip address
 no ip directed-broadcast
 shutdown
!
router isis
 net 49.0002.0000.0000.000b.00
!
ip classless
```

continues

Example S-14 Router B's Configuration File (Continued)

```

!
!
alias exec c conf t
alias exec s sh run
alias exec ii sh ip int brie
!
line con 0
  exec-timeout 30 0
  privilege level 15
  logging synchronous
  transport input none
line aux 0
line vty 0 4
!
end

```

Example S-15 The show clns interface Command from Router B

```

B# show clns interface
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 15 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 678 milliseconds
    Next IS-IS LAN Level-2 Hello in 1 seconds
Serial0 is up, line protocol is up
  CLNS protocol processing disabled
Serial0.1 is up, line protocol is up
  CLNS protocol processing disabled
Serial0.2 is up, line protocol is up
  Checksums enabled, MTU 1500, Encapsulation FRAME-RELAY
  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 disabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 43 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x2, local circuit ID 0x101

```

Example S-15 *The show clns interface Command from Router B (Continued)*

```

Level-1 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
Number of active level-2 adjacencies: 1
Next IS-IS Hello in 2 seconds
Serial0/3 is up, line protocol is up
Checksums enabled, MTU 1500, Encapsulation FRAME-RELAY
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 disabled
CLNS SSE switching disabled
DEC compatibility mode OFF for this interface
Next ESH/ISH in 11 seconds
Routing Protocol: IS-IS
Circuit Type: level-1-2
Interface number 0x1, local circuit ID 0x100
Level-1 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000C.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.00
Number of active level-2 adjacencies: 1
Next IS-IS Hello in 3 seconds

```

- 1 Identify the DIS on the Ethernet segment? How was this information apparent?
- 2 If Router A died, what would the effect be on the network?
- 3 Is summarization possible only on the routers entering the WAN cloud, or is it possible on the networks not shown in the figure, but on the other side of the routers? Give reasons for your answers.

Scenario Answers

The answers are in **bold**. The answers provided in this section are not necessarily the only possible answers to the questions. The questions are designed to test your knowledge and to give practical exercise in certain key areas. This section is intended to test and exercise skills and concepts detailed in the body of this chapter.

If your answer is different, ask yourself whether it follows the tenets explained in the answers provided. Your answer is correct not if it matches the solution provided in the book, but rather if it has included the principles of design laid out in the chapter.

In this way, the testing provided in these scenarios is deeper: It examines not only your knowledge, but also your understanding and ability to apply that knowledge to problems.

If you do not get the correct answer, refer back to the text and review the subject tested. Be certain to also review your notes on the question to ensure that you understand the principles of the subject.

Scenario S-1 Answers

- 1 Issue the commands that will allow Router A to use Integrated IS-IS routing across the NBMA cloud as if the cloud were a broadcast medium. Refer to Figure S-16 for the addressing scheme.

The highlighted portion of the configuration file in Example S-16 shows the configuration of IS-IS across the Frame Relay cloud, using the broadcast technology and LAN Hellos. The frame-relay map ip command maps the IP destination address to the outgoing DLCI and defines the interface as a broadcast interface.

The frame-relay map clns command maps to the clns process on the destination router. Without the second command, no routes appear in the IP routing table because CLNS does not receive the frames to populate the IP routing table. Remember that these are IP routes carried in the IS-IS routing protocol. It is IS-IS that updates the IP routing table.

Example S-16 *Configuration for Router A: Integrated IS-IS Broadcast Across an NBMA Frame Relay Cloud*

```
ROUTER A
Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname A
!
!
network-clock base-rate 56k
ip subnet-zero
ip tcp synwait-time 5
no ip domain-lookup
!
!
cns event-service server
!
!
controller T1 0
!
!
interface Ethernet0
 ip address 10.1.128.1 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
interface Serial0
 ip address 10.100.100.34 255.255.255.0
 ip router isis
 encapsulation frame-relay
 clockrate 56000
 frame-relay map clns 629 broadcast
 frame-relay map clns 931 broadcast
 frame-relay map ip 10.100.100.33 931 broadcast
```

Example S-16 Configuration for Router A: Integrated IS-IS Broadcast Across an NBMA Frame Relay Cloud (Continued)

```

frame-relay map ip 10.100.100.35 629 broadcast
isis circuit-type level-2-only
!
!
router isis
net 49.0001.0000.0000.000a.00
!
ip classless

```

- 2 The WAN is a Frame Relay cloud, and Router G has a point-to-point link with Router C. Issue the commands for router C that configure the link for Integrated IS-IS as a point-to-point link.

The highlighted portion of the configuration file in Example S-17 shows the configuration of IS-IS across the Frame Relay cloud, using the point-to-point technology and point-to-point Hellos. Because the link is point-to-point, there is no need to configure frame-relay map commands (because there is no choice of destination to define). The point-to-point link is just a pipe that goes to one destination. As shown in the configuration, it is only necessary to configure the interface as point-to-point, start Frame Relay, and define the DLCI. As well as configuring Frame Relay, you must also start the IS-IS process for the interface.

Example S-17 Configuration for Router C: Integrated IS-IS Point-to-Point Across an NBMA Frame Relay Cloud

```

ROUTER C
Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname C
!
!
network-clock base-rate 56k
ip subnet-zero
ip tcp synwait-time 5
no ip domain-lookup
!
cns event-service server
!
!
controller T1 0
!
interface Ethernet0
no ip address
ip router isis
shutdown
!
interface Serial0

```

continues

Example S-17 Configuration for Router C: Integrated IS-IS Point-to-Point Across an NBMA Frame Relay Cloud (Continued)

```

no ip address
encapsulation frame-relay
no fair-queue
clockrate 56000
!
interface Serial0.1 point-to-point
ip address 10.100.100.66 255.255.255.0
ip router isis
frame-relay interface-dlci 333
!
interface Serial1
ip address 10.100.100.35 255.255.255.0
ip router isis
encapsulation frame-relay
clockrate 56000
frame-relay map clns 629 broadcast
frame-relay map clns 631 broadcast
frame-relay map ip 10.100.100.33 631 broadcast
frame-relay map ip 10.100.100.34 629 broadcast
isis circuit-type level-2-only
!
router isis
net 49.0003.0000.0000.000c.00
!
ip classless

```

- 3 To reduce bandwidth consumption and to hide some network detail, summarization has been suggested as a solution over the WAN links. Issue the commands for Router A that will summarize the networks behind this router with a prefix of /16 across the WAN.

The highlighted portion of the configuration file in Example S-18 shows the summary-address command that is used to hide the routes within area 0001 from the other areas. This configuration is possible on router A because it sits on the boundary between areas. Summarizing routes reduces the network resources required by the network.

Example S-18 Configuration for Router A: Summarizing Networks from Router A in Area 0001 to All Other Areas

```

ROUTER A
! A#sh run
Building configuration...

Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!

```

Example S-18 *Configuration for Router A: Summarizing Networks from Router A in Area 0001 to All Other Areas (Continued)*

```
hostname A
!
!
network-clock base-rate 56k
ip subnet-zero
ip tcp synwait-time 5
no ip domain-lookup
!
!
cns event-service server
!
!
controller T1 0
!
!
interface Ethernet0
 ip address 10.1.1.1 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
interface Serial0
 no ip address
 encapsulation frame-relay
 no fair-queue
 clockrate 56000
!
interface Serial0.1 point-to-point
 ip address 10.100.100.34 255.255.255.0
 ip router isis
 frame-relay interface-dlci 629
 isis circuit-type level-2-only
!
interface Serial0.2 point-to-point
 ip address 10.100.32.1 255.255.255.0
 ip router isis
 frame-relay interface-dlci 931
 isis circuit-type level-2-only
!
router isis
 summary-address 10.1.0.0 255.255.0.0
 net 49.0001.0000.0000.000a.00
!
ip classless
```

Scenario S-2 Answers

- 1 Identify the DIS on the Ethernet segment? How was this information apparent?

On examining the show interface output screen in Example S-19, you can see the circuit ID is 0000.0000.000B.01. This is the system ID of the pseudonode, as made apparent by the last octet having a nonzero value. The system ID of the pseudonode is the system ID of the DIS plus the nonzero octet. Therefore, the DIS is 49.0002.0000.0000.000B.00.

Example S-19 *Identifying the DIS in Scenario S-2*

```

ROUTER B
B# show clns interface
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 15 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
      Interface number 0x0, local circuit ID 0x1
      Level-1 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
      Number of active level-1 adjacencies: 1
      Level-2 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
      Number of active level-2 adjacencies: 1
      Next IS-IS LAN Level-1 Hello in 678 milliseconds
      Next IS-IS LAN Level-2 Hello in 1 seconds

```

- 2 If Router A died, what would the effect be on the network?

If Router A died, Routers B and C would not be able to communicate with Router A or with Router D. However, Routers B and C would be able to communicate with each other. The network behind Router A would function, but it would be isolated from the others.

The neighbor tables would fail to hear the Hellos from Router A, and Routers B, and C would time out all routes that they had heard from this router. Routers behind Router A would time out Router A from the neighbor table. All the former neighbors of Router A would send LSPs. The LSPs from Router A would be purged from all the databases, new LSPs would flood the network, and the SPF algorithm would be run. Router A and the network behind it would be annexed from the larger network of Routers B and C.

- 3 Is summarization possible only on the routers entering the WAN cloud, or is it possible on the networks not shown in the figure, but on the other side of the routers? Give reasons for your answers.

Summarization is only possible on the Level 1-2 routers, acting as an ABR, and then only if the addressing scheme allows for it to be implemented. The Level 1-2 router would summarize the routes and inject them into the Level 2 network.

Suggested Further Reading

RFC 1195, "Use of OSI IS-IS for Routing TCP/IP and Dual Environments." R. Callon. December 1990.

RFC 1629, "Guidelines for OSI NSAP Allocation in the Internet" (Obsoletes RFC 1237.) R. Colella, R. Callon, E. Gardner, Y. Rekhter. May 1994.

RFC 1142, "OSI IS-IS Intra-domain Routing Protocol." D. Oran. February 1990.

ISO 10589, "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 30 April 1992

Doyle, Jeff. *Routing TCP/IP*, Volume I. Cisco Press. September 1998. ISBN: 1-57870-041-8

Martey, Abe. *IS-IS Network Design Solutions*. Cisco Press. February 2002. ISBN: 1-57870-220-8

Answers to the Quiz Questions

"Do I Know This Already?" Quiz

- 1 What portion of the address does the initialism IDP identify?

Answer: The initial domain part (IDP) is used for external routing. The IDP is used to route to the domain or autonomous system. The IDP is given by the ISO and identifies the organization responsible for assigning the format for the rest of the address by defining the DSP structure.

The IDP of the address identifies the prefix of the address, which RFC 1629 identifies as the area. This may include many fields stating the authority governing the address distribution and the unique domain and areas identifiers.

RFC 1629 states: "For routing purposes, ISO addresses are subdivided by IS-IS into the area address, the system identifier (ID), and the NSAP selector (SEL). The area address identifies both the routing domain and the area within the routing domain. Generally, the area address corresponds to the IDP plus a high-order part of the DSP (HO-DSP)."

- 2 If the NSEL is set to 0x00, what does this indicate?

Answer: Network selector (NSEL), sometimes referred to as the SEL field, describes the service at the network layer to which the packet is to be sent. This is similar to the Protocol field in IP. If the NSEL is set to 0x00, this indicates that there is no protocol defined and that this is a NET address or network layer address of a router that will forward the data to the network service access point (NSAP) of the destination.

- 3 How long is the AFI in a NSAP address and what function does it perform?

Answer: The authority and format identifier (AFI) field is 1 octet in length, and it states the governing body managing and distributing the unique addresses. The governing body has the authority to define the format of the rest of the network service access point (NSAP) address.

- 4 How long is the system ID defined in the NET on Cisco equipment?

Answer: The system ID portion of the NSAP or NET is six octets in length on all Cisco equipment.

- 5 How often, by default, does Integrated IS-IS send out Hello packets?

Answer: Integrated IS-IS sends out Hellos every 10 seconds. This is true of every interface and level of routing. On multiaccess links, the DIS sends out Hellos every 3.3 seconds.

- 6 What is a neighbor in Integrated IS-IS?

Answer: A neighbor is a router on the same link with which an adjacency is formed and routing information is then exchanged.

- 7 What is an adjacency in Integrated IS-IS?

Answer: An adjacency is the local routing information that shows the reachability of an ES or IS over a single circuit. A separate adjacency is created for each neighbor on a circuit and for each level of routing (that is, Level 1 and Level 2) on a broadcast circuit.

- 8 If the network is stable and sees no changes, how often will it send LSPs? Why are these updates sent out periodically?

Answer: A stable network sends out LSPs approximately every 15 minutes. These updates are sent out periodically to ensure that the link-state databases are synchronized throughout the autonomous system.

- 9 What command is used to configure the Integrated IS-IS router process?

Answer: The router isis EXEC command starts the routing process for integrated IS-IS.

- 10 What is the default routing level on a Cisco router?

Answer: By default the routing process for Integrated IS-IS runs as a Level 1-2 router.

- 11 What command is used to configure Integrated IS-IS routing on the interface?

Answer: The ip router isis interface command is used to start Integrated IS-IS routing on the interface.

12 How is the NET address configured on the router?

Answer: Under the command that started the routing process for Integrated IS-IS, enter the net command followed by the full network address.

13 What command is used to show the state of adjacencies on the router?

Answer: The show clns neighbor and show clns interface commands both display the adjacencies. The show clns neighbor command gives information as to the state of the link, the type of routing performed on the link, the data-link address (SNPA) of the link as well as the holdtime, or how long since it received the last Hello.

14 What command identifies the designated router for your LAN?

Answer: The designated intermediate system (DIS) is identified in the output screen of the show clns interface command or any show command that shows the pseudonode, such as show isis database or show clns neighbors. If the medium is multiaccess broadcast, the Circuit ID field shows the ID of the pseudonode. The pseudonode is identified by the nonzero value in the octet following the system ID of the DIS (for example, R2.01).

15 Explain briefly what show isis database reveals.

Answer: The show isis database command displays the LSPs in the link-state database. This database should be identical on every router in the area.

16 What command reveals the trigger for the last SPF calculation on the router?

Answer: The show isis spf-log command shows the trigger for the last 20 occurrences that the SPF calculation was run.

Q&A

1 Which system generates the pseudonode?

Answer: The designated intermediate system (DIS) generates the pseudonode, using its own system ID and setting the following octet to be a nonzero number.

2 Which timers are tracked in the neighbor table?

Answer: The holdtime, which states the amount of time since the last Hello, is held in the neighbor table for Integrated IS-IS.

3 What is a CSNP? When is it used?

Answer: The CSNP describes every link in the link-state database. These packets are sent on point-to-point links when the link comes up to synchronize the link-state databases. The designated intermediate system (DIS) on a multicast network sends out CSNPs every 10 seconds.

- 4 What is a PSNP? When is it used?

Answer: Partial sequence number packets are sent on point-to-point links to explicitly acknowledge each LSP they receive. A router on a broadcast subnetwork sends a PSNP requesting the LSPs it needs to synchronize its link-state database.

- 5 Every Integrated IS-IS PDU has the same fixed 8-octet header. What does it contain?

Answer: The following table shows the fields in the fixed header.

Field	Description
Intradomain Routing Protocol Discriminator	All Integrated IS-IS PDUs have a value of 0x83. This identifies the packet.
Length Indication	States the length of the fixed header.
Version/Protocol ID	Set to 1.
ID Length	The size of the system ID in the NSAP. This can be an integer between 1 and 8. Cisco set it to be 6, which is represented by 0.
Reserved/Packet Type	The first three bits are reserved, set to 0, and ignored. The Packet Type indicates whether this is a Hello, LSP, or SNP.
Field	Description
Version	Set to 1.
Reserved	Set to 0 and ignored.
Maximum Area Addresses	States the number of area addresses permitted for this area. Multiple area addresses are configured during transitions. Cisco set the maximum address size to 3, which is represented by 0.

- 6 When does Integrated IS-IS recalculate the SPT?

Answer: The SPT is calculated whenever an LSP is received that is either not in the database or that has a higher sequence number. This indicates that the LSP is fresher or more recent than the one held by the router. If there is a change in the status of an adjacency on the router itself, it generates its own LSP, which leads to the SPF algorithm being run and a new SPT being created.

- 7 What happens in the event of the DIS dying?

Answer: There is no backup designated router in Integrated IS-IS. Therefore, if the DIS meets an untimely death, a new DIS is elected, based on priority or highest MAC address. If another router comes online with a higher priority, it dislodges the existing DIS and rules in its place. This behavior differs from that of OSPF. When a new DIS is elected, the link-state databases are purged and new LSPs flooded.

8 Integrated IS-IS may be used to send information about which routed protocols?

Answer: Integrated IS-IS is capable of carrying both IP and CLNS.

9 How long is the holddown timer, by default?

Answer: The holddown timer is three times the Hello timer. The default Hello timer is 10 seconds, making the holddown timer 30 seconds.

10 What link-state algorithm is run to create entries for the routing table?

Answer: The name of the link-state algorithm is the Dijkstra algorithm.

11 When is an adjacency said to be INIT?

Answer: An adjacency is in the INIT state if it has sent a Hello and is waiting to receive a Hello from the other router, with the sent Hello registered in the other router's Hello.

12 For Frame Relay, when would you configure the **map** command with the broadcast parameter?

Answer: The Frame Relay map command with the broadcast parameter is used in a fully meshed environment and the network is multiaccess. This allows a DIS to be elected, streamlining of the adjacencies, and the efficient use of IP subnets.

13 Which command is used to display all update packets that are both received and sent by a router?

Answer: The command that shows all the Integrated IS-IS packets both sent and received by the router is `debug isis update-packets`. These packets are the sequence number PDUs (CSNPs and PSNPs) and LSPs that are detected by the router.

14 In the system ID 49.0000.0000.000B.00.03, what does the 03 of the last octet indicate?

Answer: The last octet is used to show fragmentation and to identify which fragment this PDU is in the LSP. This is the third fragment.

15 What is the relationship between the Hello timer and when the path is considered to have died?

Answer: The default timer is three times that of the Hello timer; therefore, the path waits for 30 seconds before declaring the path died and flushing the LSPs from the link-state database.

16 What is the purpose of the ATT bit in the Level 2 LSP?

Answer: The purpose of the ATT bit in the LSPs is to inform all Level 1 routers about the presence of a Level 2 router and where it sits.

- 17 Integrated IS-IS areas are similar to which type of areas in OSPF?

Answer: Integrated IS-IS areas are similar to stub areas in OSPF because they are unaware of any routes external to the area.

- 18 What is the purpose of the OL bit and how is it sent to other routers?

Answer: The Overload (OL) bit is set on an LSP if the router cannot store the entire link-state database. When other routers receive LSPs with this bit set, they will not send the router any transit traffic for fear that its routing table is incomplete. If the router is making decisions using incomplete data, its decisions may result in suboptimal paths or even routing loops. Traffic destined for the router can still be sent to the directly connected interfaces of a router transmitting the OL bit in its LSPs.

- 19 Describe one design restriction in configuring Level 2 routing.

Answer: There is only one hard-and-fast rule for the design of a Level 2 network: Level 2 routers must be contiguous; that is, the area cannot be fractured.

- 20 Given the following address:

49.0001.2222.2222.2222.00

Is this a NET or NSAP address? Give reasons for your choice.

Answer: The address is a NET address because the last octet is set to 0x00. There is no network service defined. This is the address of a router not an ES.

- 21 On which router may you configure route summarization?

Answer: Level 1 routes may be summarized at the Level 1-2 router at the edge of the area so that the summary addresses can be propagated throughout the other areas.

- 22 What is a pseudonode and what is its purpose?

Answer: The pseudonode is the LAN identifier for a broadcast subnetwork. The pseudonode is the system ID of the DIS plus the circuit ID. The pseudonode has links to each of the ISs, and each IS has a single link to the pseudonode. The use of the pseudonode reduces the number of links required. Instead of $n-1$ links to each of the other ISs, there is one link per IS. The DIS generates link-state PDUs on behalf of the pseudonode. These LSPs are sent to all the connected ISs.

- 23 State two reasons that a router may not be able to find a neighbor

Answer: For an adjacency to be formed and maintained, both interfaces must agree on the following:

- **The same MTU must be configured on both interfaces.**
- **If one of the routers is configured as Level 1, both routers must be Level 1-capable.**
- **The system IDs must be unique to the router.**

- **Level 1 routers must be in the same area.**
- **If the routers are in different areas, at least one of the routers must be Level 1-2.**
- **The authentication configuration must be the same on both routers.**
- **The Hello timers (including the holddown timer) must match. If one router has a Hello timer of 40 seconds, the defaults on the other router would time out the holddown timer and purge the LSP, resulting in a flapping link and endless SPF calculations.**

24 Which command shows the LSPs in detail?

Answer: The command show isis database detail shows the LSPs in detail.

25 Explain briefly why two routers cannot have the same system ID within the area.

Answer: System ID is a unique identifier for the device. Area routing (Level-1) is performed only on the system ID. Therefore, the system ID must be unique to the area to uniquely identify a particular device within the area.

26 What does TLV stand for? Briefly explain its purpose.

Answer: TLV is the same as a CLV, but some literature refers to the variable-length fields as Type/Length/Value in accordance with the IP terminology. Although the Integrated IS-IS PDUs are fixed, the TLV fields are variable-length and can expand as needed. This design allows great flexibility and movement to develop in step with technological advances. It is through the development of TLV code 128 that IS-IS was extended to carry IP, creating Integrated IS-IS.

27 How many link-state databases does a Level 1-2 router have?

Answer: A Level 1-2 router has two link-state databases, one for the Level 1 routes and the other for the Level 2 routes. A separate SPF algorithm is run for each database.

28 Why does the design of Level 1 routers storing the nearest Level 2 router as the exit point from the area sometimes result in suboptimal routing?

Answer: Level 1 areas have knowledge of networks only within their own areas. To reach another area, packets are sent to the nearest Level 2 router. Without additional configuration, this is literally the nearest router. The metrics used are the default metric of 10 on each outbound interface; therefore, the best route translates to that with the lowest hop count. As you know, the router two hops away may include a 16-Mbps Token Ring and a 56-kbps link as opposed to the three hops of Fast Ethernet and ATM. The routing decision made by Integrated IS-IS would be to take the lower bandwidth because it has the lowest hop count.

29 Integrated IS-IS packets run directly on top of which layer?

Answer: Integrated IS-IS packets run directly on top of the data link layer.

30 Where is the NET address configured?

Answer: The NET address is configured under the router process, which is started by the router isis command. It is not configured at the interface level like many routing protocols.

CD-ROM Test Questions

The following questions are included within the test engine of *CCNP BSCI Exam Certification Guide*. They are provided here for additional self-assessment and review.

Question 1

Which of the following characteristics are attributable to IS-IS? (Select all that apply.)

- A: It is an EGP.
- B: It is stable and converges quickly.
- C: It is an IGP.
- D: It was originally designed for the OSI protocols.
- E: It was designed for IPv4.

Question 2

What is the purpose of ES-IS?

- A: IS-IS area discovery
- B: Allows an ES to take over if the IS crashes
- C: Registers ES names
- D: Allows ES discovery so that an IS can pass traffic within the area

Question 3

Which of the following is an advantage of IS-IS over OSPF? (Select all that apply.)

- A: An IS-IS router produces one LSP if it is routing at Level 1 and one LSP if it is routing at Level 2; OSPF produces many LSAs.
- B: IS-IS is more extensible through its use of the TLV advertisement-encoding mechanism.
- C: IS-IS uses the Dijkstra algorithm.
- D: IS-IS has more adjustable timers.
- E: IS-IS, as implemented on Cisco IOS software, produces more accurate metrics by taking into consideration three QoS parameters.

Question 4

Which of the following accurately portrays the NSAP address structure used by Cisco IOS software?

- A: AFI (2B), IDI, HODSP, System ID (6B), NSEL (2B). Maximum total length 32 bytes.
- B: AFI (1B), IDI, HODSP, System ID (6B), NSEL (1B). Maximum total length 20 bytes.
- C: AFI (2B), IDI, HODSP, System ID (arbitrary size), NSEL (1B). Maximum total length 20 bytes.
- D: AFI (1B), FBI, System ID (arbitrary size), NSEL (1B). Maximum total length 20 bytes.
- E: AFI (2B), IDP, DSP, HODSP, LODSP, System ID (6B), NSEL (2B). Maximum total length 32 bytes.

Question 5

Which of the following statements are true about the structure of IS-IS NETs? (Select all that apply.)

- A: All routers within an area must use the same system ID.
- B: All routers within an area must use a unique system ID.
- C: All routers within an area must use a unique NSEL and the same AFI.
- D: All routers within an area must use the same prefix.
- E: All routers within an area must use a unique prefix.

Question 6

How does an L1 router accomplish routing to a destination NSAP? (Select all that apply.)

- A: If the area ID of the destination is the same as the area ID of the router, the router uses its L1 database to route by system ID.
- B: If the area ID of the destination is the same as the area ID of the router, the router uses its L1 database to route by DSP.
- C: If the area ID of the destination is the same as the area ID of the router, the router is the destination.
- D: If the area ID of the destination differs from the area ID of the router, the router uses its L1 database to route by system ID.
- E: If the area ID of the destination differs from the area ID of the router, the router uses its L2 database to route by area ID.
- F: If the area ID of the destination differs from the area ID of the router, the router forwards the traffic to the nearest L1/L2 router.

Question 7

Which of the following are IS-IS PDUs? (Select all that apply.)

- A: CLNS
- B: IP
- C: Hello
- D: DIS
- E: LSP
- F: PSNP
- G: CSNP

Question 8

What is the function of a CSNP?

- A: It is used to acknowledge and request link-state information.
- B: It uses ESH, ISH, and IIH to establish and maintain adjacencies.
- C: It is used to distribute a router's complete link-state database.
- D: It is the CLNP in the OSI suite.

Question 9

Which of the following are IS-IS network types? (Select all that apply.)

- A: Point-to-multipoint NBMA mode for incomplete mesh designs
- B: Broadcast mode for LAN
- C: Point-to-point mode for point-to-point circuits, dial interfaces, and NBMA point-to-point subinterfaces
- D: Point-to-multipoint broadcast for LAN
- E: NBMA mode for Frame Relay and ATM

Question 10

How does IS-IS represent a broadcast network in its link-state database?

- A: As $n(n-1)/2$ point-to-point links
- B: As n point-to-point links to a pseudonode
- C: As n point-to-point links to a DIS and n point-to-point links to a BDS
- D: As n^2 links, simulating a full-mesh environment
- E: As one entry

Question 11

Cisco IOS software supports which of the following IS-IS metrics?

- A: Delay
- B: Default
- C: Expense
- D: Error

Question 12

What is the size of the default “narrow” metrics? (Select all that apply.)

- A: 6 bytes for total path cost
- B: 10 bits for total path cost
- C: 6 bits for interface metric
- D: 6 bytes for interface cost
- E: 24 bits for interface cost
- F: 32 bits for total path metric

Question 13

How is the IS-IS metric of an interface determined? (Select all that apply.)

- A: The default is 10 regardless of bandwidth or delay.
- B: It is determined by $\log_{10} bw$, where bw is defined as the bandwidth of the line in kilobits per second.
- C: The default is 10 for LAN links and 1 for WAN links.
- D: It is determined by $100Mb/bw$.
- E: It can be administratively set.

Question 14

How is adjacency maintained between two intermediate systems?

- A: Hellos are multicast periodically, separately for L1 and L2.
- B: Routing tables are broadcast periodically, separately for L1 and L2.
- C: CSNPs are multicast periodically, separately for L1 and L2.
- D: The DIS pings each neighbor periodically.
- E: ICMP error messages are used to learn of a lost neighbor.

Question 15

Which of the following commands explicitly display Layer 2 routing information on a L1/L2 router? (Select all that apply.)

A: **show isis topology**

B: **show clns route**

C: **show isis route**

D: **which-route**

Question 16

Which command enables you to troubleshoot adjacency formation?

A: **show clns topology**

B: **show clns route**

C: **show clns neighbors**

D: **show isis route**

E: **show isis topology**

Question 17

How is the CLNS address specified for use by IS-IS?

A:

```
Router (config)#Router isis
Router (config-router)#Nsap 49.0001.0000.0c12.3456.00
```

B:

```
Router (config)#Router isis
Router (config-router)#Net 49.0001.0000.0c12.3456.00
```

C:

```
Router (config)# Nsap 49.0001.0000.0c12.3456.00
```

D:

```
Router (config)# Net 49.0001.0000.0c12.3456.00
```

Question 18

A router has been configured as follows:

```
Router isis
  Net 49.0001.0000.0c12.3456.00
Int e0
  Ip router isis
Int s0
  Ip router isis
```

What must be done to configure this router to limit e0 to Level 1? (Select all that apply.)

A: Nothing (because level 1 routing is the default)

B:

```
Router(config-router)#is-type level-1
```

C:

```
Router(config)#int e0
Router(config-if)#isis circuit-type level-1
```

D:

```
Router(config)#int e0
Router(config-if)#is-type level-1
```

E:

```
Router(config)#is-type level-1
```

Question 19

Which of the following commands enable you to identify IS-IS neighbors?

A: **show ip protocols**

B: **show clns route**

C: **show isis route**

D: **which-route**

E: **show isis topology**

Question 20

What is the recommended mode for IS-IS across a Frame Relay network?

- A: Point-to-point subinterfaces
- B: Broadcast
- C: NBMA
- D: General topology
- E: Point-to-multipoint

Question 21

How is IS-IS OSI routing accomplished? (Select all that apply.)

- A: Level 0 routing is the process of ES passing traffic to a directly connected IS.
- B: Level 1 routing is the process of the intermediate system routing traffic through other intermediate systems within an area based on system ID.
- C: Level 2 routing is the process of IS routing traffic between areas in a domain based on the prefix.
- D: Level 3 routing is the process of routing between domains.
- E: Level 0 routing involves interdomain routing.
- F: Level 1 routing describes interarea routing.
- G: Level 2 routing describes intra-area routing.
- H: Level 3 describes routing from an ES to a directly connected IS.

Answers:

Question 1 answers: B, C, and D

Question 2 answer: D

Question 3 answers: A, B, D

Question 4 answer: B

Question 5 answers: B, D

Question 6 answers: A, F

Question 7 answers: C, E, G

Question 8 answer: C

Question 9 answers: B, C

Question 10 answer: B

Question 11 answer: B

Question 12 answers: B, C

Question 13 answers: A, E

Question 14 answer: A

Question 15 answers: A, B

Question 16 answer: C

Question 17 answer: B

Question 18 answers: B, C

Question 19 answers: A, C, E

Question 20 answer: A

Question 21 answers: A, B, C, D