Introduction to **Intermediate System-to-Intermediate** System Protocol

In recent years, the IS-IS routing protocol has become increasingly popular, with widespread usage among Service Providers. It is a link state protocol, which enables very fast convergence with large scalability. It is also a very flexible protocol and has been extended to incorporate leading edge features such as MPLS Traffic Engineering.

The IS-IS routing protocol is a link-state protocol, as opposed to distance-vector protocols such as Interior Gateway Routing Protocol (IGRP) and Routing Information Protocol (RIP). Link-state offers several advantages over distance-vector protocols. It is faster converging, supports much larger internetworks, and is less susceptible to routing loops. Features of IS-IS include:

- Hierarchical routing
- Classless behavior
- Rapid flooding of new information
- Fast Convergence
- Very scalable
- Flexible timer tuning
- Cisco IOS implementation of multi-area routing
- Cisco IOS implementation of route-leaking
- Cisco IOS implementation of overload-bit

Intermediate System-to-Intermediate System (IS-IS) Protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in International Organization for Standardization (ISO) 10589. The protocol is designed to operate in OSI Connectionless Network Service (CLNS). Data is carried using the protocol specified in ISO 8473.

A two-level hierarchy is used to support large routing domains. A large domain may be administratively divided into areas. Each system resides in exactly one area.[1] Routing within an area is referred to as Level 1 routing. Routing between areas is referred to as Level 2 routing. A Level 2 Intermediate System (IS) keeps track of the paths to destination areas. A Level 1 IS keeps track of the routing within its own area. For a packet destined for another area, a Level 1 IS sends the packet to the nearest Level 2 IS in its own area, regardless of what the destination area is. Then the packet travels via Level 2 routing to the destination area, where it may travel via Level 1 routing to the destination. It should be noted that selecting an exit from an area based on Level 1 routing to the closest Level 2 IS might result in suboptimal routing.[2]

[1] Multiarea support per Intermediate System was introduced in later implementations of IS-IS to accommodate OSI telecommunications management networks needs, but this functionality is generally not useful for IP network design.

[2] A new extension proposed in *Domain-Wide Prefix Distribution with Multi-Level IS-IS* [RFC 2966], allows routing information available at Level 2 to be leaked into a stub area relying on Level 1 routing. This mechanism addresses the suboptimal routing problem in an Integrated IS-IS environment at the expense of scalability. However, it does not apply to a CLNS environment.

On broadcast multiaccess media (LAN), a Designated Intermediate System (DIS) is elected and will conduct the flooding over the media. The DIS is analogous to the designated router in Open Shortest Path First (OSPF) Protocol, even though the details including election process and adjacencies within a multiaccess media differ significantly.The DIS is elected by priority. The highest priority becomes the DIS. This is configurable on an interface basis. In the case of a tie, the router with the highest SNPA (MAC) address will become the DIS.

CLNS

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

In the OSI architecture there are "systems": Routers are ISs, and hosts are End Systems (ESs).

ESs themselves have no routing information; they discover ISs (routers) by listening to Intermediate System Hellos (ISHs) and sending traffic to any random router. ESs send End System Hellos (ESHs); they do not choose a designated router to handle all traffic, and optimal routing is accomplished via redirects.

ISs discover ESs by listening to ESHs, and ISs send ISHs to ESs.

There is no Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP) or Interdomain Routing Protocol (IDRP) for CLNS, but End System-to-Intermediate System (ES-IS) Protocol provides the same kind of reporting functions for ISs and ESs. The ES-IS Protocol is defined in ISO 9542.

IS-IS is an Interior Gateway Protocol (IGP) for routing OSI. IS-IS packets are not encapsulated in CLNS or IP but are encapsulated directly in the data-link layer. The IS-IS protocol family is OSI, and values such as 0xFE and 0xFEFE are used by the data-link protocol to identify the Layer 3 protocol as OSI.

Table 1 shows a basic comparison of IP and OSI services.

Table 1 Comparing IP and OSI Services

Basic connectionless service:	
IP (RFC 791)	CLNS (ISO 8473)
Neighbor greeting and error	
Reports to source about packet delivery:	
ICMP (RFC 792)	CLNP (ISO 8743)
ARP (RFC 826), IRDP (RFC 1256)	ES-IS (ISO 9542)
Routing:	
Integrated IS-IS (RFC 1195)	IS-IS (ISO 10589)
Participants are routers and hosts	Participants are ISs and ESs
IP autonomous system	ISO routing domain
Interior Gateway Protocol (IGP)	Intradomain Routing Protocol
Exterior Gateway Protocol (EGP)	Interdomain Routing Protocol (IDRP)
Border Gateway Protocol (BGP) for IP (RFC 1105)	ISO IDRP (proposal)
Static IP routes	Static CLNS routes

Integrated or Dual IS-IS

The IS-IS Routing Protocol may be used as an IGP to support IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments, and dual environments. Integrated IS-IS is deployed extensively in an IP-only environment in the top-tier Internet service provider (ISP) networks. The IS-IS working group of the Internet Engineering Task Force (IETF) developed the specification for Integrated IS-IS (RFC 1195).

Two primary methods are available for routing protocols to support dual OSI and IP routers. One method, known as "Ships in the Night," makes use of completely independent routing protocols for each of the two protocol suites. This specification presents an alternative approach, which makes use of a single integrated protocol for interior routing (that is, for calculating routes within a routing domain) for both protocol suites.

By supporting both IP and OSI traffic, this integrated protocol design supports traffic to IP hosts, OSI end systems, and dual end systems. The IS-IS Protocol can be used to support pure-IP environments, pure-OSI environments, and dual environments. IS-IS allows the interconnection of dual (IP and OSI) routing domains with other dual domains, with IP-only domains, and with OSI-only domains.

IS-IS Operations

From a high level, IS-IS operates as follows:

- Routers running IS-IS will send hello packets out all IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- Routers sharing a common data link will become IS-IS neighbors if their hello packets contain information that meets the criteria for forming an adjacency. The criteria differ slightly depending on the type of media being used (p2p or broadcast). The main criteria are matching authentication, IS-type and MTU size).
- Routers may build a link-state packet (LSP) based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- Generally, routers flood LSPs to all adjacent neighbors except the neighbor from which they received the same LSP. However, there are different forms of flooding and also a number of scenarios in which the flooding operation may differ.
- All routers will construct their link-state database from these LSPs.
- A shortest-path tree (SPT) is calculated by each IS, and from this SPT the routing table is built.

IS-IS Data-Flow Diagram

In IS-IS, routers may have adjacencies with other routers on point-to-point links. In a LAN environment, routers report their adjacencies to a Designated Intermediate System (DIS), which generates an additional LSP, commonly known as the pseudonode LSP. The DIS is responsible for conducting flooding over the LAN and also for maintaining synchronization.

The flow of information within the IS-IS routing function is represented by the IS-IS data-flow diagram (Figure 1), which consists of four processes and a Routing Information Base (RIB). The RIB consists of the link-state database and the forwarding database. The four processes in the IS-IS data-flow diagram are: *receive, update, decision,* and *forward*.

Figure1 IS-IS Data-Flow Diagram

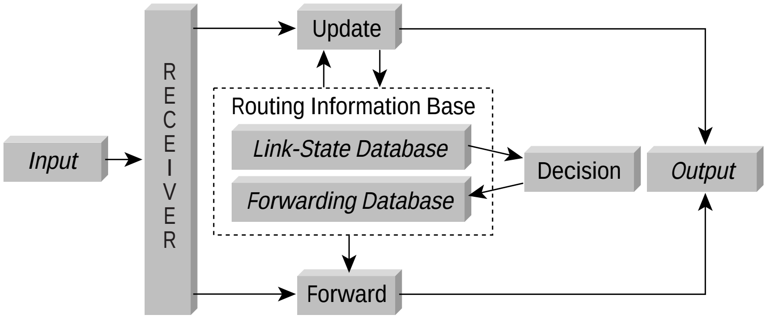

The *receive* process is the entry point for all data, including user data, error reports, routing information, and control packets. It passes user data and error reports to the forward process and passes routing information and control packets (hellos, LSPs, and sequence number packets) to the update process.

The *update* process generates local link information that is flooded to adjacent routers; in addition, the update process receives, processes, and forwards link information received from adjacent routers. This process manages the Level 1 and Level 2 link-state databases and floods Level 1 and Level 2 LSPs throughout an area.

Each LSP that resides in the link-state database has a remaining lifetime, a checksum, and a sequence number.

The LSP remaining lifetime counts down from 1200 seconds (20 minutes) to 0. The LSP originator must periodically refresh its LSPs to prevent the remaining lifetime from reaching 0. The refresh interval is 15 minutes, with a random jitter of up to 25 percent. If the remaining lifetime reaches 0, the expired LSP will be kept in the database for an additional 60 seconds (known as ZeroAgeLifetime) before it is purged.

If a router receives an LSP with an incorrect checksum, the router will cause a purge of the LSP by setting the remaining lifetime value to 0, removing the body and reflooding it. This triggers the LSP originator to send a new LSP. This behavior is different from that of OSPF, where only the originating router can purge an LSP. IS-IS can be configured so that LSPs with incorrect checksums are not purged, but the router that originated the LSP will not know that the LSP was not received.

The *decision* process runs shortest-path-first (SPF) algorithm on the link-state database, and creates the forwarding database. It computes next-hop information and computes sets of equal-cost paths, creating an adjacency set that is used for load balancing. On a Cisco router, IS-IS supports load balancing over and up to six equal-cost paths.

The *forward* process gets its input from the *receive* process and uses the forwarding database to forward data packets toward their destination. It also redirects load sharing and generates error reports.

Areas and the Routing Domain

An IS-IS routing domain is similar to a BGP autonomous system. A routing domain is a collection of areas under an administration that implements routing policies within the domain.

Backbone

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area (Figure 2).

Figure 2 IS-IS Backbone

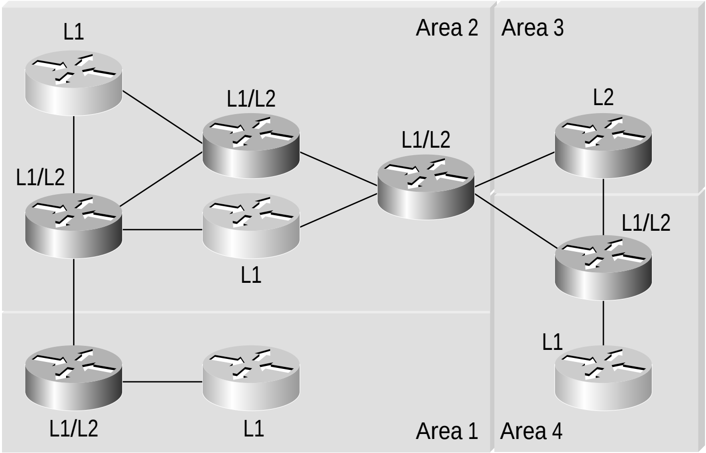

Areas

With IS-IS, an individual router is in only one area, and the border between areas is on the link that connects two routers that are in different areas (Figure 3). This is in contrast to OSPF, in which the area borders are within the Area Border Routers (ABRs) (Figure 4). The reason for this difference is that an IS-IS router generally has one network service access point (NSAP) address, and an IP router generally has multiple IP addresses.

Figure 3 OSPF Areas: Area Borders Are Within Routers

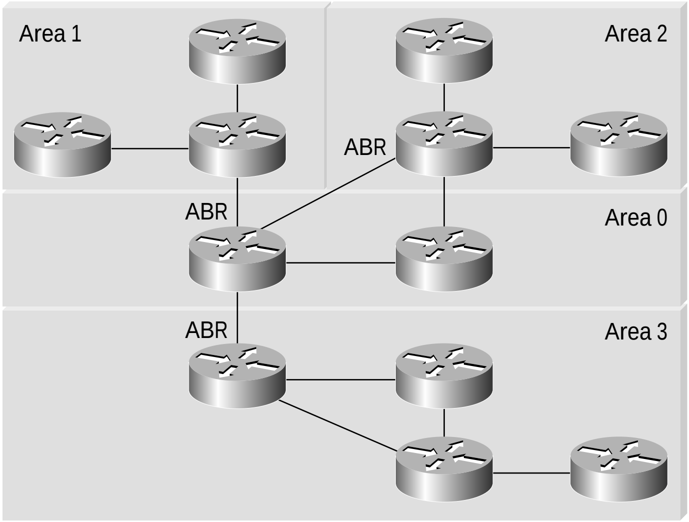

Figure 4 IS-IS Areas: Area Borders Are on Links Between Routers

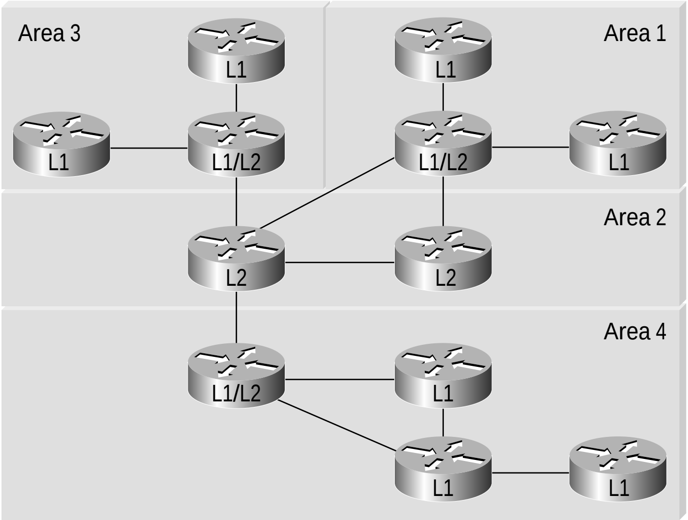

IS-IS has a two-level hierarchy. Contiguous Level 2-capable routers form the backbone. Both Level 2 and Level 1 routers live in areas. Routers can be Level 1 (L1), Level 2 (L2), or both (L1/L2). Within Cisco IOS® Software, the default configuration is both Level 1 and Level 2 at the same time which allows an IS-IS network to run with minimal configuration in a plug-and-play fashion. Level 2-capable routers connect all areas within a routing domain. Level 2 routers advertise their own area addresses (NSAP) to the other Level 2 routers in the backbone. All Level 1 routers and hosts in an area must have an NSAP with the same area address.

Level 1 Router

A Level 1 router knows the topology only of its own area and has Level 1 or Level1/Level2 neighbors in this area. It has a Level 1 link-state database with all the information for intra-area routing. It uses the closest Level 2-capable router in its own area to send packets out of the area, a scenario that may result in suboptimal routing.

Figure 5 provides an example of sub-optimal routing: The cost on all links is 10. Router A (L1) in Area X will send all traffic destined for outside Area X to Router B (L1/L2) because Router B is the closest L1/L2 neighbor. Router B is directly connected to Area Y. Router C, also L1/L2, is in Area X and is directly connected to Area Z. Router A will send packets destined for Area Z to Router B, and because Router B, Router E, and Router C are backbone routers, Router B will send this packet to Router C through Router E for delivery into Area Z. The more optimal path would be for Router A to send the packet directly to Router C through Router D.

Figure 5 Example of Sub-optimal Routing

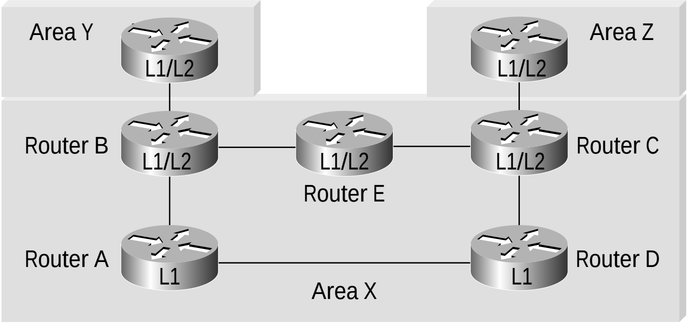

This mechanism is used to inform the Level 1 router about the closest Level 2-capable router:

A Level 1/Level 2 router that is attached to another area will set the "attached bit" in its Level 1 LSP; all the Level 1 ISs in an area will get a copy of this LSP and know where to forward packets to destinations outside the area. If the routers are running Integrated IS-IS, a default IP route will automatically be installed in the Level 1 routers pointing toward the nearest Level 1/Level 2 router that originally set the attached bit in its Level 1 LSP. A Level 1/Level 2 router that is not attached to another area can also detect that a Level 2-only neighbor is attached to another area and set the "attached bit" on behalf of this Level 2-only neighbor.

If there is more than one point to exit the area (multiple Level 2-capable routers), the closest Level 1/Level 2 router is selected based on the cost. If there are two equal cost paths then the traffic may load balance over the two paths.

Level 2 Router

A Level 2 router may have neighbors in the same or in different areas, and it has a Level 2 link-state database with all information for inter-area routing. Level 2 routers know about other areas but will not have Level 1 information from its own area. In the OSI world, a router must know the topology of its own area; so a Level 2 router should not be configured when only OSI traffic is being routed. If the traffic in an area is IP-only, all the routers can be configured as Level 2.

Level 1/Level 2 Router

A Level 1/Level 2 router may have neighbors in any area. It has two link-state databases: a Level 1 link-state database for intra-area routing and a Level 2 link-state database for inter-area routing. A Level 1/Level 2 router runs two SPFs and may require more memory and processing as a result.

A Level 1/Level 2 router running Integrated IS-IS will leak all the IP subnets from Level 1 into Level 2; these subnets can be summarized where this is desirable.

When designing a network (see figure 6), care should be taken to choose the correct setting, Level 1, Level 2, or Level 1/Level 2. When IS-IS is configured on a Cisco router, the default setting is Level 1/Level 2.

All IS-IS areas are "stub" areas, although with Cisco IOS Software Release 12.0T, it has become possible to leak Level 2 routes into Level 1, creating a sort of IS-IS not-so-stubby area.

Figure 6 A simple network running IS-IS as a routing protocol

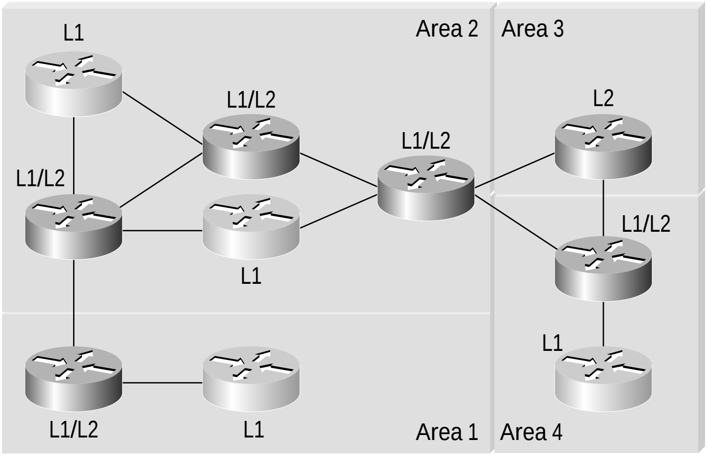

NSAP Addresses

NSAP is the network-layer address for CLNS packets. An NSAP describes an attachment to a particular service at the network layer of a node, similar to the combination of IP destination address and IP protocol number in an IP packet. NSAP encoding and format are specified by ISO 8348/Ad2.

ISO 8348/Ad2 uses the concept of hierarchical addressing domains. The global domain is the highest level. This global domain is subdivided into sub-domains, and each sub-domain is associated with an addressing authority that has a unique plan for constructing NSAP addresses.

An NSAP address (figure 7) has two major parts: the initial domain part (IDP) and the domain specific part (DSP) (Figure 7). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.

Figure 7 The NSAP address

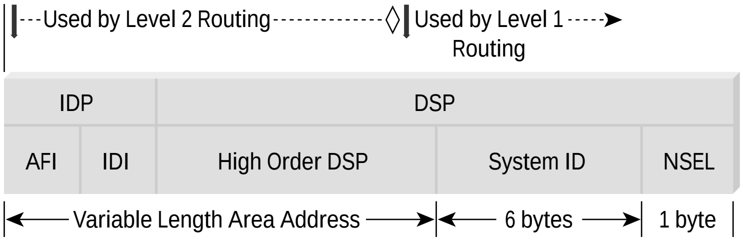

- An NSAP Consists of 3 Parts:
  Area Address, System ID, and NSAP Selector
- Total Length Between 8 and 20 Bytes
  Example: 49.0001.0000.0000.0007.00

A network entity title (NET) is an NSAP with an n-selector of zero. All router NETs have an n-selector of zero, implying the network layer of the IS itself (0 means no transport layer). For this reason, the NSAP of a router is always referred to as a NET. The NSEL (NSAP selector) is like a TCP port number: It indicates the transport layer.

Routers are identified with NETs of 8 to 20 bytes. ISO/IEC 10589 distinguishes only three fields in the NSAP address format: a variable-length area address beginning with a single octet, a system ID, and a 1-byte n-selector. Cisco implements a fixed length of 6 bytes for the system ID, which is like the OSPF router ID.

The LSP identifier is derived from the system ID (along with the pseudonode ID and LSP number). Each IS is usually configured with one NET and in one area; each system ID within an area must be unique.

The big difference between NSAP style addressing and IP style addressing is that, in general, there will be a single NSAP address for the entire router, whereas with IP there will be one IP address per interface. All ISs and ESs in a routing domain must have system IDs of the same length. All routers in an area must have the same area address. All Level 2 routers must have a unique system ID domain-wide, and all Level 1 routers must have a unique system ID area-wide. All ESs in an area will form an adjacency with a Level 1 router on a shared media segment if they share the same area address. If multiple NETs are configured on the same router, they must all have the same system ID.

There are several techniques for creating unique system IDs:

- Start numbering 1, 2, 3, 4, and so on.
- Use Media Access Control (MAC) addresses.
- Convert and use the loopback IP address: 192.168.11.1 --> 192.168.011.001--> 1921.6801.1001.

The practice of using a modified loopback IP address as the system ID may now be considered outdated because of the dynamic hostname feature. This feature uses a new Type Length Value (TLV 137) to map the router's hostname to the system ID.

Three NSAP formats are illustrated in Figure 8. The first is a simple 8-octet area ID/system ID format. The second is an OSI NSAP format, and the third is a Government OSI Profile (GOSIP) NSAP format.

Figure 8 Three NSAP Formats

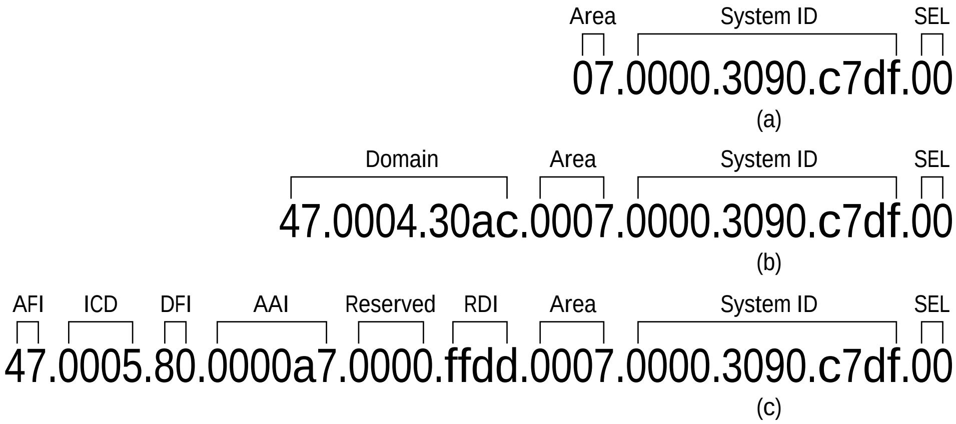

AFI: Authority and Format Identifier
ICD: International Code Designator
DFI: Domain Specific Part (DSP) Format Identifier
AAI: Administrative Authority Identifier
RDI: Routing Domain Identifier (Autonomous System Number)
SEL: Network Service Access Point (NSAP) Selector

If an official prefix is not required, you can use AFI 49, which denotes private address space, like IP address space for private Internets as defined in RFC 1918.

The AFI has a binary value between 0 and 99; this value specifies the IDI format and the DSP syntax. The DSP is in binary format, and the IDP is in decimal. ISO recognizes seven top-level network addressing authorities; each authority has its own addressing format represented by the IDI format field.

Cisco technology supports all NSAP formats that are defined by ISO 8348/Ad2.

It is possible to configure multiple NETs on a router, but no router is ever in more than one area. Configuring multiple NETs causes the areas to merge into a common area, leaking the Level 1 databases into each other. The only reasons to have multiple NETs are for splitting, merging, or renumbering areas; this method should only be used in periods of transition. Cisco Systems limits the number of configurable NETs to three per router.

Packet Types, Neighbors, Flooding, and Adjacencies

Packet Types

There are four general types of packets, and each type can be Level 1 or Level 2.

- Intermediate System-to-Intermediate System Hello (IIH)—Used by routers to detect neighbors and form adjacencies. In addition to the IIH, which is an IS-IS protocol data unit (PDU), there is an ISH and an ESH, which are End System-to-Intermediate System (ES-IS) PDUs.

- Link-state packet (LSP)—There are four types of LSPs: Level 1 pseudonode, Level 1 nonpseudonode, Level 2 pseudonode, and Level 2 nonpseudonode.
- Complete sequence number PDU (CSNP)—CSNPs contain a list of all LSPs from the current database. CSNPs are used to inform other routers of LSPs that may be outdated or missing from their own database. This ensures that all routers have the same information and are synchronized. The packets are similar to an OSPF database description packet.
- Partial sequence number PDU (PSNP)—PSNPs are used to request an LSP (or LSPs) and acknowledge receipt of an LSP (or LSPs).

The following information is included in IIH PDUs:

- Whether the PDU is a point-to-point (WAN) PDU or a LAN PDU.
- Source ID—System ID of the sending router.
- Holding time—Time period to wait to hear a hello before declaring the neighbor dead. Similar to the OSPF dead interval, the default value is three times the hello interval but can be changed with the IS-IS hello-multiplier command.
- Circuit type indicating whether the interface on which the PDU was sent is Level 1, Level 2, or Level 1/Level 2.
- PDU length.
- Local circuit ID on the sending interface (in point-to-point hello PDUs).
- LAN ID—System ID of the DIS plus the pseudonode ID (circuit ID) to differentiate LAN IDs on the same DIS.
- Priority—Higher is better. Used in DIS election (in LAN hello PDUs, there is no DR (Designated Router) election on a point-to-point link).

By default, IS-IS hellos are padded to the full maximum transmission unit (MTU) size.

The benefit of padding IIHs to the full MTU is early detection of errors caused by transmission problems with large frames or MTU mismatched on adjacent interfaces.

The drawbacks of IIH padding are that on high-speed interfaces it could be a strain on huge buffers, and on low-speed interfaces large hello PDUs waste bandwidth. This could affect time-sensitive applications such as voice over IP (VoIP).

The padding of IS-IS hellos can be turned off (in Cisco IOS Software Release 12.0[5]T and 12.0[5]S) for all interfaces on a router with the no hello padding command in router configuration mode for the IS-IS routing process, or selectively turned off for point-to-point or multipoint interfaces with the no hello padding multipoint or no hello padding point-to-point command in router configuration mode for the IS-IS routing process. Hello padding can also be turned off on an individual interface basis using the *no isis hello padding* interface configuration command.

Table 2 Hello Message Types

1	Point-to-Point IIH	Part of the IS-IS spec 10589; used by routers to form adjacencies
2	Point-to-Point IIH	Part of the IS-IS spec 10589; used by routers to form adjacencies
3	Point-to-Point IIH	Part of the IS-IS spec 10589; used by routers to form adjacencies
4	ESH	Part of the ES-IS spec 9542; similar to ICMP Router Discovery Protocol (IRDP) in TCP/IP; used for routers (ISs) and End Systems (ESs) to detect each other
5	ISH	Part of the ES-IS spec 9542; similar to IRDP in TCP/IP; used for ISs and ESs to detect each other

Even though a router sends IIH to a router it knows is connected to a link, it will also send an ISH.

Figure 9 IS-IS LAN Level 1 and Level 2 IS-IS Hello Packet Format

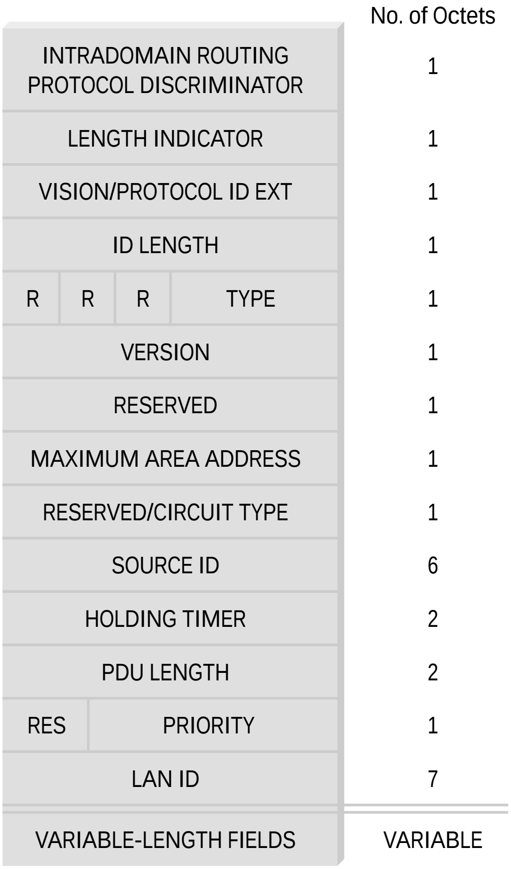

IS-IS LAN hello fields:

- Intradomain Routing Protocol discriminator—The network-layer identifier assigned to IS-IS in ISO 9577; its binary value is 10000011 (0x83).
- Length indicator—This is the length of the fixed header in octets.
- Version—It currently has value of 1.
- ID length—Length of the system ID field; must be the same for all nodes in the domain. If this is set to zero, it implies 6 octets.
- PDU Types—Values are 15 and 16 for Level 1 and Level 2 LSPs, respectively.
- Version—Value is 1.
- Maximum area addresses—Number of area addresses permitted for this IS area. Values are between 1 and 254 for actual number; 0 implies maximum of three.
- Reserved/circuit type—Top 6 bits reserved; bottom 2 bits value = 0 indicates reserved; value = 1 indicates Level 1; value = 2 indicates Level 2; value = 3 indicates Level 1 and 2.
- Source ID—System ID of transmitting router.
- Holding time—Holding time as configured on this router.
- PDU length—Length of the entire PDU, fixed header, and TLVs.

- Reserved/priority—Bit 8 reserved; bit 1 is used for priority for being the Level 1 or Level 2 DIS. Value is copied from the IIH of the DIS.
- LAN ID—A field composed of the system ID of the DIS (1–8 octets) plus a low-order octet assigned by the LAN Level 1 DIS.

Notice the variable-type length fields at the bottom of the packet. This is where the TLV information is stored. Different types of PDUs have a set of currently defined codes; any codes that are not recognized are supposed to be ignored and passed through unchanged.

The nine different types of IS-IS PDUs and the code value for the options that are valid for each of them are defined in Table 3. Codes 1 through 10 are defined in ISO 10589, and 128 through 133 are defined in RFC 1195. TLV Code 133 for authentication information is specified in RFC 1195, but Cisco technology uses the ISO Code of 10 instead. TLV Code 4 is used for partition repair and is not supported by Cisco technology.

Table 3 Valid Code Values for Nine Types of IS-IS PDUs

Level 1 LAN IS-to-IS Hello PDU Type 15	Code 1 = Area addresses Code 6 = IS neighbors (hellos) Code 8 = Padding Code 10 = Authentication information Code 129 = Protocols supported Code 132 = IP interface address
Level 2 LAN IS-to-IS Hello PDU Type 16	Code 1 = Area addresses Code 6 = IS neighbors (hellos) Code 8 = Padding Code 10 = Authentication information Code 129 = Protocols supported Code 132 = IP interface address
Point-to-point IS-IS Hello PDU Type 17	Code 1 = Area addresses Code 8 = Padding Code 10 = Authentication information Code 129 = Protocols supported Code 132 = IP interface address
Level 1 Link-State PDU Type 18	Code 1 = Area addresses Code 2 = IS neighbors (LSPs) Code 3 = ES neighbors Code 6 = IS neighbors Code 10 = Authentication information Code 128 = IP internal reachability information Code 129 = Protocols supported Code 132 = IP interface address

Table 3 Valid Code Values for Nine Types of IS-IS PDUs

Level 2 Link-State PDU Type 20	Code 1 = Area addresses Code 2 = IS neighbors (LSPs) Code 2 = IS neighbors Code 4 = Partition designated Level 2 IS Code 5 = Prefix neighbors Code 10 = Authentication information Code 128 = IP internal reachability information Code 129 = Protocols supported Code 130 = IP external reachability information Code 131 = IDRP information Code 132 = IP interface address
Level 1 Complete Sequence Number PDU Type 24	Code 9 = LSP entries Code 10 = Authentication information
Level 2 Complete Sequence Number PDU Type 25	Code 9 = LSP entries Code 10 = Authentication information
Level 1 Partial Sequence Number PDU Type 26	Code 9 = LSP entries Code 10 = Authentication information
Level 2 Partial Sequence Number PDU Type 27	Code 9 = LSP entries Code 10 = Authentication information

Figure 10 IS-IS Level 1 and Level 2 LSP Packet Format

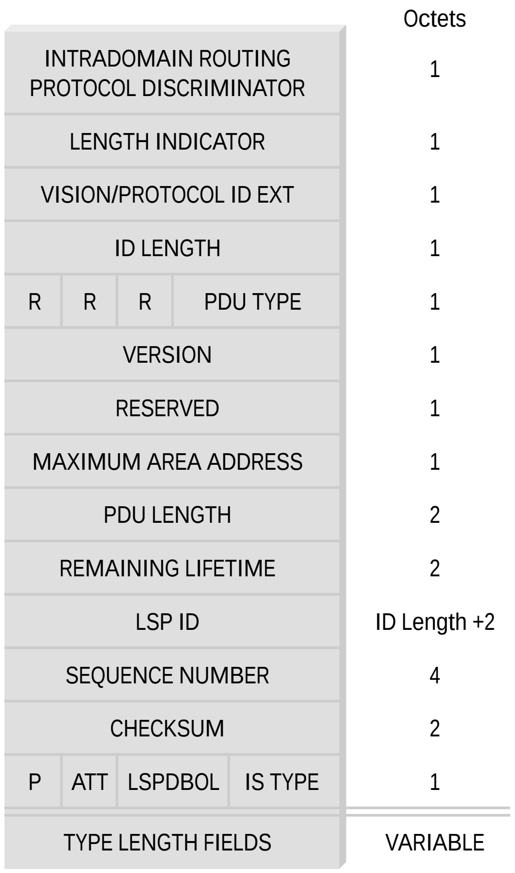

	Octets
INTRADOMAIN ROUTING PROTOCOL DISCRIMINATOR	1
LENGTH INDICATOR	1
VISION/PROTOCOL ID EXT	1
ID LENGTH	1
R R R PDU TYPE	1
VERSION	1
RESERVED	1
MAXIMUM AREA ADDRESS	1
PDU LENGTH	2
REMAINING LIFETIME	2
LSP ID	ID Length +2
SEQUENCE NUMBER	4
CHECKSUM	2
P ATT LSPDBOL IS TYPE	1
TYPE LENGTH FIELDS	VARIABLE

LSP format fields:

- Intradomain Routing Protocol discriminator—This is the network-layer identifier assigned to IS-IS in IS O9577; its binary value is 10000011.
- Length indicator—Length of the fixed header in octets.
- Version—Currently has value of 1.
- ID length—Length of the system ID field. Must be the same for all nodes in the domain; if set to zero, it implies 6 octets.
- PDU types—Decimal values are 18 and 20 for Level 1 and Level 2 LSPs, respectively.
- Version—Value is 1.
- Maximum area addresses—Number of area addresses permitted for this IS area. Values are between 1 and 254 for actual number; 0 implies a maximum of three.
- PDU length—Length of the entire PDU, fixed header, and TLVs.
- Remaining lifetime—Time in seconds before LSP expires.
- LSPID—Consists of three components: system ID, pseudonode ID, and LSP fragmentation number. Length is ID length + 2 bytes.
- Checksum—Checksum is computed from Source ID to end of PDU.
- Partition (P)—Bit 8 of the octet; when set, means originator of LSP supports partition repair.
- Attached (ATT)—Bits 4–7 of the octet. When any of these bits is set, it indicates the originator is attached to another area using the referred metric. For example, bit 4 set implies attached using the default metric.
- LSPDBOL—Bit 3. When set, it indicates the originator's LSP database is overloaded and should be circumvented in path calculations to other destinations.
- IS type—Bits 1 and 2. Only bit 1 set indicates Level 1 IS. If both are set, it indicates Level 2 IS.

Section 9 of RFC 1142 (a rewrite of ISO 10589) gives the detail about the packet layouts for each type of IS-IS PDU as well as the TLV information supported for each type. The first eight octets of all IS-IS PDUs are header fields that are common to all PDU types. As Figure 11 indicates, the Level 1 and Level 2 PDUs are identical, except for the PDU type, which differentiates them as either Level 1 or Level 2.

The lengths for the various ID fields in the PDUs: "LSP ID, SOURCE ID, START LSP ID and END LSP ID" all assume that the length of the system ID is fixed at 6 bytes. Under the column for the number of octets, an "8" means "ID LENGTH + 2," a "7" means "ID LENGTH + 1," and a "6" means "ID LENGTH." The protocol allows the system ID to vary from three to eight bytes, but in practice a six-byte system ID is always used. If the ID LENGTH field is 0, it means that the system ID is using the default length of six bytes.

Figure 11   IS-IS LAN and Point-to-Point Hello Packet Format

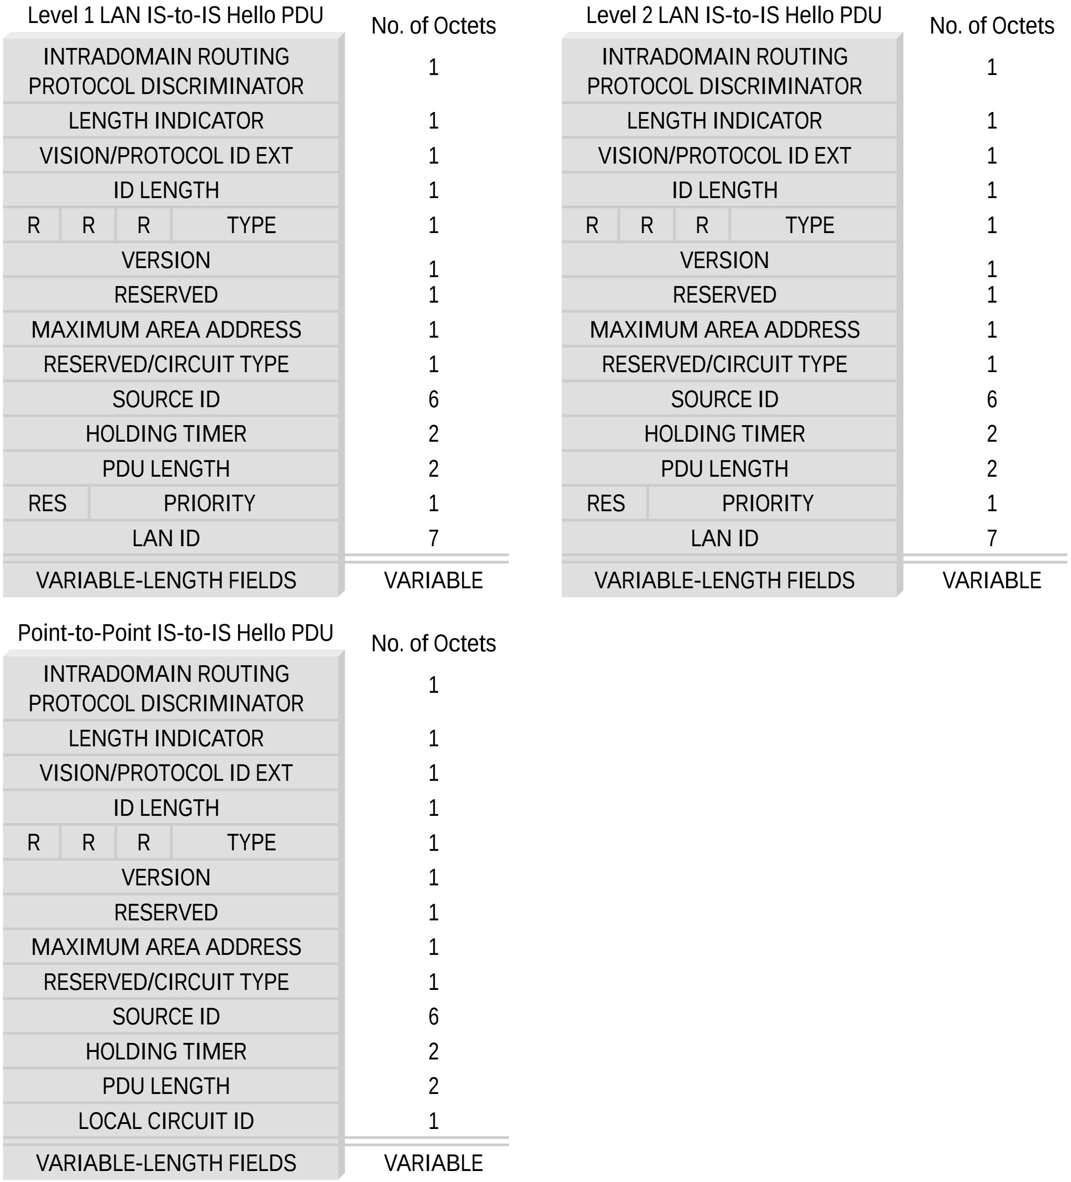

Level 1 LAN IS-to-IS Hello PDU

Level 1 LAN IS-to-IS Hello PDU	No. of Octets
INTRADOMAIN ROUTING PROTOCOL DISCRIMINATOR	1
LENGTH INDICATOR	1
VISION/PROTOCOL ID EXT	1
ID LENGTH	1
R | R | R | TYPE	1
VERSION	1
RESERVED	1
MAXIMUM AREA ADDRESS	1
RESERVED/CIRCUIT TYPE	1
SOURCE ID	6
HOLDING TIMER	2
PDU LENGTH	2
RES | PRIORITY	1
LAN ID	7
VARIABLE-LENGTH FIELDS	VARIABLE

Level 2 LAN IS-to-IS Hello PDU	No. of Octets
INTRADOMAIN ROUTING PROTOCOL DISCRIMINATOR	1
LENGTH INDICATOR	1
VISION/PROTOCOL ID EXT	1
ID LENGTH	1
R | R | R | TYPE	1
VERSION	1
RESERVED	1
MAXIMUM AREA ADDRESS	1
RESERVED/CIRCUIT TYPE	1
SOURCE ID	6
HOLDING TIMER	2
PDU LENGTH	2
RES | PRIORITY	1
LAN ID	7
VARIABLE-LENGTH FIELDS	VARIABLE

Point-to-Point IS-to-IS Hello PDU	No. of Octets
INTRADOMAIN ROUTING PROTOCOL DISCRIMINATOR	1
LENGTH INDICATOR	1
VISION/PROTOCOL ID EXT	1
ID LENGTH	1
R | R | R | TYPE	1
VERSION	1
RESERVED	1
MAXIMUM AREA ADDRESS	1
RESERVED/CIRCUIT TYPE	1
SOURCE ID	6
HOLDING TIMER	2
PDU LENGTH	2
LOCAL CIRCUIT ID	1
VARIABLE-LENGTH FIELDS	VARIABLE

Well-Known MAC Addresses

ISH packets are sent out to all IS-IS-enabled interfaces. On LANs they are sent out periodically, destined to a special multicast address. Routers will become neighbors when they see themselves in their neighbor's hello packets and link authentication information matches.

On LANs, IS-IS PDUs are forwarded to the following well-known MAC addresses:

- AllL1ISs 01-80-C2-00-00-14—The multidestination address "All Level 1 Intermediate Systems"
- AllL2ISs 01-80-C2-00-00-15—The multidestination address "All Level 2 Intermediate Systems"
- AllIntermediateSystems 09-00-2B-00-00-05—The multidestination address "All Intermediate Systems" used by ISO 9542
- AllEndSystems 09-00-2B-00-00-04—The multidestination address "All End Systems" used by ISO 9542

Adjacency Building

Neighbors on point-to-point networks always become adjacent unless they do not see themselves in their neighbors' hello PDU and match on certain parameters. On broadcast networks and nonbroadcast multiaccess (NBMA) networks, the DIS (Designated Intermediate System) will become adjacent with its neighbors.

Two routers will become neighbors if the following parameters are agreed upon:

- Level 1—The two routers sharing a common network segment must have their interfaces configured to be in the same area if they are to have a Level 1 adjacency.
- Level 2—The two routers sharing a common network segment must be configured as Level 2 if they are in different areas and want to become neighbors.
- Authentication—IS-IS allows for configuration of a password for a specified link, for an area, or for an entire domain.

The Link-State Database and Reliable Flooding

Link-State Database

All valid LSPs received by a router are stored in a link-state database. These LSPs describe the topology of an area. Routers use this link-state database to calculate its shortest-path tree.

LSPs and Reliable Flooding

Each router floods its LSPs to adjacent neighbors, and the LSPs are passed along unchanged to other adjacent routers until all the routers in the area have received them. All the Level 1 LSPs received by one router in an area describe the topology of the area.

The IS-IS link-state database consists of all the LSPs the router has received. Each node in the network maintains an identical link-state database. A change in the topology means a change in one or more of the LSPs. The router that has experienced a link going up or down will resend its LSP to inform the other routers of the change.

The LSP sequence number is increased by one to let the other routers know that the new LSP supersedes the older LSP. When a router first originates an LSP, the LSP sequence number is 1. If the sequence number increases to the maximum (oxFFFFFFFF), the IS-IS process must shut down for at least 21 minutes (MaxAge + ZeroAgeLifetime) to allow the old LSPs to age out of all the router databases.

Flooding is the process by which these new LSPs are sent throughout the network to ensure that the databases in all routers remain identical.

When a router receives a new LSP, it floods this LSP to its neighbors, except the neighbor that sent the new LSP. On point-to-point links, the neighbors acknowledge the new LSP with a PSNP, which holds the LSP ID, sequence number, checksum, and remaining lifetime. When the acknowledgment PSNP is received from a neighbor, the originating router stops sending the new LSP to that particular neighbor although it may continue to send the new LSP to other neighbors that have not yet acknowledged it. On LANs there is no explicit acknowledgement with a PSNP. Missing LSPs are detected when a CSNP is received and the list of LSPs within the CSNP is compared with the LSPs in a router's own database. If any LSPs are missing or outdated, the router will send a request for these in the form of a PSNP.

If a router receives an LSP that has an older sequence number than the one in its IS-IS database, it sends the newer LSP to the router that sent the old LSP and keeps resending it until it receives an acknowledgment PSNP from the originator of the old LSP.

LSPs must be flooded throughout an area for the databases to synchronize and for the SPF tree to be consistent within an area. It is not possible to control which LSPs are flooded by using a distribute list, although it is possible to use a routemap to control which routes are redistributed into IS-IS from another routing protocol.

Network Types

The types of networks that IS-IS defines include:

- Point-to-point networks
- Broadcast networks

Point-to-point networks, such as serial lines, connect a single pair of routers. A router running IS-IS will form an adjacency with the neighbor on the other side of a point-to-point interface. A DIS is not elected on this type of link. The basic mechanism defined in the standard is that each side of a point-to-point link declares the other side to be reachable if a hello packet is received from it. When this occurs, each side then sends a CSNP to trigger database synchronization.

Broadcast networks, such as Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI), are multiaccess in that they are able to connect more than two devices; all connected routers will receive a packet sent by one router. On broadcast networks, one IS will elect itself the DIS. Hello packets on broadcast networks are sent to the AllL1ISs or AllL2ISs MAC-layer broadcast addresses. The DIS is responsible for flooding; it will create and flood a new pseudonode LSP for each routing level it is participating in (Level 1 or Level 2) and for each LAN to which it is connected. A router can be the DIS for all connected LANs or a subset of connected LANS, depending on the configured priority or, if no priority is configured, the Layer 2 address. The DIS will also create and flood a new pseudonode LSP when a neighbor adjacency is established or torn down or the refresh interval timer for this LSP expires. The DIS mechanism reduces the amount of flooding on LANs.

NBMA networks, such as Frame Relay, Asynchronous Transfer Mode (ATM), and X.25, can connect multiple devices but have no broadcast capability. All the other routers attached to the network will not receive a packet sent by a router. Special consideration should be taken when configuring IS-IS over NBMA networks. IS-IS considers these media to be just like any other broadcast media such as Ethernet or Token Ring. In general, it is better to configure point-to-point networks on WAN interfaces and subinterfaces.

Unlike OSPF, no configuration is necessary to tell IS-IS what the network type is.

Designated Intermediate System and Pseudonodes

The idea behind the DIS is similar to that behind the designated router in OSPF. The DIS creates a pseudonode (a virtual node), and all the routers on a LAN, including the DIS, form an adjacency with the pseudonode instead of forming n*(n-1) order adjacencies with each other in a full mesh.

On a LAN, one of the routers will elect itself the DIS based on interface priority (the default is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. MAC addresses are the SNPA on LANs. On Frame Relay networks, the local data-link connection identifier (DLCI) is the SNPA. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID (in the NSAP address) will become the DIS.

A pseudonode LSP represents a LAN, including all ISs attached to that LAN, just as a non-pseudonode LSP represents a router, including all ISs and LANs connected with the router.

Figure 12 The pseudonode concept

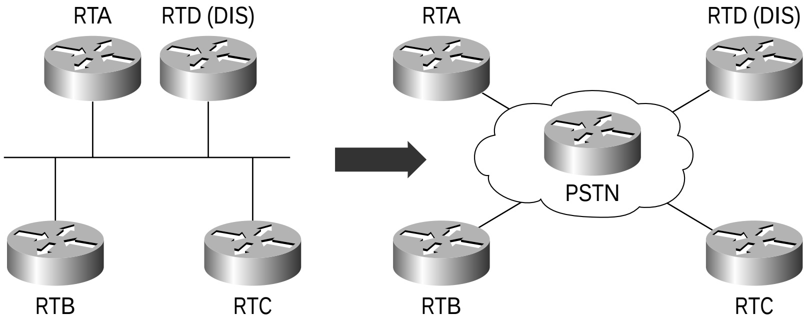

The DIS election is pre-emptive (unlike with OSPF). If a new router boots on the LAN with a higher interface priority, it becomes the DIS, purges the old pseudonode LSP, and a new set of LSPs will be flooded.

The DIS sends CSNPs describing all the LSPs in the database every 3 seconds. If a router needs an LSP because it is older than the LSP advertised by the DIS in its CSNP or it is missing an LSP that is listed in the CSNP, it will send a PSNP to the DIS and receive the LSP in return. This mechanism can work both ways: If a router sees that it has a newer version of an LSP, or it has an LSP that the DIS does not advertise in its CSNP, the router will send the newer or missing LSP to the DIS.

Metrics

The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.

The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.

Default Metric

While some routing protocols calculate the link metric automatically based on bandwidth (OSPF) or bandwidth/ delay (Enhanced Interior Gateway Routing Protocol [EIGRP]), there is no automatic calculation for IS-IS. Using old-style metrics, an interface cost is between 1 and 63 (6 bit metric value). All links use the metric of 10 by default. The total cost to a destination is the sum of the costs on all outgoing interfaces along a particular path from the source to the destination, and the least-cost paths are preferred.

The total path metric was limited to 1023 (the sum of all link metrics along a path between the calculating router and any other node or prefix). This small metric value proved insufficient for large networks and provided too little granularity for new features such as Traffic Engineering and other applications, especially with high bandwidth links. Wide metrics are also required if route-leaking is used.

Extended Metric

Cisco IOS Software addresses this issue with the support of a 24-bit metric field, the so-called “wide metric”. Using the new metric style, link metrics now have a maximum value of 16777215 (224-1) with a total path metric of 4261412864 (254 x 2^{24}).

Deploying IS-IS in the IP network with wide metrics is recommended to enable finer granularity and to support future applications such as Traffic Engineering.

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

Shortest-Path First and Link-State Database

SPF Algorithm

This section discusses the SPF algorithm for calculating routes with the IS-IS routing protocol, for support of both TCP/IP and OSI. This is based on an extension to the algorithm specified in ISO/IEC 10589.

An algorithm proposed by Edsger Dijkstra known as shortest-path first (SPF) is the basis for the route calculation. This algorithm computes the shortest paths from a single source vertex to all other vertices in a weighted, directed graph. Within the Cisco IOS implementation, weight assigned to branches of the tree is a configurable metric and spans 2^{24} per individual link and 2^{32} per path (root to leaf).

IS-IS is a link-state protocol. A distinguishing feature of a link-state protocol as it relates to distance vector protocols is that using a link-state protocol provides full visibility of the network topology, whereas distance vector protocols use information based on hearsay to build forwarding tables. This visibility provided by the link-state protocol is attained through a flooding mechanism that ensures that each router in the network (more precisely in an area) receives information that can be used to build the network map. In IS-IS, this information is flooded via link-state protocol data units,[3] and each Intermediate System or router advertises information pertaining to itself and its links.

Once the information is flooded and all routers obtain the same information, the SPF algorithm is run separately per router to process the topology and extract the shortest path from the router itself (automatically assigned at the root of the tree) to all the leaves of the tree. This is like putting a jigsaw puzzle together (Figure 13). The information derived from this process is used to populate the forwarding table on the router.

[3] Also known as link-state packets (RFC 1195)

Figure 13 LSP advertised by routers can be modeled by the jigsaw puzzle pieces

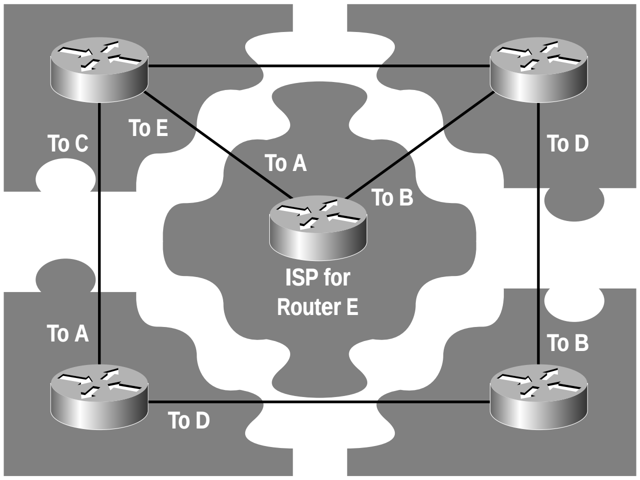

Link-State Database

- PATHS—This represents an acyclic directed graph of shortest paths from the system S performing the calculation. It is stored as a set of triples of the form <N, d(N), {Adj(N)}>, where:

  N is a system identifier. In the Level 1 algorithm, N is a 6-octet ID for OSI End Systems, a 7-octet ID for routers, or an 8-octet IP internal reachability information entry. For a router that is not a pseudonode, it is the 6-octet system ID, with a 0-appended octet. For a pseudonode, it is a true 7-octet quantity, composed of the 6-octet DIS ID and the extra octet assigned by the designated router. The IP internal reachability information entries consist of a 4-octet IP address plus a 4-octet subnet mask and will always be a leaf, that is, “End System” in PATHS.

  In the Level 2 algorithm, N is either a 7-octet router or pseudonode ID (as in the Level 1 algorithm); a variable length OSI address prefix; an 8-octet IP internal reachability information entry, or an 8-octet IP external reachability information entry. The variable-length OSI address prefixes and 8-octet IP reachability information entries will always be a leaf, that is, “End System” in PATHS. As above, the IP reachability information entries consist of an IP address-subnet mask combination.

  d(N) is N’s distance from S (that is, the total metric value from N to S).

  {Adj(N)} is a set of valid adjacencies that S may use for forwarding to N.

  When a system is placed on PATHS, the path(s) designated by its position in the graph is guaranteed to be a shortest path, and this will be reflected in the routing table, given that the route is not learned through another routing protocol with a lower administrative distance.

- TENT—This is a list of triples of the form <N, d(N), {Adj(N)}>, where N, d(N), and {Adj(N)} are as defined above for PATHS.

  TENT can intuitively be thought of as a tentative placement in PATHS. In other words, the triple <N, x, {A}> in TENT means that if N were placed in PATHS, d(N) would be x, but N cannot be placed on PATHS until is guaranteed that no path shorter than x exists.

  Similarly, the triple <N, x, {A, B}> in TENT means that if N were placed in PATHS, then d(N) would be x via either adjacency A or B.

Note: It is suggested that the implementation maintain the database TENT as a set of list of triples of the form <*,Dist,*>, sorted by Dist. In addition, it is necessary to be able to process those systems, which are pseudonodes before any nonpseudonodes at the same distance Dist.

The 8-octet system identifiers that specify IP reachability entries must always be distinguishable from other system identifiers. Two IP reachability entries that differ only in the subnet mask are still considered to be separate and will therefore have distinct system identifiers N. The SPF algorithm will therefore calculate routes to each such entry, and the correct entry will be selected in the forwarding process.

Algorithm

The function of the algorithm is to calculate the shortest path to all other nodes. To do this, a spanning tree needs to be constructed. This necessitates a logical model of nodes interconnected by point-to-point links.[4] The basic algorithm, which builds PATHS from scratch, starts out by putting the system doing the computation on PATHS (no shorter path to SELF can possibly exist). TENT is then preloaded from the local adjacency database.

For each step a node is moved into the PATHS list

- The node among all nodes on TENT that has the lowest cost from computing node is identified and moved from TENT to PATHS.
- All prefixes advertised by this node are installed in the RIB.
- All neighbors reachable from through node and moved to TENT list.

Considerations:

- If a node is directly connected to computing node, the first-hop information is obtained from the adjacency database.
- If a node is not directly connected to computing node, the first-hop information is inherited from parent node(s).
- For each node on TENT, the cost to get there from the root, and the first-hop information, is maintained.

Note that a system is not placed on PATHS unless it is the shortest path to that system. When a system N is placed on PATHS, the path to a neighbor of N, system M, is examined, as the path to N plus the link from N to M. If <M, *, *> is in PATHS, this new path will be longer and thus ignored.

If <M, *, *> is in TENT and the new path is shorter, the old entry is removed from TENT and the new path is placed in TENT. If the new path is the same length as the one in TENT, then the set of potential adjacencies {Adj(M)} is set to the union of the old set (in TENT) and the new set {Adj(N)}. If M is not in TENT, then the path is added to TENT.

Next, the algorithm finds the triple <N, x, {Adj(N)}> in TENT, with minimal x.

N is placed in PATHS. No path to N can be shorter than x at this point because all paths through systems already in PATHS have already been considered, and paths through systems in TENT still have to be greater than x because x is minimal in TENT.

When TENT is empty, PATHS is complete.

Note that external metrics can only occur in IP external reachability information entries, which correspond to a leaf (that is, “End System” in PATHS). Any route using an entry with an external metric will always be considered less desirable than any entry using an internal metric. This implies that in the addition of systems to PATHS, all systems reachable via internal routes are always added before any system reachable via external routes.

[4] Multiaccess links violate this model because multiple nodes are connected to the same link. The pseudonode was introduced to address this issue. In the shortest-path tree, a multiaccess link is modeled as a pseudonode with point-to-point links to each IS connected to the multiaccess link.

Fast Convergence at Adjacency Setup

An IS that runs out of system resources (memory, CPU, and so on) to store and process the whole link-state database is supposed to discard LSPs and to alert other routers within its area by setting the overload-bit in originated LSPs. Other routes within this area will still route packets to the overloaded router's directly connected networks, but they will not use this router for transit traffic.

When a router reloads, packets forwarded to the router could be lost, and the router would effectively behave like a black hole, dropping all the received packets for a short while. This could happen because IS-IS considers an adjacency to be established and valid before CSNP packets have been exchanged and thus before the link-state databases of neighbors have been fully synchronized.

Figure 14 Fast Convergence

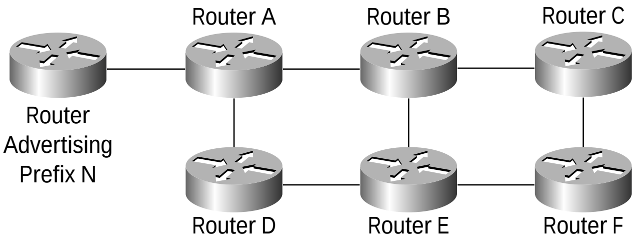

\Consider Figure 14 in the event of router B reloading. Router B would drop traffic if its neighbors include router B in their LSP while router B would not have a fully populated routing table. Thus router B would attract traffic but then not have enough information to be able to forward the traffic. This would result in packets being lost even though a valid alternate path (as used when router B is dead) exists.

Assume that the shortest path from router C to network N is through router B and router A. Router B crashes, and the network reconverges. The new path from router C to prefix N is through router F->router E ->router D->router A. However, the LSP originated from router B does not expire from the link-state databases on all routers (except router B itself). Once router B completes the reload and is back up, adjacencies with router C, router E, and Router A are established. However, before the CSNP, PSNP and full link-state database exchange is complete, router C runs SPF and uses the old version of router B LSP, hence rediscovering the old route to prefix N. Thus traffic to N is forwarded to Router B from Router C before Router B has received all LSPs, completed the SPF algorithm, and imported the prefix N into its routing table. This results in dropped packets. When router B runs an SPF based on a complete link-state database, Router B recovers full visibility and traffic to prefix N is forwarded appropriately.

In a recent IS-IS implementation in IS-IS, a router will immediately flood its own LSP even before sending CSNP packets. This does not eliminate the black-hole problem previously illustrated. However, given that the LSP is flooded immediately, the overload bit can be set to advise the rest of the network not to try routing transit traffic through the newly reloaded router. The router can be configured to advertise its LSP with the overload bit for a specific amount of time after reload.

Once the specified timer triggers, by which time the router would have a complete link-state database, the overload bit is cleared and the LSP is reflooded.

Default Routing

Default routing is achieved in two distinct ways with Integrated IS-IS:

- *Attached bit*—Set by a Level 1/Level 2 router in its own Level 1 LSP and used to indicate to all Level 1 routers (within the area) that this router is a potential exit point of the area. Level 1-only routers will default to the nearest attached Level 2 router.
- *Default information originate*—Can be configured in Level 1 as well as Level 2. The default route (0.0.0.0/0) is inserted in the router LSP (Level 1 or Level 2, according to the configuration command) and the LSP is flooded according to the router type (Level 1 or Level 2). A Level 2 router doesn't need to have a default route to originate a default route.

Level 1 routers will always prefer the explicit default route (0.0.0.0/0) found in an LSP before considering the attached bit.

Redistribution

The Integrated IS-IS specification describes the way routers are allowed to redistribute external prefixes. The Cisco implementation differs from the specification and allows more flexibility.

Redistribution from any other routing protocol, static configuration, or connected interfaces is allowed in any type of router (Level 1 and Level 2). By default the metric type will be set as *internal*, which means that the metric of the route will compete with all other internal routes. Metric type may be set to external. In that way the prefix will have a metric equal to the cost specified in the redistribution command plus a value of 64.[5] Although the metric is increased if the metric is flagged as external on redistribution, the internal/external bit used to increase the metric is actually ignored when calculating routes unless the use of external metric is specified in the configuration.

It is recommended that the metric type internal always be used when redistributing. Wide-metric TLVs do not use and discriminate between internal and external.

Summarization

Summarization is one of the key factors affecting scalability of a routing protocol. Summarization will reduce the amount of routing updates that will be flooded across the areas and the routing domain.

The use of IS-IS (especially with area routing) requires a good addressing scheme to aid summarization and avoid having a huge Level 2 database (populated with updates originated from Level 1 areas).

For IP we can summarize only native IS-IS routes into Level 2 from the Level 1 database. It is not possible to summarize IS-IS internal routes at Level 1, although it is possible to summarize external (redistributed) routes at Level 1.

[5] In earlier releases of Cisco IOS Software, the reserved bit (bit 8) was set as opposed to the internal/external bit 7. Thus external metrics were in increments of 128 rather than 64.

The key concepts for summarization of IP routes are as follows:

- Internal routes can be redistributed and summarized from Level 1 into Level 2. Summarization should be configured on the Level 1/Level 2 router, which injects the Level 1 routes into Level 2.
- If summarization is being used, all Level 1/Level 2 routers in an area must be summarizing into the backbone. If one of them is advertising more specific routes, all the traffic from the backbone will be sent toward this router because of longest-match routing.
- Internal routes cannot be summarized at Level 1 because this is not permitted by the protocol.

If the configured metric type is external (the default) or the metric type is not specified (and therefore defaults to external), redistribution will not take place. This can be confusing because the metric type does not appear in the router configuration, regardless of whether it is specified.

According to RFC 1195, the external reachability TLV is not permitted in Level 1 LSPs. Therefore, it is not possible to express an external metric type in a Level 1 LSP. This constraint will be relaxed in the future to allow Level 1 external routes to have either an internal or external metric type.

Even though the configured metric type is internal, the LSP will show that the route is “IP external” because the external IP reachability TLV is used to hold the information.

Route Leaking

Level 1 routers within an IS-IS area by default do not carry any routing information external to the area they belong to. They use a default route to exit the area. While this setup is desirable for scalability reasons, it interferes with BGP routing and Multiprotocol Label Switching (MPLS) and MPLS-VPN where all BGP next-hop addresses must be present in the local routing table. IS-IS now supports[6] a feature called “route leaking,” in which selected Level 2 routes can be advertised by a Level 1/Level 2 router into Level 1. Those leaked routes are specially tagged so they will not be re-advertised into Level 2 by another Level 1/Level 2 router.

Route-leaking imposes some risks if used in an unstable environment. An IS-IS area concept with summarization usually prevents the instabilities in one area (link or route flapping) to have an effect on other areas. Routes leaked from Level 2 into Level 1 are generally not summarized. Each time a topology change occurs in an area, all Provider Edge router addresses of this area may change metric (because of the change), and therefore the Level 1/Level 2 router that will have to propagate the PE addresses into the Level 2 core will have to recreate its Level 2 LSP. This means that leaking will have to recur and lead to a situation where for any topology change in one area you will have to re-compute (via Partial Route Calculation) many routes in (all) other areas as well. Route-leaking should therefore be planned and deployed carefully.

Security

The Cisco IS-IS implementation offers an authentication mechanism to prevent unauthorized routers from forming adjacencies or injecting TLVs. Currently, only plain-text authentication is available where the configured password is transmitted inside the IS-IS PDUs unencrypted in plain text. As such, the password can be determined by sniffing the packets. Future Cisco IOS Software releases will also contain Hashed Message Authentication Codes with MD5 (HMAC-MD5) with encrypted passwords as specified in the corresponding IETF draft.[7]

[6] This feature is also supported on OSPF ABRs to leak Type 3 link-state advertisements (LSAs) into an OSPF area.

IS-IS authentication is configured independently for adjacency establishment (hello) and for LSP authentication. The next sections describe both cases. If only LSP authentication is used, an unauthorized neighbor can still form an adjacency, but LSP packets cannot be exchanged. The ISIS database will not contain any entries for this neighbor.

LSP Authentication (Area- or Network-Wide)

Domain and area authentication can be configured separately. It is also possible to run either one or both types of authentication.

The area password is inserted and checked for Level 1 LSPs, the domain password for Level 2 LSPs.

Interface Password

This authentication can be used to make sure that only authorized neighbors form an IS-IS adjacency over a given interface. The password is configured on a per-interface base and must be specified for Level 1 and Level 2 independently.

Note: Over point-to-point interfaces, only one hello packet is sent for Level 1/Level 2 adjacencies. On those interfaces, the same password must be configured as a Level 1 and Level 2 password.

[7] http://www.ietf.org/internet-drafts/draft-ietf-isis-hmac-03.txt

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: 65 317 7777
Fax: 65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe