

# OSI Routing

---

## Background

Several routing protocols have been or are being developed under the auspices of the *International Organization for Standardization (ISO)*. ISO refers to the *Intermediate System to Intermediate System Intradomain Routing Exchange Protocol (IS-IS)* as ISO 10589. The American National Standards Institute (ANSI) X3S3.3 (network and transport layers) committee was the motivating force behind ISO standardization of IS-IS. Other ISO protocols associated with routing include ISO 9542 (*End System-to-Intermediate System*, or *ES-IS*) and ISO 10747 (*IS-IS Interdomain Routing Protocol*, or *IDRP*). Both of these protocols are discussed briefly in this chapter, but the focus is on the intradomain version of IS-IS.

IS-IS is based on work originally done at Digital Equipment Corporation for Phase V DECnet. Although IS-IS was created to route in ISO *Connectionless Network Protocol (CLNP)* networks, a version has since been created to support both CLNP and *Internet Protocol (IP)* networks. This variety of IS-IS is usually referred to as *Integrated IS-IS* and has also been called *Dual IS-IS*. Integrated IS-IS is also discussed briefly.

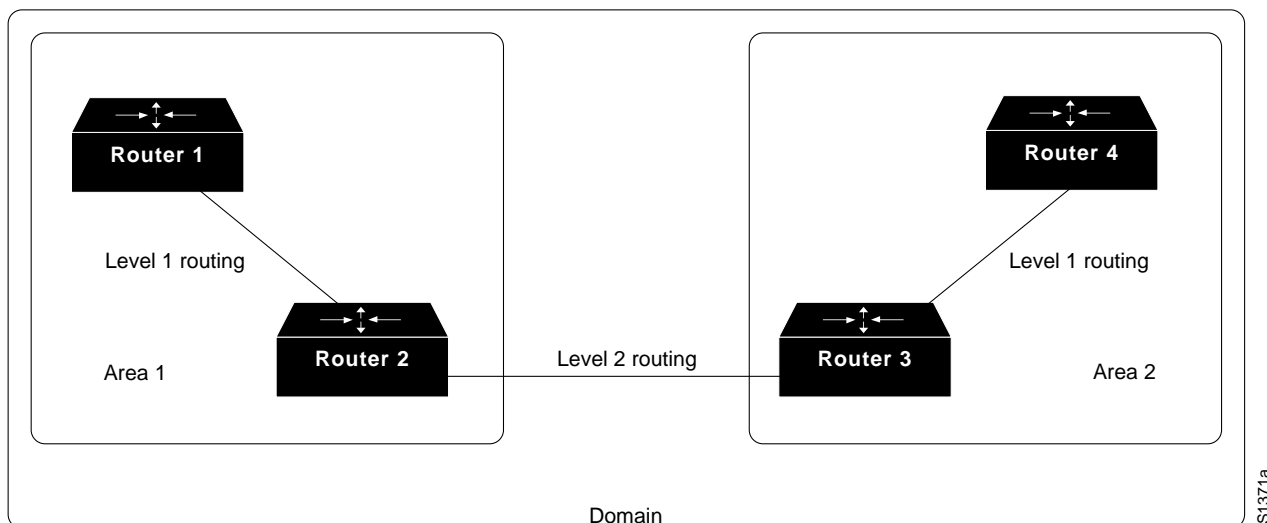
## Terminology

The world of Open Systems Interconnection (OSI) internetworking has a unique terminology. The term *end system (ES)* refers to any nonrouting network node; the term *intermediate system (IS)* refers to a router. These terms are the basis for the OSI protocols ES-IS (which allows ESs and ISs to discover each other) and IS-IS (which provides routing between ISs). Several other important OSI internetworking terms are defined as follows:

- *Area*—A group of contiguous networks and attached hosts that are specified to be an area by a network administrator or manager.
- *Domain*—A collection of connected areas. Routing domains provide full connectivity to all end systems within them.
- *Level 1 routing*—Routing within a Level 1 area.
- *Level 2 routing*—Routing between Level 1 areas.

Figure 28-1 shows the relationship between these terms.

Figure 28-1 Hierarchies in OSI Internetworks



S1371a

From a purely technological standpoint, IS-IS is quite similar to the *Open Shortest Path First* (OSPF) routing protocol. (For more information about OSPF, see Chapter 25, “Open Shortest Path First.”) Both are link state protocols. Both offer a variety of features not provided by the *Routing Information Protocol* (RIP), including *routing hierarchies*, *path splitting*, *type-of-service* (TOS) support, *authentication*, support for multiple network-layer protocols, and (with Integrated IS-IS) support for *variable length subnet masks*.

## ES-IS

ES-IS is more a discovery protocol than a routing protocol. Through ES-IS, ESs, and ISs learn about each other. This process is known as *configuration*. And, because configuration must happen before routing between ESs can occur, ES-IS is discussed here.

ES-IS distinguishes between three different types of subnetworks:

- *Point-to-point subnetworks*—Provide a point-to-point link between two systems. Many wide-area network (WAN) serial links are point-to-point networks.
- *Broadcast subnetworks*—Direct a single physical message to all nodes on the subnetwork. Ethernet and IEEE 802.3 are examples of broadcast subnetworks. For more information about Ethernet and IEEE 802.3, see Chapter 5, “Ethernet/IEEE 802.3.”
- *General-topology subnetworks*—Support an arbitrary number of systems. However, unlike broadcast subnetworks, the cost of an n-way transmission scales directly with the subnetwork size on a general-topology subnetwork. X.25 is an example of a general-topology subnetwork. For more information about X.25, see Chapter 12, “X.25.”

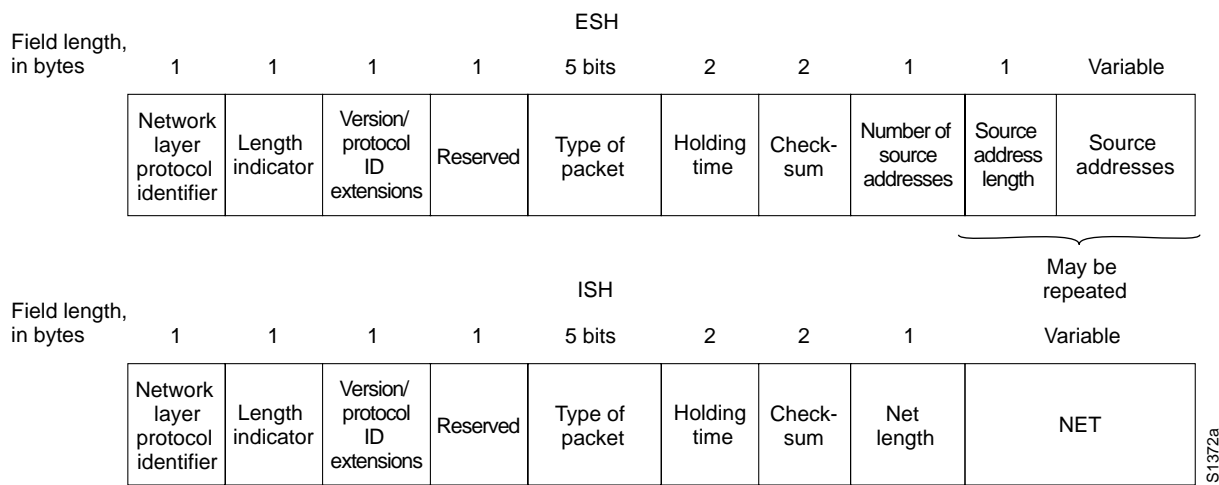
Configuration information is transmitted at regular intervals through two types of messages. *ES hello messages* (ESHs) are generated by ESs and sent to every IS on the subnetwork. *IS hello messages* (ISHs) are generated by ISs and sent to all ESs on the subnetwork. These hello messages are primarily intended to convey the subnetwork and network-layer addresses of the systems that generate them.

Where possible, ES-IS attempts to send configuration information to many systems simultaneously. On broadcast subnetworks, ES-IS hello messages are sent to all ISs through a special multicast address. ISs send hello messages to a special multicast address designating all end systems. When operating on a general-topology subnetwork, ES-IS generally does not transmit configuration information because of the high cost of multicast transmissions.

ES-IS conveys both network-layer addresses and subnetwork addresses. OSI network-layer addresses identify either the *network service access point* (NSAP), which is the interface between Layer 3 and Layer 4, or the *network entity title* (NET), which is the network layer entity in an OSI IS. OSI subnetwork addresses (sometimes called *subnetwork point of attachment* addresses, or *SNPAs*) are the points at which an ES or IS is physically attached to a subnetwork. The SNPA address uniquely identifies each system attached to the subnetwork. In an Ethernet network, for example, the SNPA is the 48-bit *Media Access Control* (MAC) address. Part of the configuration information transmitted by ES-IS is the NSAP-to-SNPA or NET-to-SNPA mapping.

Figure 28-2 shows the frame formats of both ESH and ISH packets.

**Figure 28-2 ESH and ISH Packet Formats**



## IS-IS

IS-IS is a link state routing protocol. It floods the network with link state information in order to build a complete, consistent picture of network topology.

### Routing Hierarchy

To simplify router design and operation, IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs know how to communicate with other Level 1 ISs in the same area. Level 2 ISs know how to communicate with ISs in other areas. To summarize, Level 1 ISs form Level 1 areas; Level 2 ISs route between Level 1 areas.

Level 2 ISs form an intradomain routing backbone. In other words, Level 2 ISs can get to other Level 2 ISs by traversing only Level 2 ISs. The backbone simplifies design because Level 1 ISs now only need to know how to get to the nearest Level 2 IS. The backbone routing protocol can also change without impacting the intra-area routing protocol.

Inter-ES Communication

OSI routing is accomplished as follows. Each ES lives in a particular area. ESs discover the nearest IS by listening to ISH packets. When an ES wants to send a packet to another ES, it sends the packet to one of the ISs on its directly attached network. The router looks up the destination address and forwards the packet along the best route. If the destination ES is on the same subnetwork, the local IS will know this from listening to ESHs, and will forward the packet appropriately. In this case, the IS may also provide a *redirect* (RD) message back to the source to tell it that a more direct route is available. If the destination address is an ES on another subnetwork in the same area, the IS will know the correct route and will forward the packet appropriately. If the destination address is an ES in another area, the Level 1 IS sends the packet to the nearest Level 2 IS. Forwarding through Level 2 ISs continues until the packet reaches a Level 2 IS in the destination area. Within the destination area, ISs forward the packet along the best path until the destination ES is reached.

Each IS generates an update specifying the ESs and ISs to which it is connected, as well as the associated metrics. The update is sent to all neighboring ISs, which forward (flood) it to their neighbors, and so on. Sequence numbers terminate the flood and distinguish old updates from new ones. Because all ISs receive link state updates from all other ISs, each IS can build a complete full topology database. When the topology changes, new updates are sent.

Metrics

IS-IS uses a single required default metric with a maximum path value of 1,024. The metric is arbitrary and is typically assigned by a network administrator. Any single link can have a maximum value of 64. Path lengths are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types, while at the same time ensuring that the shortest path algorithm used for route computation would be reasonably efficient.

IS-IS also defines three additional metrics (costs) as an option for those administrators who feel they are necessary. The *delay* cost reflects the amount of delay on the link. The *expense* cost reflects the communications cost associated with using the link. The *error* cost reflects the error rate of the link.

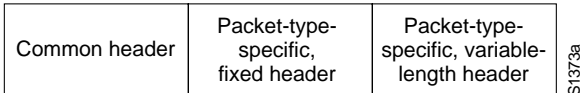
IS-IS maintains a mapping of these four metrics to the *quality of service* (QOS) option in the CLNP packet header. Using these mappings, IS-IS can compute routes through the internetwork.

Packet Formats

- IS-IS uses three basic packet formats:
- *IS-IS hello packets*
  - *Link state packets (LSPs)*
  - *Sequence numbers packets (SNPs)*


Each of the three IS-IS packets has a complex format with three different logical parts. The first part is an 8-byte fixed header shared by all three packet types. The second part is a packet-type-specific portion with a fixed format. The third logical part is also packet-type-specific, but is of variable length. The logical format of IS-IS packets is shown in Figure 28-3.

Figure 28-3 IS-IS Logical Packet Format



Each of the three packet types shares a common header, as shown in Figure 28-4.

**Figure 28-4 IS-IS Common Header Format**



The fields of the IS-IS common header are as follows:

- *Protocol identifier*—Identifies the IS-IS protocol. This field contains a constant (131).
- *Header length*—Contains the fixed header length. The length is always equal to 8 bytes, but is included so that IS-IS packets do not differ significantly from CLNP packets.
- *Version*—Contains a value of 1 in the current IS-IS specification.
- *ID length*—Specifies the size of the ID portion of an NSAP address. If the field contains a value between 1 and 8 inclusive, the ID portion of an NSAP address is that number of bytes. If the field contains a value of zero, the ID portion of an NSAP address is 6 bytes. If the field contains a value of 255 (all ones), the ID portion of an NSAP address is zero bytes.
- *Packet type*—Specifies the type of IS-IS packet (hello, LSP, or SNP).
- *Version*—Repeated after the packet type field.
- *Reserved*—Ignored by the receiver and is equal to zero.
- *Maximum area addresses*—Specifies the number of addresses permitted in this area.

Following the common header, each packet type has a different additional fixed portion, followed by a variable portion.

## Integrated IS-IS

Integrated IS-IS is a version of IS-IS that uses a single routing algorithm to support more network-layer protocols than just CLNP. Integrated IS-IS is sometimes called Dual IS-IS, after a version designed for IP and CLNP networks.

Several fields are added to IS-IS packets to allow IS-IS to support additional network layers. These fields inform routers about the following:

- The reachability of network addresses from other protocol suites
- Which protocols are supported by which routers
- Other information required by a specific protocol suite

Integrated IS-IS represents one of two ways of supporting multiple network-layer protocols in a router, the other being the *ships-in-the-night* approach. Ships-in-the-night advocates the use of a completely separate and distinct routing protocol for each network protocol, so that the multiple routing protocols essentially exist independently (with different types of routing information passing like ships in the night). The ability to route multiple network-layer protocols through tables calculated by a single routing protocol saves some router resources.

## Interdomain Routing Protocol (IDRP)

IDRP is the OSI protocol designed to move information between routing domains. As such, it is designed to operate seamlessly with CLNP, ES-IS, and IS-IS. IDRP is based on the *Border Gateway Protocol* (BGP), an interdomain routing protocol that originated in the IP community. For more information on BGP, see Chapter 27, “Border Gateway Protocol.”

IDRP introduces several new terms, including the following:

- *Border intermediate system* (BIS)—An IS that participates in interdomain routing. As such, it uses IDRP.
- *Routing domain* (RD)—A group of ESs and ISs operating under the same set of administrative rules, including the sharing of a common routing plan.
- *Routing domain identifier* (RDI)—A unique RD identifier.
- *Routing information base* (RIB)—The routing database used by IDRP. RIBs are built by each BIS from information received from within the RD and from other BISs. A RIB contains the set of routes chosen for use by a particular BIS.
- *Confederation*—A group of RDs. The confederation appears to RDs outside the confederation as a single RD. A confederation’s topology is not visible to RDs outside the confederation. Confederations help reduce network traffic by acting as internetwork firewalls and may be nested within one another.

An IDRP route is a sequence of RDIs. Some of these RDIs can be confederations. Each BIS is configured to know the RD and confederations to which it belongs, and learns about, other BISs, RDs, and confederations through information exchanges with each neighbor. As with distance vector routing, routes to a particular destination accumulate outward from the destination. Only routes that satisfy a BIS’s local policies and have been selected for use will be passed on to other BISs. Route recalculation is partial and occurs when one of three events occurs: an incremental routing update with new routes is received; a BIS neighbor goes down; or a BIS neighbor comes up.

IDRP features include the following:

- Support for CLNP QOS
- Loop suppression by keeping track of all RDs traversed by a route
- Reduction of route information and processing by using confederations, the compression of RD path information, and other means
- Reliability by using a built-in reliable transport
- Security by using cryptographic signatures on a per-packet basis
- Route servers
- RIB refresh packets