

Open Systems Interconnection (OSI) Routing Protocol

Background

The International Organization for Standardization (ISO) developed a complete suite of routing protocols for use in the Open Systems Interconnection (OSI) protocol suite. These include *Intermediate System-to-Intermediate Systems (IS-IS)*, *End System-to-Intermediate System (ES-IS)*, and *Interdomain Routing Protocol (IDRP)*. This chapters addresses the basic operations of each of these protocols.

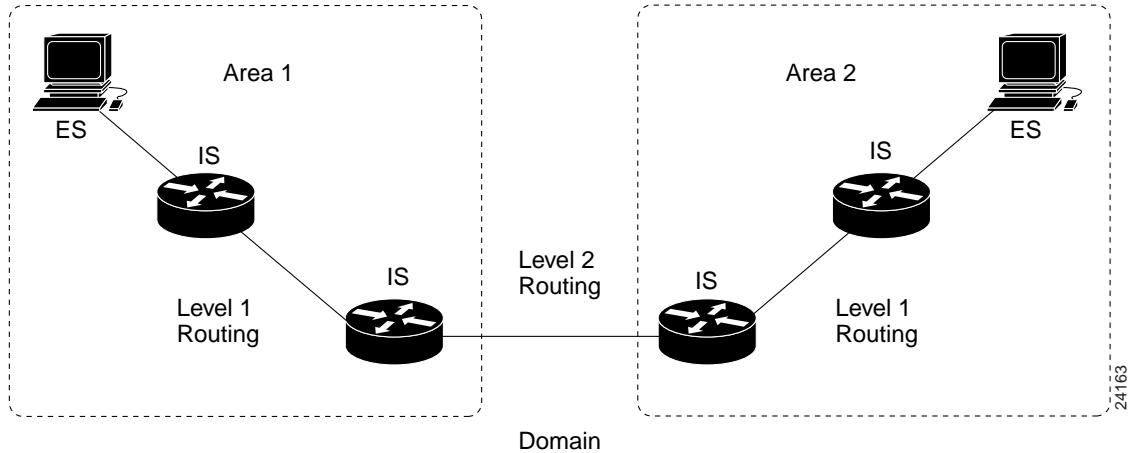
IS-IS is based on work originally done at Digital Equipment Corporation for DECnet/OSI (DECnet Phase V). IS-IS originally was developed to route in ISO *Connectionless Network Protocol (CLNP)* networks. A version has since been created that supports both CLNP and *Internet Protocol (IP)* networks; this version usually is referred to as *Integrated IS-IS* (it also has been called *Dual IS-IS*).

OSI routing protocols are summarized in several ISO documents, including ISO 10589, which defines IS-IS. 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 documents include ISO 9542 (which defines ES-IS) and ISO 10747 (which defines IDRP).

OSI Networking Terminology

The world of OSI networking uses some specific terminology, such as *end system (ES)*, which refers to any nonrouting network nodes, and *intermediate system (IS)*, which refers to a router. These terms form the basis for the ES-IS and IS-IS OSI protocols. The ES-IS protocol enables ESs and ISs to discover each other. The IS-IS protocol provides routing between ISs. Other important OSI networking terms include *area*, *domain*, *Level 1 routing*, and *Level 2 routing*. An area is a group of contiguous networks and attached hosts that is specified to be an area by a network administrator or manager. A domain is a collection of connected areas. Routing domains provide full connectivity to all end systems within them. Level 1 routing is routing within a Level 1 area, while Level 2 routing is routing between Level 1 areas. Figure 41-1 illustrates the relationship between areas and domains and depicts the levels of routing between the two.

Figure 41-1 Areas exist within a larger domain and use Level 2 routing to communicate.




End System-to-Intermediate System (ES-IS)

End System-to-Intermediate System (ES-IS) is an OSI protocol that defines how end systems (hosts) and intermediate systems (routers) learn about each other, a process known as *configuration*. Configuration must happen before routing between ESs can occur.

ES-IS is more of a *discovery* protocol than a routing protocol. It distinguishes between three different types of subnetworks: *point-to-point subnetworks*, *broadcast subnetworks*, and *general-topology subnetworks*. Point-to-point subnetworks, such as WAN serial links, provide a point-to-point link between two systems. Broadcast subnetworks, such as Ethernet and IEEE 802.3, direct a single physical message to all nodes on the subnetwork. General-topology subnetworks, such as X.25, support an arbitrary number of systems. Unlike broadcast subnetworks, however, the cost of an *n*-way transmission scales directly with the subnetwork size on a general-topology subnetwork. Figure 41-2 illustrates the three types of ES-IS subnetworks.

Figure 41-2 ES-IS can be deployed in point-to-point, broadcast, and general topology subnetworks.



ES-IS Configuration

ES-IS configuration is the process whereby ESs and ISs discover each other configuration that routing between ESs can occur. ES-IS configuration information is transmitted at regular intervals through two types of messages: *ES hello messages* (ESHs) and *IS hello messages* (ISHs). ESHs are generated by ESs and sent to every IS on the subnetwork. ISHs are generated by ISs and sent to all ESs on the subnetwork. These hello messages primarily are intended to convey the subnetwork and network-layer addresses of the systems that generate them. Where possible, ES-IS attempts to send configuration information simultaneously to many systems. On broadcast subnetworks, ES-IS hello messages are sent to all ISs through a special multicast address that designates 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 Addressing Information

The ES-IS configuration protocol conveys both OSI network-layer and OSI subnetwork addresses. OSI network-layer addresses identify either the *network service-access point* (NSAP), which is the interface between OSI Layer 3 and Layer 4, or the *network entity title* (NET), which is the network-layer entity in an OSI IS. OSI subnetwork addresses, or subnetwork *point-of-attachment* addresses (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.

Intermediate System-to-Intermediate System (IS-IS)

Intermediate System-to-Intermediate System (IS-IS) is an OSI link-state hierarchical routing protocol that floods the network with link-state information to build a complete, consistent picture of network topology. To simplify router design and operation, IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Hierarchical routing simplifies backbone design because Level 1 ISs only need to know how to get to the nearest Level 2 IS. The backbone routing protocol also can change without impacting the intra-area routing protocol.

OSI Routing Operation

Each ES lives in a particular area. OSI routing begins when the ESs discovers 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 then looks up the destination address and forwards the packet along the best route. If the destination ES is on the same subnetwork, the local IS will know this from listening to ESHs and will forward the packet appropriately. The IS also might 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.

Link-state update messages help ISs learn about the network topology. First, each IS generates an update specifying the ESs and ISs to which it is connected, as well as the associated metrics. The update then 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.) Using these updates topology of the network. When the topology changes, new updates are sent.

IS-IS Metrics

IS-IS uses a single required default metric with a maximum path value of 1,024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links 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 will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error. The delay cost metric reflects the amount of delay on the link. The expense cost metric reflects the communications cost associated with using the link. The error cost metric 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. IS-IS uses these mappings to compute routes through the internetwork.

IS-IS packets consist of eight fields.

IS-IS uses three basic packet formats: *IS-IS hello packets*, *link-state packets (LSPs)*, and *sequence-numbers packets (SNPs)*. Each of the three IS-IS packets has a complex format with the following three different logical parts. The first part consists of 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 part is also packet-type-specific but of variable length. Figure 41-3 illustrates the logical format of IS-IS packets. Figure 41-4 shows the common header fields of the IS-IS packets.

Figure 41-3 IS-IS packets consist of three logical headers.

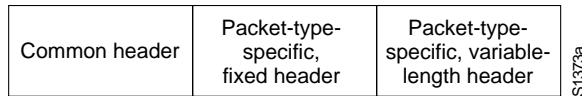



Figure 41-4 IS-IS packets consist of eight fields.



The following descriptions summarize the fields illustrated in Figure 41-4:

- *Protocol Identifier*—Identifies the IS-IS protocol and contains the constant 131.
- *Header Length*—Contains the fixed header length. The length always is 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*—Repeats after the packet type field.
- *Reserved*—Is ignored by the receiver and is equal to 0.
- *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 the OSI IS-IS routing protocol that uses a single routing algorithm to support more network-layer protocols than just CLNP. Integrated IS-IS is sometimes called Dual IS-IS, named 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 reachability of network addresses from other protocol suites and other information required by a specific protocol suite. Integrated IS-IS implementations send only one set of routing updates, which is more efficient than two separate implementations.

Integrated IS-IS represents one of two ways of supporting multiple network-layer protocols in a router; the other is the *ships-in-the-night* approach. Ships-in-the-night routing advocates the use of a completely separate and distinct routing protocol for each network protocol so that the multiple routing protocols essentially exist independently. Essentially, the different types of routing information pass like ships in the night. Integrated routing has the capability to route multiple network-layer protocols through tables calculated by a single routing protocol, thus saving some router resources. Integrated IS-IS uses this approach.

Interdomain Routing Protocol (IDRP)

The Interdomain Routing Protocol is an OSI protocol that specifies how routers communicate with routers in different domains. IDRP 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. IDRP features include the following:

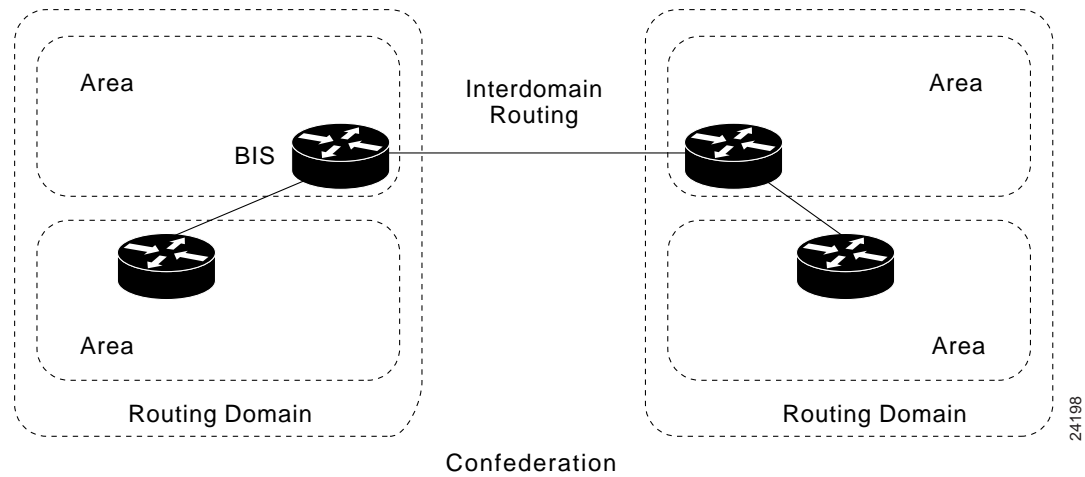
- Support for CLNP quality of service (QoS)
- Loop suppression by keeping track of all RDs traversed by a route
- Reduction of route information and processing by using confederations, the compression of RD path information, and other means
- Reliability by using a built-in reliable transport
- Security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

IDRP Terminology

IDRP introduces several environment-specific terms. These include *border intermediate system* (BIS), *routing domain* (RD), *routing-domain identifier* (RDI), *routing-information base* (RIB), and *confederation*. A BIS is an IS that participates in interdomain routing and, as such, uses IDRP. An RD is a group of ESs and ISs that operate under the same set of administrative rules and share a common routing plan. An RDI is a unique RD identifier. A RIB is a routing database used by IDRP that is 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. A confederation is a group of RDs that appears to RDs outside the confederation as a single RD. The confederation's topology is not visible

to RDs outside the confederation. Confederations must be nested within one another and help reduce network traffic by acting as internetwork firewalls. Figure 41-5 illustrates the relationship between IDRP entities.

Figure 41-5 Domains communicate via border intermediate systems (BISs).



IDRP Routing

An IDRP route is a sequence of RDIs, some of which can be confederations. Each BIS is configured to know the RD and the confederations to which it belongs. It 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.