Overview of the ISO IS-IS Intra-Domain Routing Protocol

Ross Callon

Dave Oran

Radia Perlman

November, 1991.

1 ISO IS-IS

The IS-IS Routing Protocol [1] has been developed in ISO to provide routing for pure OSI environments. In particular, IS-IS is designed to work in conjunction with ISO 8473 (The ISO Connectionless Network Layer Protocol [2]), and ISO 9542 (The ISO End System to Intermediate System Protocol [3]). This paper briefly describes the manner in which IS-IS is used to support pure OSI environments. Enhancements for support of IP and multi-protocol environments have been specified [4,5], but are beyond the scope of this paper.

NOTE: The ISO term for a router is an "intermediate system". IS-IS is so named because it is a protocol which runs between routers. In other words, it is an "intermediate system to intermediate system" (or IS to IS) protocol. In this paper, the terms "router", "intermediate system", and "IS" are used interchangeably.

IS-IS provides an "intra-domain" routing protocol. That is, it provides a standard for routing within a routing domain. The boundaries of routing domains are defined by network management, by setting some links to be “external domain”. If a link is marked as “external domain”, no IS-IS routing messages are sent on that link.

ISO has not yet completed a standard for inter-domain routing. The OSI Inter-domain routing protocol is currently at the "commitee draft" stage [6], and is expected to become a full standard before the end of 1992. However, the inter-domain routing protocol is not yet widely available in products. Instead, manual configuration is used. The link is statically configured with the set of address prefixes reachable via that link, and with the method by which they can be reached (such as the DTE address to be dialed to reach that address, or the fact that the DTE address should be extracted from the IDP portion of the ISO address).

OSI IS-IS routing makes use of two-level hierarchical routing. A routing domain is partitioned into “areas”. Level 1 routers know the topology in their area, including all routers and end systems in their area. However, level 1 routers do not know the identity of routers or destinations outside of their area. Level 1 routers forward all traffic for destinations outside of their area to a level 2 router in their area. Similarly, level 2 routers know the level 2 topology, and know which addresses are reachable via each level 2 router. However, level 2 routers do not need to know the topology within any level 1 area, except to the extent that a level 2 router may also be a level 1 router within a single area. Only level 2 routers can exchange data packets or routing information directly with external routers located outside of the routing domains.

ISO provides for flexible, variable length addresses (known as Network Service Access Point Addresses, or NSAPs [7]), which allow for multi-level hierarchical address assignment. These addresses provide the flexibility needed to simultaneously solve two critical problems: (i) How to administer a worldwide address space; and (ii) how to assign addresses in a manner which makes routing feasible in a worldwide internet. However, assignment of addresses needs to be done with great care, if the potential advantages of this addressing flexibility are to be realized in actual networks [8]. The ISO NSAP address format is illustrated in figure 1.

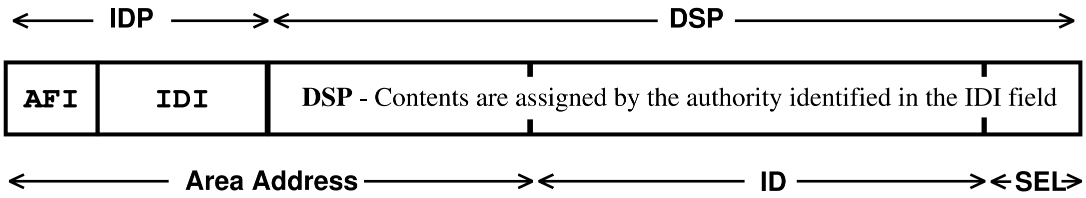

Figure 1: ISO Hierarchical Address Structure

For administrative purposes, ISO addresses are subdivided into the Initial Domain Part (IDP), and the Domain Specific Part (DSP). The IDP is the part which is standardized by ISO, and specifies the format and authority responsible for assigning the rest of the address. The IDP is further subddivided into the Authority and Format Identifier (AFI) and the Initial Domain Identifier (IDI). The DSP is assigned by whatever addressing authority is specified by the IDP.

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

Usually, all nodes in an area have the same area address. However, sometimes an area might have multiple addresses. Motivations for allowing this are:

- It might be desirable to change the address of an area. The most graceful way of changing an area from having address A to having address B is to first allow it to have both addresses A and B, and then after all nodes in the area have been modified to recognize both addresses, then one by one the nodes can be modified to “forget” address A.

- It might be desirable to merge areas A and B into one area. The method for accomplishing this is to, one by one, add knowledge of address B into the A partition, and similarly add knowledge of address A into the B partition.

- It might be desirable to partition an area C into two areas, A and B (where “A” might equal “C”, in which case this example becomes one of removing a portion of an area). This would be accomplished by first introducing knowledge of address A into the appropriate nodes (those destined to become area A), and knowledge of address B into the appropriate nodes, and then one by one removing knowledge of address C.

Since OSI addressing explicitly identifies the area, it is very easy for level 1 routers to identify packets going to destinations outside of their area, which need to be forwarded to level 2 routers.

In IS-IS, there are two types of routers:

- Level 1 intermediate systems — these nodes route based on the ID portion of the ISO address. They route within an area. They recognize, based on the destination address in a packet, whether the destination is within the area. If so, they route towards the destination. If not, they route to the nearest level 2 router.

- Level 2 intermediate systems — these nodes route based on the area address. They route towards areas, without regard to the internal structure of an area. A level 2 IS may also be a level 1 IS in one area.

A level 1 router will have the area portion of its address manually configured. It will refuse to become a neighbor with a node whose area addresses do not overlap its area addresses. However, if level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, then the level 1 router will accept the other node as a neighbor.

A level 2 router will accept another level 2 router as a neighbor, regardless of area address. However, if the area addresses do not overlap, the link would be considered by both routers to be "level 2 only", and only level 2 LSPs would flow on the link. External links (to other routing domains) must be from level 2 routers.

IS-IS provides an optional partition repair function. In the unlikely case that a level 1 area become partitioned, this function, if implemented, allows the partition to be repaired via use of level 2 routes.

IS-IS requires that the set of level 2 routers be connected. Should the level 2 backbone become partitioned, there is no provision for use of level 1 links to repair a level 2 partition.

In unusual cases, a single level 2 router may lose connectivity to the level 2 backbone. In this case the level 2 router will indicate in its level 1 LSPs that it is not "attached" (i.e., does not have connectivity to any destination outside of the area), thereby allowing level 1 routers in the area to route traffic for outside of the area to a different level 2 router. Level 1 routers therefore route traffic to destinations outside of their area only to level 2 routers which indicate in their level 1 LSPs that they are "attached".

An end system may autoconfigure the area portion of its address by extracting the area portion of a neighboring router's address. If this is the case, then an endnode will always accept a router as a neighbor. Since the standard does not specify that the end system MUST autoconfigure its area address, an end system may be configured with an area address. In this case the end system would ignore router neighbors with non-matching area addresses.

Special treatment is necessary for broadcast subnetworks, such as LANs. This solves two sets of issues: (i) In the absence of special treatment, each router on the subnetwork would announce a link to every other router on the subnetwork, resulting in n-squared links reported; (ii) Again, in the absence of special treatment, each router on the LAN would report the same identical list of end systems on the LAN, resulting in substantial duplication.

These problems are avoided by use of a "pseudonode", which represents the LAN. Each router on the LAN reports that it has a link to the pseudonode (rather than reporting a link to every other router on the LAN). One of the routers on the LAN is elected "designated router". The designated router then sends out an LSP on behalf of the pseudonode, reporting links to all of the routers on the LAN. This reduces the potential n-squared links to n links. In addition, only the pseudonode LSP includes the list of end systems on the LAN, thereby eliminating the potential duplication (for further information on designated routers and pseudonodes, see [1]).

The IS-IS provides for optional Quality of Service (QOS) routing, based on throughput (the default metric), delay, expense, or residual error probability. This is described in detail in [1].

IS-IS has a provision for authentication information to be carried in all IS-IS PDUs. Currently the only form of authentication which is defined is simple passwords. A password may be associated with each link, each area, and with the level 2 subdomain. A router not in possession of the appropriate password(s) is prohibited from participating in the corresponding function (i.e., may not initialize a link, be a member of the area, or a member of the level 2 subdomain, respectively). Procedures are provided to allow graceful migration of passwords without disrupting operation of the routing protocol. The authentication functions are extensible so that a stronger, cryptographically-based security scheme may be added in an upwardly compatible fashion at a future date.

2 IS-IS Packet Types

The packets used in IS-IS routing protocol fall into three main classes: (i) Hello Packets; (ii) Link State Packets (LSPs); and (iii) Sequence Number Packets (SNPs).

Hello packets are used to initialize and maintain adjacencies between neighboring routers. There are three types of IS-IS Hello packets: (i) “Level 1 LAN IS to IS Hello PDUs” are used by level 1 routers on broadcast LANs. (ii) “Level 2 LAN IS to IS Hello PDUs” are used by level 2 routers on broadcast LANs. (iii) “Point-to-Point IS to IS Hello PDUs” are used on non-broadcast media, such as point-to-point links, or general topology subnetworks.

On point-to-point links, the exchange of ISO 9542 ISHs (intermediate system Hellos) is used to initialize the link, and to allow each router to know if there is a router on the other end of the link, before IS-IS Hellos are exchanged. All routers implementing IS-IS (whether IP-only, OSI-only, or dual), if they have any interfaces on point-to-point links, must therefore be able to transmit ISO 9542 ISHs on their point-to-point links.

Link State Packets (LSPs) are used to exchange link state information. There are two types of LSPs: (i) “Level 1 Link State PDUs” are transmitted by level 1 routers. (ii) “Level 2 Link State PDUs” are transmitted by level 2 routers. Note that level 2 routers will, in most cases, also be level 1 routers, and will therefore transmit both sorts of LSPs.

Sequence number PDUs are used to ensure that neighboring routers have the same notion of what is the most recent LSP from each other router. The sequence number PDUs therefore serve a similar function to acknowledgement packets, but allow more efficient operation. There are four types of sequence number packets: (i) “Level 1 Complete Sequence Numbers PDU”; (ii) “Level 2 Complete Sequence Numbers PDU”; (iii) “Level 1 Partial Sequence Numbers PDU”; and (iv) “Level 2 Partial Sequence Numbers PDU”. A partial sequence number packet lists the most recent sequence number of one or more LSPs, and operates much like an acknowlegement. A partial sequence number packet differs from an conventional acknowledgement in the sense that it may acknowlege multiple LSPs at once, and in the sense that it may act as a request for information. A complete sequence number packet contains the most recent sequence number of all LSPs in the database. A complete sequence number packet may therefore be used to ensure synchronization of the database between adjacent routers either periodically, or when a link first comes up.

3 Symbols and Abbreviations

AA	Administrative Authority (a three octet field in the GOSIP version 2.0 NSAP address format)
AFI	Authority and Format Identifier (first octet of all OSI NSAP addresses — identifies format of the rest of the address)

CLNP	Connection-Less Network Protocol (ISO 8473, the OSI connectionless network layer protocol — very similar to IP)
DFI	DSP Format Identifier (a one octet field in the GOSIP version 2.0 NSAP address format)
ES	End System (The OSI term for a host)
ES-IS	End System to Intermediate System Routeing Exchange Protocol (ISO 9542 — OSI protocol between routers and end systems)
ICD	International Code Designator (ISO standard for identifying organizations)
IP	Internetwork Protocol (an Internet Standard Network Layer Protocol)
IS	Intermediate System (The OSI term for a router)
IS-IS	Intermediate System to Intermediate System Routeing Exchange Protocol (the ISO protocol for routing within a single routing domain)
IS-IS Hello	An Hello packet defined by the IS-IS protocol (a type of packet used by the IS-IS protocol)
ISH	An Hello packet defined by ISO 9542 (ES-IS protocol). (not the same as IS-IS Hello)
ISO	International Organization for Standardization (an international body which is authorized to write standards of many kinds)
LSP	Link State Packet (a type of packet used by the IS-IS protocol)
NLPID	Network Layer Protocol ID (A one-octet field identifying a network layer protocol)
NSAP	Network Service Access Point (a conceptual interface point at which the network service is made available)
SEL	NSAP Selector (the last octet of NSAP addresses, also called NSEL)
OSI	Open Systems Interconnection (an international standard protocol architecture)
RD	Routing Domain (the set of routers and end systems using a single instance of a routing protocol such as IS-IS)
SNPA	Subnetwork Point of Attachment (a conceptual interface at which a subnetwork service is provided)
TCP	Transmission Control Protocol (an Internet Standard Transport Layer Protocol)
TCP/IP	The protocol suite based on TCP, IP, and related protocols (the Internet standard protocol architecture)

4 References

[1] "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO 10589, forthcoming.

[2] “Protocol for Providing the Connectionless-Mode Network Service”, ISO 8473, March 1987.

[3] “End System to Intermediate System Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473)”, ISO 9542, March 1988.

[4] "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", RFC 1195, Ross Callon, December 1990.

[5] "Digital Network Architecture (Phase V) Network Routing Layer Functional Specification", Digital Equipment Corporation No. EK-DNA03-FS-001.

[6] "Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs", ISO/CD 10747, July 1991.

[7] "Network Service Definition, Addendum 2: Network Layer Addressing", ISO 8348/Addendum 2, 1988.

[8] "Guidelines for OSI NSAP Allocation in the Internet", R.Colella, E.Gardner, and R.Callon, RFC 1237, July 1991.