Network Working Group
Request for Comments: 1195

Ross W. Callon
Digital Equipment Corporation
December, 1990

Use of OSI IS-IS for Routing in TCP/IP and Dual Environments

Status of this Memo

This RFC specifies a protocol on the IAB Standards Track for the internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This RFC specifies an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/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. This specification was developed by the IS-IS working group of the Internet Engineering Task Force.

The OSI IS-IS protocol has reached a mature state, and is ready for implementation and operational use. The most recent version of the OSI IS-IS protocol is contained in ISO DP 10589 [1]. The proposed standard for using IS-IS for dual routing will therefore make use of this version (with a minor bug correction, as discussed in Annex B). We expect that future versions of this proposed standard will upgrade to the final International Standard version of IS-IS when available.

Comments should be sent to “isis@merit.edu”.

Contents

Figures

1 Introduction: Overview of the Protocol

The TCP/IP protocol suite has been growing in importance as a multi-vendor communications architecture. With the anticipated emergence of OSI, we expect coexistence of TCP/IP and OSI to continue for an extended period of time. There is a critical need for routers to support both IP traffic and OSI traffic in parallel.

There are two main methods that 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 alternate approach, which makes use of a single integrated protocol for interior routing (i.e., for calculating routes within a routing domain) for both protocol suites.

This integrated protocol design is based on the OSI Intra-domain IS-IS routing protocol [1], with IP-specific functions added. This RFC is considered a companion to the OSI IS-IS Routing spec, and will only describe the required additional features.

By supporting both IP and OSI traffic, this integrated protocol design supports traffic to IP hosts, OSI end systems, and dual end systems. This approach is “integrated” in the sense that the IS-IS protocol can be used to support pure-IP environments, pure-OSI environments, and dual environments. In addition, this approach allows interconnection of dual (IP and OSI) routing domains with other dual domains, with IP-only domains, and with OSI-only domains.

The protocol specified here is based on the work of the IETF IS-IS working group.

1.1 What the Integrated IS-IS offers

The integrated IS-IS provides a single routing protocol which will simultaneously provide an efficient routing protocol for TCP/IP, and for OSI. This design makes use of the OSI IS-IS routing protocol, augmented with IP-specific information. This design provides explicit support for IP subnetting, variable subnet masks, TOS-based routing, and external routing. There is provision for authentication information, including the use of passwords or other mechanisms. The precise form of authentication mechanisms (other than passwords) is outside of the scope of this document.

Both OSI and IP packets are forwarded “as is” — i.e., they are transmitted directly over the underlying link layer services without the need for mutual encapsulation. The integrated IS-IS is a dynamic routing protocol, based on the SPF (Dijkstra) routing algorithm.

The protocol described in this specification allows for mixing of IP-only, OSI-only, and dual (IP and OSI) routers, as defined below.

An IP-only IS-IS router (or “IP-only” router) is defined to be a router which: (i) Uses IS-IS as the routing protocol for IP, as specified in this report; and (ii) Does not otherwise support OSI protocols. For example, such routers would not be able to forward OSI CLNP packets.

An OSI-only router is defined to be a router which uses IS-IS as the routing protocol for OSI, as specified in [1]. Generally, OSI-only routers may be expected to conform to OSI standards, and may be implemented independent of this specification.

A dual IS-IS router (or "dual" router) is defined to be a router which uses IS-IS as a single integrated routing protocol for both IP and OSI, as specified in this report.

This approach does not change the way that IP packets are handled. IP-only and dual routers are required to conform to the requirements of Internet Gateways [4]. The integrated IS-IS protocol described in this report outlines an Interior Gateway Protocol (IGP) which will provide routing within a TCP/IP routing domain (i.e., autonomous system). Other aspects of router functionality (e.g., operation of ICMP, ARP, EGP, etc.) are not affected by this proposal.

Similarly, this approach does not change the way that OSI packets are handled. There will be no change at all to the contents nor to the handling of ISO 8473 Data packets and Error Reports, nor to ISO 9542 Redirects and ES Hellos. ISO 9542 IS Hellos transmitted on LANs are similarly unchanged. ISO 9542 IS Hellos transmitted on point-to-point links are unchanged except for the addition of IP-related information. Similarly, other OSI packets (specifically those involved in the IS-IS intra-domain routing protocol) remain unchanged except for the addition of IP-related information.

This approach makes use of the existing IS-IS packets, with IP-specific fields added. Specifically: (i) authentication information may be added to all IS-IS packets; (ii) the protocols supported by each router, as well as each router's IP addresses, are specified in ISO 9542 IS Hello, IS-IS Hello and Link State Packets; (iii) internally reachable IP addresses are specified in all Link State Packets; and (iv) externally reachable IP addresses, and external routing protocol information, may be specified in level 2 Link State Packets. The detailed encoding and interpretation of this information is specified in sections 3, 4, and 5 of this RFC.

The protocol described in this report may be used to provide routing in an IP-only routing domain, in which all routers are IP-only. Similarly, this protocol may be used to provide routing in a pure dual domain, in which all routers are dual. Finally, this protocol may be used to provide routing in a mixed domain, in which some routers are IP-only, some routers are OSI-only, and some routers are dual. The specific topological restrictions which apply in this latter case are described in detail in section 1.4 ("Support of Mixed Routing Domains"). The use of IS-IS for support of pure OSI domains is specified in [1].

This protocol specification does not constrain which network management protocol(s) may be used to manage IS-IS-based routers. Management information bases (MIBs) for managing IP-only, OSI-only, and dual routers, compatible with CMIP, CMOT, and/or SNMP, are the subject of a separate, companion document [8].

1.2 Overview of the ISO IS-IS Protocol

The IS-IS Routing Protocol 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 section briefly describes the manner in which IS-IS is used to support pure OSI environments. Enhancements for support of IP and dual environments are specified elsewhere in this report.

In IS-IS, the network is partitioned into “routing domains”. The boundaries of routing domains are defined by network management, by setting some links to be “exterior links”. If a link is marked as “exterior”, no IS-IS routing messages are sent on that link.

Currently, ISO does not have a standard for inter-domain routing (i.e., for routing between separate autonomous routing domains). 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.

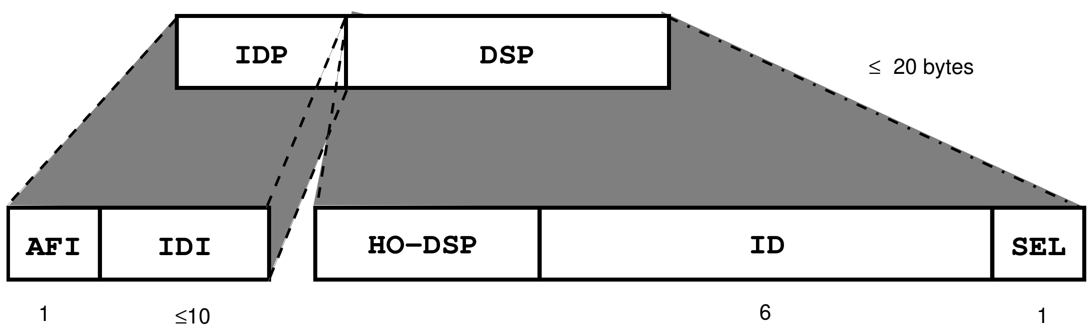

Figure 1 - ISO Hierarchical Address Structure

As illustrated in figure 1, 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 DSP is assigned by whatever addressing authority is specified by the IDP. The DSP is further subdivided into a “High Order Part of DSP” (HO-DSP), a system identifier (ID), and an NSAP selector (SEL). The HO-DSP may use any format desired by the authority which is identified by the IDP. Together, the combination of [IDP, HO-DSP] identify both the routing domain and the area within the routing domain. The combination of [IDP,HO-DSP] may therefore be referred to as the “Area Address”.

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 (i.e., on the combination of [IDP, HO-DSP]). 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”, thereby allowing level 1 routers in the area to route traffic for outside of the domain 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 greater detail in section 3.5, and in [1].

1.3 Overview of the Integrated IS-IS

The integrated IS-IS allows a single routing protocol to be used to route both IP and OSI packets. This implies that the same two-level hierarchy will be used for both IP and OSI routing. Each area will be specified to be either IP-only (only IP traffic can be routed in that particular area), OSI-only (only OSI traffic can be routed in that area), or dual (both IP and OSI traffic can be routed in the area).

This proposal does not allow for partial overlap of OSI and IP areas. For example, if one area is OSI-only, and another area is IP-only, then it is not permissible to have some routers be in both areas. Similarly, a single backbone is used for the routing domain. There is no provision for independent OSI and IP backbones.

Similarly, within an IP-only or dual area, the amount of knowledge maintained by routers about specific IP destinations will be as similar as possible as for OSI. For example, IP-capable level 1 routers will maintain the topology within the area, and will be able to route directly to IP destinations within the area. However, IP-capable level 1 routers will not maintain information about destinations outside of the area. Just as in normal OSI routing, traffic to destinations outside of the area will be forwarded to the nearest level 2 router. Since IP routes to subnets, rather than to specific end systems, IP routers will not need to keep nor distribute lists of IP host identifiers (note that routes to hosts can be announced by using a subnet mask of all ones).

The IP address structure allows networks to be partitioned into subnets, and allows subnets to be recursively subdivided into smaller subnets. However, it is undesireable to require any specific relationship between IP subnet addresses and IS-IS areas. For example, in many cases, the dual routers may be installed into existing environments, which already have assigned IP and/or OSI addresses. In addition, even if IP addresses are not already pre-assigned, the address limitations of IP constrain what addresses may be assigned. We therefore will not require any specific relationship between IP addresses and the area structure. The IP addresses can be assigned completely independently of the OSI addresses and IS-IS area structure. As will be described in section 3.2 (“Hierarchical Abbreviation of IP Reachability Information”), greater efficiency and scaling of the routing algorithm can be achieved if there is some correspondence between the IP address assignment structure and the area structure.

Within an area, level 1 routers exchange link state packets which identify the IP addresses reachable by each router. Specifically, zero or more [IP address, subnet mask, metric] combinations may be included in each Link State Packet. Each level 1 router is manually configured with the [IP address, subnet mask, metric] combinations which are reachable on each interface. A level 1 router routes as follows:

- If a specified destination address matches an [IP address, subnet mask, metric] reachable within the area, the packet is routed via level 1 routing.

- If a specified destination address does not match any [IP address, subnet mask, metric] combination listed as reachable within the area, the packet is routed towards the nearest level 2 router.

Flexible use of the limited IP address space is important in order to cope with the anticipated growth of IP environments. Thus an area (and by implication a routing domain) may simultaneously make use of a variety of different address masks for different subnets in the area (or domain). Generally, if a specified destination address matches more than one [IP address, subnet mask] pair, the more specific address is the one routed towards (the one with more "1" bits in the mask — this is known as "best match" routing).

Level 2 routers include in their level 2 LSPs a complete list of [IP address, subnet mask, metric] specifying all IP addresses reachable in their area. As described in section 3, this information may be obtained from a combination of the level 1 LSPs (obtained from level 1 routers in the same area), and/or by manual configuration. In addition, Level 2 routers may report external reachability information, corresponding to addresses which can be reached via routers in other routing domains (autonomous systems)

Default routes may be announced by use of a subnet mask containing all zeroes. Default routes should be used with great care, since they can result in “black holes”. Default routes are permitted only at level 2 as external routes (i.e., included in the "IP External Reachability Information" field, as explained in sections 3 and 5). Default routes are not permitted at level 1.

The integrated IS-IS provides optional Type of Service (TOS) routing, through use of the QOS feature from IS-IS.

1.4 Support of Mixed Routing Domains

The integrated IS-IS proposal specifically allows for three types of routing domains:

- Pure IP

- Pure OSI

- Dual

In a pure IP routing domain, all routers must be IP-capable. IP-only routers may be freely mixed with dual routers. Some fields specifically related to OSI operation may be included by dual routers, and will be ignored by IP-only routers. Only IP traffic will be routed in an pure IP domain. Any OSI traffic may be discarded (except for the IS-IS packets necessary for operation of the routing protocol).

In a pure OSI routing domain, all routers must be OSI-capable. OSI-only routers may be freely mixed with dual routers. Some fields specifically related to IP operation may be included by dual routers, and will be ignored by OSI-only routers. Only OSI traffic will be routed in a pure OSI domain. Any IP traffic may be discarded.

In a dual routing domain, IP-only, OSI-only, and dual routers may be mixed on a per-area basis. Specifically, each area may itself be defined to be pure IP, pure OSI, or dual.

In a pure IP area within a dual domain, IP-only and dual routers may be freely mixed. Only IP traffic can be routed by level 1 routing within a pure-IP area.

In a pure-OSI area within a dual domain, OSI-only and dual routers may be freely mixed. Only OSI traffic can be routed by level 1 routing within a pure OSI area.

In a dual area within a dual routing domain only dual routers may be used. Both IP and OSI traffic can be routed within a dual area.

Within a dual domain, if both IP and OSI traffic are to be routed between areas then all level 2 routers must be dual.

1.5 Advantages of Using Integrated IS-IS

Use of the integrated IS-IS protocol, as a single protocol for routing both IP and OSI packets in a dual environment, has significant advantages over using separate protocols for independently routing IP and OSI traffic.

An alternative approach is known as “Ships In the Night” (S.I.N.). With the S.I.N. approach, completely separate routing protocols are used for IP and for OSI. For example, OSPF [5] may be used for routing IP traffic, and IS-IS [1] may be used for routing OSI traffic. With S.I.N., the two routing protocols operate more or less independently. However, dual routers will need to implement both routing protocols, and therefore there will be some degree of competition for resources.

Note that S.I.N. and the integrated IS-IS approach are not really completely separate options. In particular, if the integrated IS-IS is used within a routing domain for routing of IP and OSI traffic, it is still possible to use other independent routing protocols for routing other protocol suites.

In the future, optional extensions to IS-IS may be defined for routing other common protocol suites. However, such future options are outside of the scope of this document. This section will compare integrated IS-IS and S.I.N. for routing of IP and OSI only.

A primary advantage of the integrated IS-IS relates to the network management effort required. Since the integrated IS-IS provides a single routing protocol, within a single coordinated routing domain using a single backbone,, this implies that there is less information to configure. This combined with a single coordinated MIB simplifies network management.

Note that the operation of two routing protocols with the S.I.N. approach are not really independent, since they must share common resources. However, with the integrated IS-IS, the interactions are explicit, whereas with S.I.N., the interactions are implicit. Since the interactions are explicit, again it may be easier to manage and debug dual routers.

Another advantage of the integrated IS-IS is that, since it requires only one routing protocol, it uses fewer resources. In particular, less implementation resources are needed (since only one protocol needs to be implemented), less CPU and memory resources are used in the router (since only one protocol needs to be run), and less network resources are used (since only one set of routing packets need to be transmitted). Primarily this translates into a financial savings, since each of these three types of resources cost money. This implies that dual routers based on the integrated IS-ginal Finger protocol, this murchase and operate than dual routers based on S.I.N.

Note that the operation of two routing protocols with the S.I.N. approach are not really independent, since they must share common resources. For example, if one routing protocol becomes unstable and starts to use excessive resources, the other protocol is likely to suffer. A bug in one protocol could crash the other. However, with the integrated IS-IS, the interactions are explicit and are defined into the protocol and software interactions. With S.I.N., the interactions are implicit.

The use of a single integrated routing protocol similarly reduces the likely frequency of software upgrades. Specifically, if you have two different routing protocols in your router, then you have to upgrade the software any time EITHER of the protocols change. If you make use of a single integrated routing protocol, then software changes are still likely to be needed, but less frequently.