Internet Engineering Task Force A. Bansal, T. Przygienda, A. Patel
INTERNET DRAFT Fore, Bell Labs/Lucent
20 Jan 1998

IS-IS over IPv4
<draft-ietf-isis-wg-over-ip-00.txt>

Status of This Memo

Abstract

This draft describes an optional implementation technique within IS-IS [ISO90, Cal90a, Cal90b] used today by several ISPs for routing within their clouds. IS-IS is an interior gateway routing protocol developed originally by OSI and used with IP extensions as IGP. This draft describes how to encapsulate IS-IS packets in IPv4 [Pos81] format. Such an encapsulation has many advantages, one of those being the possibility to run integrated IS-IS on anything that understands IPv4, including avian carriers [Wai90].

1. Introduction

Encapsulation of IS-IS as defined in ISO 10589 [ISO90, Cal90a, Cal90b] uses directly over Link Layer protocols as opposed to IP routing protocols [Moy97] that are encapsulated over IP directly.

By defining an encapsulation of IS-IS in IP, we save on special OSI

encapsulation on several media types. Such an encapsulation would solve fragmentation problems of large LSPs and remove the necessity for OSI PPP extensions [Kat92] when IS-IS is run over negotiated PPP links. Additionally, on certain media, such as P2P ATM links, no LLC/SNAP encapsulation is necessary to provide multi-protocol routing, allowing for gains in efficiency.

2. Encapsulation of IS-IS and ISO 9542 packets over IP

IS-IS is encapsulated directly over the Internet Protocol's network layer. IS-IS packets are therefore encapsulated by IP and local data-link headers. Within this encapsulation, IS-IS packets are propagated as usual, however without the appropriate link-layer fields but starting at NLPI.

IS-IS does not normally provide a way to transmit packets larger than MTU size. This proposal allows to use IP fragmentation when transmitting such packets. If necessary, the length of IS-IS packets over IP can be up to 65,535 bytes (including the IP header). The IS-IS packet types that are likely to be large (LSPs, CSNPs, PSNPs) can usually be split into several separate protocol packets, without loss of functionality. This is recommended; IP fragmentation SHOULD be avoided whenever possible since it can lead to different problems, such as loss of fragments causing the retransmission of complete IP packets. Following rules apply:

- IIHs MUST have the size of [InterfaceMTU - IP headersize] (1) and have the DF bit set. IIH's TTL MUST be set to 1 to prevent them from traveling multiple hops.

- SNPs on IP encapsulated interfaces MUST NOT be larger than the minimum of [InterfaceMTU - IP headersize] and respective originatingLSPBufferSize.

- LSP fragments MAY BE built with a size up to the value of corresponding originatingLSPBufferSize.

- LSP fragments MUST NOT be larger than the respective originatingLSPBufferSize.

---

1. not of maximum of Buffer and DataLinkSize used normally

- LSPs MUST be sent allowing IP fragmentation (DF bit not set).

- LSP fragments SHOULD NOT exceed the size of the minimum of dataLinkBlockSize and respective originatingLSPBufferSize.

This set of rules allows to configure a network with respective originatingLSPBufferSize larger than some interfaces' MTUs.
In a mixed environment, care must be taken that respective originatingLSPBufferSize does net exceed the MTU size of interfaces without IP encapsulation.
The other important features of IS-IS in IP's IP encapsulation are:

- Use of IP multicast. Some IS-IS in IP messages are multicast, when sent over broadcast networks. Three distinct IP multicast addresses are used. Packets sent to these multicast addresses should never be forwarded; they are meant to travel a single hop only. To ensure that these packets will not travel multiple hops, their IP TTL must be set to 1.

  IPAllL1ISs

  OSI multicast value of this address was 01-80-C2-00-00-14. This multicast address has been assigned the IP address value 224.0.0.? for IP encapsulated IS-IS. All routers running L1 IS-IS in IP should be prepared to receive packets sent to this address. Hello packets are always sent to this destination. Also, certain IS-IS in IP protocol packets are sent to this address during the flooding procedure.

  IPAllL2ISs

  OSI multicast value of this address was 01-80-C2-00-00-15. This multicast address has been assigned the IP address value 224.0.0.? for IP encapsulated IS-IS. All routers running L2 IS-IS in IP should be prepared to receive packets sent to this address. Hello packets are always sent to this destination. Also, certain IS-IS in IP protocol packets are sent to this address during the flooding procedure.

IPAllIntermediateSystems

OSI multicast value of this address was 09-00-2B-0O-00-05. This multicast address has been assigned the IP address value 224.0.0.? for IP encapsulated IS-IS. All routers running IS-IS in IP should be prepared to receive packets sent to this address. ISO 9542 is using this address.

- IS-IS in IP is IP protocol number ??. This number has been requested with the Network Information Center. IP protocol number assignments are documented in [RP94].

  Note: For development purposes, IP Protocol number 9 (Private IGP protocol number) [RP94] will be used until an official number is granted.

- All IS-IS in IP routing protocol packets are sent using the normal service TOS value of binary 0000 defined in [Alm92].

- Routing protocol packets are sent with IP precedence set to Internetwork Control. IS-IS in IP protocol packets should be given precedence over regular IP data traffic, in both sending and receiving. Setting the IP precedence field in the IP header to Internetwork Control [Pos81] may help implement this objective.

3. Internal Encapsulation

On point-to-point links no MAC addresses are used by IS-IS. Therefore the Intradomain Routeing Protocol Discriminator or ISO 9542 Network Layer Protocol Identifier starts directly after the IP header. For P2P link running PPP, the Payload format will consist of PPP header, followed by the IP header and NLPID as indicated in figure 1.

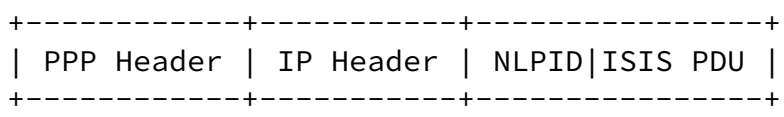

Figure 1: Encapsulation of ISIS frames over PPP

For P2P ATM links using VC muxing, the payload format must not include the PPP header as indicated in figure 2.

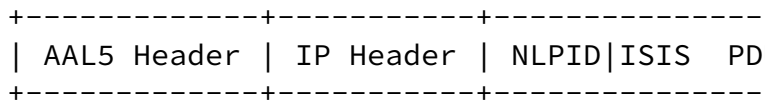

Figure 2: Encapsulation of ISIS frames over ATM

In the case of broadcast media, MAC addresses are used for adjacency identification. It is rather painful from the implementation perspective to assume that an IS-IS must have access to MAC headers when receiving frames. On the other hand, introducing an artificial `bridging' encapsulation header preceding the Intradomain Routeing Protocol Discriminator within the routed frame was perceived as a very bulky solution. As yet another approach, based on the observation that ethernets always have at least one IP address, MAC address necessary for adjacency maintenance could be built algorithmically using the source address of the IP packet received. However, this would necessitate e.g. the allocation of a 2-byte prefix within the 802.2/3 MAC space. Given the difficulties surrounding each of these solutions, no special encapsulation is defined for the ethernet case and the protocol implementation is either required to received the complete frame, including layer 2 encapsulation or obtain the appropriate MAC address from the source IP address using an interface to the lower layers of the IP stack.

4. Interoperability with Devices without IP Encapsulation

An interoperability solution for devices using IP encapsulation and OSI encapsulation of ISIS frames would be only useful if it could significantly ease the migration path in the existing networks. Given the fact that graceful migration is not a paramount issue for existing networks and any solution would invariable lead to the problem of partitioning of broadcast media, such a solution is not defined.

5. Acknowledgments

The encapsulation description part has been "borrowed" from a well-known RFC [Moy97] with the author's consent. Tony Li, Mike Shand, Henk Smit, Rajesh Varadarajan, Jeff Swinton, Stacy Smith gave many constructive comments.

6. Security Consideration

ISIS security applies to the work presented. No specific security issues with the proposed solutions are known. Things like IPSec may influence the work in strange and unknown ways ;-)

References

[Alm92] P. Almquist. Type of Service in the Internet Protocol Suite. INTERNET-RFC, Internet Engineering Task Force, July 1992.

[Cal90a] R. Callon. OSI ISIS Intradomain Routing Protocol. INTERNET-RFC, Internet Engineering Task Force, February 1990.

[Cal90b] R. Callon. Use of OSI ISIS for Routing in TCP/IP and Dual Environments. INTERNET-RFC, Internet Engineering Task Force, December 1990.

[Hei93] J. Heinanen. Rfc 1483, Multiprotocol Encapsulation over ATM Adaptation Layer 5. Internet Engineering Task Force, July 1993.

[ISO90] ISO. Information Technology - Telecommunications and Information Exchange between Systems - Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service. ISO, 1990.

[Kat92] D. Katz. Rfc 1377, The PPP OSI Network Layer Control Protocol (OSINLCP). Internet Engineering Task Force, November 1992.

[MD90] J. Mogul and S. Deering. Rfc 1191, Path MTU Discovery. Internet Engineering Task Force, November 1990.

[Moy97] J. Moy. OSPFv2, RFC 2178. Internet Engineering Task Force, July 1997.

[Pos81] J. Postel. Rfc 791, rfc Internet Protocol. Internet Engineering Task Force, September 1981.

[RP94] J. Reynolds and J. Postel. Rfc 1700, Assigned Numbers. Internet Engineering Task Force, October 1994.

[Wai90] D. Waitzman. Rfc 1149, standard for the Transmission of IP Datagrams on Avian Carriers. Internet Engineering Task Force, April 1990.

Authors' Addresses

Tony Przygienda
Bell Labs, Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733-3030
prz@dnrc.bell-labs.com

Ajay Patel
Bell Labs, Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733-3030
ajayp@dnrc.bell-labs.com

Atul Bansal
FORE Systems,
1595 Spring Hill Rd
Vienna, VA 22181
bansal@fore.com