Internet Engineering Task Force A. Bansal, T. Przygienda, A. Patel INTERNET DRAFT Siara

Fore,

1999

0ct

IS-IS over IPv4 <draft-ietf-isis-wg-over-ip-02.txt>

Status of This Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Abstract

This draft describes an optional implementation technique within IS-IS [IS090 , 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]. Encapsulation of IS-IS PDUs in IPv6 is outside the scope of this document.

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.

| Przygienda et al. | Expires Mar 2000 | [Page 1] |
|-------------------|------------------|----------|
|-------------------|------------------|----------|

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 not exceed the size of [InterfaceMT U - IP headersize] (1) and have the DF bit set. MTU size padding rules should be followed as described in ISO 10589.

- SNPs MUST NOT be larger than the respective originatingLSPBufferSize.

- SNPs MUST be sent allowing IP fragmentation (DF bit not set).
- SNPs SHOULD NOT be larger than the minimum of [InterfaceMT U IP headersize] and respective originatingLSPBufferSize.
- LSP fragments MAY BE built with a size up to the value of corresponding originatingLSPBufferSize.

1. not of maximum of Buffer and DataLinkSize used normally

Przygienda et al.Expires Mar 2000[Page2]Internet DraftIS-IS over IPv4Oct1999199119911991

- LSP fragments MUST NOT be larger than the respective originatingLSPBufferSize.
- 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. A default originatingLSPBufferSize of 1472 is recommended.

The other important features of IS-IS in IP's IP encapsulation are:

- IS-IS in IP has been assigned IP protocol number 124. IP protocol number assignments are documented in [RP94].

Historical note: For development purposes, IP Protocol number 9 (Private IGP protocol number) [RP94] has been used until an official number was granted.

- 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 224.0.0.19 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.

| Przygienda et al. 3] | Expires Mar 2000 | [Page | |
|-------------------------|------------------|-------|--|
| Internet Draft 1999 | IS-IS over IPv4 | 0ct | |

IPAllL2ISs

OSI multicast value of this address was 01-80-C2-00-00-15.

This multicast address has been assigned the IP address value 224.0.0.20 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-00-00-05. This multicast address has been assigned the IP address value 224.0.0.21 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.

- 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 to 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 the usual NLPID as indicated in figure 1.

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

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

| Przygienda et al. 4] | Expires Mar 2000 | [Page |
|-------------------------|----------------------|-------|
| Internet Draft 1999 | IS-IS over IPv4 | Oct |
| ++ | + | |
| PPP Header | IP Header ISIS PDU | |
| ++ | + | |

Figure 1: Encapsulation of ISIS frames over PPP

+----+ | AAL5 Header | IP Header | ISIS PDU | +----+

Figure 2: Encapsulation of ISIS frames over ATM

when receiving frames. However, based on the observation that ethernets have at least one IP address assigned, MAC address necessary for adjacency maintenance can be built algorithmically using the source address of the IP packet received. To prevent collisions with universally administered addresses, the addresses are designated by having bit one in byte zero of the MAC set to 1 to indicate locally administered address as specified by the according IEEE standard. When receiving an IP encapsulated ISIS PDU, the artificial MAC address is assumed to consist (in network order) of 4 bytes of source IP address aligned as least significant bytes, ``locally administered'' bit being set and all remaining bits being zero. Figure 3 visualizes such an address for a source IP address 128.127.128.1.

Figure 3: Artificial MAC for IP address 128.127.128.1

Expires Mar 2000

Przygienda et al. 5]

IS-IS over IPv4

[Page

0ct

Internet Draft 1999

> When operating on an interface encapsulating IS-IS within IP, encapsulated PDUs MUST be sent with a constant IP source address. Any change of the IP address on the interface MUST be considered equivalent to change of MAC address in ISO mode. In all places in which MAC addresses are being used in ISO mode, source IP address MUST be used to compute artificial MACs when sending and parsing received PDUs.

Any PDU received on IP encapsulated broadcast interface and containing MACs with either the ``locally administered'' bit not being set or any of the remaining bits in the first two bytes being set SHOULD be discarded.

<u>4</u>. 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 invariably 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, Dave Katz, Christian Hopps, Mike Shand, Henk Smit, Rajesh Varadarajan, Jeff Swinton, Stacy Smith and others provided constructive comments.

<u>6</u>. 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 o Suite. INTERNET-RFC 1992. | f Service in the Internet Protocol , Internet Engineering Task Force, | July | |
|------------------------|---|--|----------------|--|
| Przygienda et 6] | al. | Expires Mar 2000 | [Page | |
| Internet Draft 1999 | : | IS-IS over IPv4 | 0ct | |
| [Cal90a] | R. Callon. OSI ISIS INTERNET-RFC, Intern 1990. | Intradomain Routing Protocol. et Engineering Task Force, February | ý | |
| [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. <u>Rfc 14</u> Adaptation Layer 5. 1993. | 83, Multiprotocol Encapsulation ove Internet Engineering Task Force, . | er ATM July | |
| [IS090] | ISO. Information Te Information Exchange to Intermediate Syst Use in Conjunction w Connectionless-Mode | chnology - Telecommunications and between Systems - Intermediate Sys em Routing Exchange Protocol for ith the Protocol for Providing the Network Service. ISO, 1990. | stem | |
| [Kat92] | D. Katz. <u>Rfc 1377</u> , Protocol (OSINLCP). November 1992. | The PPP OSI Network Layer Control Internet Engineering Task Force, | | |
| [MD90] | J. Mogul and S. Deer Internet Engineering | ing. <u>Rfc 1191</u> , Path MTU Discovery Task Force, November 1990. | | |

| [Moy97] | J. Moy. OSPFv2, <u>RF0</u> July 1997. | <u>C 2178</u> . | Internet | Engineer | ing Task | Force, |
|---|--|------------------------|------------------------|--------------|----------|--------|
| [Pos81] | J. Postel. <u>Rfc 791</u> , Engineering Task For | , rfc Int rce, Sept | ernet Pro ember 198 | tocol. 1. | Internet | |
| [RP94] | J. Reynolds and J. Postel. <u>Rfc 1700</u> , Assigned Numbers. Internet Engineering Task Force, October 1994. | | | | | |
| [Wai90] | D. Waitzman. <u>Rfc 1149</u> , standard for the Transmission of IP Datagrams on Avian Carriers. Internet Engineering Task Force, April 1990. | | | | | |
| Authors' Addre | esses | | | | | |
| Tony Przygieno Siara Systems | da | | | | | |
| Przygienda et 7] | al. | Expires | Mar 2000 | | | [Page |
| Internet Draft 1999 | t | IS-IS ov | ver IPv4 | | | Oct |
| <mark>300</mark> Ferguson I Mountain View, prz@siara.com | Drive , CA 94043 | | | | | |
| Ajay Patel Siara Systems <mark>300</mark> Ferguson I Mountain View, ajayp@siara.co | Drive , CA 94043 Dm | | | | | |
| Atul Bansal FORE Systems, <u>1595</u> Spring H : Vienna, VA 22: bansal@fore.co | ill Rd 181 Dm | | | | | |
| Przygienda et 8] | al. | Expires | Mar 2000 | | | [Page |