Network Working Group
Internet Draft
Expiration Date: December 1999

Henk Smit
cisco Systems

IS-IS and IP over the same ATM VC with aal5mux encapsulation

draft-hsmit-isis-aal5mux-00.txt

1. Status of this Memo

2. Abstract

This document describes a way to send and receive both IP and IS-IS packets over one and the same ATM Virtual Circuit without using an extra layer of encapsulation.

3. Introduction

IS-IS is a routing protocol used for IP and CLNS routing. A number of large ISP networks use IS-IS as their IP IGP. The IS-IS protocol is specified in [1], with extensions for supporting IPv4 specified in [2]. IS-IS packets are encapsulated directly in the datalink layer, as part of the OSI protocol family. IS-IS packets are not encapsulated in an IP header, nor in a CLNS header.

ATM is a technology used to transmit data, voice and video. When data packets are sent over ATM, they are encapsulated in an AAL5 frame header, and the AAL5 frame is divided in a number of fixed size

cells. Cell headers are 5 bytes. The AAL5 headers are 8 bytes. Cells have a payload of 48 bytes.

3.1. ATM encapsulation methods

When data packets are sent over an ATM VC, the receiver must have a way to recognize each packet. Currently there are 3 ways to do this:

- Reserve a separate VC for each layer-3 protocol.
  This is called 'Null Encapsulation', 'VC multiplexing',
  or 'aal5mux'. This encapsulation is described in [3].

- Add a 8-byte LLC/SNAP header in front of the layer-3 header.
  This is called 'LLC/SNAP Encapsulation' or aal5snap.
  This encapsulation is also described in [3].

- Add two bytes with the NLPID in front of the data packet.
  This is called aal5nlpid encapsulation.

It is common practice to send both routing protocol packets and user traffic over the same physical or logical path. Doing so prevents certain failure modes. Thus IS-IS packets and IP packets must be sent over the same ATM VC. Because IS-IS packets are not encapsulated in IP packets, users of IS-IS for IP routing must use aal5snap or aal5nlpid encapsulation.

3.2. Motivation

The current estimate is that 40% of the packet on the Internet are small packets, of 40 bytes or less. Most of these are TCP acknowledgements. They consist of a 20 byte IP header, and a 20 byte TPC header, which makes them exactly 40 bytes long.

The ATM cell payload is 48 bytes. The 48 bytes of AAL5, IP and TCP headers will exactly fill one ATM cell. If one uses aal5nlpid or aal5snap, we need two cells to store all 50 or 56 bytes of AAL5, SNAP/NLPID, IP and TCP headers. Because cells are fixed in length, the biggest part of the second cell will be unused.

Bandwidth used by ATM, which is not used to transport user traffic, is called cell-tax. A quick calculation will show that using aal5snap increases the cell-tax by 5% to 7% of the total bandwidth.

4. Sending IS-IS and IP over the same aal5mux Virtual Circuit

Unfortunately aal5mux supports only one layer-3 protocol per VC. The reason is that the receiver must be able to de-multiplex packets to different layer-3 stacks. Without extra information in the AAL5 frame, indicating the layer-3 protocol, layer-2 at the receiver can

not distinguish packets from different layer-3 protocols.

But if there is no extra information in the AAL5 frame, can the receiver not make the distinction between IP and IS-IS by just looking at the bare layer-3 packet ? Fortunately the answer is yes. By looking at the first byte of each packet, the receiver can determine if a packet is IP or IS-IS.

The first two fields in the IP header are the 4-bit version number and the 4-bit header length. The value of the first byte is normally 0x45. If there are IP header options attached to the IP header, the first byte can be between 0x46 and 0x4F.

The first byte in an IS-IS packet is always 0x83.

Thus by looking at the first byte of an incoming packet, the receiver can separate IP and IS-IS packets. Because of this feature we do not depend on the ATM layer anymore to help us with the demultiplexing. Routers can now send and receive both IS-IS and IP packets in a non-ambiguous way over the same aal5mux VC.

5. Remaining issues

If IS-IS is used in Dual-mode, to build both IP and CLNS routing tables, it is necessary that CLNS packets are forwarded over the same path as IP packets. Implementations supporting IP and CLNS routing might want to consider extending the check for the first byte in the packet. CLNS packets always start with 0x81. ES-IS packets start with 0x82.

This proposal depends on the fact that all IP packets in the global Internet are version 4. The method of looking at the first byte can be extended to distinquish between IS-IS and any version of the IP protocol, except version 8.

6. Security Considerations

The extension described in this document might have the surprising effect to some network administrators that IS-IS or IP packets can now be sent over VCs which did not carry those type of packets in the past.

7. Acknowledgements

The author would like to thank Jonathan Gardner for the discussion

that initiated this proposal.

8. References

[1] ISO 10589, "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)" [Also republished as RFC 1142]

[2] RFC 1195, "Use of OSI IS-IS for routing in TCP/IP and dual environments", R.W. Callon, Dec. 1990

[3] RFC 1483, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", Juha Heinanen, July 1993

9. Author's Address

Henk Smit
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
e-mail: hsmit@cisco.com