

IPv4 over MAPOS Version 1

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Authors' Note

This memo documents a mechanism for supporting Version 4 of the Internet Protocol (IPv4) on Version 1 of the Multiple Access Protocol over SONET/SDH. This document is NOT the product of an IETF working group nor is it a standards track document. It has not necessarily benefited from the widespread and in-depth community review that standards track documents receive.

Abstract

This document describes a protocol for transmission of the Internet Protocol Version 4 (IPv4) over Multiple Access Over SONET/SDH (MAPOS) version 1. MAPOS is a link layer protocol and provides multiple access capability over SONET/SDH links. IP runs on top of MAPOS. This document explains IP datagram encapsulation in HDLC frame of MAPOS, and the Address Resolution Protocol (ARP).

1. Introduction

Multiple Access Protocol over SONET/SDH (MAPOS) [1] is a high-speed link-layer protocol that provides multiple access capability over SONET/SDH. Its frame format is based on the HDLC-like framing [2] for PPP. A component called "Frame Switch" [1] allows multiple nodes to be connected together in a star topology to form a LAN. Using long haul SONET/SDH links, the nodes on such a "SONET-LAN" can span over a wide geographical area. The Internet Protocol (IP) [3] datagrams are transmitted in MAPOS HDLC frames [1].

This document describes a protocol for transmission of IP datagrams over MAPOS version 1 [1]. It explains IP datagram encapsulation in HDLC frame of MAPOS, and ARP (Address Resolution Protocol) for mapping between IP address and HDLC address.

2. Frame Format for Encapsulating IP Datagrams

An IP datagram is transmitted in a MAPOS HDLC frame. The protocol field of the frame must contain the value 0x0021 (hexadecimal) as defined by the "MAPOS Version 1 Assigned Numbers" [4]. The information field contains the IP datagram.

The information field may be one to 65,280 octets in length; the MTU(Maximum Transmission Unit) of MAPOS is 65,280 octets. Although the large MTU size can suppress the overhead of IP header processing, it may cause fragmentation anywhere along the path from the source to the destination and result in performance degradation. To cope with the issue, Path MTU discovery [5] may be used.

3. Address Mapping

This section explains MAPOS ARP and the mapping of special addresses.

3.1 ARP cache

Each node on a MAPOS network maintains an "ARP cache" that maps destination IP addresses to their corresponding 8-bit HDLC addresses. Entries are added to this cache either manually or by the Address Resolution Protocol described below. Entries are removed from this cache manually, by the UNARP mechanism, or by ARP cache validation mechanism. An implementation must provide a mechanism for manually adding or removing arbitrary ARP cache entries.

3.2 Address Resolution Protocol (ARP)

This subsection describes MAPOS ARP protocol and its packet format.

3.2.1 Overview

The MAPOS ARP is similar to that for ethernet. Prior to sending an IP datagram, the node must know the destination HDLC address corresponding to the destination IP address. When its ARP cache does not contain the corresponding entry, it uses ARP to translate the IP address to the HDLC address. That is, it broadcasts an ARP request containing the destination IP address. In response to the request, the node which has the IP address sends an ARP reply containing the HDLC address. The returned HDLC address is stored in the ARP cache.

3.2.2 ARP Frame Format

The protocol field for an ARP frame must contain 0xFE01 (hexadecimal) as defined by the "MAPOS Version 1 Assigned Numbers" [4]. The information field contains the ARP packet as shown below.

| | | |
|---|------------------------------------|---|
| Hardware Address Space (25:MAPOS) 16 bits | | Protocol Address Space (2048 in Dec) 16 bits |
| Hard Addr Length (4) 8 bits | Proto Addr Length (4) 8 bits | Operation Code (1:Request 2:Response) 16 bits |
| Sender HDLC Address | | 32 bits |
| Sender IP Address | | 32 bits |
| Target HDLC Address | | 32 bits |
| Target IP Address | | 32 bits |

Figure 5 ARP packet format

Hardware Address Space (16 bits)

The hardware address space for MAPOS ARP is 25 in Decimal as assigned by IANA [6].

Protocol Address Space (16 bits)

The protocol address space for IP is 2048 in Decimal.

Hardware Address Length (8 bits)

The hardware address length is 4.

Protocol Address Length (8 bits)

The protocol address length for IP is 4.

Operation Code (16 bits)

The operation code is 1 for request and 2 for response.

Sender hardware (HDLC) Address (32 bits)

Contains the sender's HDLC address in an ARP request, and the target HDLC address in an ARP response. The 8-bit HDLC address is placed in the least significant place of the 32-bit field. The remaining bits should be zero.

Sender Protocol (IP) Address (32 bits)

Contains the sender's IP address in an ARP request, and the target IP address in an ARP response.

Target hardware (HDLC) Address (32 bits)

Contains unknown target HDLC address (all zeros) in an ARP request, and sender's HDLC address in an ARP response. The 8-bit HDLC address is placed in the least significant place of the 32-bit field. The remaining bits should be zero.

Target Protocol (IP) Address (32 bits)

Contains the target IP address in an ARP request, and the sender's IP address in an ARP response.

3.3 UNARP

An implementation MUST provide an UNARP mechanism to flush obsolete ARP cache entries. The mechanism is similar to the ARP extension described in [7]. When a node detects that its port has come up, it MUST broadcast an UNARP packet. It forces every other node to clear the obsolete ARP entry which was created by the node previously connected to the switch port. An UNARP is an ARP clear request with the following values:

| | | |
|--------------------------------|---|-----------------------------|
| Hardware Address Space | : | 25 |
| Protocol Address Space | : | 2048 |
| Hardware Address Length | : | 4 |
| Protocol Address Length | : | 4 |
| Operation Code | : | 23 (MAPOS-UNARP) |
| Sender hardware (HDLC) Address | : | HDLC address of the node |
| Sender Protocol (IP) Address | : | IP address of the node |
| Target hardware (HDLC) Address | : | all 1 |
| Target Protocol (IP) Address | : | 255.255.255.255 (broadcast) |

Hardware Address Space (16 bits)

The hardware address space for MAPOS ARP is 25 in Decimal as assigned by IANA [6].

Operation Code (16 bits)

The operation code is 23 for MAPOS-UNARP in Decimal as assigned by IANA [6].

The node MUST send three UNARP packets at 30 seconds intervals. The receiving node of the packet MUST clear the ARP cache entry associated with the Sender Protocol (IP) Address, if and only if the corresponding Hardware (HDLC) Address is not equal to that contained in the UNARP packet. That is, if both the Sender Hardware (HDLC) Address and the Sender Protocol(IP) Address match those of the cached entry, the entry is left unchanged.

3.4 ARP Cache Validation

An implementation MUST provide a mechanism to remove out-of-date cache entries and it SHOULD provide options to configure the timeout value [8]. One approach is to periodically time-out the cache entries, even if they are in use. This approach involves ARP cache timeouts in the order of a minute or less.

Furthermore, when the link is lost on an interface, all ARP cache entries associated with the interface MUST be removed immediately. Causes for link loss includes conditions such as loss of carrier and out-of-synchronization.

3.5 IP Broadcast and multicast

In broadcast and multicast frames, the most significant bit of the HDLC address must be 1 [1]. In addition, the least significant bit must always be 1 to indicate the end of the field [1].

In the case of IP broadcast, the remaining six bits of the HDLC address must be all 1s. That is, it should be mapped to the HDLC broadcast address 0xFF (hexadecimal).

In the case of IP multicast, the remaining six bits of the HDLC address must contain the lowest-order six bits of the IP multicast group address. It resembles IP multicast extension for ethernet described in [9]. Exceptions arise when these six bits are either all zeros or all ones, in which case they should be altered to the bit sequence 111110.

4. Security Considerations

Security issues are not discussed in this memo.

References

- [1] Murakami, K. and M. Maruyama, "MAPOS - Multiple Access Protocol over SONET/SDH, Version 1," RFC-2171, June 1997.
- [2] Simpson, W., editor, "PPP in HDLC-like Framing," RFC-1662, July 1994.
- [3] Postel, J., "Internet Protocol (IP)," RFC-791, September 1981.
- [4] Maruyama, M. and K. Murakami, "MAPOS Version 1 Assigned Numbers," RFC-2172, June 1997.
- [5] Mogul, J. and S. Deering, "Path MTU Discovery," RFC-1191, Nov. 1990.
- [6] IANA, "IANA-Assignments",
<http://www.iana.org/iana/assignments.html>
- [7] Malkin, G., "ARP Extension - UNARP," RFC-1868, November 1995.
- [8] Braden, R., "Requirements for Internet Hosts - Communication Layers," RFC-1122, October 1989.
- [9] Deering, S., "Host Extensions for IP Multicasting," RFC-1112, August 1989.

Acknowledgements

The authors would like to acknowledge the contributions and thoughtful suggestions of John P. Mullaney, Clark Bremer, Masayuki Kobayashi, Paul Francis, Toshiaki Yoshida, and Takahiro Sajima.

Author's Address

Ken Murakami
NTT Software Laboratories
3-9-11, Midori-cho
Musashino-shi
Tokyo-180, Japan
E-mail: murakami@ntt-20.ecl.net

Mitsuru Maruyama
NTT Software Laboratories
3-9-11, Midori-cho
Musashino-shi
Tokyo-180, Japan
E-mail: mitsuru@ntt-20.ecl.net

