

Network Working Group
Request for Comments: 3422
Category: Informational

O. Okamoto
M. Maruyama
NTT Laboratories
T. Sajima
Sun Microsystems
November 2002

Forwarding Media Access Control (MAC) Frames over Multiple
Access Protocol over Synchronous Optical Network/Synchronous Digital
Hierarchy (MAPOS)

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

IESG Note

This memo documents a way of tunneling Ethernet frames over MAPOS networks. 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 memo describes a method for forwarding media access control (MAC) frames over Multiple Access Protocol over Synchronous Optical Network/Synchronous Digital Hierarchy (MAPOS), thus providing a way to unify MAPOS network environment and MAC-based Local Area Network (LAN) environment.

1. Network Model

In the Network model assumed in this memo, MAC-based LAN traffic is forwarded by a MAPOS switched network. This model allows distant LANs to be interconnected to form a single LAN segment. Transparent LAN Service (TLS) is provided by encapsulating MAC frames in MAPOS frames and by mapping MAC addresses to MAPOS addresses.

This network model is shown in figure 1. "MAPOS network" is composed of MAPOS switches, SONET/SDH leased lines and optical fiber cables. A LAN is connected to a MAPOS network by a Network Adapter (NA) which has a MAPOS interface and an ethernet interface. A unique MAPOS address is assigned to each NA by NSP (Node-Switch Protocol) [2].




Figure 1. VPN network service model with LANs N1 and N2

Host H1 in LAN N1 and host H2 in LAN N2 are connected to distinct MAC-based LANs. Transparent LAN service is provided by MAPOS network N0 exchanging MAC frames between Host H1 and Host H2.

Using this mechanism, a single VLAN segment can be setup from multiple LANs that may be geographically located far away from each other.

The use of a switched technology is recommended for building a MAC-based LAN. In some cases, however, this becomes a requirement. A likely example is the situation where a MAC-based LAN having two network adapters, both attached to the same MAPOS network (for redundancy). If the LAN is built using shared (non-switched) technology, then this loop configuration is bound to be stormed by incessant broadcast traffic. This can only be circumvented by using switched technology with support for broadcast spanning tree [7].

2. Forwarding a MAC Frame

This section describes the MAC frame forwarding mechanism in the MAPOS network.

2.1. Outline

In figure 2, LANs N1 and N2 communicates via MAPOS network N0. NAs B1 and B2 are gateways into Network N0, and they each have a MAPOS interface and an ethernet interface.

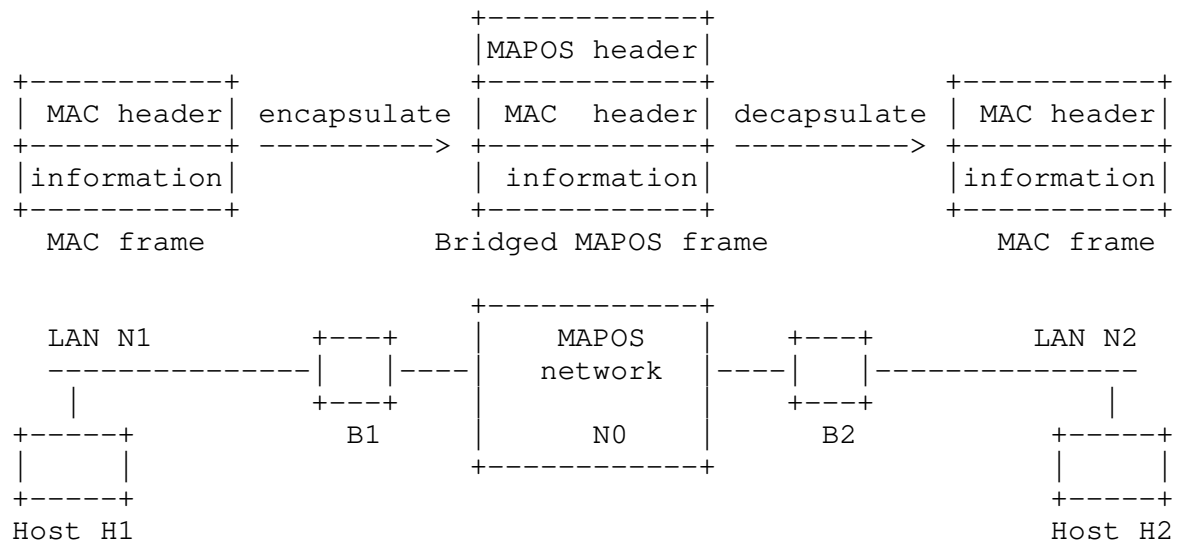


Figure 2. Forwarding a MAC frame from H1 to H2 over the VPN

The process of forwarding a MAC frame transparently from host H1 to host H2 is also shown in figure 2. NA B1 encapsulates a MAC frame from host H1, and forwards it to MAPOS network N0. NA B2 decapsulates the MAPOS frame, then forwards the MAC frame to host H2.

2.2. MAPOS encapsulation format

To transmit a MAC frame into MAPOS network, the NA encapsulates the frame as shown in the following figures. This frame format is based on Bridged LAN Traffic for PPP [4]; only the fields with semantics specific to this document are described below. The fields are transmitted from left to right.

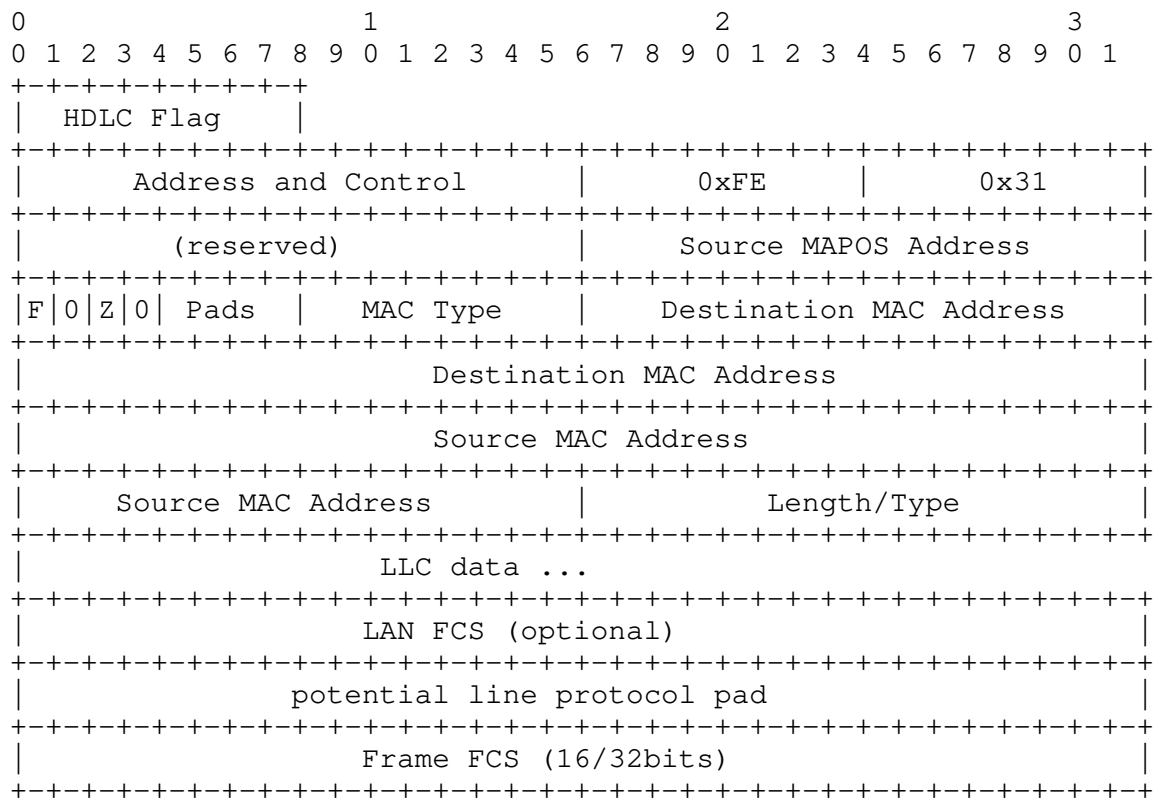


Figure 3. 802.3 Frame format (IEEE 802 Un-tagged Frame)

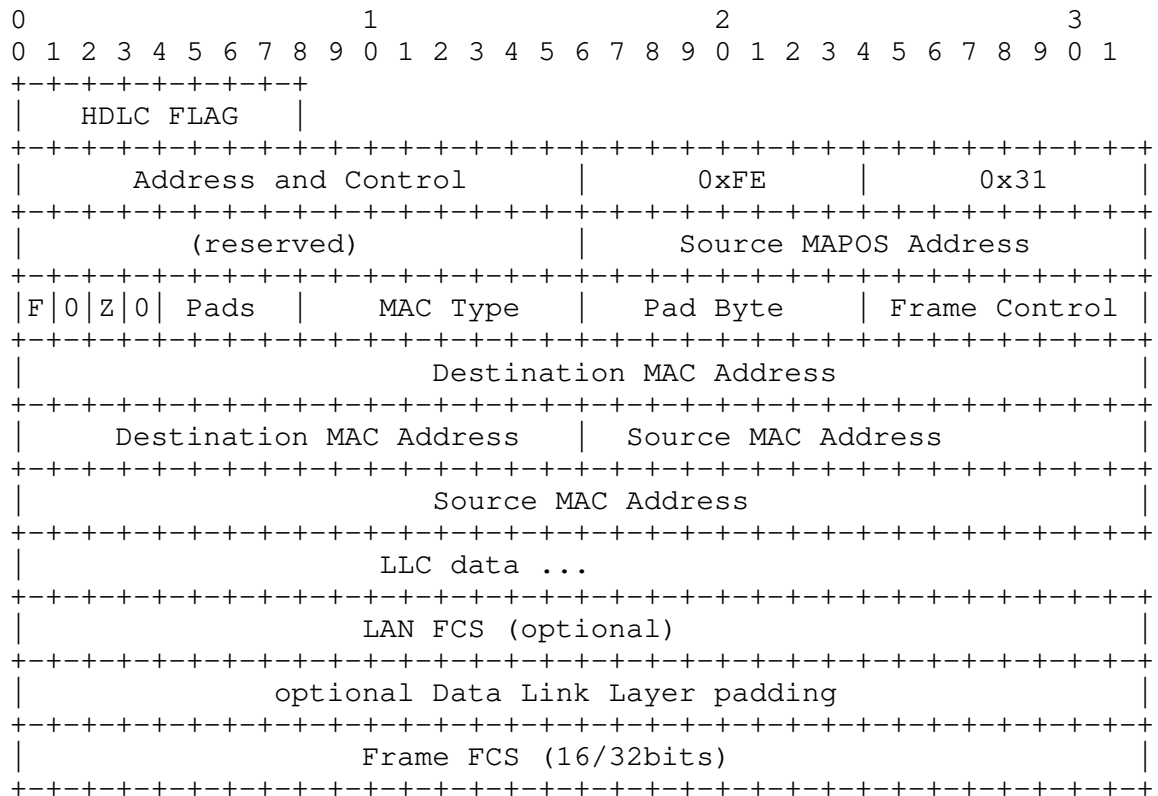


Figure 4. 802.4/802.5/FDDI Frame format (IEEE 802 Un-tagged Frame)

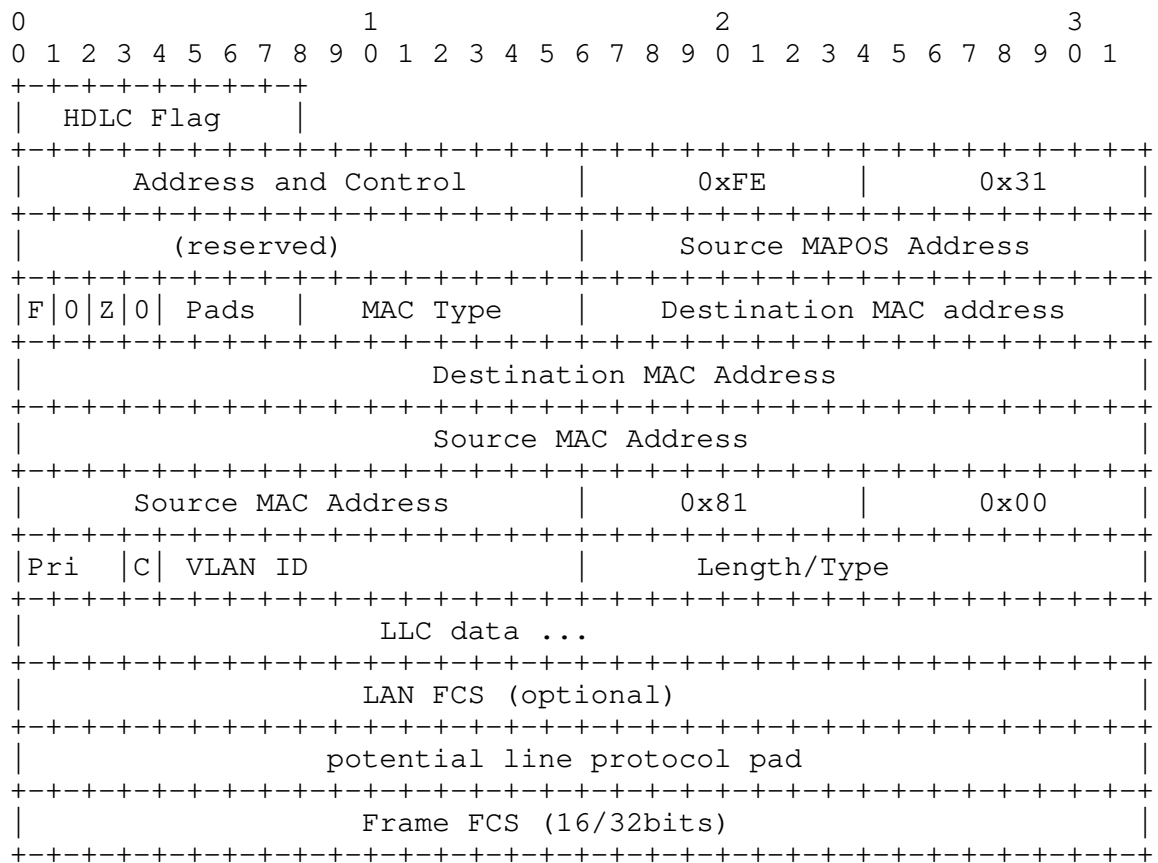


Figure 5. 802.3 Frame format (IEEE 802 Tagged Frame)

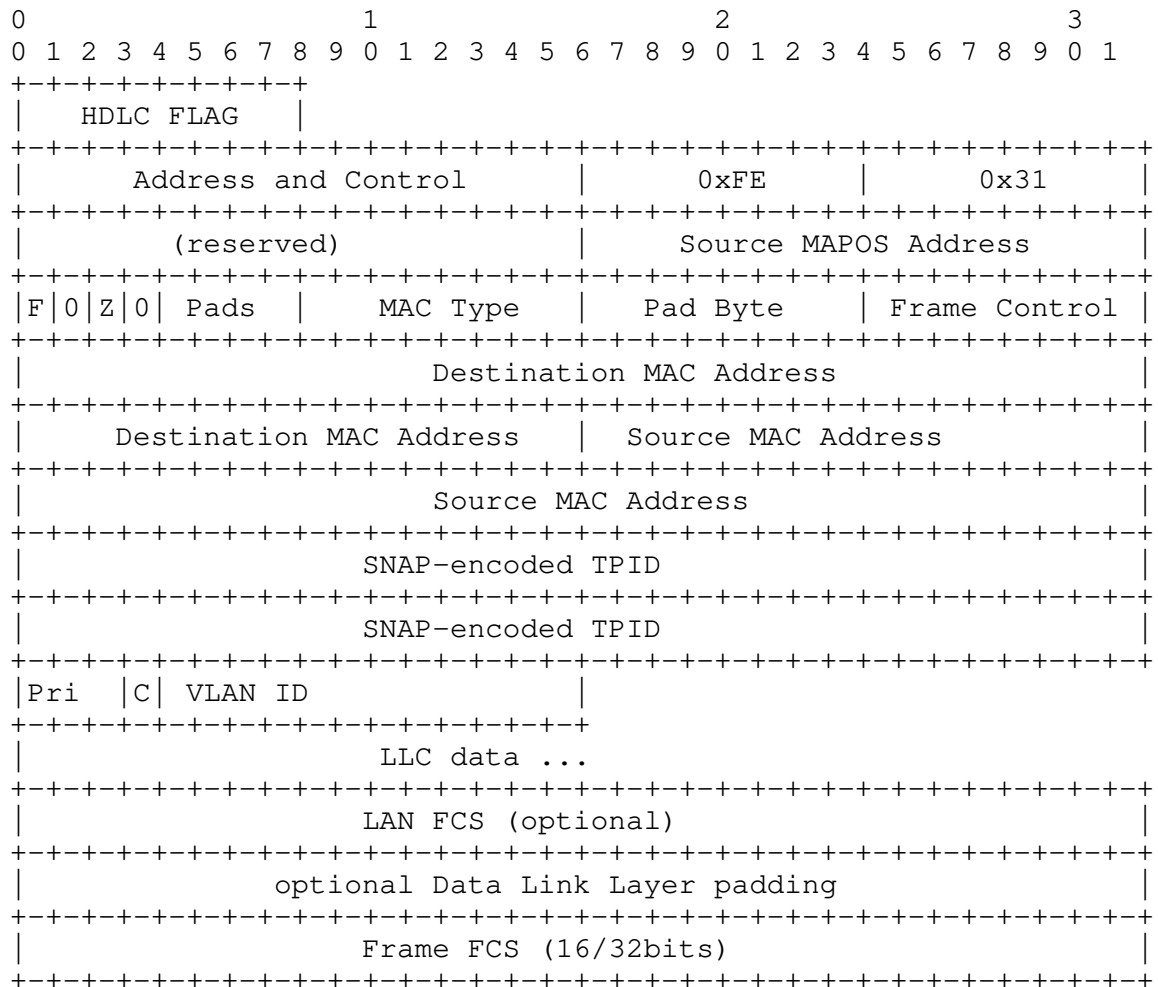


Figure 6. 802.4/802.5/FDDI Frame format (IEEE 802 Tagged Frame)

Address and Control

These fields contain the destination HDLC address as defined by MAPOS Version 1 [1] and MAPOS 16 [3].

Protocol Field

0xFE31 for bridged LAN traffic for MAPOS. NA should only accept NSP (0xFE03) and bridged MAPOS frames (0xFE31) frames; others should be silently discarded.

Source MAPOS address

Contains the MAPOS address of the sending NA. For MAPOS version 1 [1] the 8-bit HDLC address is placed in the least significant place of the 16-bit field and the upper eight bits must be zero.

3. Determination of the Destination MAPOS Address

The destination MAPOS address for a MAC frame to be bridged is determined by searching the address table composed of entries of the form

{destination MAC address, destination MAPOS address}

during the encapsulation phase.

For example, in figure 2, when a MAC frame to be sent to host H2 is encapsulated, the destination MAPOS address corresponding to NA B2 is used.

Determination of the destination MAPOS address for forwarding a MAC unicast frame is described in 3.1. The way for forwarding a MAC broadcast or multicast frame is described in 3.2. Methods for populating the address table are explained in 3.3.

3.1. Destination MAPOS address for forwarding a MAC unicast frame

In NA, entries of the form

{destination MAC address, destination MAPOS address}

are held in its address table. When a MAC frame is received by the ethernet interface, the address table is searched using the destination MAC address as the key. If a matching entry is found, the corresponding MAPOS address is used as the destination MAPOS address. If no matching entry exists, MAC broadcast forwarding (3.2) is used.

3.2. Forwarding a MAC broadcast or multicast frame

All MAC broadcast or multicast frames must be duplicated for transmission (via MAPOS unicast) to each of the peer network adapters in the same VLAN as the sending network adapter.

Consider an example shown in figure 7 where six LANs N1 through N6 are connected to the MAPOS network via network adapters B1 through B6.




Figure 7. Six networks connected to the MAPOS network

If a VLAN is configured with LANs N1, N2, and N3, a MAC broadcast or multicast frame originating from LAN N1 must not be forwarded to LAN N4, N5, or N6 but only to LANs N1, N2, and N3. It is duplicated twice for encapsulation and delivery to B2 and B3 via MAPOS unicast.

A set of network adapters that belongs to the same VLAN defines the broadcast scope of the VLAN. Before a VLAN is put to use, each NA in the VLAN must be configured with the MAPOS addresses of its peer NAs. A NA should silently discard bridged MAPOS frames with a MAPOS source address that is not among the peers that the NA knows about.

The use of MAPOS multicast for forwarding MAC broadcast frames is under further study.

3.3. Methods for configuring the address table

This section describes two methods for setting up an address table: static and dynamic. NA must implement the static method described in 3.3.1. The dynamic method (3.3.2) is optional, but an implementation must provide an option to disable this feature.

3.3.1. Static setup of address table

The address table can be set up statically. Before using a VLAN, address table entries for each NA in the VLAN must be populated manually.

These entries are considered permanent until they are manually removed, and must not be "aged" or overwritten by the dynamic procedure described in 3.3.2.

3.3.2. Dynamic setup of address table

The address table can also be set up dynamically. A NA discovers entries for its address table from incoming encapsulated MAPOS frames.

The NA adds the pair

```
{source MAC address, source MAPOS address}
```

to its address table when it receives an encapsulated MAPOS frame.

Entries discovered this way are subject to aging timer (should be configurable with the default of 300 seconds). Once the timer for an entry expires, the entry is removed from the address table. The timer is reset each time an encapsulated MAPOS frame with the same source MAC address is received.

There must be at most one entry for a source MAC address. If a discovered MAPOS address for a MAC address differs from the previously discovered address, the new one takes precedence and the address table entry must be overwritten. Under no circumstance may a discovered entry overwrite a statically created entry (3.3.1).

Discovery process using ARP [6] packets between host H1 (the MAC address is h1) in LAN N1 and host H2 (the MAC address is h2) in LAN N2 is shown below.

The MAPOS addresses of NAs B1, B2, B3 are b1, b2, b3 respectively.

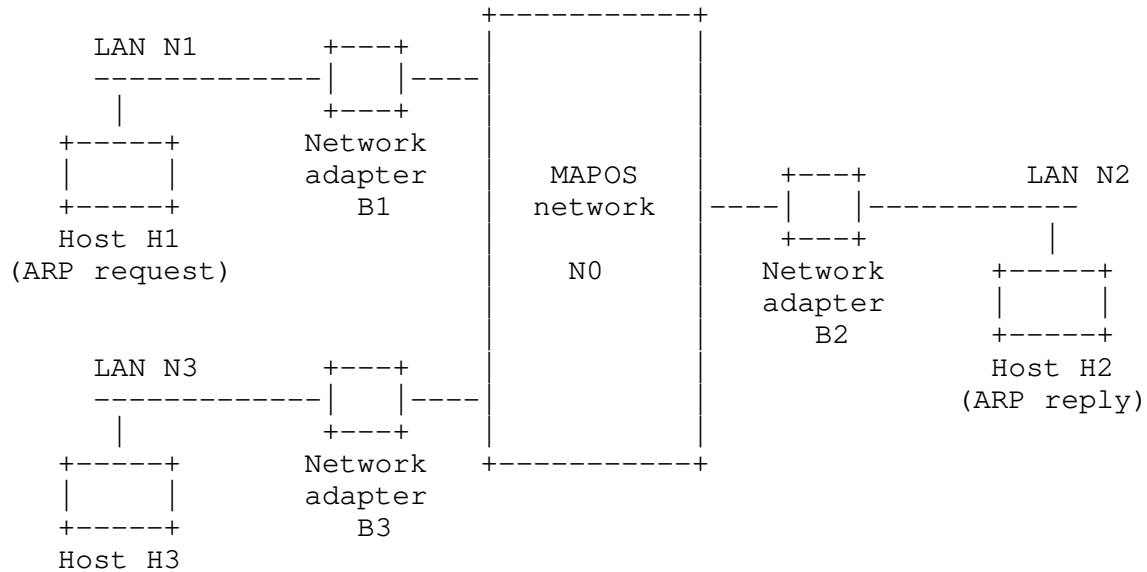


Figure 8. Three networks connected to the MAPOS network

- (1) Host H1 transmits an ARP request frame. An ARP request frame is a MAC broadcast Frame.
- (2) At NA B1, ARP request frame is received and is encapsulated. Because the VPN is composed of LANs N1, N2, and N3, the NA B1 must send a MAPOS frame that has destination MAPOS address b2 and another MAPOS frame that has destination MAPOS address b3. MAPOS address b1 is stored in the source MAPOS address field of each frame.
- (3) The bridged MAPOS frame arrives at NAs B2 and B3 from the MAPOS network.
- (4) NAs B2 and B3 receive the bridged MAPOS frame, and the pair {h1, b1} is added to their address tables.
- (5) In NA B2, the received MAPOS frame is decapsulated, and the MAC frame is forwarded to LAN N2. Similarly, in NA B3, the received MAPOS frame is decapsulated, and the MAC frame is forwarded to LAN N3.
- (6) At host H2, which exists in LAN N2, an ARP reply frame is transmitted to host H1.

- (7) Via the ethernet interface on NA B2, the ARP reply frame is received, and MAPOS encapsulation is done.

Because the entry

{h1, b1}

is registered in the address table, b1 is determined to be the destination MAPOS address. The bridged frame is forwarded to the MAPOS network.

- (8) MAPOS network delivers the bridged MAPOS frame to NA B1.
- (9) NA B1 decapsulates the bridged MAPOS frame, and forwards the MAC frame to LAN N1. At the same time, the entry {h2 , b2} is registered into NA B1 address table.
- (10) Host H1 receives the ARP reply frame.

4. Connecting a MAPOS Host to the VLAN

In order for a native MAPOS host to connect to a VLAN, it must have its own unique MAC address and implement all the features of a network adapter appropriate for the MAC framing that it wishes to use.


5. Security Considerations

This section discusses some of the security factors that need to be considered when planning a transparent LAN service described in section 1, "Network Model."

5.1 Management boundaries

In a large network, different parts of the network are managed by different organizations, and it is essential to clearly define the boundaries of management responsibilities.

A probable scenario is that a common carrier provides transparent LAN service to a variety of customers. Each customer is a distinct organization, expecting virtual private network service. In such a case, the common carrier should take management responsibility for the MAPOS network, optical cables to customer sites, and the network adapters that reside in customer premises.



In essence, the customer is allowed to do no more than connecting the cable from their MAC-based LAN to the network adapters. Common carrier should be very careful to monitor and protect their assets, including SONET/SDH connections and network adapters. In particular, network adapters serve as the primary line of defense against attacks and should be closely guarded.

5.2 Risks

Privacy of every customer connected to the carrier's MAPOS network may be compromised.

5.3 Attack against network adapters

A network adapter should be a dedicated device. This makes the device simple and easier to harden against break-in attempts. In the worst case, the device may crash causing network outage that only affects the customer that the failed network adapter serves. At this point, the privacy of other customers is still safe.

A more meaningful attack would be to replace a network adapter with some other intelligent agent that knows how network adapters work. This is possible because network adapters are customer premise equipment. Using such a device, an attacker can infiltrate the networks of other customers. Filtering based on source MAPOS address in bridging traffic is ineffective because this field is filled-in by network adapters -- MAPOS networks do not forward source addresses.

5.4 Filtering at network adapters and MAPOS switches

Network adapters should have the following frame filtering functions.

- Each NA in a VLAN is configured with the MAPOS addresses of its peer NAs that belongs to the same VLAN. A NA should only accept bridged MAPOS frames with a source MAPOS address of one of its VLAN peers.
- A NA should never import discovered address table entries with a MAPOS address that is not the address of one of its VLAN peers.
- If a NA detects that the amount of broadcast traffic from a host on MAC-base LAN exceeds a predefined threshold, the NA should stop forwarding traffic from that host.

By default, frame filtering by MAPOS switches is optional. It is desirable for a MAPOS switch to implement the following filtering features.

- A line interface of a MAPOS switch is made aware of the MAPOS addresses in the VLAN to which the interface participates. The interface discards all incoming bridged traffic (from the NA) that is destined to addresses outside of the VLAN's set.
- MAPOS switch assigns a MAPOS address to a NA using NSP. The switch discards all incoming bridged traffic (from the NA) with the source MAPOS address different from the one that is assigned by NSP.

5.5 Additional protection measures

A common carrier can implement additional protective measures such as the following.

- SONET/SDH connection is closely monitored. Once a network adapter is detected to have gone down, subsequent attempts at re-connecting to the MAPOS network are refused until manually re-enabled.
- Above method is effective against real attacks, but it also hinders timely recovery from accidents such as power outages. A reasonable trade-off solution is to implement an authentication mechanism between the MAPOS network and network adapters. Much like Challenge Handshake Authentication Protocol (CHAP) [8] used in PPP connection. Something similar may be implemented by defining additional message types to NSP.

6. References

- [1] Murakami, K. and M. Maruyama, "MAPOS - Multiple Access Protocol over SONET/SDH, Version 1", RFC 2171, June 1997.
- [2] Murakami, K. and M. Maruyama, "A MAPOS version 1 Extension - Node-Switch Protocol", RFC 2173, June 1997.
- [3] Murakami, K. and M. Maruyama, "MAPOS16 - Multiple Access Protocol over SONET/SDH with 16 Bit Addressing", RFC 2175, June 1997.
- [4] Higashiyama, M. and F.Baker, "PPP Bridging Control Protocol (BCP)", RFC 2878, July 2000.
- [5] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, January 2002.

- [6] Plummer, D.C., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.
- [7] IEEE 802.1D-1993, "Media Access Control (MAC) Bridges," ISO/IEC 15802-3:1993 ANSI/IEEE Std 802.1D, 1993 edition, July 1993.
- [8] Simpson, W., "PPP Challenge Handshake Authentication Protocols", RFC 1994, August 1996.

7. Acknowledgements

The authors would like to acknowledge the contributions and thoughtful suggestions of Naohisa Takahashi, Tetsuo Kawano and Tsuyoshi Ogura.

Appendix - Validation of the MAC Frame Forwarding Mechanism

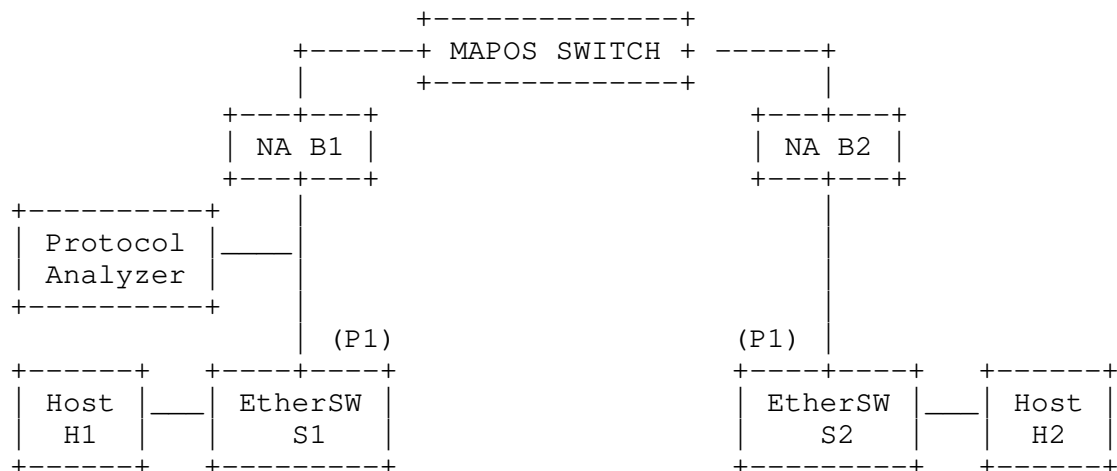
This appendix describes the configuration and procedure used to validate the soundness of the mechanism described in this document. The key points are:

- MAC frames are correctly forwarded by MAPOS network, and
- Even if a network contains loops, broadcast packets do not storm the network. MAC-based networks must use broadcast spanning tree technology in order for this to work.

(1) Verification of MAC frame forwarding on MAPOS network

Hosts H1 and H2, Ethernet switches S1 and S2, network adapters B1 and B2, and a MAPOS switch are connected as shown below. An ethernet protocol analyzer is placed between S1 and B1 for traffic monitoring.

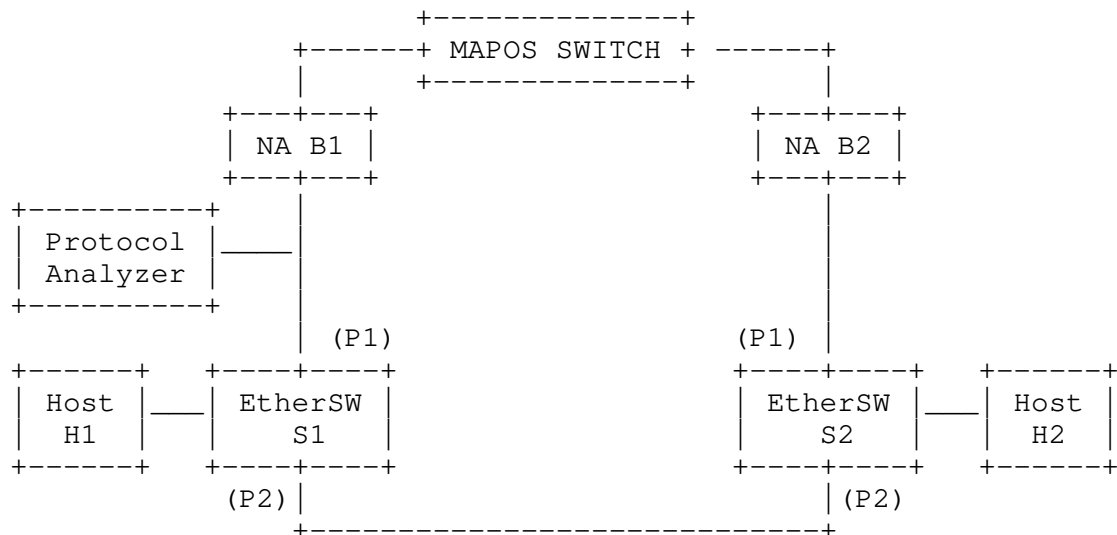
In the diagrams that follow, the hosts are x86 PC running FreeBSD 4.4-RELEASE, ethernet switches are Extreme Summit5i, network adapters are OKI Electric MA-1, and the MAPOS switch is CSR CoreSwitch80.



Correct forwarding of unicast MAC frames (ping) are observed between H1 and H2 through path (P1).

(2) Verification of spanning tree operation

- Enable spanning tree on S1 and S2.
- Connect S1 and S2 via path (P2) for redundancy.



It is observed that broadcast packets are correctly exchanged between S1 and S2, and that broadcast forwarding loop does not exist.

(3) Verification of spanning tree fail over

- H1 and H2 communication takes place through path (P1).
Spanning tree is configured such that Path (P2) is blocked.

It is observed that severing the link at any point along path (P1) makes the spanning tree configure itself to use path (P2).

It is also observed that restoring path (P1) makes the spanning tree configures itself to use path (P1).

Authors' Addresses

Osamu Okamoto
NTT Network Service System Laboratories
3-9-11, Midori-cho Musashino-shi
Tokyo 180-8585, Japan

EEmail: okamoto.osamu@lab.ntt.co.jp

Mitsuru Maruyama
NTT Network Innovation Laboratories
3-9-11, Midori-cho Musashino-shi
Tokyo 180-8585, Japan

EEmail: mitsuru@core.ecl.net

Takahiro Sajima
Sun Microsystems, K.K.
4-10-1, Yoga Setagaya-ku
Tokyo 158-8633, Japan

EEmail: tjs@sun.com

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

