

Network Working Group
Request for Comments: 3572
Category: Informational

T. Ogura
M. Maruyama
NTT Network Innovation Labs
T. Yoshida
Werk Mikro Systems
July 2003

Internet Protocol Version 6 over MAPOS
(Multiple Access Protocol Over SONET/SDH)

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 (2003). All Rights Reserved.

IESG Note

This memo documents a way of carrying IPv6 packets 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

Multiple Access Protocol over SONET/SDH (MAPOS) is a high-speed link-layer protocol that provides multiple access capability over a Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH).

This document specifies the frame format for encapsulating an IPv6 datagram in a MAPOS frame. It also specifies the method of forming IPv6 interface identifiers, the method of detecting duplicate addresses, and the format of the Source/Target Link-layer Addresses option field used in IPv6 Neighbor Discovery messages.

Table of Contents

- 1. Introduction 2
- 2. Frame Format for Encapsulating IPv6 Datagrams. 3
 - 2.1. Frame Format 3
 - 2.2. Maximum Transmission Unit (MTU). 3
 - 2.3. Destination Address Mapping. 4
 - 2.3.1. Unicast. 4
 - 2.3.2. Multicast 4
- 3. Interface Identifier 6
- 4. Duplicate Address Detection. 8
- 5. Source/Target Link-layer Address Option. 9
- 6. Security Considerations. 10
 - 6.1. Issues concerning Link-layer Addresses 10
 - 6.1.1. Protection against fraudulent reception of traffic 10
 - 6.1.2. Protection against improper traffic. 11
 - 6.2. Uniqueness of Interface Identifiers. 11
- 7. References. 12
- 8. Authors' Addresses 13
- 9. Full Copyright Statement 14

1. Introduction

Multiple Access Protocol over SONET/SDH (MAPOS) [1][2] is a high-speed link-layer protocol that provides multiple access capability over SONET/SDH. Its frame format is based on the HDLC-like (High Level Data Link Control) framing [3] for PPP. A component called a "Frame Switch" [1] allows multiple nodes (hosts and routers) 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 a wide geographical area.

This document specifies the frame format for encapsulating an Internet Protocol version 6 (IPv6) [4] datagram in a MAPOS frame, the method of forming IPv6 interface identifiers, the method of detecting duplicate addresses, and the format of the Source/Target Link-layer Addresses option field used in Neighbor Discovery messages such as Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement, and Redirect messages.

In the remainder of this document, the term "MAPOS" is used unless the distinction between MAPOS version 1 [1] and MAPOS 16 [2] is required.

2. Frame Format for Encapsulating IPv6 Datagrams

2.1. Frame Format

MAPOS uses the same HDLC-like framing as PPP-over-SONET, described in [3]. The MAPOS frame begins and ends with a flag sequence 01111110 (0x7E), and the MAPOS frame header contains address, control, and protocol fields. The address field contains a destination HDLC address. In MAPOS 16, the address field is extended to 16 bits, and the control field of MAPOS version 1 is omitted. The frame check sequence (FCS) field is 16 bits long by default, but a 32-bit FCS may be used optionally. Details of the MAPOS frame format are described in [1][2].

An IPv6 datagram is encapsulated in the MAPOS frame. In the case of encapsulating an IPv6 datagram, the protocol field must contain the value 0x0057 (hexadecimal). The IPv6 datagram is stored in the information field which follows immediately after the protocol field. That is, this field contains the IPv6 header followed immediately by the payload. Figure 1 shows the frame format. The fields are transmitted from left to right.




Figure 1. Frame format.

2.2. Maximum Transmission Unit (MTU)

The length of the information field of the MAPOS frame may vary, but shall not exceed 65,280 (64K - 256) octets [1][2]. The default maximum transmission unit (MTU) is 65,280 octets.

However, the MTU size may be reduced by a Router Advertisement [5] containing an MTU option that specifies a smaller MTU, or by manual configuration of each node. If a Router Advertisement received on a MAPOS interface has an MTU option specifying an MTU larger than 65,280, or larger than a manually configured value, that MTU option may be logged for the system management but must be otherwise ignored.

2.3. Destination Address Mapping

This section specifies the method of mapping an IPv6 destination address to the address field in the MAPOS frame header.

2.3.1. Unicast

In unicasting, the address field of a MAPOS frame contains the HDLC address that has been assigned via NSP (Node Switch Protocol) [6] to the MAPOS interface, which has the IPv6 unicast destination address.

In order to determine the destination HDLC address that corresponds to an IPv6 unicast destination address, the sender uses Link-layer Address Resolution described in [5].

2.3.2. Multicast

Address resolution is never performed on IPv6 multicast addresses. An IPv6 multicast destination address is mapped to the address field in the MAPOS frame header as described below for MAPOS version 1 and MAPOS 16.

MAPOS version 1:

The address field of the MAPOS version 1 frame header contains an 8-bit-wide destination HDLC address [1]. The least significant bit (LSB) of the field must always be 1 to indicate the end of the field. The most significant bit (MSB) is used to indicate whether the frame is a unicast or a multicast frame.

In the case of an IPv6 multicast, the MSB of the address field is 1 to indicate that the frame is multicast. As described above, the LSB of the address field is 1. The other six bits of the address field must contain the lowest-order six bits of the IPv6 multicast address. Figure 2 shows the address field of the MAPOS version 1 frame header in the case of an IPv6 multicast, where D(1) through D(6) represent the lowest-order six bits of the IPv6 multicast address. Exceptions arise when these six bits are either all zeros or all ones. In these cases, they should be altered to the bit sequence 111110. That is, the address field should be 0xFD (hexadecimal).




Figure 2. Address mapping in multicasting (MAPOS version 1).

MAPOS 16:

The address field of the MAPOS 16 frame header contains the 16-bit-wide destination HDLC address [2]. The LSB of the first octet must always be 0 to indicate the continuation of this field, and the LSB of the second octet must always be 1 to indicate the end of this field. The MSB of the first octet is used to indicate whether the frame is a unicast or a multicast frame.

In the case of an IPv6 multicast, the MSB of the first octet is 1 to indicate that the frame is multicast. As described above, the LSB of the first octet is 0 and the LSB of the second octet is 1. The other 13 bits of the address field must contain the lowest-order 13 bits of the IPv6 multicast address. Figure 3 shows the address field of the MAPOS 16 frame header in the case of an IPv6 multicast, where D(1) through D(13) represent the lowest-order 13 bits of the IPv6 multicast address. Exceptions arise when these 13 bits are either all zeros or all ones. In these cases, the address field should be 0xFEFD (hexadecimal).




Figure 3. Address mapping in multicasting (MAPOS 16).

3. Interface Identifier

This section specifies the method of forming the interface identifier [7].

A node that has one or more MAPOS interfaces must create one or more EUI-64 [8] based interface identifiers. Here, it should be noted that deriving interface identifiers from HDLC addresses of MAPOS interfaces is undesirable for the following reasons.

1. When a node is connected to a frame switch, an HDLC address is assigned to the interface of the node from the frame switch via NSP [6]. (In the remainder of this document, the term "MAPOS address" is used to refer to the address.) The value of the MAPOS address assigned to the interface depends on the combination of the switch number of the frame switch and the port number of the frame switch to which the interface is connected. The switch number is required to be unique only within a MAPOS multi-switch environment [6]; that is, there can be frame switches that have the same switch number in different MAPOS multi-switch environment separated by IP routers. Therefore, the uniqueness of a MAPOS address is guaranteed only within a MAPOS multi-switch environment.

Furthermore, if an implementation ensures that the link between the interface of the node and the port of the frame switch is hot-swappable, the port number of the frame switch or the frame switch connected to the interface of the node can be changed, so the MAPOS address assigned to the interface can also be changed without performing a system re-start of the node.

In short, the global uniqueness of a MAPOS address is not guaranteed, and a MAPOS address is not a built-in address but can be changed without performing a system re-start. Thus, if an interface identifier were derived from a MAPOS address, it could also be changed without a system re-start. This would not follow the recommendation in [7].

2. In the case of a point-to-point connection between two nodes, the same MAPOS address is assigned to each interface. Specifically, in the case of MAPOS version 1, the assigned address is 0x03 [6], and in the case of MAPOS 16, the assigned address is 0x0003 [2]. It is not easy to achieve link-locality of the interface identifier in a strict manner using the same Link-layer address.

For the above reasons, nodes with MAPOS interfaces must not derive their interface identifiers from their MAPOS addresses.

The following are methods of forming an interface identifier in the order of preference. These are almost the same as the methods described in [9] except that a MAPOS address must not be used as a source of uniqueness when an IEEE global identifier is unavailable.

- 1) If an IEEE global identifier (EUI-48 or EUI-64) is available anywhere on the node, it should be used to construct the interface identifier due to its uniqueness. When extracting an IEEE global identifier from another device on the node, care should be taken to ensure that the extracted identifier is presented in canonical ordering [10].

The only transformation from an EUI-64 identifier is to invert the "u" bit (universal/local bit in IEEE EUI-64 terminology). For example, for a globally unique EUI-64 identifier as shown in Figure 4:




Figure 4. Globally unique EUI-64 identifier.

where "c" are the bits of the assigned company_id, "0" is the value of the universal/local bit to indicate global scope, "g" is the group/individual bit, and "e" are the bits of the extension identifier, the IPv6 interface identifier would be as shown in Figure 5. The only change is inverting the value of the universal/local bit.




Figure 5. IPv6 interface identifier derived from a globally unique EUI-64 identifier.

In the case of an EUI-48 identifier, it is first converted to the EUI-64 format by inserting two octets, with hexadecimal values of 0xFF and 0xFE, in the middle of the 48-bit MAC (between the company_id and extension-identifier portions of the EUI-48 value).

For example, for a globally unique 48-bit EUI-48 identifier as shown in Figure 6:




Figure 6. Globally unique EUI-48 identifier.

where "c" are the bits of the assigned company_id, "0" is the value of the universal/local bit to indicate global scope, "g" is the group/individual bit, and "e" are the bits of the extension identifier, the IPv6 interface identifier would be as shown in Figure 7.




Figure 7. IPv6 interface identifier derived from a globally unique EUI-48 identifier.

- 2) If an IEEE global identifier is not available, a different source of uniqueness should be used. Suggested sources of uniqueness include machine serial numbers, etc. MAPOS addresses must not be used.

In this case, the "u" bit of the interface identifier must be set to 0.

- 3) If a good source of uniqueness cannot be found, it is recommended that a random number be generated. In this case the "u" bit of the interface identifier must be set to 0.

4. Duplicate Address Detection

Immediately after the system start-up, the MAPOS address has not yet been assigned to a MAPOS interface. The assignment is not completed until the adjacent frame switch, or adjacent node in the case of a point-to-point connection between two nodes, has delivered the MAPOS address to the interface via NSP [6]. Until then, no data transmission can be performed on the interface. Thus, a node must


conduct duplicate address detection [11] on all unicast addresses of MAPOS interfaces after the MAPOS address assignment has been completed by NSP.

5. Source/Target Link-layer Address Option

As specified in [5], the Source/Target Link-layer Address option is one of the options included in Neighbor Discovery messages. In [5], the length of the Source/Target Link-layer Address option field is specified in units of 8 octets. However, in the case of MAPOS, the length of the address field is 2 octets (MAPOS 16) or 1 octet (MAPOS version 1)[1][2]. Thus, if the exact form of the address field is embedded in the Link-layer Address field of the Source/Target Link-layer Address option field, the total length of the option field is 4 octets (MAPOS 16) or 3 octets (MAPOS version 1), both of which are shorter than 8 octets.

For the above reason, in the case of MAPOS, the Link-layer Address field of the Source/Target Link-layer Address option must be extended with zeros in order to extend the length of the option field to 8 octets, and the Length field must be set to 1 as shown below.

MAPOS version 1:



Fields:


Type: 1 for Source link-layer address.
 2 for Target link-layer address.

Length: 1 (in units of 8 octets).

Address: MAPOS version 1 8-bit address.

Figure 8. Format of the Source/Target Link-layer Address option field (MAPOS version 1).

MAPOS 16:



Fields:

Type: 1 for Source link-layer address.
 2 for Target link-layer address.

Length: 1 (in units of 8 octets).

Link-layer Address: MAPOS 16 16-bit address.

Figure 9. Format of the Source/Target Link-layer Address option field (MAPOS 16).

6. Security Considerations

In MAPOS, a link-layer address (MAPOS address) is assigned to a network interface by a frame switch via NSP; unlike other link-layer protocols such as Ethernet that use a built-in address on a network interface. Security considerations derived from this are described in 6.1 and 6.2. Because there is no link-layer security in MAPOS, the same security considerations as those of other link-layer protocols would be applied to other points.

6.1. Issues concerning Link-layer Addresses

6.1.1. Protection against fraudulent reception of traffic

In MAPOS, a MAPOS address is assigned by a frame switch, and it consists of the switch number and the port number of the switch to which the network interface is connected. (In the case of a point-to-point connection between two nodes, a fixed address is assigned to their network interfaces.) This brings the following advantages.

1. The value of the MAPOS address of a MAPOS network interface indicates the location of the interface in the MAPOS network. In other words, the value itself of the destination address of a MAPOS frame defines the actual location of the network interface to which the frame should be finally delivered. Therefore, as long as MAPOS addresses of network interfaces of nodes that have

been connected to the network through proper administrative process are held and frames are delivered only to those addresses, other nodes cannot receive frames unless their network interfaces are connected to the same ports of frame switches as those to which network interfaces of properly administered nodes are connected. This makes fraudulent reception of traffic difficult.

2. In the case where MAPOS addresses are not administered as mentioned above, it is possible that a malicious node could hijack traffic by spoofing its IPv6 address in a response to an IPv6 Neighbor Discovery. Even in this case, the node must advertise the true MAPOS address of its network interface in the response so that it can receive successive frames. This makes it easy to pinpoint the location of the host.

6.1.2. Protection against improper traffic

A MAPOS frame does not have a field for including its sender's address. Therefore, in the case where a node sends one-way improper traffic maliciously or accidentally, there is no way to obtain the sender's MAPOS address from the traffic and this leads to difficulty in identifying the node (because source IP addresses might be forged).

An effective way to alleviate the difficulty is to moderate the size of MAPOS multi-switch environment [6]. A common approach is to separate it using IP routers. This makes it easy to identify the node sending improper traffic within the multi-switch environment. To secure the environment against improper traffic from outside it, boundary IP routers need to block it using packet filtering based on IP layer information.

6.2. Uniqueness of Interface Identifiers

Global uniqueness of a MAPOS address is not guaranteed, and a MAPOS address is not a built-in address but can be changed without performing a system re-start if an implementation ensures that the link between the network interface of the node and the port of the frame switch is hot-swappable. Thus, an interface identifier must not be derived from a MAPOS address in order to ensure that the interface identifier is not changed without a system re-start.

As a consequence, in IP Version 6 over MAPOS, the existence of network interfaces other than MAPOS that have IEEE global identifier based addresses has great importance in creating interface identifiers. However, it may be common for there to be no such interfaces on a node, so a different source of uniqueness must be used. Therefore, sufficient care should be taken to prevent

duplication of interface identifiers. At present, there is no protection against duplication through accident or forgery.

7. 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, "MAPOS 16 - Multiple Access Protocol over SONET/SDH with 16 Bit Addressing", RFC 2175, June 1997.
- [3] Simpson, W., Ed., "PPP in HDLC-like Framing", STD 51, RFC 1662, July 1994.
- [4] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [5] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [6] Murakami, K. and M. Maruyama, "A MAPOS version 1 Extension - Node Switch Protocol", RFC 2173, June 1997.
- [7] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [8] IEEE, "Guidelines of 64-bit Global Identifier (EUI-64) Registration Authority", <http://standards.ieee.org/db/oui/tutorials/EUI64.html>, March 1997.
- [9] Haskin, D. and E. Allen, "IP Version 6 over PPP", RFC 2472, December 1998.
- [10] Narten, T. and C. Burton, "A Caution On The Canonical Ordering Of Link-Layer Addresses", RFC 2469, December 1998.
- [11] Thompson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

8. Authors' Addresses

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

EMail: ogura@core.ecl.net

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

EMail: mitsuru@core.ecl.net

Toshiaki Yoshida
Werk Mikro Systems
250-1, Mikajiri
Kumagaya
Saitama 360-0843, Japan

EMail: yoshida@peta.arch.ecl.net

9. Full Copyright Statement

Copyright (C) The Internet Society (2003). 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 assignees.

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.

