

Network Working Group
Request for Comments: 2171
Category: Informational

K. Murakami
M. Maruyama
NTT Laboratories
June 1997

MAPOS - Multiple Access Protocol over SONET/SDH 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 multiple access protocol for transmission of network-protocol datagrams, encapsulated in High-Level Data Link Control (HDLC) frames, 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 the protocol MAPOS, Multiple Access Protocol over SONET/SDH, for transmitting network-protocol datagrams over SONET/SDH. It focuses on the core protocol -- other documents listed in the bibliography may be referenced in conjunction with this document to provide support and services for protocols at higher layers.

1. Introduction

1.1 SONET/SDH

The Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) [1][2][3][4] family of ITU-T standard protocols are designed to provide common, simple, and flexible interface for broadband optical fiber transmission systems. It enables direct octet-synchronous multiplexing of lower rate tributaries. SONET/SDH-compliant transmission systems are widely deployed by telephone carriers world wide.


This document defines the MAPOS protocol -- a method for transmitting HDLC frames over SONET/SDH. The protocol provides multiple access capability to SONET/SDH, an inherently point-to-point link. This enables construction of seamless networking environment using SONET/SDH as transmission media for both LAN and WAN.

1.2 Possible Configurations


The MAPOS protocol provides multiple access, broadcast / multicast-capable switched LAN environment using SONET/SDH lines as transmission media. Possible configurations of MAPOS system are shown in the following diagrams. In (a), two end nodes are connected to each other. Figure (b) shows a star-topology "SONET-LAN" where multiple end nodes are connected to an HDLC frame switch. The frame switch forwards packets between nodes and provides multiple access capability. In (c), multiple frame switches are linked together, creating a switching cluster.



(a) Point-to-Point configuration



(b) Point-to-Multipoint configuration



(c) Switching cluster configuration

Figure 1. Possible configurations

Each port on a switch has an unique identifier within the switch. A node connected to a switch port must inherit the address of the port. That is, the node address is equal to the port identifier and is unique within the switch.

In a switch cluster, a node address is subnetted. The high-order bits, the part where the corresponding bits in the "subnet mask" are 1, indicate the switch address. The remaining low-order bits indicate the unique node address within the switch. The two fields form an unique address for a given node.

In either case, the address may be configured manually into a node interface, or automatically by the address assignment mechanism described in [5].

Note that any two components may be connected either directly, or via a long-haul SONET/SDH leased line.

1.3 Packet Transmission

The protocol is connection-less -- when a node wish to communicate with some other node, it simply fills-in the destination address of an HDLC frame, places it in one or more SONET/SDH payloads, and sends it over a SONET/SDH link.

The switch forwards the frame to its destination based on the destination address. In a switch cluster, the frame may be forwarded by multiple switches and is eventually delivered to the specified node. Broadcast and multicast are also supported. Frames with an invalid destination address are silently discarded.

Like ethernet, the multiple access capability is provided by a switch or a switch cluster. Since MAPOS is a link layer protocol, it is independent of the upper layer protocols. That is, it can support any network layer protocols such as IP. MAPOS IPv4 support is described in [6].

2. Physical Layer

This protocol treats the underlying end-to-end SONET/SDH transmission link as if it was a plain, transparent channel. It sends HDLC frames in SONET/SDH payloads, and expects them to arrive at the other end unaltered.

Each node and switch should terminate SONET/SDH overhead such as section overhead, line overhead, and path overhead according to the specification of SONET/SDH. Unfortunately, SONET and SDH overhead interpretations are not identical. In addition, some SONET/SDH implementations utilize some overhead bytes in proprietary manner.

The detail of the interpretation is beyond the scope of this document. Appendix A describes some of the most significant differences among SONET, SDH, and their implementations that often causes interoperability problems. Implementors of SONET/SDH interfaces are strongly encouraged to be aware of such differences, and provide workaround options in their products.

3. Data Link Layer

3.1 HDLC Frame Format

MAPOS uses the same HDLC-like framing as used in PPP-over-SONET, described in RFC-1662[7]. Figure 2 shows the frame format. Logical Link Control (LLC), and Sublayer/Sub-Network Access Protocol (SNAP) are not used. It does not include the bytes for transparency. The fields are transmitted from left to right.




Figure 2. Frame format

Flag Sequence

Flag sequence is used for frame synchronization. Each frame begins and ends with a flag sequence 01111110 (0x7E). If a frame immediately follows another, one flag sequence may be treated as the end of the preceding frame and the beginning of the immediately following frame. When the line is idle, the flag sequence is to be transmitted continuously on the line.

Address

The address field contains the destination HDLC address. A frame is forwarded by a switch based on this field. It is 8 bits wide. The LSB indicates the end of this field, and must always be 1. The MSB is used to indicate if the frame is a unicast or a multicast frame. The MSB of 0 means unicast, with the remaining six bits indicating the destination node address. MSB of 1 means multicast, with the remaining six bits indicating the group address. The address 11111111 (0xFF) means that the frame is a broadcast frame. The address 00000001 (0x01) is reserved to identify the control processor inside a switch. Frames with an invalid address should be silently discarded.

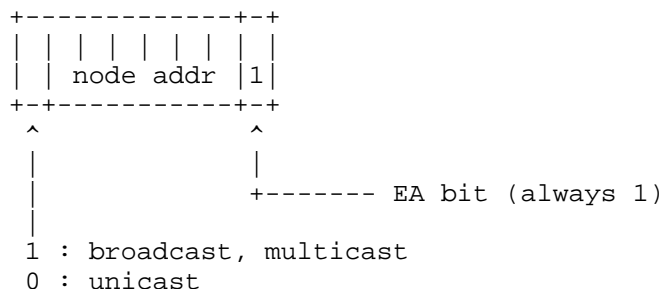


Figure 3 Address format

Control

The control field contains single octet 00000011 (0x03) which, in HDLC nomenclature, means that the frame is an Unnumbered Information (UI) with the Poll/Final (P/F) bit set to zero. Frames with any other control field values should be silently discarded.

Protocol

The protocol field indicates the protocol to which the datagram encapsulated in the information field belongs. It conforms to the ISO 3309 extension mechanism, and the value for this field may be obtained from the most recent "Assigned Numbers" [8] and "MAPOS Version 1 Assigned Numbers" [9].

Information

The information field contains the datagram for the protocol specified in the protocol field. The length of this field may vary, but shall not exceed 65,280 (64K - 256) octets.

Frame Check Sequence (FCS)

By default, the frame check sequence (FCS) field is 16-bits long. Optionally, 32 bit FCS may be used instead. The FCS is calculated over all bits of the address, control, protocol, and information fields prior to escape conversions. The least significant octet of the result is transmitted first as it contains the coefficient of the highest term.

Inter-frame fill

A sending station must continuously transmit the flag sequence as inter-frame fill after the FCS field. The inter-frame flag sequences must be silently discarded by the receiving station.

When an under-run occurs during DMA in the sending station, it must abort the frame transfer and continuously send the flag sequence to indicate the error.

3.2 Octet-Synchronous Framing

MAPOS uses an octet stuffing procedure because it treats SONET/SDH as a byte-oriented synchronous link. Since SONET/SDH provides transparency, Async-Control-Character-Map (ACCM) is not used. HDLC frames are mapped into the SONET/SDH payload as follows.

Each HDLC frame is separated from another frame by one or more flag sequence, 01111110 (0x7E). An escape sequence is defined to escape the flag sequence and itself. Prior to sending the frame, but after the FCS computation, every occurrence of 01111110 (0x7E) other than the flags is to be converted to the sequence 01111101 01011110 (0x7D 0x5E), and the sequence 01111101 (0x7D) is to be converted to the sequence 01111101 01011101 (0x7D 0x5D). Upon receiving a frame, this conversion must be reversed prior to FCS computation.

4. Further Reading

To fully utilize MAPOS protocol, it is useful to reference other documents[5][6][9][10] in conjunction with this document.

5. Security Considerations

Security issues are not discussed in this memo.

References

- [1] CCITT Recommendation G.707: Synchronous Digital Hierarchy Bit Rates (1990).
- [2] CCITT Recommendation G.708: Network Node Interface for Synchronous Digital Hierarchy (1990).
- [3] CCITT Recommendation G.709: Synchronous Multiplexing Structure (1990).
- [4] American National Standard for Telecommunications - Digital Hierarchy - Optical Interface Rates and Formats Specification, ANSI T1.105-1991.
- [5] Murakami, K. and M. Maruyama, "A MAPOS version 1 Extension - Node Switch Protocol," RFC2173, June, 1997.

- [6] Murakami, K. and M. Maruyama, "IPv4 over MAPOS Version 1," RFC2176, June, 1997.
- [7] Simpson, W., editor, "PPP in HDLC-like Framing," RFC1662, July 1994.
- [8] IANA, "IANA-Assignments,"
<http://www.iana.org/iana/assignments.html>
- [9] Maruyama, M. and K. Murakami, "MAPOS Version 1 Assigned Numbers," RFC2172, June 1997.
- [10] Murakami, K. and M. Maruyama, "A MAPOS version 1 Extension - Switch Switch Protocol," RFC2174, June, 1997.

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

APPENDIX A. Differences among SONET, SDH, and their Implementations

This section briefly describes the major differences among SONET which is an ANSI standard, SDH, an ITU-T standard, and their implementations.

AU pointer (H1, H2, H3)

The AU pointer consists of bytes H1, H2, and H3. The bits 5 and 6 of the H1 byte are called "SS bits," and are used to indicate the offset into the payload where the beginning of a SPE is located. (Note that "SPE" is a SONET term -- SDH calls it "VC.") In the case of OC-3c, SONET sets the SS bits of the second and the third H1 bytes to 0, whereas SDH sets them to 10 for AU-4, and 01 for AU-31. Although the SS bits may be ignored at the receiving station, some transmission systems discards SONET/SDH frames with SS bits that it doesn't expect -- the sending station should be aware of this, and include a configuration option to handle it.

Z1 and Z2

The Z bytes are reserved in SONET/SDH. Some transmission systems, however, use them in a proprietary manner. SONET uses Z1 for Line Error Monitoring. NTT, a carrier in Japan, utilized Z1 for Automatic Protection Switching (APS.)

DCC Bytes

The D bytes are called the Data Communication channel (DCC), and are defined for maintenance and operations. However, some carriers and vendors use them in a proprietary manner. For example, NTT's STM-1 UNI uses the D4, D5, and D6 bytes to transfer section and path maintenance information.

