

Network Working Group
Request for Comments: 3293
Category: Standards Track

A. Doria
Lulea University of Technology
J. Buerkle
Nortel Networks
T. Worster
June 2002

General Switch Management Protocol (GSMP)
Packet Encapsulations for Asynchronous Transfer Mode (ATM),
Ethernet and Transmission Control Protocol (TCP)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

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

Abstract

This memo specifies the encapsulation of GSMP (General Switch Management Protocol) packets in ATM (Asynchronous Transfer Mode), Ethernet and TCP (Transmission Control Protocol).

Specification of Requirements

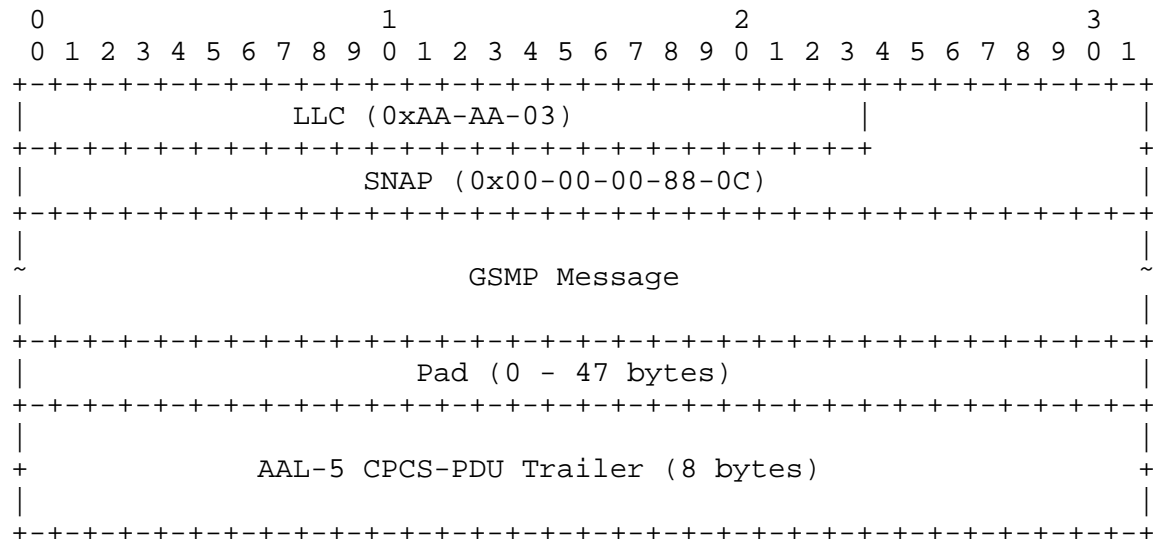
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [7].

1. Introduction

GSMP messages are defined in [1] and MAY be encapsulated in several different protocols for transport. This memo specifies their encapsulation in ATM AAL-5, in Ethernet or in TCP. Other encapsulations may be defined in future specifications.

2. ATM Encapsulation

GSMP packets are variable length and for an ATM data link layer they are encapsulated directly in an AAL-5 CPCS-PDU [3][4] with an LLC/SNAP header as illustrated:



(The convention in the documentation of Internet Protocols [5] is to express numbers in decimal. Numbers in hexadecimal format are specified by prefacing them with the characters "0x". Numbers in binary format are specified by prefacing them with the characters "0b". Data is pictured in "big-endian" order. That is, fields are described left to right, with the most significant byte on the left and the least significant byte on the right. Whenever a diagram shows a group of bytes, the order of transmission of those bytes is the normal order in which they are read in English. Whenever a byte represents a numeric quantity the left most bit in the diagram is the high order or most significant bit. That is, the bit labelled 0 is the most significant bit. Similarly, whenever a multi-byte field represents a numeric quantity the left most bit of the whole field is the most significant bit. When a multi-byte quantity is transmitted, the most significant byte is transmitted first. This is the same coding convention as is used in the ATM layer [2] and AAL-5 [3][4].)

The LLC/SNAP header contains the bytes: 0xAA 0xAA 0x03 0x00 0x00 0x00 0x88 0x0C. (0x880C is the assigned Ethertype for GSMP.)

The maximum transmission unit (MTU) of the GSMP Message field is 1492 bytes.

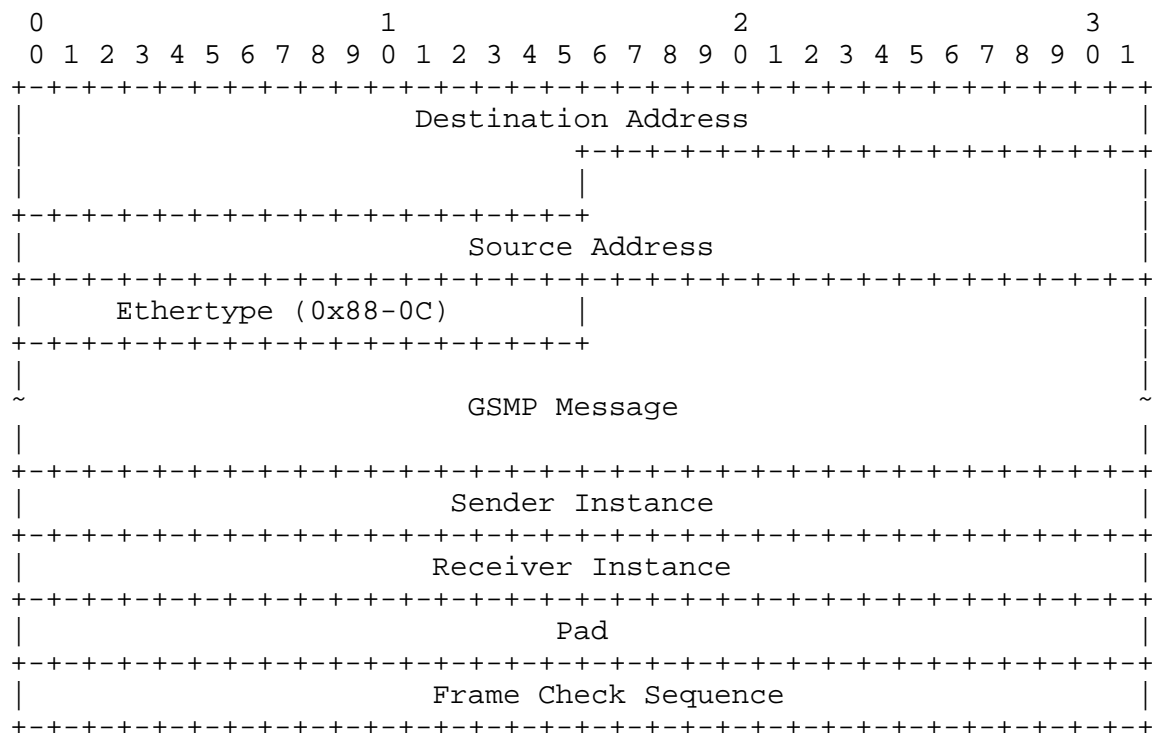
The virtual channel over which a GSMP session is established between a controller and the switch it is controlling is called the GSMP control channel. The default VPI and VCI of the GSMP control channel for LLC/SNAP encapsulated GSMP messages on an ATM data link layer is:

VPI = 0
VCI = 15.

The GSMP control channel MAY be changed using the GSMP MIB.

3. Ethernet Encapsulation

GSMP packets MAY be encapsulated on an Ethernet data link as illustrated:



Destination Address

For the SYN message of the adjacency protocol the Destination Address is the broadcast address 0xFFFFFFFFFFFF. (Alternatively, it is also valid to configure the node with the unicast 48-bit IEEE MAC address of the destination. In this case the configured unicast Destination Address is used in the SYN message.) For all other messages the Destination Address is the unicast 48-bit

IEEE. MAC address of the destination. This address may be discovered from the Source Address field of messages received during synchronisation of the adjacency protocol.

Source Address

For all messages, the Source Address is the 48-bit IEEE MAC address of the sender.

Ethertype

The assigned Ethertype for GSMP is 0x880C.

GSMP Message

The maximum transmission unit (MTU) of the GSMP Message field is 1492 bytes.

Sender Instance

The Sender Instance number for the link obtained from the adjacency protocol. This field is already present in the adjacency protocol message. It is appended to all non-adjacency GSMP messages in the Ethernet encapsulation to offer additional protection against the introduction of corrupt state.

Receiver Instance

The Receiver Instance number is what the sender believes is the current instance number for the link, allocated by the entity at the far end of the link. This field is already present in the adjacency protocol message. It is appended to all non-adjacency GSMP messages in the Ethernet encapsulation to offer additional protection against the introduction of corrupt state.

Pad

After adjacency has been established the minimum length of the data field of an Ethernet packet is 46 bytes. If necessary, padding should be added such that it meets the minimum Ethernet frame size. This padding should be bytes of zero and is not to be considered part of the GSMP message.

Frame Check Sequence

The Frame Check Sequence (FCS) is defined in IEEE 802.3 [6] as follows:

Note: This section is included for informational and historical purposes only. The normative reference can be found in IEEE 802.3 Standard [6].

"A cyclic redundancy check (CRC) is used by the transmit and receive algorithms to generate a CRC value for the FCS field. The frame check sequence (FCS) field contains a 4-byte (32-bit)

cyclic redundancy check (CRC) value. This value is computed as a function of the contents of the source address, destination address, length, LLC data and pad (that is, all fields except the preamble, SFD, FCS and extension). The encoding is defined by the following generating polynomial.

$$G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x^1."$$

The procedure for the CRC calculation can be found in [6].

After the adjacency protocol has achieved synchronisation, for every GSMP message received with an Ethernet encapsulation, the receiver must check the Source Address from the Ethernet MAC header, the Sender Instance, and the Receiver Instance. The incoming GSMP message must be discarded if the Sender Instance and the Source Address do not match the values of the Sender Instance and the Sender Name stored by the "Update Peer Verifier" operation of the GSMP adjacency protocol. The incoming GSMP message must also be discarded if it arrives over any port other than the port over which the adjacency protocol has achieved synchronisation. In addition, the incoming message must also be discarded if the Receiver Instance field does not match the current value for the Sender Instance of the GSMP adjacency protocol.

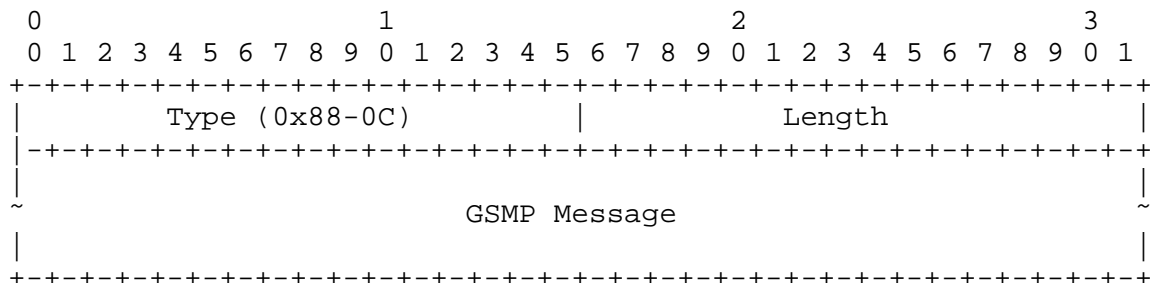
4. TCP/IP Encapsulation

When GSMP messages are transported over an IP network, they MUST be transported using the TCP encapsulation. TCP provides reliable transport, network flow control, and end-system flow control suitable for networks that may have high loss and variable or unpredictable delay.

For TCP encapsulations of GSMP messages, the controller runs the client code and the switch runs the server code. Upon initialisation, the server is listening on GSMP's TCP port number: 6068. The controller establishes a TCP connection with each switch it manages. The switch under control MUST be a multi-connection server (PORT 6068) to allow creation of multiple control sessions from N GSMP controller instances. Adjacency protocol messages, which are used to synchronise the controller and switch and maintain handshakes, are sent by the controller to the switch after the TCP connection is established. GSMP messages other than adjacency protocol messages MUST NOT be sent until after the adjacency protocol has achieved synchronisation. The actual GSMP message flow will occur on other ports.

4.1 Message Formats

GSMP messages are sent over a TCP connection. A GSMP message is processed only after it is entirely received. A four-byte TLV header field is prepended to the GSMP message to provide delineation of GSMP messages within the TCP stream.



Type

This 2-byte field indicates the type code of the following message. The type code for GSMP messages is 0x88-0C (i.e., the same as GSMP's Ethertype).

Length

This 2-byte unsigned integer indicates the total length of the GSMP message only. It does not include the 4-byte TLV header.

4.2 TCP/IP Security consideration

When GSMPv3 is implemented for use in IP networks, provisions for security between the controller and client MUST be available and MUST be provided by IP Security [IPSEC]. In this case, the IPSEC Encapsulation Security Payload (ESP) MUST be used to provide both integrity and confidentiality.

5. Security Considerations

The security of GSMP's TCP/IP control channel has been addressed in Section 4.2. For all uses of GSMP over an IP network it is REQUIRED that GSMP be run over TCP/IP using the security considerations discussed in Section 4.2. Security using ATM and Ethernet encapsulations MAY be provided at the link layer. Discussion of these methods is beyond the scope of this specification. For secure operation over any media, the IP encapsulation with IPsec SHOULD be used.

References

- [1] Doria, A., Sundell, K., Hellstrand, F. and T. Worster, "General Switch Management Protocol (GSMP) V3", RFC 3292, June 2002.
- [2] "B-ISDN ATM Layer Specification," International Telecommunication Union, ITU-T Recommendation I.361, Feb. 1999.
- [3] "B-ISDN ATM Adaptation Layer (AAL) Specification," International Telecommunication Union, ITU-T Recommendation I.363, Mar. 1993.
- [4] "B-ISDN ATM Adaptation Layer specification: Type 5 AAL", International Telecommunication Union, ITU-T Recommendation I.363.5, Aug. 1996.
- [5] Reynolds, J., Editor, "Assigned Numbers", RFC 3232, January 2002.
- [6] IEEE Std 802.3, 1998 Edition
"Information technology-Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications"
- [7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Tom Worster

Phone: +1 617 247 2624

EMail: fsb@thefsb.org

Avri Doria

Div. of Computer Communications

Lulea University of Technology

S-971 87 Lulea

Sweden

Phone: +1 401 663 5024

EMail: avri@acm.com

Joachim Buerkle

Nortel Networks Germany GmbH & Co. KG

Hahnstr. 37-39

60528 Frankfurt am Main

Germany

EMail: Joachim.Buerkle@nortelnetworks.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.

