

Network Working Group
Request for Comments: 3292
Category: Standards Track

A. Doria
Lulea University of Technology
F. Hellstrand
K. Sundell
Nortel Networks
T. Worster
June 2002

General Switch Management Protocol (GSMP) V3

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 document describes the General Switch Management Protocol Version 3 (GSMPv3). The GSMPv3 is an asymmetric protocol that allows one or more external switch controllers to establish and maintain the state of a label switch such as, an ATM, frame relay or MPLS switch. The GSMPv3 allows control of both unicast and multicast switch connection state as well as control of switch system resources and QoS features.

Acknowledgement

GSMP was created by P. Newman, W. Edwards, R. Hinden, E. Hoffman, F. Ching Liaw, T. Lyon, and G. Minshall (see [6] and [7]). This version of GSMP is based on their work.

Contributors

In addition to the authors/editors listed in the heading, many members of the GSMP group have made significant contributions to this specification. Among the contributors who have contributed materially are: Constantin Adam, Clint Bishard, Joachim Buerkle, Torbjorn Hedqvist, Georg Kullgren, Aurel A. Lazar, Mahesan Nandikesan, Matt Peters, Hans Sjostrand, Balaji Srinivasan, Jaroslav Sydir, Chao-Chun Wang.

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 [RFC2119].

Table of Contents

1. Introduction	4
2. GSMP Packet Encapsulation	6
3. Common Definitions and Procedures	6
3.1 GSMP Packet Format	7
3.1.1 Basic GSMP Message format	7
3.1.2 Fields commonly found in GSMP messages	11
3.1.3 Labels	12
3.1.4 Failure Response Messages	17
4. Connection Management Messages	18
4.1 General Message Definitions	18
4.2 Add Branch Message	25
4.2.1 ATM specific procedures:	29
4.3 Delete Tree Message	30
4.4 Verify Tree Message	30
4.5 Delete All Input Port Message	30
4.6 Delete All Output Port Message	31
4.7 Delete Branches Message	32
4.8 Move Output Branch Message	35
4.8.1 ATM Specific Procedures:	37
4.9 Move Input Branch Message	38
4.9.1 ATM Specific Procedures:	41
5. Reservation Management Messages	42
5.1 Reservation Request Message	43
5.2 Delete Reservation Message	46
5.3 Delete All Reservations Message	47
6. Management Messages	47
6.1 Port Management Message	47
6.2 Label Range Message	53
6.2.1 Labels	56
7. State and Statistics Messages	60
7.1 Connection Activity Message	61
7.2 Statistics Messages	64
7.2.1 Port Statistics Message	67
7.2.2 Connection Statistics Message	67
7.2.3 QoS Class Statistics Message	68
7.3 Report Connection State Message	68
8. Configuration Messages	73
8.1 Switch Configuration Message	73
8.1.1 Configuration Message Processing	75
8.2 Port Configuration Message	75

8.2.1 PortType Specific Data	79
8.3 All Ports Configuration Message	87
8.4 Service Configuration Message	89
9. Event Messages	93
9.1 Port Up Message	95
9.2 Port Down Message	95
9.3 Invalid Label Message	95
9.4 New Port Message	96
9.5 Dead Port Message	96
9.6 Adjacency Update Message	96
10. Service Model Definition	96
10.1 Overview	96
10.2 Service Model Definitions	97
10.2.1 Original Specifications	97
10.2.2 Service Definitions	98
10.2.3 Capability Sets	99
10.3 Service Model Procedures	99
10.4 Service Definitions	100
10.4.1 ATM Forum Service Categories	101
10.4.2 Integrated Services	104
10.4.3 MPLS CR-LDP	105
10.4.4 Frame Relay	105
10.4.5 DiffServ	106
10.5 Format and Encoding of the Traffic Parameters	106
10.5.1 Traffic Parameters for ATM Forum Services	106
10.5.2 Traffic Parameters for Int-Serv Controlled Load Service	107
10.5.3 Traffic Parameters for CRLDP Service	108
10.5.4 Traffic Parameters for Frame Relay Service	109
10.6 Traffic Controls (TC) Flags	110
11. Adjacency Protocol	111
11.1 Packet Format	112
11.2 Procedure	115
11.2.1 State Tables	117
11.3 Partition Information State	118
11.4 Loss of Synchronisation.....	119
11.5 Multiple Controllers Per Switch Partition	119
11.5.1 Multiple Controller Adjacency Process	120
12. Failure Response Codes	121
12.1 Description of Failure and Warning Response Messages	121
12.2 Summary of Failure Response Codes and Warnings	127
13. Security Considerations	128
Appendix A Summary of Messages	129
Appendix B IANA Considerations	130
References	134
Authors' Addresses	136
Full Copyright Statement	137

1. Introduction

The General Switch Management Protocol (GSMP) is a general purpose protocol to control a label switch. GSMP allows a controller to establish and release connections across the switch, add and delete leaves on a multicast connection, manage switch ports, request configuration information, request and delete reservation of switch resources, and request statistics. It also allows the switch to inform the controller of asynchronous events such as a link going down. The GSMP protocol is asymmetric, the controller being the master and the switch being the slave. Multiple switches may be controlled by a single controller using multiple instantiations of the protocol over separate control connections. Also a switch may be controlled by more than one controller by using the technique of partitioning.

A "physical" switch can be partitioned into several virtual switches that are referred to as partitions. In this version of GSMP, switch partitioning is static and occurs prior to running GSMP. The partitions of a physical switch are isolated from each other by the implementation and the controller assumes that the resources allocated to a partition are at all times available to that partition. A partition appears to its controller as a label switch. Throughout the rest of this document, the term switch (or equivalently, label switch) is used to refer to either a physical, non-partitioned switch or to a partition. The resources allocated to a partition appear to the controller as if they were the actual physical resources of the partition. For example if the bandwidth of a port were divided among several partitions, each partition would appear to the controller to have its own independent port.

GSMP controls a partitioned switch through the use of a partition identifier that is carried in every GSMP message. Each partition has a one-to-one control relationship with its own logical controller entity (which in the remainder of the document is referred to simply as a controller) and GSMP independently maintains adjacency between each controller-partition pair.

Kinds of label switches include frame or cell switches that support connection oriented switching, using the exact match-forwarding algorithm based on labels attached to incoming cells or frames. A switch is assumed to contain multiple "ports". Each port is a combination of one "input port" and one "output port". Some GSMP requests refer to the port as a whole, whereas other requests are specific to the input port or the output port. Cells or labelled frames arrive at the switch from an external communication link on

incoming labelled channels at an input port. Cells or labelled frames depart from the switch to an external communication link on labelled channels from an output port.

A switch may support multiple label types, however, each switch port can support only one label type. The label type supported by a given port is indicated by the switch to the controller in a port configuration message. Connections may be established between ports, supporting different label types. Label types include ATM, Frame Relay, MPLS Generic and FEC Labels.

A connection across a switch is formed by connecting an incoming labelled channel to one or more outgoing labelled channels. Connections are referenced by the input port on which they originate and the Label values of their incoming labelled channel.

GSMP supports point-to-point and point-to-multipoint connections. A multipoint-to-point connection is specified by establishing multiple point-to-point connections, each of them specifying the same output branch. A multipoint-to-multipoint connection is specified by establishing multiple point-to-multipoint trees each of them specifying the same output branches.

In general a connection is established with a certain quality of service (QoS). This version of GSMP includes a default QoS Configuration and additionally allows the negotiation of alternative, optional QoS configurations. The default QoS Configuration includes three QoS Models: a Service Model, a Simple Abstract Model (strict priorities) and a QoS Profile Model.

The Service Model is based on service definitions found external to GSMP such as in Integrated Services or ATM Service Categories. Each connection is assigned a specific service that defines the handling of the connection by the switch. Additionally, traffic parameters and traffic controls may be assigned to the connection depending on the assigned service.

In the Simple Abstract Model, a connection is assigned a priority when it is established. It may be assumed that for connections that share the same output port, a cell or frame on a connection with a higher priority is much more likely to exit the switch before a cell or frame on a connection with a lower priority if they are both in the switch at the same time. The number of priorities that each port of the switch supports may be obtained from the port configuration message.

The QoS Profile Model provides a simple mechanism that allows connection to be assigned QoS semantics defined externally to GSMP. The QoS Profile Model can be used to indicate pre-defined Differentiated Service Per Hop Behaviours (PHBs). Definition of QoS profiles is outside of the scope of this specification.

All GSMP switches MUST support the default QoS Configuration. A GSMP switch may additionally support one or more alternative QoS Configurations. The QoS models of alternative QoS configurations are defined outside the GSMP specification. GSMP includes a negotiation mechanism that allows a controller to select from the QoS configurations that a switch supports.

GSMP contains an adjacency protocol. The adjacency protocol is used to synchronise states across the link, to negotiate which version of the GSMP protocol to use, to discover the identity of the entity at the other end of a link, and to detect when it changes.

2. GSMP Packet Encapsulation

GSMP packets may be transported via any suitable medium. GSMP packet encapsulations for ATM, Ethernet and TCP are specified in [15]. Additional encapsulations for GSMP packets may be defined in separate documents.

3. Common Definitions and Procedures

GSMP is a master-slave protocol. The controller issues request messages to the switch. Each request message indicates whether a response is required from the switch and contains a transaction identifier to enable the response to be associated with the request. The switch replies with a response message indicating either a successful result or a failure. There are six classes of GSMP request-response message: Connection Management, Reservation Management, Port Management, State and Statistics, Configuration, and Quality of Service. The switch may also generate asynchronous Event messages to inform the controller of asynchronous events. The controller can be required to acknowledge event messages, but by default does not do so. There is also an adjacency protocol message used to establish synchronisation across the link and maintain a handshake.

For the request-response messages, each message type has a format for the request message and a format for the success response. Unless otherwise specified a failure response message is identical to the request message that caused the failure, with the Code field indicating the nature of the failure.

Switch ports are described by a 32-bit port number. The switch assigns port numbers and it may typically choose to structure the 32 bits into opaque sub-fields that have meaning to the physical structure of the switch (e.g., slot, port). In general, a port in the same physical location on the switch will always have the same port number, even across power cycles. The internal structure of the port number is opaque to the GSMP protocol. However, for the purposes of network management such as logging, port naming, and graphical representation, a switch may declare the physical location (physical slot and port) of each port. Alternatively, this information may be obtained by looking up the product identity in a database.

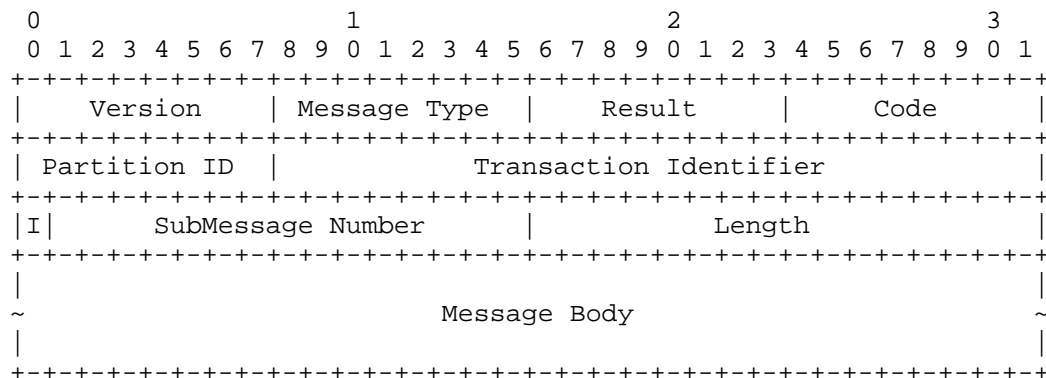
Each switch port also maintains a port session number assigned by the switch. A message, with an incorrect port session number **MUST** be rejected. This allows the controller to detect a link failure and to keep states synchronised.

Except for the adjacency protocol message, no GSMP messages may be sent across the link until the adjacency protocol has achieved synchronisation, and all GSMP messages received on a link that do not currently have state synchronisation **MUST** be discarded.

3.1 GSMP Packet Format

3.1.1 Basic GSMP Message format

All GSMP messages, except the adjacency protocol message, have the following format:



(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 [1] and AAL-5 [2][3].)

Version

The version number of the GSMP protocol being used in this session. It SHOULD be set by the sender of the message to the GSMP protocol version negotiated by the adjacency protocol.

Message Type

The GSMP message type. GSMP messages fall into the following classes: Connection Management, Reservation Management, Port Management, State and Statistics, Configuration, Quality of Service, Events and messages belonging to an Abstract or Resource Model (ARM) extension. Each class has a number of different message types. In addition, one Message Type is allocated to the adjacency protocol.

Result

Field in a Connection Management request message, a Port Management request message, or a Quality of Service request message that is used to indicate whether a response is required to the request message if the outcome is successful. A value of "NoSuccessAck" indicates that the request message does not expect a response if the outcome is successful, and a value of "AckAll" indicates that a response is expected if the outcome is successful. In both cases a failure response MUST be generated if the request fails. For State and Statistics, and Configuration request messages, a value of "NoSuccessAck" in the request message is ignored and the request message is handled as if the field was set to "AckAll". (This facility was added to reduce the control traffic in the case where the

controller periodically checks that the state in the switch is correct. If the controller does not use this capability, all request messages SHOULD be sent with a value of "AckAll".)

In a response message, the result field can have three values: "Success," "More," and "Failure". The "Success" and "More" results both indicate a success response. All messages that belong to the same success response will have the same Transaction Identifier. The "Success" result indicates a success response that may be contained in a single message or the final message of a success response spanning multiple messages.

"More" in the result indicates that the message, either request or response, exceeds the maximum transmission unit of the data link and that one or more further messages will be sent to complete the success response.

ReturnReceipt is a result field used in Events to indicate that an acknowledgement is required for the message. The default for Events Messages is that the controller will not acknowledge Events. In the case where a switch requires acknowledgement, it will set the Result Field to ReturnReceipt in the header of the Event Message.

The encoding of the result field is:

NoSuccessAck:	Result = 1
AckAll:	Result = 2
Success:	Result = 3
Failure:	Result = 4
More:	Result = 5
ReturnReceipt	Result = 6

The Result field is not used in an adjacency protocol message.

Code

Field gives further information concerning the result in a response message. It is mostly used to pass an error code in a failure response but can also be used to give further information in a success response message or an event message. In a request message, the code field is not used and is set to zero. In an adjacency protocol message, the Code field is used to determine the function of the message.

Partition ID

Field used to associate the command with a specific switch partition. The format of the Partition ID is not defined in GSMP. If desired, the Partition ID can be divided into multiple sub-identifiers within a single partition. For example: the Partition ID could be subdivided into a 6-bit partition number and a 2-bit sub-identifier which would allow a switch to support 64 partitions with 4 available IDs per partition.

Transaction Identifier

Used to associate a request message with its response message. For request messages, the controller may select any transaction identifier. For response messages, the transaction identifier is set to the value of the transaction identifier from the message to which it is a response. For event messages, the transaction identifier SHOULD be set to zero. The Transaction Identifier is not used, and the field is not present, in the adjacency protocol.

I flag

If I is set then the SubMessage Number field indicates the total number of SubMessage segments that compose the entire message. If it is not set then the SubMessage Number field indicates the sequence number of this SubMessage segment within the whole message.

SubMessage Number

When a message is segmented because it exceeds the MTU of the link layer, each segment will include a submessage number to indicate its position. Alternatively, if it is the first submessage in a sequence of submessages, the I flag will be set and this field will contain the total count of submessage segments.

Length

Length of the GSMP message including its header fields and defined GSMP message body. The length of additional data appended to the end of the standard message SHOULD be included in the Length field.

3.1.2 Fields commonly found in GSMP messages

The following fields are frequently found in GSMP messages. They are defined here to avoid repetition.

Port

Gives the port number of the switch port to which the message applies.

Port Session Number

Each switch port maintains a Port Session Number assigned by the switch. The port session number of a port remains unchanged while the port is continuously in the Available state and the link status is continuously Up. When a port returns to the Available state after it has been Unavailable or in any of the Loopback states, or when the line status returns to the Up state after it has been Down or in Test, or after a power cycle, a new Port Session Number MUST be generated. Port session numbers SHOULD be assigned using some form of random number.

If the Port Session Number in a request message does not match the current Port Session Number for the specified port, a failure response message MUST be returned with the Code field indicating, "5: Invalid port session number". The current port session number for a port may be obtained using a Port Configuration or an All Ports Configuration message.

3.1.2.1 Additional General Message Information

1. Any field in a GSMP message that is unused or defined as "reserved" MUST be set to zero by the sender and ignored by the receiver.
2. Flags that are undefined will be designated as: x: reserved
3. It is not an error for a GSMP message to contain additional data after the end of the Message Body. This is allowed to support proprietary and experimental purposes. However, the maximum transmission unit of the GSMP message, as defined by the data link layer encapsulation, MUST NOT be exceeded. The length of additional data appended to the end of the standard message SHOULD be included in the message length field.
4. A success response message MUST NOT be sent until the requested operation has been successfully completed.


3.1.3 Labels

All labels in GSMP have a common structure composed of tuples, consisting of a Type, a Length, and a Value. Such tuples are commonly known as TLV's, and are a good way of encoding information in a flexible and extensible format. A label TLV is encoded as a 2 octet field that uses 12 bits to specify a Type and four bits to specify certain behaviour specified below, followed by a 2 octet Length field, followed by a variable length Value field. Additionally, a label field can be composed of many stacked labels that together constitute the label.

A summary of TLV labels supported in this version of the protocol is listed below:

TLV Label	Type	Section Title
-----	----	-----
ATM Label	0x100	ATM TLV Labels
FR Label	0x101	Frame Relay TLV Labels
MPLS Gen Label	0x102	MPLS Generic TLV Labels
FEC Label	0x103	FEC TLV Labels

All Labels will be designated as follow:



x: Reserved Flags.
These are generally used by specific messages and will be defined in those messages.

S: Stacked Label Indicator
Label Stacking is discussed below in section 3.1.3.5

Label Type
A 12-bit field indicating the type of label.

Label Length
A 16-bit field indicating the length of the Label Value field in bytes.

Label Value

A variable length field that is an integer number of 32 bit words long. The Label Value field is interpreted according to the Label Type as described in the following sections.

3.1.3.1 ATM Labels

If the Label Type = ATM Label, the labels MUST be interpreted as an ATM labels as shown:

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9																								
x S x x				ATM Label (0x100)												Label Length																																															
x x x x				VPI												VCI																																															

For a virtual path connection (switched as a single virtual path connection) or a virtual path (switched as one or more virtual channel connections within the virtual path) the VCI field is not used.

ATM distinguishes between virtual path connections and virtual channel connections. The connection management messages apply both to virtual channel connections and virtual path connections. The Add Branch and Move Branch connection management messages have two Message Types. One Message Type indicates that a virtual channel connection is required, and the other Message Type indicates that a virtual path connection is required. The Delete Branches, Delete Tree, and Delete All connection management messages have only a single Message Type because they do not need to distinguish between virtual channel connections and virtual path connections. For virtual path connections, neither Input VCI fields nor Output VCI fields are required. They SHOULD be set to zero by the sender and ignored by the receiver. Virtual channel branches may not be added to an existing virtual path connection. Conversely, virtual path branches may not be added to an existing virtual channel connection. In the Port Configuration message each switch input port may declare whether it is capable of supporting virtual path switching (i.e., accepting connection management messages requesting virtual path connections).

3.1.3.2 Frame Relay Labels

If the TLV Type = FR Label, the labels MUST be interpreted as a Frame Relay labels as shown:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|S|x|x|   FR Label (0x101)   |           Label Length           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x x x x| Res |Len|           DLCI           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Res

The Res field is reserved in [21], i.e., it is not explicitly reserved by GSMP.

Len

The Len field specifies the number of bits of the DLCI. The following values are supported:

Len	DLCI bits
0	10
2	23

DLCI

DLCI is the binary value of the Frame Relay Label. The significant number of bits (10 or 23) of the label value is to be encoded into the Data Link Connection Identifier (DLCI) field when part of the Frame Relay data link header [13].

3.1.3.3 MPLS Generic Labels

If a port's attribute PortType=MPLS, then that port's labels are for use on links for which label values are independent of the underlying link technology. Examples of such links are PPP and Ethernet. On such links the labels are carried in MPLS label stacks [14]. If the Label Type = MPLS Generic Label, the labels MUST be interpreted as Generic MPLS labels as shown:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|S|x|x| MPLS Gen Label (0x102)|           Label Length           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x x x x x x x x x x x x|           MPLS Label           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

MPLS Label

This is a 20-bit label value as specified in [14], represented as a 20-bit number in a 4-byte field.

3.1.3.4 FEC Labels

Labels may be bound to Forwarding Equivalence Classes (FECs) as defined in [18]. A FEC is a list of one or more FEC elements. The FEC TLV encodes FEC items. In this version of the protocol only, Prefix FECs are supported. If the Label Type = FEC Label, the labels MUST be interpreted as Forwarding Equivalence Class Labels as shown:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|x|S|x|x|   FEC Label (0x103)   |   Label Length   |
+-----+-----+-----+-----+
~                               FEC Element 1       ~
+-----+-----+-----+-----+
~                               FEC Element n       ~
+-----+-----+-----+-----+

```

FEC Element

The FEC element encoding depends on the type of FEC element. In this version of GSMP only, Prefix FECs are supported.

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Element Type |   Address Family   | Prefix Length |
+-----+-----+-----+-----+
~                               Prefix                               ~
+-----+-----+-----+-----+

```

Element Type

In this version of GSMP the only supported Element Type is Prefix FEC Elements. The Prefix FEC Element is a one-octet value, encoded as 0x02.

Address Family

Two-byte quantity containing a value from ADDRESS FAMILY NUMBERS in [5], that encodes the address family for the address prefix in the Prefix field.

Prefix Length

One byte containing the length in bits of the address prefix that follows. A length of zero indicates a prefix that matches all addresses (the default destination); in this case the Prefix itself is zero bytes.

Prefix

An address prefix encoded according to the Address Family field, whose length, in bits, was specified in the Prefix Length field.

3.1.3.5 Label Stacking

Label stacking is a technique used in MPLS [14] that allows hierarchical labelling. MPLS label stacking is similar to, but subtly different from, the VPI/VCI hierarchy of labels in ATM. There is no set limit to the depth of label stacks that can be used in GSMP.

When the Stacked Label Indicator S is set to 1 it indicates that an additional label field will be appended to the adjacent label field. For example, a stacked Input Short Label could be designated as follows:

```

+-----+
|x|S|x|x|                                     |
+-----+                                     |
~                                             ~
+-----+                                     |
** |x|S|x|x|                                     |
+-----+                                     |
~                                             ~
+-----+

```

Input Label

Stacked Input Label

** Note: There can be zero or more Stacked Labels fields (like those marked **) following an Input or Output Label field. A Stacked Label follows the previous label field if and only if the S Flag in the previous label is set.

When a label is extended by stacking, it is treated by the protocol as a single extended label, and all operations on that label are atomic. For example, in an add branch message, the entire input label is switched for the entire output label. Likewise, in Move Input Branch and Move Output Branch messages, the entire label is swapped. For that reason, in all messages that designate a label field, it will be depicted as a single 64-bit field, though it might be instantiated by many 64-bit fields in practice.

3.1.4 Failure Response Messages

A failure response message is formed by returning the request message that caused the failure with the Result field in the header indicating failure (Result = 4) and the Code field giving the failure code. The failure code specifies the reason for the switch being unable to satisfy the request message.

If the switch issues a failure response in reply to a request message, no change should be made to the state of the switch as a result of the message causing the failure. (For request messages that contain multiple requests, such as the Delete Branches message, the failure response message will specify which requests were successful and which failed. The successful requests may result in changed state.)

A warning response message is a success response (Result = 3) with the Code field specifying the warning code. The warning code specifies a warning that was generated during the successful operation.

If the switch issues a failure response it MUST choose the most specific failure code according to the following precedence:

- Invalid Message
- General Message Failure
- Specific Message Failure
A failure response specified in the text defining the message type.
- Connection Failures
- Virtual Path Connection Failures
- Multicast Failures
- QoS Failures
- General Failures
- Warnings

If multiple failures match in any of the categories, the one that is listed first should be returned. Descriptions of the Failure response messages can be found in section 12.

4. Connection Management Messages

4.1 General Message Definitions

Connection management messages are used by the controller to establish, delete, modify and verify connections across the switch. The Add Branch, Delete Tree, and Delete All connection management messages have the following format, for both request and response messages:

0										1										2										3																																																											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																																										
										Version																				Message Type																				Result																				Code																			
										Partition ID																				Transaction Identifier																																																											
I										SubMessage Number																				Length																																																											
										Port Session Number																																																																															
										Reservation ID																																																																															
										Input Port																																																																															
										Input Service Selector																																																																															
										Output Port																																																																															
										Output Service Selector																																																																															
IQS OQS P x N O										Adaptation Method																																																																															
x S x x										Input Label																																																																															
~																				~																																																																					
x S x x										Output Label																																																																															
~																				~																																																																					

If a valid Reservation ID is specified and the Service Model is used (i.e., IQS or OQS=0b10) then the Traffic Parameters Block may be omitted from the Add Branch message indicating that the Traffic Parameters specified in the corresponding Reservation Request message are to be used.

Input Port

Identifies a switch input port.

Input Label

Identifies an incoming labelled channel arriving at the switch input port indicated by the Input Port field. The value in the Input Label field MUST be interpreted according to the Label Type attribute of the switch input port indicated by the Input Port field.

Input Service Selector

Identifies details of the service specification being used for the connection. The interpretation depends upon the Input QoS Model Selector (IQS).

IQS = 00: In this case, the Input Service Selector indicates a simple priority.

IQS = 01: In this case, the Input Service Selector is an opaque service profile identifier. The definition of these service profiles is outside the scope of this specification. Service Profiles can be used to indicate pre-defined Differentiated Service Per Hop Behaviours.

IQS = 10: In this case, the Input Service Selector corresponds to a Service Spec as defined in Chapter 8.2. When the value of either IQS or OQS is set to 0b10, then a Traffic Parameters Block is appended to the message.

IQS = 11: In this case the Input Service Selector corresponds to an ARM service specification. Definition of ARM service specifications is outside the scope of this specification and is determined by the MType as defined in Chapter 8.1.

Output Port

Identifies a switch output port.

Output Label

Identifies an outgoing labelled channel departing at the switch output port indicated by the Output Port field. The value in the Output Label field MUST be interpreted according to the Label Type attribute of the switch input port indicated by the Output Port field

Output Service Selector

Identifies details of the service model being used. The interpretation depends upon the Output QoS Model selector (OQS).

OQS = 00: In this case the Output Service Selector indicates a simple priority.

OQS = 01: In this case the Output Service Selector is an opaque service profile identifier. The definition of these service profiles is outside the scope of this specification. Service Profiles can be used to indicate pre-defined Differentiated Service Per Hop Behaviours.

OQS = 10: In this case the Output Service Selector corresponds to a Service Spec as defined in Chapter 8.2. When the value of either IQS or OQS is set to 0b10 then a Traffic Parameters Block is appended to the message.

OQS = 11: In this case the Output Service Selector corresponds to an ARM service specification. Definition of ARM service specifications is outside the scope of this specification and is determined by the MType as defined in Chapter 8.1.

IQS, OQS

Input and Output QoS Model Selector:

The QoS Model Selector is used to specify a QoS Model for the connection. The values of IQS and OQS determine respectively the interpretation of the Input Service Selector and the Output Service Selector, and SHOULD be interpreted as a priority, a QoS profile, a service specification, or an ARM specification as shown:

IQS/OQS	QoS Model	Service Selector
-----	-----	-----
00	Simple Abstract	Model Priority
01	QoS Profile Model	QoS Profile
10	Default Service Model	Service Specification
11	Optional ARM	ARM Specification

P Flag

If the Parameter flag is set it indicates that a single instance of the Traffic Parameter block is provided. This occurs in cases where the Input Traffic Parameters are identical to Output Traffic Parameters.

N Flag

The Null flag is used to indicate a null adaptation method. This occurs when the branch is connecting two ports of the same type.

O Flag

The Opaque flag indicates whether the adaptation fields are opaque, or whether they are defined by the protocol. See the definition of Adaptation Method below for further information.

Adaptation Method

The adaptation method is used to define the adaptation framing that may be in use when moving traffic from one port type to another port type; e.g., from a frame relay port to an ATM port. The content of this field is defined by the Opaque flag. If the Opaque flag is set, then this field is defined by the switch manufacturer and is not defined in this protocol. If the opaque flag is not set, the field is divided into two 12-bit fields as follows:

```

+++++
|IQS|OQS|P|x|N|O|      Input Adaptation      |      Output Adaptation      |
+++++

```

Input Adaptation

Adaptation framing method used on incoming connections.

Output Adaptation

Adaptation framing method used on outgoing connections.

Adaptation Types:

0x100	PPP
0x200	FRF.5
0x201	FRF.8

Input and Output TC Flags

TC (Traffic Control) Flags are used in Add Branch, Move Input Branch and Move Output Branch messages for connections using the Service Model (i.e., when IQS or OQS=0b10). The TC Flags field is defined in Section 10.6.

Input and Output Traffic Parameters Block

This variable length field is used in Add Branch, Move Input Branch and Move Output Branch messages for connections using the Service Model (i.e., when IQS or OQS=0b10). Traffic Parameters Block is defined in Section 10.5. The Traffic Parameters Block may be omitted if a valid, non-zero Reservation ID is specified, in which case the Traffic Parameters of the corresponding Reservation Request message are used. If the P flag is set, then the appended message block will only include a single traffic parameter block which will be used for both input and output traffic.

For all connection management messages, except the Delete Branches message, the success response message is a copy of the request message returned with the Result field indicating success. The Code field is not used in a connection management success response message.

The failure response message is a copy of the request message returned with a Result field indicating failure.

Fundamentally, no distinction is made between point-to-point and point-to-multipoint connections. By default, the first Add Branch message for a particular Input Port and Input Label will establish a point-to-point connection. The second Add Branch message with the same Input Port and Input Label fields will convert the connection to a point-to-multipoint connection with two branches. However, to avoid possible inefficiency with some switch designs, the Multicast Flag is provided. If the controller knows that a new connection is point-to-multipoint when establishing the first branch, it may indicate this in the Multicast Flag. Subsequent Add Branch messages with the same Input Port and Input Label fields will add further branches to the point-to-multipoint connection. Use of the Delete Branch message on a point-to-multipoint connection with two branches will result in a point-to-point connection. However, the switch may structure this connection as a point-to-multipoint connection with a single output branch if it chooses. (For some switch designs this structure may be more convenient.) Use of the Delete Branch message on a point-to-point connection will delete the point-to-point connection. There is no concept of a connection with zero output branches. All connections are unidirectional, one input labelled channel to one or more output labelled channels.

In GSMP a multipoint-to-point connection is specified by establishing multiple point-to-point connections, each of them specifying the same output branch. (An output branch is specified by an output port and output label.)

The connection management messages may be issued regardless of the Port Status of the switch port. Connections may be established or deleted when a switch port is in the Available, Unavailable, or any of the Loopback states. However, all connection states on an input port will be deleted when the port returns to the Available state from any other state, i.e., when a Port Management message is received for that port with the Function field indicating either Bring Up, or Reset Input Port.

When the value of either IQS or OQS is set to 0b10 then the following Traffic Parameters Block is appended to the above message:

```

+++++
|Input TC Flags |x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|
+++++
|
|           Input Traffic Parameters Block           |
|
+++++
|Output TC Flags|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|
+++++
|
|           Output Traffic Parameters Block           |
|
+++++

```

Note: Fields and Parameters that have been explained in the description of the general connection message will not be explained in this section. Please refer to section 4.1 for details.

M: Multicast

Multicast flags are used as a hint for point-to-multipoint or multipoint-to-point connections in the Add Branch message. They are not used in any other connection management messages and in these messages they SHOULD be set to zero. There are two instances of the M-bit in the Add Branch message; one for input branch specified by the Input Port and Input Label fields and one for the output branch specified by the Output Port and Output Label fields. If set for the input branch (in front of Input Label field), it indicates that the connection is very likely to be a point-to-multipoint connection. If zero, it indicates that this connection is very likely to be a point-to-point connection or is unknown. If set for the output branch (in front of the Output Label field), it indicates that the connection is very likely to be a multipoint-to-point connection. If zero, it indicates that this connection is very likely to be a point-to-point connection or is unknown.

If M flags are set for input as well as output branches, it indicates that the connection is very likely to be a multipoint-to-multipoint connection.

The Multicast flags are only used in the Add Branch message when establishing the first branch of a new connection. It is not required to be set when establishing subsequent branches of a point-to-multipoint or a multipoint-to-point connection and

on such connections it SHOULD be ignored by the receiver. (Except in cases where the connection replace bit is enabled and set, the receipt of the second and subsequent Add Branch messages from the receiver indicates a point-to-multipoint or a multipoint-to-point connection.) If it is known that this is the first branch of a point-to-multipoint or a multipoint-to-point connection, this flag SHOULD be set. If it is unknown, or if it is known that the connection is point-to-point, this flag SHOULD be zero. The use of the multicast flag is not mandatory and may be ignored by the switch. If unused, the flags SHOULD be set to zero. Some switches use a different data structure for multicast connections rather than for point-to-point connections. These flags prevent the switch from setting up a point-to-point structure for the first branch of a multicast connection that MUST immediately be deleted and reconfigured as point-to-multipoint or multipoint-to-point when the second branch is established.

B: Bi-directional

The Bi-directional flag applies only to the Add Branch message. In all other Connection Management messages it is not used. It may only be used when establishing a point-to-point connection. The Bi-directional flag in an Add Branch message, if set, requests that two unidirectional connections be established, one in the forward direction, and one in the reverse direction. It is equivalent to two Add Branch messages, one specifying the forward direction, and one specifying the reverse direction. The forward direction uses the values of Input Port, Input Label, Output Port and Output Label as specified in the Add Branch message. The reverse direction is derived by exchanging the values specified in the Input Port and Input Label fields, with those of the Output Port and Output Label fields respectively. Thus, a connection in the reverse direction originates at the input port specified by the Output Port field, on the label specified by the Output Label field. It departs from the output port specified by the Input Port field, on the label specified by the Input Label field.

The Bi-directional flag is simply a convenience to establish two unidirectional connections in opposite directions between the same two ports, with identical Labels, using a single Add Branch message. In all future messages the two unidirectional connections MUST be handled separately. There is no bi-directional delete message. However, a single Delete Branches message with two Delete Branch Elements, one for the forward connection and one for the reverse, may be used.

R: Connection Replace

The Connection Replace flag applies only to the Add Branch message and is not used in any other Connection Management messages. The R flag is used in cases when creation of multipoint-to-point connections is undesirable (e.g., POTS applications where fan-in is meaningless). If the R flag is set, the new connection replaces any existing connection if the label is already in use at the same Output Port.

The Connection Replace mechanism allows a single Add Connection command to function as either a Move Branch message or a combination of Delete Branch/Add Branch messages. This mechanism is provided to support existing 64k call handling applications, such as emulating 64k voice switches.

The use of R flag is optional and MUST be pre-configured in the Port Management message [see section 6.1] to activate its use. The R flag MUST NOT be set if it is not pre-configured with the Port Management message. The switch MUST then return a Failure Response message: "36: Replace of connection is not activated on switch". Information about whether the function is active or not, can be obtained by using the Port Configuration message [see section 8.2].

The R flag MUST NOT be set if either the M flag or the B flag is set. If a switch receives an Add connection request that has the R flag set with either the B or the M flag set, it MUST return a failure response message of: "37: Connection replacement mode cannot be combined with Bi-directional or Multicast mode"

If the connection specified by the Input Port and Input Label fields does not already exist, it MUST be established with the single output branch specified in the request message. If the Bi-directional Flag in the Flags field is set, the reverse connection MUST also be established. The output branch SHOULD have the QoS attributes specified by the Class of Service field.

If the connection specified by the Input Port and Input Label fields already exists and the R flag is not set, but the specified output branch does not, the new output branch MUST be added. The new output branch SHOULD have the QoS attributes specified by the Class of Service field.

If the connection specified by the Input Port and Input Label fields already exists and the specified output branch also already exists, the QoS attributes of the connection, specified by the Class of Service field, if different from the request message, SHOULD be

changed to that in the request message. A success response message MUST be sent if the Result field of the request message is "AckAll". This allows the controller to periodically reassert the state of a connection or to change its priority. If the result field of the request message is "NoSuccessAck" a success response message SHOULD NOT be returned. This may be used to reduce the traffic on the control link for messages that are reasserting a previously established state. For messages that are reasserting a previously established state, the switch MUST always check that this state is correctly established in the switch hardware (i.e., the actual connection tables used to forward cells or frames).

If the connection specified by the Input Port and Input Label fields already exists, and the Bi-directional Flag in the Flags field is set, a failure response MUST be returned indicating: "15: Point-to-point bi-directional connection already exists".

It should be noted that different switches support multicast in different ways. There may be a limit to the total number of point-to-multipoint or multipoint-to-point connections certain switches can support, and possibly a limit on the maximum number of branches that a point-to-multipoint or multipoint-to-point connection may specify. Some switches also impose a limit on the number of different Label values that may be assigned e.g., to the output branches of a point-to-multipoint connection. Many switches are incapable of supporting more than a single branch of any particular point-to-multipoint connection on the same output port. Specific failure codes are defined for some of these conditions.

4.2.1 ATM specific procedures:

To request an ATM virtual path connection the ATM Virtual Path Connection (VPC) Add Branch message is:

Message Type = 26

An ATM virtual path connection can only be established between ATM ports, i.e., ports with the "ATM" Label Type attribute. If an ATM VPC Add Branch message is received and either the switch input port specified by the Input Port field or the switch output port specified by the Output Port field is not an ATM port, a failure response message MUST be returned indicating, "28: ATM Virtual path switching is not supported on non-ATM ports".

If an ATM VPC Add Branch message is received and the switch input port specified by the Input Port field does not support virtual path switching, a failure response message MUST be returned indicating, "24: ATM virtual path switching is not supported on this input port".

If an ATM virtual path connection already exists on the virtual path specified by the Input Port and Input VPI fields, a failure response message MUST be returned, indicating "27: Attempt to add an ATM virtual channel connection branch to an existing virtual path connection". For the VPC Add Branch message, if a virtual channel connection already exists on any of the virtual channels within the virtual path specified by the Input Port and Input VPI fields, a failure response message MUST be returned indicating, "26: Attempt to add an ATM virtual path connection branch to an existing virtual channel connection".

4.3 Delete Tree Message

The Delete Tree message is a Connection Management message used to delete an entire connection. All remaining branches of the connection are deleted. A connection is defined by the Input Port and the Input Label fields. The Output Port and Output Label fields are not used in this message. The Delete Tree message is:

Message Type = 18

If the Result field of the request message is "AckAll" a success response message MUST be sent upon successful deletion of the specified connection. The success message MUST NOT be sent until the delete operation has been completed and if possible, not until all data on the connection, queued for transmission, has been transmitted.

4.4 Verify Tree Message

The Verify Tree message has been removed from this version of GSMP.

Message Type = 19

If a request message is received with Message Type = 19, a failure response MUST be returned with the Code field indicating:

"3: The specified request is not implemented on this switch.".

4.5 Delete All Input Port Message

The Delete All Input Port message is a connection management message used to delete all connections on a switch input port. All connections that originate at the specified input port MUST be deleted. On completion of the operation all dynamically assigned Label values for the specified port MUST be unassigned, i.e., there MUST be no connections established in the Label space that GSMP controls on this port. The Service Selectors, Output Port, Input

Label and Output Label fields are not used in this message. The Delete All Input Port message is:

Message Type = 20

If the Result field of the request message is "AckAll", a success response message MUST be sent upon completion of the operation. The success response message MUST NOT be sent until the operation has been completed.

The following failure response messages may be returned to a Delete All Input Port request.

3: The specified request is not implemented on this switch.

4: One or more of the specified ports does not exist.

5: Invalid Port Session Number.

If any field in a Delete All Input Port message not covered by the above failure codes is invalid, a failure response MUST be returned indicating: "2: Invalid request message". Else, the Delete All Input Port operation MUST be completed successfully and a success message returned. No other failure messages are permitted.

4.6 Delete All Output Port Message

The Delete All message is a connection management message used to delete all connections on a switch output port. All connections that have the specified output port MUST be deleted. On completion of the operation all dynamically assigned Label values for the specified port MUST be unassigned, i.e., there MUST be no connections established in the Label space that GSMP controls on this port. The Service Selectors, Input Port, Input Label and Output Label fields are not used in this message. The Delete All Output Port message is:

Message Type = 21

If the Result field of the request message is "AckAll", a success response message MUST be sent upon completion of the operation. The success response message MUST NOT be sent until the operation has been completed.

The following failure response messages may be returned to a Delete All Output Port request.

In each Delete Branch Element, a connection is specified by the Input Port and Input Label fields. The specific branch to be deleted is indicated by the Output Port and Output Label fields.

If the Result field of the Delete Branches request message is "AckAll" a success response message MUST be sent upon successful deletion of the branches specified by all of the Delete Branch Elements. The success response message MUST NOT be sent until all of the delete branch operations have been completed. The success response message is only sent if all of the requested delete branch operations were successful. No Delete Branch Elements are returned in a Delete Branches success response message and the Number of Elements field MUST be set to zero.

If there is a failure in any of the Delete Branch Elements, a Delete Branches failure response message MUST be returned. The Delete Branches failure response message is a copy of the request message with the Code field of the entire message set to "10: General Message Failure" and the Error field of each Delete Branch Element indicating the result of each requested delete operation. A failure in any of the Delete Branch Elements MUST NOT interfere with the processing of any other Delete Branch Elements.

4.8 Move Output Branch Message

The Move Output Branch message is used to move a branch of an existing connection from its current output port label to a new output port label in a single atomic transaction. The Move Output Branch connection management message has the following format for both request and response messages:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number												Length																										
Port Session Number																																							
Input Port																																							
Input Service Selector																																							
Old Output Port																																							
New Output Port																																							
Output Service Selector																																							
I	Q	S		O	Q	S		P		x		N		O		Adaptation Method																							
x		S		x		x		Input Label																															
~																																							
x		S		x		x		Old Output Label																															
~																																							
x		S		x		x		New Output Label																															
~																																							

When the value of either IQS or OQS is set to 0b10 then the following Traffic Parameters Block is appended to the above message:

```

+++++
|Input TC Flags |x x x x x x x x x x x x x x x x x x x x x|
+++++
|
|           Input Traffic Parameters Block           |
|~
|
+++++
|Output TC Flags|x x x x x x x x x x x x x x x x x x x x x|
+++++
|
|           Output Traffic Parameters Block          |
|~
|
+++++

```

Note: Fields and Parameters that have been explained in the description of the general connection message will not be explained in this section. Please refer to section 4.1 for details.

The Move Output Branch message is a connection management message used to move a single output branch of connection from its current output port and Output Label, to a new output port and Output Label on the same connection. None of the connection's other output branches are modified. When the operation is complete the original Output Label on the original output port will be deleted from the connection.

The Move Output Branch message is:

Message Type = 22

For the Move Output Branch message, if the connection specified by the Input Port and Input Label fields already exists, and the output branch specified by the Old Output Port and Old Output Label fields exists as a branch on that connection, the output branch specified by the New Output Port and New Output Label fields is added to the connection and the branch specified by the Old Output Port and Old Output Label fields is deleted. If the Result field of the request message is "AckAll", a success response message MUST be sent upon successful completion of the operation. The success response message MUST NOT be sent until the Move Branch operation has been completed.

For the Move Output Branch message, if the connection specified by the Input Port and Input Label fields already exists, but the output branch specified by the Old Output Port and Old Output Label fields

does not exist as a branch on that connection, a failure response MUST be returned with the Code field indicating, "12: The specified branch does not exist".

4.8.1 ATM Specific Procedures:

The ATM VPC Move Output Branch message is a connection management message used to move a single output branch of a virtual path connection from its current output port and output VPI, to a new output port and output VPI on the same virtual channel connection. None of the other output branches are modified. When the operation is complete the original output VPI on the original output port will be deleted from the connection.

The VPC Move Branch message is:

Message Type = 27

For the VPC Move Output Branch message, if the virtual path connection specified by the Input Port and Input VPI fields already exists, and the output branch specified by the Old Output Port and Old Output VPI fields exists as a branch on that connection, the output branch specified by the New Output Port and New Output VPI fields is added to the connection and the branch specified by the Old Output Port and Old Output VPI fields is deleted. If the Result field of the request message is "AckAll", a success response message MUST be sent upon successful completion of the operation. The success response message MUST NOT be sent until the Move Branch operation has been completed.

For the VPC Move Output Branch message, if the virtual path connection specified by the Input Port and Input VPI fields already exists, but the output branch specified by the Old Output Port and Old Output VPI fields does not exist as a branch on that connection, a failure response MUST be returned with the Code field indicating, "12: The specified branch does not exist".

If the virtual channel connection specified by the Input Port and Input Label fields; or the virtual path connection specified by the Input Port and Input VPI fields; does not exist, a failure response MUST be returned with the Code field indicating, "11: The specified connection does not exist".

If the output branch specified by the New Output Port, New Output VPI, and New Output VCI fields for a virtual channel connection; or the output branch specified by the New Output Port and New Output VPI fields for a virtual path connection; is already in use by any connection other than that specified by the Input Port and Input

Label fields, then the resulting output branch will have multiple input branches. If multiple point-to-point connections share the same output branch, the result will be a multipoint-to-point connection. If multiple point-to-multipoint trees share the same output branches, the result will be a multipoint-to-multipoint connection.

4.9 Move Input Branch Message

The Move Input Branch message is used to move a branch of an existing connection from its current input port label to a new input port label in a single atomic transaction. The Move Input Branch connection management message has the following format for both request and response messages:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number																			Length																			
Port Session Number																																							
Output Port																																							
Input Service Selector																																							
Old Input Port																																							
New Input Port																																							
Output Service Selector																																							
I	Q	S	P	x	N	O	Adaptation Method																																
x	S	x	x	Output Label																																			
~																																							
x	S	x	x	Old Input Label																																			
~																																							
x	S	x	x	New Input Label																																			
~																																							

When the value of either IQS or OQS is set to 0b10, then the following Traffic Parameters Block is appended to the above message:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|Input TC Flags |x x x x x x x x x x x x x x x x x x x x x x|
+-----+-----+-----+-----+-----+-----+-----+-----+
|
|               Input Traffic Parameters Block               |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
|Output TC Flags|x x x x x x x x x x x x x x x x x x x x x x|
+-----+-----+-----+-----+-----+-----+-----+-----+
|
|               Output Traffic Parameters Block               |
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Note: Fields and Parameters that have been explained in the description of the general connection message will not be explained in this section. Please refer to section 4.1 for details.

The Move Input Branch message is a connection management message used to move a single input branch of connection from its current input port and Input Label, to a new input port and Input Label on the same connection. None of the connection's other input branches are modified. When the operation is complete, the original Input Label on the original input port will be deleted from the connection.

The Move Input Branch message is:

Message Type = 23

For the Move Input Branch message, if the connection specified by the Output Port and Output Label fields already exists, and the input branch specified by the Old Input Port and Old Input Label fields exists as a branch on that connection, the input branch specified by the New Input Port and New Input Label fields is added to the connection and the branch specified by the Old Input Port and Old Input Label fields is deleted. If the Result field of the request message is "AckAll", a success response message MUST be sent upon successful completion of the operation. The success response message MUST NOT be sent until the Move Input Branch operation has been completed.

For the Move Input Branch message, if the connection specified by the Output Port and Output Label fields already exists, but the input branch specified by the Old Input Port and Old Input Label fields does not exist as a branch on that connection, a failure response MUST be returned with the Code field indicating, "12: The specified branch does not exist".

4.9.1 ATM Specific Procedures:

The ATM VPC Move Input Branch message is a connection management message used to move a single input branch of a virtual path connection from its current input port and input VPI, to a new input port and input VPI on the same virtual channel connection. None of the other input branches are modified. When the operation is complete, the original input VPI on the original input port will be deleted from the connection.

The VPC Move Input Branch message is:

Message Type = 28

For the VPC Move Input Branch message, if the virtual path connection specified by the Output Port and Output VPI fields already exists, and the input branch specified by the Old Input Port and Old Input VPI fields exists as a branch on that connection, the input branch specified by the New Input Port and New Input VPI fields is added to the connection and the branch specified by the Old Input Port and Old Input VPI fields is deleted. If the Result field of the request message is "AckAll" a success response message MUST be sent upon successful completion of the operation. The success response message MUST NOT be sent until the Move Input Branch operation has been completed.

For the VPC Move Input Branch message, if the virtual path connection specified by the Output Port and Output VPI fields already exists, but the input branch specified by the Old Input Port and Old Input

VPI fields does not exist as a branch on that connection, a failure response MUST be returned with the Code field indicating, "12: The specified branch does not exist".

If the virtual channel connection specified by the Output Port and Output Label fields, or if the virtual path connection specified by the Output Port and Output VPI fields does not exist, a failure response MUST be returned with the Code field indicating, "11: The specified connection does not exist".

If the input branch specified by the New Input Port, New Input VPI, and New Input VCI fields for a virtual channel connection, or the input branch specified by the New Input Port and New Input VPI fields for a virtual path connection, is already in use by any connection other than that specified by the Output Port and Output Label fields, then the resulting input branch will have multiple output branches. If multiple point-to-point connections share the same input branch, the result will be a point-to-multipoint connection. If multiple multipoint-to-point trees share the same input branches, the result will be a multipoint-to-multipoint connection.

5. Reservation Management Messages

GSMP allows switch resources (e.g., bandwidth, buffers, queues, labels, etc.) to be reserved for connections before the connections themselves are established. This is achieved through the manipulation of Reservations in the switch.

Reservations are hard state objects in the switch that can be created by the controller by sending a Reservation Request message. Each Reservation is uniquely identified by an identifying number called a Reservation ID. Reservation objects can be deleted with the Delete Reservation message or the Delete All Reservations message. A reservation object is also deleted when the Reservation is deployed by specifying a Reservation ID in a valid Add Branch message.

The reserved resources MUST remain reserved until either the reservation is deployed, in which case the resources are applied to a branch, or the reservation is explicitly deleted (with a Delete Reservation message or a Delete All Reservations message), in which case the resources are freed. Reservations and reserved resources are deleted if the switch is reset.

A Reservation object includes its Reservation ID plus all the connection state associated with a branch with the exception that the branch's input label and/or output label may be unspecified. The Request Reservation message is therefore almost identical to the Add Branch message.

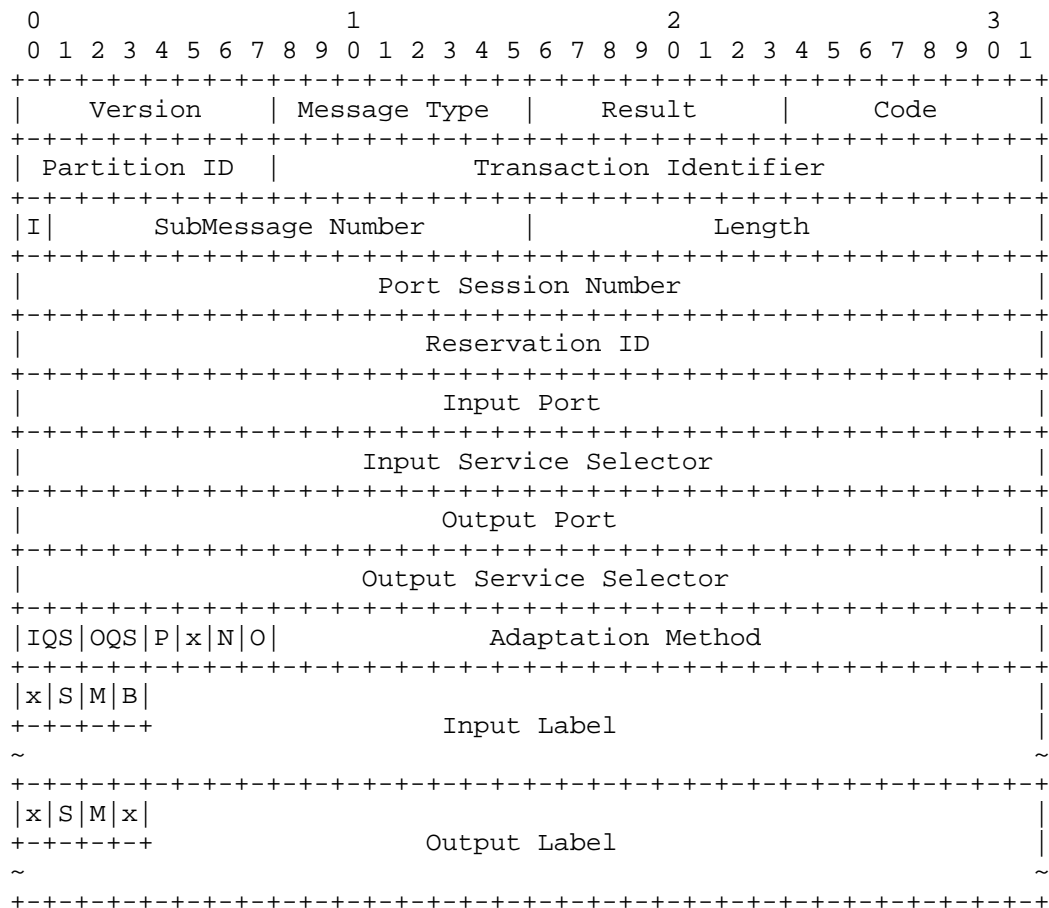
The switch establishes the maximum number of reservations it can store by setting the value of Max Reservations in the Switch Configuration response message. The switch indicates that it does not support reservations by setting Max Reservations to 0. The valid range of Reservation IDs is 1 to Max Reservations).

5.1 Reservation Request Message

The Reservation Request message creates a Reservation in the switch and reserves switch resources for a connection that may later be established using an Add Branch message. The Reservation Request Message is:

Message Type = 70

The Reservation Request message has the following format for the request message:



When the value of either IQS or OQS is set to 0b10 then the following Traffic Parameters Block is appended to the above message:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|Input TC Flags |x x x x x x x x x x x x x x x x x x x x x x x|
+-----+-----+-----+-----+-----+-----+-----+-----+
|
|               Input Traffic Parameters Block
|
+-----+-----+-----+-----+-----+-----+-----+-----+
|Output TC Flags|x x x x x x x x x x x x x x x x x x x x x x x|
+-----+-----+-----+-----+-----+-----+-----+-----+
|
|               Output Traffic Parameters Block
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Note: Fields and Parameters that have been explained in the description of the general connection message will not be explained in this section. Please refer to section 4.1 for details.

All the fields of the Reservation Request message have the same meanings as they do in the Add Branch message with the following exceptions:

Reservation ID

Specifies the Reservation ID of the Reservation. If the numerical value of the Reservation ID is greater than the value of the Max Reservations (from the Switch Configuration message), a failure response is returned indicating "20: the Reservation ID out of Range". If the value of Reservation ID matches that of an extant Reservation, a failure response is returned indicating "22: Reservation ID in use".

Input Label

If a specific input label is specified, then that label is reserved along with the required resources. If the Input Label is 0 then the switch reserves the resources, but will not bind them to a label until the add branch command is given, which references the Reservation Id. If the input label is 0, then all stacked labels MUST also be zeroed.

Output Label

If a specific Output Label is specified then that label is reserved along with the required resources. If the Output Label is 0 then the switch reserves the resources, but will not bind them to a label until the add branch command is given which references the Reservation Id. If the Output Label is 0, then all stacked labels MUST also be zeroed

When the switch receives a valid Reservation Request it reserves all the appropriate switch resources needed to establish a branch with corresponding attributes. If sufficient resources are not available, a failure response is returned indicating "18: Insufficient resources". Other failure responses are as defined for the Add Branch message.

5.2 Delete Reservation Message

The Delete Reservation message deletes a Reservation object in the switch and frees the reserved switch resources associated with that reservation object. The Reservation Request Message is:

Message Type = 71

The Delete Reservation message has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number																			Length																			
Port Session Number																																							
Reservation ID																																							

If the Reservation ID matches that of an extant Reservation then the reservation is deleted and corresponding switch resources are freed. If the numerical value of the Reservation ID is greater than the value of the Max Reservations (from the Switch Configuration message), a failure response is returned indicating "20: Reservation ID out of Range". If the value of Reservation ID does not match that of any extant Reservation, a failure response is returned indicating "23: Non-existent reservation ID".

5.3 Delete All Reservations Message

The Delete All Reservation message deletes all extant Reservation objects in the switch and frees the reserved switch resources of these reservations. The Reservation Request Message is:

Message Type = 72

The Delete All Reservation message has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number										Length																												

6. Management Messages

6.1 Port Management Message

The Port Management message allows a port to be brought into service, to be taken out of service, to be set to loop back, reset, or to change the transmit data rate. Only the Bring Up and the Reset Input Port functions change the connection state (established connections) on the input port. Only the Bring Up function changes the value of the Port Session Number. The Port Management message MAY also be used for enabling the replace connection mechanism. The Port Management message is also used as part of the Event Message flow control mechanism.

If the Result field of the request message is "AckAll", a success response message MUST be sent upon successful completion of the operation. The success response message MUST NOT be sent until the operation has been completed. The Port Management Message is:

Message Type = 32

The Port Management message has the following format for the request and success response messages:

0										1										2										3										
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1									
Version										Message Type										Result										Code										
Partition ID										Transaction Identifier																														
I	SubMessage Number															Length																								
Port																																								
Port Session Number																																								
Event Sequence Number																																								
R	x	x	x	x	x	x	x	x	x	Duration										Function																				
Event Flags															Flow Control Flags																									
Transmit Data Rate																																								

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Event Sequence Number

The success response message gives the current value of the Event Sequence Number of the switch port indicated by the Port field. The Event Sequence Number is set to zero when the port is initialised. It is incremented by one each time the port detects an asynchronous event that the switch would normally report via an Event message. If the Event Sequence Number in the success response differs from the Event Sequence Number of the most recent Event message received for that port, events have occurred that were not reported via an Event message. This is most likely to be due to the flow control that restricts the rate at which a switch can send Event messages for each port. In the request message this field is not used.

R: Connection Replace

The R flag shall only be checked when the Function field = 1 (Bring Up). If the R flag is set in the Port Management request message, it indicates that a switch controller requests the switch port to support the Connection Replace mechanism.

Connection Replace behaviour is described in chapter 4.2. If a switch does not support the Connection Replace mechanism, it MUST reply with the failure response: "45: Connection Replace mechanism not supported on switch" and reset the R-flag. Upon successful response, the R flag SHOULD remain set in the response message.

Duration

Is the length of time in seconds, that any of the loopback states remain in operation. When the duration has expired, the port will automatically be returned to service. If another Port Management message is received for the same port before the duration has expired, the loopback will continue to remain in operation for the length of time specified by the Duration field in the new message. The Duration field is only used in request messages with the Function field set to Internal Loopback, External Loopback, or Bothway Loopback.

Function

Specifies the action to be taken. The specified action will be taken regardless of the current status of the port (Available, Unavailable, or any Loopback state). If the specified function requires a new Port Session Number to be generated, the new Port Session Number MUST be returned in the success response message. The defined values of the Function field are:

Bring Up:

Function = 1. Bring the port into service. All connections that originate at the specified input port MUST be deleted and a new Port Session Number MUST be selected, preferably using some form of random number. On completion of the operation all dynamically assigned Label values for the specified input port MUST be unassigned, i.e., no connections will be established in the Label space that GSMP controls on this input port. Afterwards, the Port Status of the port will be Available.

Take Down:

Function = 2. Take the port out of service. Any data received at this port will be discarded. No data will be transmitted from this port. Afterwards, the Port Status of the port will be Unavailable.

The behaviour is undefined if the port is taken down over which the GSMP session that controls the switch is running. (In this case the most probable behaviour would be for the switch either to ignore the message or to terminate the current GSMP session and to initiate another session,

possibly with the backup controller, if any.) The correct method to reset the link over which GSMP is running is to issue an RSTACK message in the adjacency protocol.

Internal Loopback:

Function = 3. Data arriving at the output port from the switch fabric are looped through to the input port to return to the switch fabric. All of the functions of the input port above the physical layer, e.g., header translation, are performed upon the looped back data. Afterwards, the Port Status of the port will be Internal Loopback.

External Loopback:

Function = 4. Data arriving at the input port from the external communications link are immediately looped back to the communications link at the physical layer without entering the input port. None of the functions of the input port, above the physical layer are performed upon the looped back data. Afterwards, the Port Status of the port will be External Loopback.

Bothway Loopback:

Function = 5. Both internal and external loopbacks are performed. Afterwards, the Port Status of the port will be Bothway Loopback.

Reset Input Port:

Function = 6. All connections that originate at the specified input port MUST be deleted and the input and output port hardware re-initialised. On completion of the operation, all dynamically assigned Label values for the specified input port MUST be unassigned, i.e., no connections will be established in the Label space that GSMP controls on this input port. The range of labels that may be controlled by GSMP on this port will be set to the default values specified in the Port Configuration message. The transmit data rate of the output port MUST be set to its default value. The Port Session Number is not changed by the Reset Input Port function. Afterwards, the Port Status of the port will be Unavailable.

Reset Flags:

Function = 7. This function is used to reset the Event Flags and Flow Control Flags. For each bit that is set in the Event Flags field, the corresponding Event Flag in the switch port MUST be reset to 0. For each bit that is set in the Flow Control Flags field, the corresponding Flow Control Flag in the switch port MUST be toggled; i.e., flow control

for the corresponding event is turned off if it is currently on and it is turned on if it is currently off. The Port Status of the port is not changed by this function.

Set Transmit Data Rate:

Function = 8. Sets the transmit data rate of the output port as close as possible to the rate specified in the Transmit Data Rate field. In the success response message, the Transmit Data Rate MUST indicate the actual transmit data rate of the output port. If the transmit data rate of the requested output port cannot be changed a failure response MUST be returned with the Code field indicating: "43: The transmit data rate of this output port cannot be changed". If the transmit data rate of the requested output port can be changed, but the value of the Transmit Data Rate field is beyond the range of acceptable values, a failure response MUST be returned with the Code field indicating: "44: Requested transmit data rate out of range for this output port". In the failure response message, the Transmit Data Rate MUST contain the same value as contained in the request message that caused the failure. The transmit data rate of the output port is not changed by the Bring Up, Take Down, or any of the Loopback functions. It is returned to the default value by the Reset Input Port function.

Transmit Data Rate

This field is only used in request and success response messages with the Function field set to "Set Transmit Data Rate". It is used to set the output data rate of the output port. It is specified in cells/s and bytes/s. If the Transmit Data Rate field contains the value 0xFFFFFFFF the transmit data rate of the output port SHOULD be set to the highest valid value.

Event Flags

Field in the request message that is used to reset the Event Flags in the switch port indicated by the Port field. Each Event Flag in a switch port corresponds to a type of Event message. When a switch port sends an Event message, it sets the corresponding Event Flag on that port. Depending on the setting in the Flow Control Flag, a port is either subject to flow control or not. If it is subject to flow control, then it is not permitted to send another Event message of the same type before the Event Flag has been reset. To reset an event flag, the Function field in the request message is set to "Reset Flags". For each bit that is set in the Event Flags field, the corresponding Event Flag in the switch port is reset.

6.2 Label Range Message

The default label range, Min Label to Max Label, is specified for each port by the Port Configuration or the All Ports Configuration messages. When the protocol is initialised, before the transmission of any Label Range messages, the label range of each port will be set to the default label range. (The default label range is dependent upon the switch design and configuration and is not specified by the GSMP protocol.) The Label Range message allows the range of labels supported by a specified port, to be changed. Each switch port MUST declare whether it supports the Label Range message in the Port Configuration or the All Ports Configuration messages. The Label Range message is:

Message Type = 33

The Label Range message has the following format for the request and success response messages:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number												Length																										
Port																																							
Port Session Number																																							
Q	M	D	x	Range Count										Range Length																									
Label Range Block																																							

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Each element of the Label Range Block has the following format:

```

+-----+
|x|x|V|C|
+-----+
~
+-----+
|x|x|x|x|
+-----+
~
+-----+
|
+-----+

```

Min Label

Max Label

Remaining Labels

Flags

Q: Query

If the Query flag is set in a request message, the switch MUST respond with the current range of valid labels. The current label range is not changed by a request message with the Query flag set. If the Query flag is zero, the message is requesting a label change operation.

M: Multipoint Query

If the Multipoint Query flag is set the switch MUST respond with the current range of valid specialized multipoint labels. The current label range is not changed by a request message with the Multipoint Query flag set.

D: Non-contiguous Label Range Indicator

This flag will be set in a Query response if the labels available for assignment belong to a non-contiguous set.

V: Label

The Label flag use is port type specific.

C: Multipoint Capable

Indicates label range that can be used for multipoint connections.

Range Count

Count of Label Range elements contained in the Label Range Block.

Range Length

Byte count in the Label Range Block.

Min Label

The minimum label value in the range.

Max Label

The maximum label value in the range.

Remaining Labels

The maximum number of remaining labels that could be requested for allocation on the specified port.

The success response to a Label Range message requesting a change of label range is a copy of the request message with the Remaining Labels field updated to the new values after the Label Range operation.

If the switch is unable to satisfy a request to change the Label range, it MUST return a failure response message with the Code field set to: "40: Cannot support one or more requested label ranges". In this failure response message, the switch MUST use the Min Label and Max Label fields to suggest a label range that it is able to satisfy.

A Label Range request message may be issued regardless of the Port Status or the Line Status of the target switch port. If the Port field of the request message contains an invalid port (a port that does not exist or a port that has been removed from the switch) a failure response message MUST be returned with the Code field set to, "4: One or more of the specified ports does not exist".

If the Query flag is set in the request message, the switch MUST reply with a success response message containing the current range of valid labels that are supported by the port. The Min Label and Max Label fields are not used in the request message.

If the Multipoint Query flag is set in the request message and the switch does not support a range of valid multipoint labels, then the switch MUST reply with a failure response message with the Code field set to, "42: Specialised multipoint labels not supported". The Min Label and Max Label fields are not used in the Multipoint request message.

If a label range changes and there are extant connection states with labels used by the previous label range, a success response message MUST be returned with the Code field set to, "46: One or more labels are still used in the previous Label Range". This action indicates that the label range has successfully changed but with a warning that there are extant connection states for the previous label range.

6.2.1 Labels

6.2.1.1 ATM Labels

If the Label Type = ATM Label, the labels range message MUST be interpreted as an ATM Label as shown:

0										1										2										3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																		
x x V C										ATM Label (0x100)																				Label Length																			
x x x x										min VPI																				min VCI																			
x x x x										ATM Label (0x100)																				Label Length																			
x x x x										max VPI																				max VCI																			
										Remaining VPI's																				Remaining VCI's																			

V: Label

If the Label flag is set, the message refers to a range of VPI's only. The Min VCI and Max VCI fields are unused. If the Label flag is zero the message refers to a range of VCI's on either one VPI or on a range of VPI's.

Min VPI, Max VPI

Specify a range of VPI values, Min VPI to Max VPI inclusive. A single VPI may be specified with a Min VPI and a Max VPI having the same value. In a request message, if the value of the Max VPI field is less than or equal to the value of the Min VPI field, the requested range is a single VPI with a value equal to the Min VPI field. Zero is a valid value. In a request message, if the Query flag is set, and the Label flag is zero, the Max VPI field specifies a single VPI and the Min VPI field is not used. The maximum valid value of these fields for both request and response messages is 0xFFF.

Min VCI, Max VCI

Specify a range of VCI values, Min VCI to Max VCI inclusive. A single VCI may be specified with a Min VCI and a Max VCI having the same value. In a request message, if the value of the Max VCI field is less than or equal to the value of the Min VCI field, the requested range is a single VCI with a value equal to the Min VCI field. Zero is a valid value. (However, VPI=0, VCI=0 is not available as a virtual channel connection as it is used as a special value in ATM to indicate an unassigned cell.)

Remaining VPI's, Remaining VCI's

These fields are unused in the request message. In the success response message and in the failure response message these fields give the maximum number of remaining VPI's and VCI's that could be requested for allocation on the specified port (after completion of the requested operation in the case of the success response). It gives the switch controller an idea of how many VPI's and VCI's it could request. The number given is the maximum possible given the constraints of the switch hardware. There is no implication that this number of VPI's and VCI's is available to every switch port.

If the Query flag and the Label flag are set in the request message, the switch MUST reply with a success response message containing the current range of valid VPI's that are supported by the port. The Min VPI and Max VPI fields are not used in the request message.

If the Query flag is set and the Label flag is zero in the request message, the switch MUST reply with a success response message containing the current range of valid VCI's that are supported by the VPI specified by the Max VPI field. If the requested VPI is invalid, a failure response MUST be returned indicating: "13: One or more of the specified Input Labels is invalid". The Min VPI field is not used in either the request or success response messages.

If the Query flag is zero and the Label flag is set in the request message, the Min VPI and Max VPI fields specify the new range of VPI's to be allocated to the input port specified by the Port field. The range of VPI's previously allocated to this port SHOULD be increased or decreased to the specified value.

If the Query flag and the Label flag are zero in the request message, the Min VCI and Max VCI fields specify the range of VCI's to be allocated to each of the VPI's specified by the VPI range. The range of VCI's previously allocated to each of the VPI's within the specified VPI range on this port, it SHOULD be increased or decreased to the specified value. The allocated VCI range MUST be the same on each of the VPI's within the specified VPI range.

If the switch is unable to satisfy a request to change the label range, it MUST return a failure response message with the Code field set to: "40: Cannot support one or more requested label ranges". If the switch is unable to satisfy a request to change the VPI, the switch MUST use the Min VPI and Max VPI fields to suggest a VPI range that it would be able to satisfy and set the VCI fields to zero, or if the switch is unable to satisfy a request to change the VCI range

on all VPI's within the requested VPI range, the switch MUST use the Min VPI, Max VPI, Min VCI, and Max VCI fields to suggest a VPI and VCI range that it would be able to satisfy.

In all other failure response messages for the label range operation, the switch MUST return the values of Min VPI, Max VPI, Min VCI, and Max VCI from the request message.

While switches can typically support all 256 or 4096 VPI's, the VCI range that can be supported is often more constrained. Often the Min VCI MUST be 0 or 32. Typically all VCI's within a particular VPI MUST be contiguous. The hint in the failure response message allows the switch to suggest a label range that it could satisfy in view of its particular architecture.

While the Label Range message is defined to specify both a range of VPI's and a range of VCI's within each VPI, the most likely use is to change either the VPI range or the range of VCI's within a single VPI. It is possible for a VPI to be valid but to be allocated no valid VCI's. Such a VPI could be used for a virtual path connection, but to support virtual channel connections it would need to be allocated a range of VCI's.

6.2.1.2 Frame Relay Labels

If the Label Type = FR Label, the labels range message MUST be interpreted as Frame Relay Labels as shown:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|V|C|      FR Label (0x101)      |      Label Length      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x| Res |Len|      Min DLCI      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x|      FR Label (0x101)      |      Label Length      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x| Res |Len|      Max DLCI      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Remaining DLCI          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

V: Label

The Label flag is not used.

Res

The Res field is reserved in [21], i.e., it is not explicitly reserved by GSMP.

Len

The Len field specifies the number of bits of the DLCI. The following values are supported:

Len	DLCI bits
0	10
2	23

Min DLCI, Max DLCI

Specify a range of DLCI values, Min DLCI to Max DLCI inclusive. The values SHOULD be right justified in the 23-bit fields and the preceding bits SHOULD be set to zero. A single DLCI may be specified with a Min DLCI and a Max DLCI having the same value. In a request message, if the value of the Max DLCI field is less than or equal to the value of the Min DLCI field, the requested range is a single DLCI with a value equal to the Min DLCI field. Zero is a valid value.

Remaining DLCI's

This field is unused in the request message. In the success response message and in the failure response message, this field gives the maximum number of remaining DLCI's that could be requested for allocation on the specified port (after completion of the requested operation in the case of the success response). It gives the switch controller an idea of how many DLCI's it could request. The number given is the maximum possible given the constraints of the switch hardware. There is no implication that this number of DLCI's is available to every switch port.

6.2.1.3 MPLS Generic Labels

The Label Range Block for PortTypes using MPLS labels. These types of labels are for use on links for which label values are independent of the underlying link technology. Examples of such links are PPP and Ethernet. On such links the labels are carried in MPLS label stacks [14]. If Label Type = MPLS Gen Label, the labels range message MUST be interpreted as MPLS Generic Label as shown:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|V|C| MPLS Gen Label (0x102)| Label Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x|x|x|x|x|x|x|x|x| Min MPLS Label |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x|x| MPLS Gen Label (0x102)| Label Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|x|x|x|x|x|x|x|x|x|x|x| Max MPLS Label |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Remaining Labels |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

V: Label

The Label flag is not used.

Min MPLS Label, Max MPLS Label

Specify a range of MPLS label values, Min MPLS Label to Max MPLS Label inclusive. The Max and Min MPLS label fields are 20 bits each.

Remaining MPLS Labels

This field is unused in the request message. In the success response message and in the failure response message this field gives the maximum number of remaining MPLS Labels that could be requested for allocation on the specified port (after completion of the requested operation in the case of the success response). It gives the switch controller an idea of how many MPLS Labels it could request. The number given is the maximum possible given the constraints of the switch hardware. There is no implication that this number of Labels is available to every switch port.

6.2.1.4 FEC Labels

The Label Range message is not used for FEC Labels and is for further study.

7. State and Statistics Messages

The state and statistics messages permit the controller to request the values of various hardware counters associated with the switch input and output ports and connections. They also permit the controller to request the connection state of a switch input port. The Connection Activity message is used to determine whether one or

more specific connections have recently been carrying traffic. The Statistics message is used to query the various port and connection traffic and error counters.

The Report Connection State message is used to request an input port to report the connection state for a single connection, a single ATM virtual path connection, or for the entire input port.

7.1 Connection Activity Message

The Connection Activity message is used to determine whether one or more specific connections have recently been carrying traffic. The Connection Activity message contains one or more Activity Records. Each Activity Record is used to request and return activity information concerning a single connection. Each connection is specified by its input port and Input Label which are specified in the Input Port and Input Label fields of each Activity Record.

Two forms of activity detection are supported. If the switch supports per connection traffic accounting, the current value of the traffic counter for each specified connection MUST be returned. The units of traffic counted are not specified but will typically be either cells or frames. The controller MUST compare the traffic counts returned in the message with previous values for each of the specified connections to determine whether each connection has been active in the intervening period. If the switch does not support per connection traffic accounting, but is capable of detecting per connection activity by some other unspecified means, the result may be indicated for each connection using the Flags field. The Connection Activity message is:

Message Type = 48

The Connection Activity request and success response messages have the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number												Length																										
Number of Records										x x x x x x x x x x x x x x x x x x																													
Activity Records																																							

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Number of Records

Field specifies the number of Activity Records to follow. The number of Activity records in a single Connection Activity message MUST NOT cause the packet length to exceed the maximum transmission unit defined by the encapsulation.

Each Activity Record has the following format:

V C A x	TC Count		TC Block Length	
Input Port				
Traffic Count				
x S x x	Input Label			

Flags

V: Valid Record

In the success response message the Valid Record flag is used to indicate an invalid Activity Record. The flag MUST be zero if any of the fields in this Activity Record are invalid, if the input port specified by the Input Port field does not exist, or if the specified connection does not exist. If the Valid Record flag is zero in a success response message, the Counter flag, the Activity flag, and the Traffic Count field are undefined. If the Valid Record flag is set, the Activity Record is valid, and the Counter and Activity flags are valid. The Valid Record flag is not used in the request message.

C: Counter

In a success response message, if the Valid Record flag is set, the Counter flag, if zero, indicates that the value in the Traffic Count field is valid. If set, it indicates that the value in the Activity flag is valid. The Counter flag is not used in the request message.

A: Activity

In a success response message, if the Valid Record and Counter flags are set, the Activity flag, if set, indicates that there has been some activity on this connection since the last Connection Activity message for this connection. If zero, it indicates that there has been no activity on this connection since the last Connection Activity message for this connection. The Activity flag is not used in the request message.

TC Count

In cases where per connection traffic counting is supported, this field contains the count of Traffic Count entries.

TC Block Length

In cases where per connection traffic counting is supported, this field contains the Traffic Count block size in bytes.

Input Port

Identifies the port number of the input port on which the connection of interest originates in order to identify the connection (regardless of whether the traffic count for the connection is maintained on the input port or the output port).

Input Label

Fields identify the specific connection for which statistics are being requested.

Traffic Count

Field is not used in the request message. In the success response message, if the switch supports per connection traffic counting, the Traffic Count field MUST be set to the value of a free running, connection specific, 64-bit traffic counter counting traffic flowing across the specified connection. The value of the traffic counter is not modified by reading it. If per connection traffic counting is supported, the switch MUST report the Connection Activity result using the traffic count rather than using the Activity flag.

The format of the failure response is the same as the request message with the Number of Records field set to zero and no Connection Activity records returned in the message. If the switch is incapable of detecting per connection activity, a failure response MUST be returned indicating, "3: The specified request is not implemented on this switch".

7.2 Statistics Messages

The Statistics messages are used to query the various port, connection and error counters.

The Statistics request messages have the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number																			Length																			
Port																																							
x	S	x	x	Label																																			
~																																							

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Label

The Label Fields identifies the specific connection for which statistics are being requested.

The success response for the Statistics message has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number											Length																											
Port																																							
x	S	x	x																																				
										Label																													
~																																							
+																																							
+																																							
Input Cell Count																																							
+																																							
+																																							
Input Frame Count																																							
+																																							
+																																							
Input Cell Discard Count																																							
+																																							
+																																							
Input Frame Discard Count																																							
+																																							
+																																							
Header Checksum Error Count																																							
+																																							
+																																							
Input Invalid Label Count																																							

Output Cell Count
Output Frame Count
Output Cell Discard Count
Output Frame Discard Count

Note: Field and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Input Cell Count, Output Cell Count

Give the value of a free running 64-bit counter counting cells arriving at the input or departing from the output respectively. These fields are relevant for label type = ATM, for all other label types these fields SHOULD be set to zero by the sender and ignored by the receiver.

Input Frame Count, Output Frame Count

Give the value of a free running 64-bit counter counting frames (packets) arriving at the input or departing from the output respectively. These fields are relevant for label types = FR and MPLS, for all other label types these fields SHOULD be set to zero by the sender and ignored by the receiver.

Input Cell Discard Count, Output Cell Discard Count

Give the value of a free running 64-bit counter counting cells discarded due to queue overflow on an input port or on an output port respectively. These fields are relevant for label type = ATM, for all other label types these fields SHOULD be set to zero by the sender and ignored by the receiver.

Input Frame Discard Count, Output Frame Discard Count

Give the value of a free running 64-bit counter counting frames discarded due to congestion running on an input port or on an output port respectively. These fields are relevant for label

types = FR and MPLS, for all other label types these fields SHOULD be set to zero by the sender and ignored by the receiver.

Header Checksum Error Count

Gives the value of a free running 64-bit counter counting cells or frames discarded due to header checksum errors on arrival at an input port. For an ATM switch this would be the HEC count.

Invalid Label Count

Gives the value of a free running 64-bit counter counting cells or frames discarded because their Label is invalid on arrival at an input port.

7.2.1 Port Statistics Message

The Port Statistics message requests the statistics for the switch port specified in the Port field. The contents of the Label field in the Port Statistics request message is ignored. All of the count fields in the success response message refer to per-port counts regardless of the connection to which the cells or frames belong. Any of the count fields in the success response message not supported by the port MUST be set to zero. The Port Statistics message is:

Message Type = 49

7.2.2 Connection Statistics Message

The Connection Statistics message requests the statistics for the connection specified in the Label field that originates on the switch input port specified in the Port field. All of the count fields in the success response message refer only to the specified connection. The Header Checksum Error Count and Invalid Label Count fields are not connection specific and MUST be set to zero. Any of the other count fields not supported on a per connection basis MUST be set to zero in the success response message. The Connection Statistics message is:

Message Type = 50

7.2.3 QoS Class Statistics Message

The QoS Class Statistics message is not supported in this version of GSMP.

Message Type = 51 is reserved.

7.3 Report Connection State Message

The Report Connection State message is used to request an input port to report the connection state for a single connection or for the entire input port. The Report Connection State message is:

Message Type = 52

The Report Connection State request message has the following format:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Version  | Message Type |   Result   |   Code   |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Partition ID | Transaction Identifier |
+-----+-----+-----+-----+-----+-----+-----+-----+
| I | SubMessage Number | Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Input Port |
+-----+-----+-----+-----+-----+-----+-----+-----+
| x | S | A | V |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Input Label |
+-----+-----+-----+-----+-----+-----+-----+-----+
~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Note: Field and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Input Port

Identifies the port number of the input port for which the connection state is being requested.

Flags

A: All Connections

If the All Connections flag is set, the message requests the connection state for all connections that originate at the input port specified by the Input Port field. In this case the Input Label field and the Label flag are unused.

V: ATM VPI

The ATM VPI flag may only be set for ports with PortType=ATM. If the switch receives a Report Connection State message in which the ATM VPI flag is set and in which the input port specified by the Input Port field does not have PortType=ATM, the switch MUST return a Failure response "28: ATM Virtual Path switching is not supported on non-ATM ports".

If the All Connections flag is zero and the ATM VPI flag is also zero, the message requests the connection state for the connection that originates at the input port specified by the Port and Input Label fields.

ATM specific procedures:

If the All Connections flag is zero and the ATM VPI flag is set and the input port specified by the Input Port field has LabelType=ATM, the message requests the connection state for the virtual path connection that originates at the input port specified by the Input Port and Input VPI fields. If the specified Input VPI identifies an ATM virtual path connection (i.e., a single switched virtual path) the state for that connection is requested. If the specified Input VPI identifies a virtual path containing virtual channel connections, the message requests the connection state for all virtual channel connections that belong to the specified virtual path.

Input Label

Field identifies the specific connection for which the connection state is being requested. For requests that do not require a connection to be specified, the Input Label field is not used.

response to the same request message MUST have the same Input Port and Transaction Identifier fields as the request message. A single Connection Record MUST NOT be split across multiple success response messages. "More" in the Result field of a response message indicates that one or more further success response messages should be expected in response to the same request message. "Success" in the Result field indicates that the response to the request has been completed. The Result values are defined in chapter 3.1.1.

Each Connection Record has the following format:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|A|V|P|      Record Count      |      Record Length      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|S|x|x|
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               Input Label                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               ~
~                               Output Branch Records                       ~
|                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Flags

A: All Connections

V: ATM VPI

For the first Connection Record in each success response message, the All Connections and the ATM VPI flags MUST be the same as those of the request message. For successive Connection Records in the same success response message, these flags are not used.

P: ATM VPC

The ATM VPC flag may only be set for ports with PortType=ATM. The ATM VPC flag, if set and only if set, indicates that the Connection Record refers to an ATM virtual path connection.

Input Label

The input label of the connection specified in this Connection Record.

Record Count

Count of Output Branch Records included in a response message.

Record Length

Length in bytes of Output Branch Records field

Output Branch Records

Each Connection Record MUST contain one or more Output Branch Records. Each Output Branch Record specifies a single output branch belonging to the connection identified by the Input Label field of the Connection Record and the Input Port field of the Report Connection State message. A point-to-point connection will require only a single Output Branch Record. A point-to-multipoint connection will require multiple Output Branch Records. If a point-to-multipoint connection has more output branches than can fit in a single Connection Record contained within a single success response message, that connection may be reported using multiple Connection Records in multiple success response messages.

Each Output Branch Record has the following format:

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Output Port                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|x|S|x|x|                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Output Label                                     |
~                                                                                     ~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Output Port

The output port of the switch to which this output branch is routed.

Output Label

The output label of the output branch specified in this Output Branch Record.

ATM specific procedures:

If this Output Branch Record is part of a Connection Record that specifies a virtual path connection (the ATM VPC flag is set) the Output VCI field is unused.

A Report Connection State request message may be issued regardless of the Port Status or the Line Status of the target switch port.

If the Input Port of the request message is valid, and the All Connections flag is set, but there are no connections established on that port, a failure response message MUST be returned with the Code field set to, "10: General Message Failure". For the Report Connection State message, this failure code indicates that no

connections matching the request message were found. This failure message SHOULD also be returned if the Input Port of the request message is valid, the All Connections flag is zero, and no connections are found on that port matching the specified connection.

8. Configuration Messages

The configuration messages permit the controller to discover the capabilities of the switch. Three configuration request messages have been defined: Switch, Port, and All Ports.

8.1 Switch Configuration Message

The Switch Configuration message requests the global (non port-specific) configuration for the switch. The Switch Configuration message is:

Message Type = 64

The Port field is not used in the switch configuration message.

The Switch Configuration message has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number										Length																												
MType					MType					MType					MType																								
Firmware Version Number										Window Size																													
Switch Type																																							
										Switch Name																													
										Max Reservations																													

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

MType

Represents an alternative QoS Configuration type. In the request message the requested MType is in the most significant (leftmost) MType byte; the other three MType bytes are unused. The reply message will either accept the MType request by including the requested MType in the leftmost MType field of the response message or it will reject the MType request by responding with MType=0, the default MType, in the first MType field. Optionally, in the case of a rejection, the switch reply can include up to 3 additional MType values, each of which indicates an available alternative QoS Configuration. A switch that supports only the default QoS Configuration always returns MType=0 in all four MType fields. MType negotiation is discussed in section 8.1.1.

- | | | |
|---------|---|---|
| 0 | - | Indicates use of the default GSMP model |
| 1-200 | - | Reserved |
| 201-255 | - | Experimental |

Firmware Version Number

The version number of the switch control firmware installed.

Window Size

The maximum number of unacknowledged request messages that may be transmitted by the controller without the possibility of loss. This field is used to prevent request messages being lost in the switch because of overflow in the receive buffer. The field is a hint to the controller. If desired, the controller may experiment with higher and lower window sizes to determine heuristically the best window size.

Switch Type

A 16-bit field allocated by the manufacturer of the switch. (For these purposes, the manufacturer of the switch is assumed to be the organisation identified by the OUI in the Switch Name field.) The Switch Type identifies the product. When the Switch Type is combined with the OUI from the Switch Name the product is uniquely identified. Network Management may use this identification to obtain product related information from a database.

Switch Name

A 48-bit quantity that is unique within the operational context of the device. A 48-bit IEEE 802 MAC address, if available, may be used as the Switch Name. The most significant 24 bits

of the Switch Name MUST be an Organisationally Unique Identifier (OUI) that identifies the manufacturer of the switch.

Max Reservations

The maximum number of Reservations that the switch can support (see Chapter 5). A value of 0 indicates that the switch does not support Reservations.

8.1.1 Configuration Message Processing

After adjacency between a controller and after a switch is first established the controller that opts to use a QoS Configuration model other than the default would send the Switch Configuration request including the requested QoS Configuration's MType value in the request message. This request MUST be sent before any connection messages are exchanged. If the switch can support the requested QoS configuration, then the switch includes the requested MType value in the response message as an indication that it accepts the request. If the switch cannot support the requested QoS Configuration, it replaces the MType value in the request message with that of the default QoS Configuration, i.e., MType=0.

The switch configuration response messages may additionally include the MType values of up to three alternative QoS Configurations that the switch supports and that the controller may choose between.

The exchange continues until the controller sends a requested MType that the switch accepts or until it sends a connection request message. If the exchange ends without confirmation of an alternate switch model, then the default Mtype=0 is be used.

Once an MType has been established for the switch, it cannot be changed without full restart, that is the re-establishment of adjacency with the resetting of all connections.

8.2 Port Configuration Message

The Port Configuration message requests the switch for the configuration information of a single switch port. The Port field in the request message specifies the port for which the configuration is requested. The Port Configuration message is:

Message Type = 65.

The Port Configuration success response message has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number															Length																							
Port																																							
Port Session Number																																							
Event Sequence Number																																							
Event Flags															Port Attribute Flags																								
PortType										S x x x x x x x x										Data Fields Length																			
PortType Specific Data																																							
x x x x x x x x x x x x x x x x										Number of Service Specs																													
Service Specs List																																							

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Port

The switch port to which the configuration information refers. Configuration information relating to both the input and the output sides of the switch port is given. Port numbers are 32 bits wide and allocated by the switch. The switch may choose to structure the 32 bits into subfields that have meaning to the physical structure of the switch hardware (e.g., physical slot and port). This structure may be indicated in the Physical Slot Number and Physical Port Number fields.

Event Sequence Number

The Event Sequence Number is set to zero when the port is initialised. It is incremented by one each time the port detects an asynchronous event that the switch would normally report via an Event message. The Event Sequence Number is explained in section 9.

Event Flags

Event Flags in a switch port corresponds to a type of Event message.

Port Attribute Flags

Port Attribute Flags indicate specific behaviour of a switch port. The format of the Port Attribute Flags field is given below:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
  |R|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|x|
  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

R: Connection Replace flag

If set, indicates that connections being established by an Add Branch message with a corresponding R-bit set will replace any previously established connection if a clash between the established output branch and the requested output branch occurs [see chapter 4.2].

x: Unused.

PortType

- 1: PortType is ATM
- 2: PortType is FR
- 3: PortType is MPLS

S: Service Model

If set, indicates that Service Model data follows the PortSpecific port configuration data.

Data Fields Length

The total length in bytes of the combined PortType Specific Data and Service Model Data fields. The length of each of these fields may be derived from the other data so the value of Data Fields Length serves primarily as a check and to assist parsing of the All Ports Configuration message success response.

PortType Specific Data

This field contains the configuration data specific to the particular port type as specified by the PortType field. The field format and length also depends on the value of the PortType. PortType Specific Data is defined below.

Number of Service Specs

Field contains the total number of Service Specs following in the remainder of the Port Configuration message response or Port Configuration Record.

Service Specs List

The Service Specs correspond to the Input and Output Service selectors used in Connection Management and Reservation messages. Specifically they define the possible values used when the Service Selector (IQS or OQS) is set to 0b10 indicating the use of the default service specification model defined in Chapter 10.

Service Spec

The format of each service spec is given below:

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|           Service ID           |           Capability Set ID           |
+-----+-----+-----+-----+
```

Each Service Spec identifies a Service supported by the switch together with the Capability Set ID that identifies the parameters of that instance of the Service. The Service Spec List may contain more than one Service Spec sharing the same Service ID. However, each Service Spec in the Service Specs List MUST be unique.

Service ID

Field contains the Service ID of a Service supported on the port. Service ID values are defined as part of the Service definition in Chapter 9.6.

Capability Set ID

Field identifies a Capability Set ID of the Service specified by the Service ID that is supported on the port. Capability Set ID values are defined by the Switch in the Service Configuration response message (see Section 8.4). The switch MUST NOT return a {Service ID, Capability Set ID} pair that is not reported in a Service Configuration response message.

8.2.1 PortType Specific Data

The length, format and semantics of the PortType Specific Data field in the Port Configuration message success response and in the Port Records of the All Port Configuration message success response all depend on the PortType value of the same message or record respectively. The specification of the PortType Specific Data field is given below. For each defined PortType value the Min and Max Label fields are given in the subsequent subsections.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
P M L R Q										Label Range Count										Label Range Length																			
										Default Label Range Block																													
										Receive Data Rate																													
										Transmit Data Rate																													
Port Status										Line Type										Line Status										Priorities									
Physical Slot Number																				Physical Port Number																			

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Where each of the ranges in the Default Label Range Blocks will have the following format:

x x V C																																							
										Min Label																													
x x x x																																							
										Max Label																													

Flags

P: VP Switching

The ATM VPC flag may only be set for ports with PortType=ATM. The VP Switching flag, if set, indicates that this input port is capable of supporting virtual path switching. Else, if zero, it indicates that this input port is only capable of virtual channel switching.

M: Multicast Labels

The Multicast Labels flag, if set, indicates that this output port is capable of labelling each output branch of a point-to-multipoint tree with a different label. If zero, it indicates that this output port is not able to label each output branch of a point-to-multipoint tree with a different label.

L: Logical Multicast

The Logical Multicast flag, if set, indicates that this output port is capable of supporting more than a single branch from any point-to-multipoint connection. This capability is often referred to as logical multicast. If zero, it indicates that this output port can only support a single output branch from each point-to-multipoint connection.

R: Label Range

The Label Range flag, if set, indicates that this switch port is capable of reallocating its label range and therefore accepts the Label Range message. Else, if zero, it indicates that this port does not accept Label Range messages.

Q: QoS

The QoS flag, if set, indicates that this switch port is capable of handling the Quality of Service messages defined in section 9 of this specification. Else, if zero, it indicates that this port does not accept the Quality of Service messages.

V: Label

The Label flag is port type specific.

C: Multipoint Capable

This flag indicates that the label range may be used for multipoint connections.

Label Range Count

The total number of Default Label Range elements contained in the Default Label Range Block.

Label Range Length

Byte count in the Default Label Range Block.

Min Label

The specification of the Min Label field for each defined PortType value is given in the subsequent subsections. The default minimum value of a dynamically assigned incoming label that the connection table on the input port supports and that may be controlled by GSMP. This value is not changed as a result of the Label Range message.

Max Label

The specification of the Max Label field for each defined PortType value is given in the subsequent subsections. The default maximum value of a dynamically assigned incoming label that the connection table on the input port supports and that may be controlled by GSMP. This value is not changed as a result of the Label Range message.

Receive Data Rate

The maximum rate of data that may arrive at the input port in;

cells/s	for PortType = ATM
bytes/s	for PortType = FR
bytes/s	for PortType = MPLS

Transmit Data Rate

The maximum rate of data that may depart from the output port in;

cells/s	for PortType = ATM
bytes/s	for PortType = FR
bytes/s	for PortType = MPLS

(The transmit data rate of the output port may be changed by the Set Transmit Data Rate function of the Port Management message.)

Port Status

Gives the administrative state of the port. The defined values of the Port Status field are:

Available:

Port Status = 1. The port is available to both send and receive cells or frames. When a port changes to the Available state from any other administrative state, all dynamically assigned connections MUST be cleared and a new Port Session Number MUST be generated.

Unavailable:

Port Status = 2. The port has intentionally been taken out of service. No cells or frames will be transmitted from this port. No cells or frames will be received by this port.

Internal Loopback:

Port Status = 3. The port has intentionally been taken out of service and is in internal loopback: cells or frames arriving at the output port from the switch fabric are looped through to the input port to return to the switch fabric. All of the functions of the input port above the physical layer, e.g., header translation, are performed upon the looped back cells or frames.

External Loopback:

Port Status = 4. The port has intentionally been taken out of service and is in external loopback: cells or frames arriving at the input port from the external communications link are immediately looped back to the communications link at the physical layer without entering the input port. None of the functions of the input port above the physical layer are performed upon the looped back cells or frames.

Bothway Loopback:

Port Status = 5. The port has intentionally been taken out of service and is in both internal and external loopback.

The Port Status of the port over which the GSMP session controlling the switch is running MUST be declared Available. The controller will ignore any other Port status for this port. The Port Status of switch ports after power-on initialisation is not defined by GSMP.

Line Type

The type of physical transmission interface for this port. The values for this field are defined by the IANAifType's specified in [17].

The following values are identified for use in this version of the protocol.

```
PortType = Unknown: other(1)
PortType = MPLS:    ethernetCsmacd(6),
                    ppp(23)
PortType = ATM:     atm(37)
PortType = FR:      frameRelayService(44)
```

Line Status

The status of the physical transmission medium connected to the port. The defined values of the Line Status field are:

Up:

Line Status = 1. The line is able to both send and receive. When the Line Status changes to Up from either the Down or Test states, a new Port Session Number MUST be generated.

Down:

Line Status = 2. The line is unable either to send or receive or both.

Test:

Line Status = 3. The port or line is in a test mode, for example, power-on test.

Priorities

The number of different priority levels that this output port can assign to connections. Zero is invalid in this field. If an output port is able to support "Q" priorities, the highest priority is numbered zero and the lowest priority is numbered "Q-1". The ability to offer different qualities of service to different connections based upon their priority is assumed to be a property of the output port of the switch. It may be assumed that for connections that share the same output port, a cell or frame on a connection with a higher priority is much more likely to exit the switch before a cell or frame on a connection with a lower priority if they are both in the switch at the same time.

Physical Slot Number

The physical location of the slot in which the port is located. It is an unsigned 16-bit integer that can take any value except 0xFFFF. The value 0xFFFF is used to indicate "unknown". The Physical Slot Number is not used by the GSMP protocol. It is provided to assist network management in functions such as logging, port naming, and graphical representation.

Physical Port Number

The physical location of the port within the slot in which the port is located. It is an unsigned 16-bit integer that can take any value except 0xFFFF. The value 0xFFFF is used to indicate "unknown". The Physical Port Number is not used by the GSMP protocol. It is provided to assist network management in functions such as logging, port naming, and graphical representation.

There MUST be a one to one mapping between the Port Number and the Physical Slot Number and Physical Port Number combination. Two different Port Numbers MUST NOT yield the same Physical Slot Number and Physical Port Number combination. The same Port Number MUST yield the same Physical Slot Number and Physical Port Number within a single GSMP session. If both Physical Slot Number and Physical Port Number indicate "unknown" the physical location of switch ports may be discovered by looking up the product identity in a database to reveal the physical interpretation of the 32-bit Port Number.

8.2.1.1 PortType Specific data for PortType=ATM

If PortType=ATM, the Default Label Range Block has the following format:

0										1										2										3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9										
x x V x										ATM Label (0x100)																				Label Length																			
x x x x										VPI																				VCI																			

V: Label

If the Label flag is set, the message refers to a range of VPI's only. The Min VCI and Max VCI fields are unused. If the Label flag is zero the message refers to a range of VCI's on either one VPI or on a range of VPI's.

Min VPI

The default minimum value of dynamically assigned incoming VPI that the connection table on the input port supports and that may be controlled by GSMP.

Max VPI

The default maximum value of dynamically assigned incoming VPI that the connection table on the input port supports and that may be controlled by GSMP.

At power-on, after a hardware reset, and after the Reset Input Port function of the Port Management message, the input port MUST handle all values of VPI within the range Min VPI to Max VPI inclusive and GSMP MUST be able to control all values within this range. It should be noted that the range Min VPI to Max VPI refers only to the incoming VPI range that can be supported by the associated port. No restriction is placed on the values of outgoing VPI's that may be written into the cell header. If the switch does not support virtual paths it is acceptable for both Min VPI and Max VPI to specify the same value, most likely zero.

Use of the Label Range message allows the range of VPI's supported by the port to be changed. However, the Min VPI and Max VPI fields in the Port Configuration and All Ports Configuration messages always report the same default values regardless of the operation of the Label Range message.

Min VCI

The default minimum value of a dynamically assigned incoming VCI that the connection table on the input port can support and may be controlled by GSMP. This value is not changed as a result of the Label Range message.

Max VCI

The default maximum value of a dynamically assigned incoming VCI that the connection table on the input port can support and may be controlled by GSMP.

At power-on, after a hardware reset, and after the Reset Input Port function of the Port Management message, the input port MUST handle all values of VCI within the range Min VCI to Max VCI inclusive, for each of the virtual paths in the range Min VPI to Max VPI inclusive, and GSMP MUST be able to control all values within this range. It should be noted that the range Min VCI to Max VCI refers only to the incoming VCI range that can be supported by the associated port on each of the virtual paths in the range Min VPI to Max VPI. No restriction is placed on the values of outgoing VCI's that may be written into the cell header. Use of the Label Range message allows the range of VCI's to be changed on each VPI supported by the port. However, the Min VCI and Max VCI fields in the Port Configuration and All Ports Configuration messages always report the same default values regardless of the operation of the Label Range message.

For a port over which the GSMP protocol is operating, the VCI of the GSMP control channel may or may not be reported as lying within the range Min VCI to Max VCI. A switch should honour a connection request message that specifies the VCI value of the GSMP control channel even if it lies outside the range Min VCI to Max VCI

8.2.1.2 PortType Specific data for PortType=FR

If PortType=FR, the Default Label Range Block has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x|S|x|x|      FR Label (0x101)      |      Label Length      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|x x x x|Res|Len|      DLCI      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Res

The Res field is reserved in [21], i.e., it is not explicitly reserved by GSMP.

Len

This field specifies the number of bits of the DLCI. The following values are supported:

Len	DLCI bits
0	10
2	23

Min DLCI, Max DLCI

Specify a range of DLCI values, Min DLCI to Max DLCI inclusive. The values SHOULD be right justified in the 23-bit fields and the preceding bits SHOULD be set to zero. A single DLCI may be specified with a Min DLCI and a Max DLCI having the same value. In a request message, if the value of the Max DLCI field is less than or equal to the value of the Min DLCI field, the requested range is a single DLCI with a value equal to the Min DLCI field. Zero is a valid value.

8.2.1.3 PortType Specific data for PortType=MPLS

The Default Label Range Block for PortTypes using MPLS labels. These types of labels are for use on links for which label values are independent of the underlying link technology. Examples of such links are PPP and Ethernet. On such links the labels are carried in MPLS label stacks [14]. Ports of the Type MPLS have the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x|x|x|x| MPLS Gen Label (0x102)|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x|x|x|x|x|x|x|x|x|x|x|x|                               MPLS Label |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Min MPLS Label, Max MPLS Label

Specify a range of MPLS label values, Min MPLS Label to Max MPLS Label inclusive. The Max and Min MPLS label fields are 20 bits each.

8.2.1.4 PortType Specific data for PortType=FEC

The Default Label Range Block for PortTypes using FEC labels is not used. The Label Range Count and Label Range Length fields defined in [8.2.1] should be set to 0.

8.3 All Ports Configuration Message

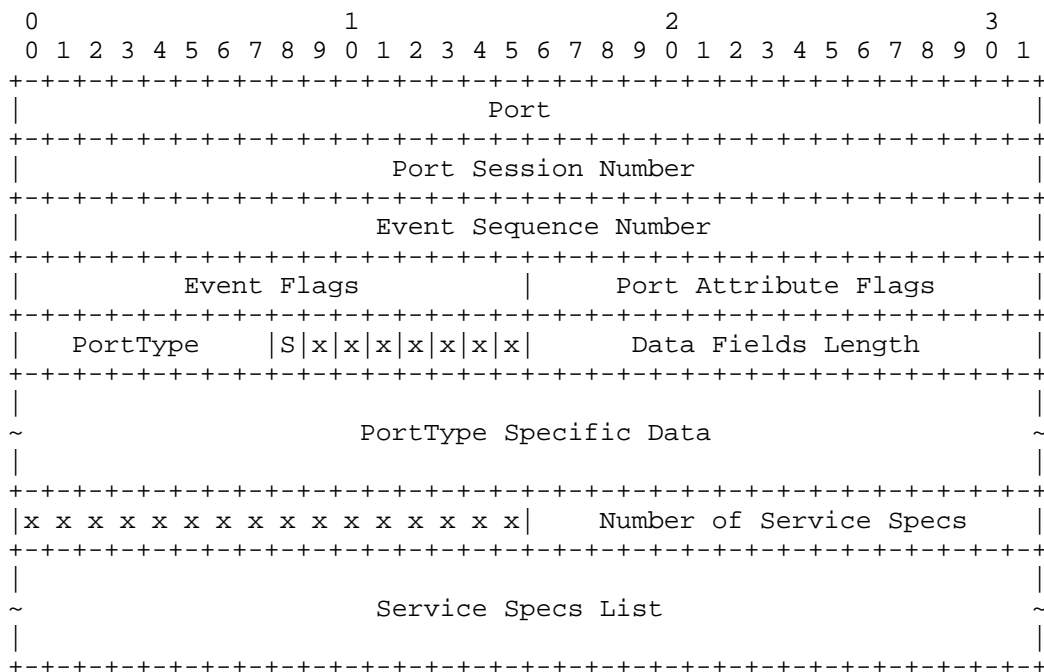
The All Ports Configuration message requests the switch for the configuration information of all of its ports. The All Ports Configuration message is:

Message Type = 66

The Port field is not used in the request message.

Port Records

Follow in the remainder of the message. Each port record has the following format:



The definition of the fields in the Port Record is exactly the same as that of the Port Configuration message [section 8.2].

8.4 Service Configuration Message

The Service Configuration message requests the switch for the configuration information of the Services that are supported. The Service Configuration message is:

Message Type = 67

The Service Configuration success response message has the following format:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|   Version   | Message Type |   Result   |   Code   |
+-----+-----+-----+-----+
| Partition ID | Transaction Identifier |
+-----+-----+-----+-----+
| I | SubMessage Number | Length |
+-----+-----+-----+-----+
| x x x x x x x x x x x x x x x | Number of Service Records |
+-----+-----+-----+-----+
|                                     |
~                               Service Records                               ~
|                                     |
+-----+-----+-----+-----+

```

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Number of Service Records

Field gives the total number of Service Records to be returned in the Service Records field.

Service Records

A sequence of zero or more Service Records. The switch returns one Service Record for each Service that it supports on any of its ports. A Service record contains the configuration data of the specified Service. Each Service Record has the following format:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|   Service ID   | Number of Cap. Set. Records |
+-----+-----+-----+-----+
|                                     |
~                               Capability Set Records                               ~
|                                     |
+-----+-----+-----+-----+

```

Service ID

The Service ID Field identifies the Service supported by the port. The Services are defined with their Service ID values as described in section 10.2.

Number of Cap. Set. Records

Field gives the total number of Capability Set Records to be returned in the Service Record field.

Capability Set Records

The switch returns one or more Capability Set Records in each Service Record. A Capability Set contains a set of parameters that describe the QoS parameter values and traffic controls that apply to an instance of the Service. Each Capability Set record has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Cap. Set ID										Traffic Controls																													
CLR										CTD																													
Frequency										CDV																													

Capability Set ID

The value in this Field defines a Capability Set ID supported by the switch. The values of a Capability Set ID are assigned by the switch and used in Port Configuration messages to identify Capability Sets supported by individual ports. Each Capability Set Record within a Service Record MUST have a unique Capability Set ID.

Traffic Controls

Field identifies the availability of Traffic Controls within the Capability Set. Traffic Controls are defined as part of the respective Service definition, see Chapter 10. Some or all of the Traffic Controls may be undefined for a given Service, in which case the corresponding Flag is ignored by the controller. The Traffic Controls field is formatted into Traffic Control Sub-fields as follows:

0										1										
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5					
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+					
	U		D		I		E		S		V		x	x	x	x				
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+					

Traffic Control Sub-fields have the following encoding:

0b00 Indicates that the Traffic Control is not available in the Capability Set.

- 0b01 Indicates that the Traffic Control is applied to all connections that use the Capability Set.
- 0b10 Indicates that the Traffic Control is available for application to connections that use the Capability Set on a per connection basis.
- 0b11 Reserved

Traffic Control Sub-fields:

- U: Usage Parameter Control
The Usage Parameter Control sub-field indicates the availability of Usage Parameter Control for the specified Service and Capability Set.
- D: Packet Discard
The Packet Discard sub-field indicates the availability of Packet Discard for the specified Service and Capability Set.
- I: Ingress Shaping
The Ingress Shaping sub-field indicates the availability of Ingress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance for the specified Service and Capability Set.
- E: Egress Shaping, Peak Rate
The Egress Shaping, Peak Rate sub-field indicates the availability of Egress Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance for the specified Service and Capability Set.
- S: Egress Traffic Shaping, Sustainable Rate
The Egress Shaping, Sustainable Rate sub-field, if set, indicates that Egress Traffic Shaping to the Sustainable Cell Rate and Maximum Burst Size is available for the specified Service and Capability Set.
- V: VC Merge
The VC Merge sub-field indicates the availability of ATM Virtual Channel Merge (i.e., multipoint to point ATM switching with a traffic control to avoid AAL5 PDU interleaving) capability for the specified Service and Capability Set.

QoS Parameters

The remaining four fields in the Capability Set Record contain the values of QoS Parameters. QoS Parameters are defined as part of the respective Service definition, see Chapter 9.6. Some or all of the QoS Parameters may be undefined for a given Service, in which case the corresponding field is ignored by the controller.

CLR: Cell Loss Ratio

The Cell Loss Ratio parameter indicates the CLR guaranteed by the switch for the specified Service. A cell loss ratio is expressed as an order of magnitude n , where the CLR takes the value of ten raised to the power of $-n$, i.e., $\log(\text{CLR}) = -n$. The value n is coded as a binary integer, having a range of $1 \leq n \leq 15$. In addition, the value 0b1111 1111 indicates that no CLR guarantees are given.

Frequency

The frequency field is coded as an 8 bit unsigned integer. Frequency applies to the MPLS CR-LDP Service (see Section 10.4.3). Valid values of Frequency are:

- 0 - Very frequent
- 1 - Frequent
- 2 - Unspecified

CTD: Cell Transfer Delay

The CTD value is expressed in units of microseconds. It is coded as a 24-bit integer.

CDV: Peak-to-peak Cell Delay Variation

The CDV value is expressed in units of microseconds. It is coded as a 24-bit integer.

9. Event Messages

Event messages allow the switch to inform the controller of certain asynchronous events. By default the controller does not acknowledge event messages unless ReturnReceipt is set in the Result field. The Code field is only used in case of Adjacency Update message, otherwise it is not used and SHOULD be set to zero. Event messages are not sent during initialisation. Event messages have the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
I	SubMessage Number											Length																											
										Port																													
										Port Session Number																													
										Event Sequence Number																													
x S x x																																							
										Label																													
~																																							

Note: Fields and Parameters that have been explained in the description of the general messages will not be explained in this section. Please refer to section 3.1 for details.

Event Sequence Number

The current value of the Event Sequence Number for the specified port. The Event Sequence Number is set to zero when the port is initialised. It is incremented by one each time the port detects an asynchronous event that the switch would normally report via an Event message. The Event Sequence Number MUST be incremented each time an event occurs even if the switch is prevented from sending an Event message due to the action of the flow control.

Label

Field gives the Label to which the event message refers. If this field is not required by the event message it is set to zero.

Each switch port MUST maintain an Event Sequence Number and a set of Event Flags, one Event Flag for each type of Event message. When a switch sends an Event message it MUST set the Event Flag for that port corresponding to the Event type. If Flow Control is activated for this Event type for this Port then the switch MUST NOT send another Event message of the same type for that port until the Event Flag has been reset. Event Flags are reset by the "Reset Event Flags" function of the Port Management message. This is a simple flow control preventing the switch from flooding the controller with

event messages. The Event Sequence Number of the port MUST be incremented every time an event is detected on that port even if the port is prevented from reporting the event due to the action of the flow control. This allows the controller to detect that it has not been informed of some events that have occurred on the port due to the action of the flow control.

9.1 Port Up Message

The Port Up message informs the controller that the Line Status of a port has changed from, either the Down or Test state to the Up state. When the Line Status of a switch port changes to the Up state from either the Down or Test state a new Port Session Number MUST be generated, preferably using some form of random number. The new Port Session Number is given in the Port Session Number field. The Label field is not used and is set to zero. The Port Up message is:

Message Type = 80

9.2 Port Down Message

The Port Down message informs the controller that the Line Status of a port has changed from the Up state or Test state to the Down state. This message will be sent to report link failure if the switch is capable of detecting link failure. The port session number that was valid before the port went down is reported in the Port Session Number field. The Label field is not used and is set to zero. The Port Down message is:

Message Type = 81

9.3 Invalid Label Message

The Invalid Label message is sent to inform the controller that one or more cells or frames have arrived at an input port with a Label that is currently not allocated to an assigned connection. The input port is indicated in the Port field, and the Label in the Label field. The Invalid Label message is:

Message Type = 82

9.4 New Port Message

The New Port message informs the controller that a new port has been added to the switch. The port number of the new port is given in the Port field. A new Port Session Number MUST be assigned, preferably using some form of random number. The new Port Session Number is given in the Port Session Number field. The state of the new port is undefined so the Label field is not used and is set to zero. The New Port message is:

Message Type = 83

9.5 Dead Port Message

The Dead Port message informs the controller that a port has been removed from the switch. The port number of the port is given in the Port field. The Port Session Number that was valid before the port was removed is reported in the Port Session Number field. The Label fields are not used and are set to zero. The Dead Port message is:

Message Type = 84

9.6 Adjacency Update Message

The Adjacency Update message informs the controller when adjacencies, i.e., other controllers controlling a specific partition, are joining or leaving. When a new adjacency has been established, the switch sends an Adjacency Update message to every controller with an established adjacency to that partition. The Adjacency Update message is also sent when adjacency is lost between the partition and a controller, provided that there are any remaining adjacencies with that partition. The Code field is used to indicate the number of adjacencies known by the switch partition. The Label field is not used and SHOULD be set to zero. The Adjacency Update message is:

Message Type = 85

10. Service Model Definition

10.1 Overview

In the GSMP Service Model a controller may request the switch to establish a connection with a given Service. The requested Service is identified by including a Service ID in the Add Branch message or the Reservation Message. The Service ID refers to a Service Definition provided in this chapter of the GSMP specification.

A switch that implements one or more of the Services, as defined below, advertises the availability of these Services in the Service Configuration message response (see Section 8.4). Details of the switch's implementation of a given Service that are important to the controller (e.g., the value of delay or loss bounds or the availability of traffic controls such as policers or shapers) are reported in the form of a Capability Set in the Service Configuration message response.

Thus a switch's implementation of a Service is defined in two parts: the Service Definition, which is part of the GSMP specification, and the Capability Set, which describes attributes of the Service specific to the switch. A switch may support more than one Capability Set for a given Service. For example if a switch supports one Service with two different values of a delay bound it could do this by reporting two Capability Sets for that Service.

The Service Definition is identified in GSMP messages by the Service ID, a sixteen-bit identifier. Assigned numbers for the Service ID are given with the Service Definitions in Section 10.4. The Capability Set is identified in GSMP messages by the Capability Set ID, a sixteen-bit identifier. Numbers for the Capability Set ID are assigned by the switch and are advertised in the Service Configuration message response.

The switch reports all its supported Services and Capability Sets in the Service Configuration message response. The subset of Services and Capability Sets supported on an individual port is reported in the Port Configuration message response or in the All Ports Configuration message response. In these messages the Services and Capability Sets supported on the specified port are indicated by a list of {Service ID, Capability Set ID} number pairs.

10.2 Service Model Definitions

Terms and objects defined for the GSMP Service Model are given in this section.

10.2.1 Original Specifications

Services in GSMP are defined largely with reference to Original Specifications, i.e., the standards or implementation agreements published by organisations such as ITU-T, IETF, and ATM Forum that originally defined the Service. This version of GSMP refers to 4 original specifications: [8], [9], [10] and [11].

10.2.2 Service Definitions

Each Service Definition in GSMP includes definition of:

Traffic Parameters

Traffic Parameter definitions are associated with Services while Traffic Parameter values are associated with connections.

Traffic Parameters quantitatively describe a connection's requirements on the Service. For example, Peak Cell Rate is a Traffic Parameter of the Service defined by the ATM Forum Constant Bit Rate Service Category.

Some Traffic Parameters are mandatory and some are optional, depending on the Service.

Semantics of Traffic Parameters are defined by reference to Original Specifications.

QoS Parameters

QoS Parameters and their values are associated with Services.

QoS Parameters express quantitative characteristics of a switch's support of a Service. They include, for example, quantitative bounds on switch induced loss and delay.

Some QoS Parameters will be mandatory and some will be optional.

Semantics of QoS Parameters are defined by reference to Original Specifications.

Traffic Controls

The implementation of some Services may include the use of Traffic Controls. Traffic Controls include, for example functions such as policing, input shaping, output shaping, tagging and marking, frame vs. cell merge, frame vs. cell discard.

Switches are not required to support Traffic Controls. Any function that is always required in the implementation of a Service is considered part of the Service and is not considered a Traffic Control.

If a switch supports a Traffic Control then the control may be applied either to all connections that use a given Capability Set (see below) or to individual connections.

The definition of a Traffic Control is associated with a Service. Traffic Controls are defined, as far as possible, by reference to Original Specifications.

10.2.3 Capability Sets

For each Service that a switch supports the switch MUST also support at least one Capability Set. A Capability Set establishes characteristics of a switch's implementation of a Service. It may be appropriate for a switch to support more than one Capability Set for a given Service.

A Capability Set may contain, depending on the Service definition, QoS Parameter values and an indication of availability of Traffic Controls.

If a switch reports QoS Parameter values in a Capability Set then these apply to all the connections that use that Capability Set.

For each Traffic Control defined for a given Service the switch reports availability of that control as one of the following:

Not available in the Capability Set,

Applied to all connections that use the Capability Set, or

Available for application to connections that use the Capability Set on a per connection basis. In this case, a controller may request application of the Traffic Control in connection management messages.

10.3 Service Model Procedures

A switch's Services and Capability Sets are reported to a controller in a Service Configuration message. A Service Configuration message response includes the list of Services defined for GSMP that the switch supports and, for each Service, a specification of the Capability Sets supported for the Service. Services are referred to by numbers standardised in the GSMP specification. Capability Sets are referred to by a numbering system reported by the switch. Each Capability Set within a given Service includes a unique identifying number together with the switch's specification of QoS Parameters and Traffic Controls.

A switch need not support all the defined Services and Capability Sets on every port. The supported Services and Capability Sets are reported to the controller on a per port basis in port configuration messages. Port configuration response messages list the supported

Services using the standardised identifying numbers and the Capability Sets by using the identifying numbers established in the switch Service configuration messages.

GSMP does not provide a negotiation mechanism by which a controller may establish or modify Capability Sets.

When a controller establishes a connection, the connection management message includes indication of the Service and the Capability Set. Depending on these the connection management message may additionally include Traffic Parameter values and Traffic Control flags.

A connection with a given Service can only be established if both the requested Service and the requested Capability Set are available on all of the connection's input and output ports.

Refresh of an extant connection is permitted but the add branch message requesting the message MUST NOT include indication of Service, Capability Sets or Traffic Parameters.

An extant connection's Traffic Parameters may be changed without first deleting the connection. The Service and Capability Sets of an extant connection cannot be changed.

Move branch messages may be refused on the grounds of resource depletion.

10.4 Service Definitions

This section sets forth the definition of Services. The following Service Identifiers are defined:

ID	Service Type
1	CBR= 1
2	rt-VBR.1
3	rt-VBR.2
4	rt-VBR.3
5	nrt-VBR.1
6	nrt-VBR.2
7	nrt-VBR.3
8	UBR.1
9	UBR.2
10-11	Reserved
12	GFR.1
13	GFR.2
14-19	Reserved
20	Int-Serv Controlled Load

21-24	Reserved
25	MPLS CR-LDP QoS
26-29	Reserved
30	Frame Relay Service
31-49	Reserved
50-69	Reserved GMPLS
70-65535	Reserved

Each Service will be defined in its own subsection. Each Service definition includes the following definitions:

Service Identifier

The reference number used to identify the Service in GSMP messages.

Service Characteristics

A definition of the Service.

Traffic Parameters

A definition of the Traffic Parameters used in connection management messages.

QoS Parameters

A definition of the QoS Parameters that are included in the Capability Set for instances of the Service.

Traffic Controls

A definition of the Traffic Controls that may be supported by an instance of the Service.

Descriptive text is avoided wherever possible in order to minimise any possibility of semantic conflict with the Original Specifications.

10.4.1 ATM Forum Service Categories

10.4.1.1 CBR

Service Identifier:

CBR.1 - Service ID = 1

Service Characteristics:

Equivalent to ATM Forum CBR.1 Service, see [8].

Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance

QoS Parameters:

- Cell Loss Ratio
- Maximum Cell Transfer Delay
- Peak-to-peak Cell Delay Variation

Traffic Controls:

- (U) Usage Parameter Control
- (I) Ingress Traffic Shaping to the Peak Cell Rate
- (E) Egress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance
- (D) Packet Discard

10.4.1.2 rt-VBR

Service Identifier:

- rt-VBR.1 - Service ID = 2
- rt-VBR.2 - Service ID = 3
- rt-VBR.3 - Service ID = 4

Service Characteristics:

Equivalent to ATM Forum rt-VBR Service, see [8].

Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance
- Sustainable Cell Rate
- Maximum Burst Size

QoS Parameters:

- Cell Loss Ratio
- Maximum Cell Transfer Delay
- Peak-to-peak Cell Delay Variation

Traffic Controls:

- (U) Usage Parameter Control
- (I) Ingress Traffic Shaping to the Peak Cell Rate
- (E) Egress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance
- (S) Egress Traffic Shaping to the Sustainable Cell Rate and Maximum Burst Size
- (P) Packet Discard
- (V) VC Merge

10.4.1.3 nrt-VBR

Service Identifier:

- nrt-VBR.1 - Service ID = 5
- nrt-VBR.2 - Service ID = 6
- nrt-VBR.3 - Service ID = 7

Service Characteristics:

Equivalent to ATM Forum nrt-VBR Service, see [8].

Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance
- Sustainable Cell Rate
- Maximum Burst Size

QoS Parameter:

- Cell Loss Ratio

Traffic Controls:

- (U) Usage Parameter Control
- (I) Ingress Traffic Shaping to the Peak Cell Rate
- (E) Egress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance
- (S) Egress Traffic Shaping to the Sustainable Cell Rate and Maximum Burst Size
- (P) Packet Discard
- (V) VC Merge

10.4.1.4 UBR

Service Identifier:

- UBR.1 - Service ID = 8
- UBR.2 - Service ID = 9

Service Characteristics:

Equivalent to ATM Forum UBR Service, see [8].

Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance

QoS Parameter:

None

Traffic Controls:

- (U) Usage Parameter Control
- (I) Ingress Traffic Shaping to the Peak Cell Rate

- (E) Egress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance
- (P) Packet Discard
- (V) VC Merge

10.4.1.5 ABR

ABR is not supported in this version of GSMP.

10.4.1.6 GFR

Service Identifier:

GFR.1 - Service ID = 12

GFR.2 - Service ID = 13

Service Characteristics:

Equivalent to ATM Forum GFR Service, see [8].

Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance
- Minimum Cell Rate
- Maximum Burst Size
- Maximum Frame Size

QoS Parameter:

- Cell Loss Ratio

Traffic Controls:

- (U) Usage Parameter Control
- (I) Ingress Traffic Shaping to the Peak Cell Rate
- (E) Egress Traffic Shaping to the Peak Cell Rate and Cell Delay Variation Tolerance
- (V) VC Merge

10.4.2 Integrated Services

10.4.2.1 Controlled Load

Service Identifier:

Int-Serv Controlled Load - Service ID = 20

Service Characteristics:

See [9].

Traffic Parameters:

- Token bucket rate (r)
- Token bucket depth (b)
- Peak rate (p)
- Minimum policed unit (m)
- Maximum packet size (M)

QoS Parameter:

None.

Traffic Controls:

None.

10.4.3 MPLS CR-LDP

Service Identifier:

MPLS CR-LDP QoS - Service ID = 25

Service Characteristics:

See [10].

Traffic Parameters:

- Peak Data Rate
- Peak Burst Size
- Committed Data Rate
- Committed Burst Size
- Excess Burst Size
- Weight

QoS Parameter:

- Frequency

Traffic Controls:

None currently defined.

10.4.4 Frame Relay

Service Identifier:

Frame Relay Service - Service ID = 30

Service Characteristics:

Equivalent to Frame Relay Bearer Service, see [11].

Traffic Parameters:

- Committed Information Rate
- Committed Burst Rate
- Excess Burst Rate

QoS Parameters:

None.

Traffic Controls:

- Usage Parameter Control
- Egress Traffic Shaping to the Committed Information Rate and Committed Burst Size

10.4.5 DiffServ

DiffServ is not supported in this version of GSMP.

10.5 Format and encoding of the Traffic Parameters

Connection management messages that use the GSMP Service Model (i.e., those that have IQS or OQS set to 0b10) include the Traffic Parameters Block that specifies the Traffic Parameter values of a connection. The required Traffic Parameters of a given Service are given in Section 10.4. The format and encoding of these parameters are given below.

10.5.1 Traffic Parameters for ATM Forum Services

The Traffic Parameters:

- Peak Cell Rate
- Cell Delay Variation Tolerance
- Sustainable Cell Rate
- Maximum Burst Size
- Minimum Cell Rate
- Maximum Frame Size

are defined in [8]. These Parameters are encoded as 24-bit unsigned integers. Peak Cell Rate, Sustainable Cell Rate, and Minimum Cell Rate are in units of cells per second. Cell Delay Variation Tolerance is in units of microseconds. Maximum Burst Size and Maximum Frame Size are in units of cells. In GSMP messages, the individual Traffic Parameters are encoded as follows:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|x x x x x x x x|                24 bit unsigned integer          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The format of the Traffic Parameters Block in connection management messages depends on the Service. It is a sequence of the 32 bit words (as shown above) corresponding to the Traffic Parameters as specified in the Service Definitions given in Section 10.4.1 in the order given there.

10.5.2 Traffic Parameters for Int-Serv Controlled Load Service

The Traffic Parameters:

- Token bucket rate (r)
- Token bucket size (b)
- Peak rate (p)

are defined in [9]. They are encoded as 32-bit IEEE single-precision floating point numbers. The Traffic Parameters Token bucket rate (r) and Peak rate (p) are in units of bytes per seconds. The Traffic Parameter Token bucket size (b) is in units of bytes.

The Traffic Parameters:

- Minimum policed unit (m)
- Maximum packet size (M)

are defined in [9]. They are encoded as 32 integer in units of bytes.

The Traffic Parameters Block for the Int-Serv Controlled Load Service is as follows:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Token bucket rate (r)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Token bucket size (b)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Peak rate (p)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Minimum policed unit (m)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Maximum packet size (M)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

10.5.3 Traffic Parameters for CRLDP Service

The Traffic Parameters:

- Peak Data Rate
- Peak Burst Size
- Committed Data Rate
- Committed Burst Size
- Excess Burst Size

are defined in [10] to be encoded as a 32-bit IEEE single-precision floating point number. A value of positive infinity is represented as an IEEE single-precision floating-point number with an exponent of all ones (255) and a sign and mantissa of all zeros. The values Peak Data Rate and Committed Data Rate are in units of bytes per second. The values Peak Burst Size, Committed Burst Size and Excess Burst Size are in units of bytes.

The Traffic Parameter

- Weight

is defined in [10] to be an 8-bit unsigned integer indicating the weight of the CRLSP. Valid weight values are from 1 to 255. The value 0 means that weight is not applicable for the CRLSP.

The Traffic Parameters Block for the CRLDP Service is as follows:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
										Peak Data Rate																													
										Peak Burst Size																													
										Committed Data Rate																													
										Committed Burst Size																													
										Excess Burst Size																													
										Weight																													

10.5.4 Traffic Parameters for Frame Relay Service

The Traffic Parameters:

- Committed Information Rate
- Committed Burst Size
- Excess Burst Size

are defined in [11]. Format and encoding of these parameters for frame relay signalling messages are defined in [12]. (Note that in [12] the Committed Information Rate is called "Throughput".) GSMP uses the encoding defined in [12] but uses a different format.

The format of the Traffic Parameters Block for Frame Relay Service is as follows:

0										1										2										3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																		
x x x x x x x x x x x x x x										Mag										x x x x x										CIR Multiplier																			
x x x x x x x x x x x x x x										Mag										x x										CBS Multiplier																			
x x x x x x x x x x x x x x										Mag										x x										EBS Multiplier																			

Mag

This field is an unsigned integer in the range from 0 to 6. The value 7 is not allowed. Mag is the decimal exponent for the adjacent multiplier field (which itself functions as a mantissa).

CIR Multiplier

This field is an unsigned integer. It functions as the mantissa of the Committed Information Rate Traffic Parameter.

CBS Multiplier**EBS Multiplier**

These fields are unsigned integers. They function as the mantissas of the Committed Burst Size and Excess Burst Size Traffic Parameters respectively.

The Traffic Parameter Values are related to their encoding in GSMP messages as follows:

Committed Information Rate = $10^{(\text{Mag})} * (\text{CIR Multiplier})$

Committed Burst Size = $10^{(\text{Mag})} * (\text{CBS Multiplier})$

Excess Burst Size = $10^{(\text{Mag})} * (\text{EBS Multiplier})$

10.6 Traffic Controls (TC) Flags

The TC Flags field in Add Branch messages for connections using the Service Model are set by the controller to indicate that specific traffic controls are requested for the requested connection. The TC Flags field is shown below:

```

  0 1 2 3 4 5 6 7
+-----+
|U|D|I|E|S|V|P|x|
+-----+
```

U: Usage Parameter Control

When set, this flag indicates that Usage Parameter Control is requested.

D: Packet Discard

When set, this flag indicates that Packet Discard is requested.

I: Ingress Shaping

When set, this flag indicates the availability of Ingress Traffic Shaping to the Peak Rate and Delay Variation Tolerance is requested.

E: Egress Shaping, Peak Rate

When set, this flag indicates that Egress Shaping to the Peak Rate and Delay Variation Tolerance is requested.

S: Egress Traffic Shaping, Sustainable Rate

When set, this flag indicates that Egress Traffic Shaping to the Sustainable Rate and Maximum Burst Size is requested.

V: VC Merge

When set, this flag indicates that ATM Virtual Channel Merge (i.e., multipoint to point ATM switching with a traffic control to avoid AAL5 PDU interleaving) is requested.

P: Port

When set indicates that traffic block pertains to Ingress Port.

x: Reserved

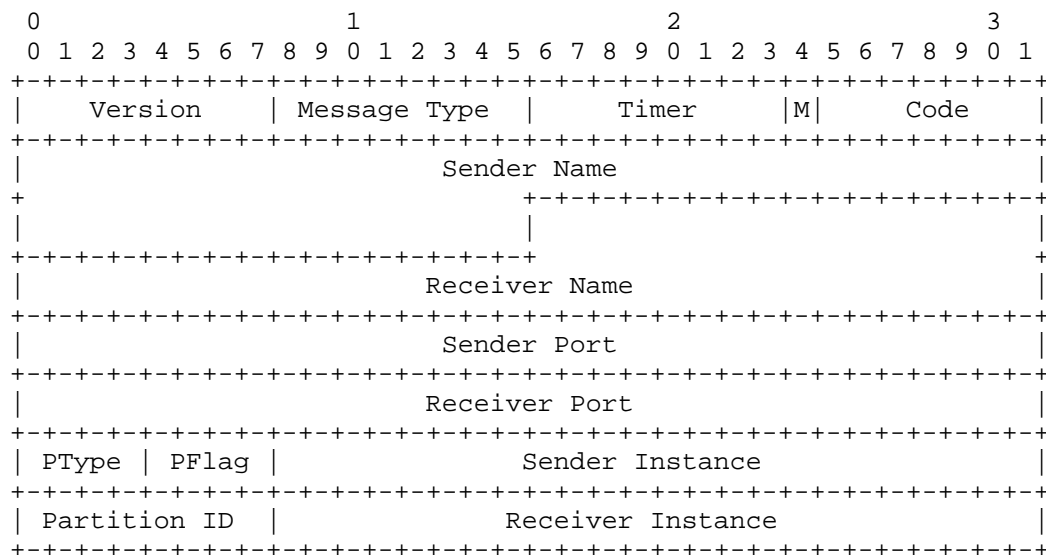
The controller may set (to one) the flag corresponding to the requested Traffic Control if the corresponding Traffic Control has been indicated in the Service Configuration response message (Section 8.4) as available for application to connections that use the requested Capability Set on a per connection basis. (The requested Capability Set is indicated by the Capability Set ID the least significant byte of the Service Selector field of the Add Branch message.) If the Traffic Control has been indicated in the Service Configuration response message as either not available in the Capability Set or applied to all connections that use the Capability Set then the controller sets the flag to zero and the switch ignores the flag.

11. Adjacency Protocol

The adjacency protocol is used to synchronise state across the link, to agree on which version of the protocol to use, to discover the identity of the entity at the other end of a link, and to detect when it changes. GSMP is a hard state protocol. It is therefore important to detect loss of contact between switch and controller, and to detect any change of identity of switch or controller. No GSMP messages other than those of the adjacency protocol may be sent across the link until the adjacency protocol has achieved synchronisation.

11.1 Packet Format

All GSMP messages belonging to the adjacency protocol have the following structure:



Version

In the adjacency protocol the Version field is used for version negotiation. The version negotiation is performed before synchronisation is achieved. In a SYN message the Version field always contains the highest version understood by the sender. A receiver receiving a SYN message with a version higher than understood will ignore that message. A receiver receiving a SYN message with a version lower than its own highest version, but a version that it understands, will reply with a SYNACK with the version from the received SYN in its GSMP Version field. This defines the version of the GSMP protocol to be used while the adjacency protocol remains synchronised. All other messages will use the agreed version in the Version field.

The version number for the version of the GSMP protocol defined by this specification is Version = 3.

Message Type

The adjacency protocol is:

Message Type = 10

Timer

The Timer field is used to inform the receiver of the timer value used in the adjacency protocol of the sender. The timer specifies the nominal time between periodic adjacency protocol messages. It is a constant for the duration of a GSMP session. The timer field is specified in units of 100ms.

M-Flag

The M-Flag is used in the SYN message to indicate whether the sender is a master or a slave. If the M-Flag is set in the SYN message, the sender is a master. If zero, the sender is a slave. The GSMP protocol is asymmetric, the controller being the master and the switch being the slave. The M-Flag prevents a master from synchronising with another master, or a slave with another slave. If a slave receives a SYN message with a zero M-Flag, it MUST ignore that SYN message. If a master receives a SYN message with the M-Flag set, it MUST ignore that SYN message. In all other messages the M-Flag is not used.

Code

Field specifies the function of the message. Four Codes are defined for the adjacency protocol:

```
SYN:      Code = 1
SYNACK:   Code = 2
ACK:      Code = 3
RSTACK:   Code = 4.
```

Sender Name

For the SYN, SYNACK, and ACK messages, is the name of the entity sending the message. The Sender Name is a 48-bit quantity that is unique within the operational context of the device. A 48-bit IEEE 802 MAC address, if available, may be used for the Sender Name. If the Ethernet encapsulation is used the Sender Name MUST be the Source Address from the MAC header. For the RSTACK message, the Sender Name field is set to the value of the Receiver Name field from the incoming message that caused the RSTACK message to be generated.

Receiver Name

For the SYN, SYNACK, and ACK messages, is the name of the entity that the sender of the message believes is at the far end of the link. If the sender of the message does not know the name of the entity at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Name field is set to the value of the Sender Name field from the incoming message that caused the RSTACK message to be generated.

Sender Port

For the SYN, SYNACK, and ACK messages, is the local port number of the link across which the message is being sent. For the RSTACK message, the Sender Port field is set to the value of the Receiver Port field from the incoming message that caused the RSTACK message to be generated.

Receiver Port

For the SYN, SYNACK, and ACK messages, is what the sender believes is the local port number for the link, allocated by the entity at the far end of the link. If the sender of the message does not know the port number at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Port field is set to the value of the Sender Port field from the incoming message that caused the RSTACK message to be generated.

PType

PType is used to specify if partitions are used and how the Partition ID is negotiated.

Type of partition being requested.

- 0 No Partition
- 1 Fixed Partition Request
- 2 Fixed Partition Assigned

PFlag

Used to indicate the type of partition request.

- 1 - New Adjacency.
In the case of a new adjacency, the state of the switch will be reset.
- 2 - Recovered Adjacency.
In the case of a recovered adjacency, the state of the switch will remain, and the Switch Controller will be responsible for confirming that the state of the switch matches the desired state.

Sender Instance

For the SYN, SYNACK, and ACK messages, is the sender's instance number for the link. It is used to detect when the link comes back up after going down or when the identity of the entity at the other end of the link changes. The instance number is a 24-bit number that is guaranteed to be unique within the recent past and to change when the link or node comes back up after going down. Zero is not a valid instance number. For the RSTACK message, the Sender Instance field is set to the value

of the Receiver Instance field from the incoming message that caused the RSTACK message to be generated.

Partition ID

Field used to associate the message with a specific switch partition.

Receiver Instance

For the SYN, SYNACK, and ACK messages, is what the sender believes is the current instance number for the link, allocated by the entity at the far end of the link. If the sender of the message does not know the current instance number at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Instance field is set to the value of the Sender Instance field from the incoming message that caused the RSTACK message to be generated.

11.2 Procedure

The adjacency protocol is described by the following rules and state tables.

The rules and state tables use the following operations:

- o The "Update Peer Verifier" operation is defined as storing the values of the Sender Instance, Sender Port, Sender Name and Partition ID fields from a SYN or SYNACK message received from the entity at the far end of the link.
- o The procedure "Reset the link" is defined as:
 1. Generate a new instance number for the link
 2. Delete the peer verifier (set to zero the values of Sender Instance, Sender Port, and Sender Name previously stored by the Update Peer Verifier operation)
 3. Send a SYN message
 4. Enter the SYNSENT state.
- o The state tables use the following Boolean terms and operators:
 - A The Sender Instance in the incoming message matches the value stored from a previous message by the "Update Peer Verifier" operation.
 - B The Sender Instance, Sender Port, Sender Name and Partition ID fields in the incoming message match the values stored from a previous message by the "Update Peer Verifier" operation.

- C The Receiver Instance, Receiver Port, Receiver Name and Partition ID fields in the incoming message match the values of the Sender Instance, Sender Port, Sender Name and Partition ID currently sent in outgoing SYN, SYNACK, and ACK messages.

"&&" Represents the logical AND operation

"||" Represents the logical OR operation

"!" Represents the logical negation (NOT) operation.

- o A timer is required for the periodic generation of SYN, SYNACK, and ACK messages. The value of the timer is announced in the Timer field. The period of the timer is unspecified but a value of one second is suggested.

There are two independent events: the timer expires, and a packet arrives. The processing rules for these events are:

Timer Expires: Reset Timer
 If state = SYNSENT Send SYN
 If state = SYNRCVD Send SYNACK
 If state = ESTAB Send ACK

Packet Arrives:
 If incoming message is an RSTACK:
 If (A && C && !SYNSENT) Reset the link
 Else discard the message.
 If incoming message is a SYN, SYNACK, or ACK:
 Response defined by the following State Tables.
 If incoming message is any other GSMP message and
 state != ESTAB:
 Discard incoming message.
 If state = SYNSENT Send SYN (Note 1)
 If state = SYNRCVD Send SYNACK (Note 1)

Note 1: No more than two SYN or SYNACK messages should be sent within any time period of length defined by the timer.

- o State synchronisation across a link is considered to be achieved when the protocol reaches the ESTAB state. All GSMP messages, other than adjacency protocol messages, that are received before synchronisation is achieved, will be discarded.

11.2.1 State Tables

State: SYNSENT

Condition	Action	New State
SYNACK && C	Update Peer Verifier; Send ACK	ESTAB
SYNACK && !C	Send RSTACK	SYNSENT
SYN	Update Peer Verifier; Send SYNACK	SYNRCVD
ACK	Send RSTACK	SYNSENT

State: SYNRCVD

Condition	Action	New State
SYNACK && C	Update Peer Verifier; Send ACK	ESTAB
SYNACK && !C	Send RSTACK	SYNRCVD
SYN	Update Peer Verifier; Send SYNACK	SYNRCVD
ACK && B && C	Send ACK	ESTAB
ACK && !(B && C)	Send RSTACK	SYNRCVD

State: ESTAB

Condition	Action	New State
SYN SYNACK	Send ACK (note 2)	ESTAB
ACK && B && C	Send ACK (note 3)	ESTAB
ACK && !(B && C)	Send RSTACK	ESTAB

Note 2: No more than two ACKs should be sent within any time period of length defined by the timer. Thus, one ACK MUST be sent every time the timer expires. In addition, one further ACK may be sent between timer expirations if the incoming message is a SYN or SYNACK. This additional ACK allows the adjacency protocol to reach synchronisation more quickly.

Note 3: No more than one ACK should be sent within any time period of length defined by the timer.

11.3 Partition Information State

Each instance of a [switch controller-switch partition] pair will need to establish adjacency synchronisation independently.

Part of the process of establishing synchronisation when using partition will be to establish the assignment of partition identifiers. The following scenarios are provided for:

- A controller can request a specific partition ID by setting the PType to Fixed Partition Request.
- A controller can let the switch decide whether it wants to assign a fixed partition ID or not, by setting the PType to No Partition.
- A switch can assign the specific Partition ID to the session by setting the PType to Fixed Partition Assigned. A switch can specify that no partitions are handled in the session by setting the PType to No Partition.

The assignment is determined by the following behaviour:

- An adjacency message from a controller with PType = 1 and Code = SYN SHOULD be treated as a partition request.
- An adjacency message from a switch with PType = 2 and Code = SYN SHOULD be treated as a partition assignment.
- An adjacency message from a controller or a switch with PType = 2 and Code = (SYNACK || ACK) SHOULD be treated as a success response, the partition is assigned.
- An adjacency message from a controller with PType = 0 and Code = SYN indicates that the controller has not specified if it requests partitions or not.

- An adjacency message from a switch with PType = 0 and Code = SYN indicates that the switch does not support partitions.
- An adjacency message from a controller or a switch with PType = 0 and Code = (SYNACK || ACK) indicates that the session does not support partitions.
- An adjacency message from a controller or a switch with PType = (1 || 2) and Code = RSTACK indicates that requested Partition ID is unavailable.
- An adjacency message from a controller or a switch with PType = 0 and Code = RSTACK indicates that an unidentified error has occurred. The session SHOULD be reset.

All other combinations of PType and Code are undefined in this version of GSMP.

11.4 Loss of Synchronisation

If after synchronisation is achieved, no valid GSMP messages are received in any period of time in excess of three times the value of the Timer field announced in the incoming adjacency protocol messages, loss of synchronisation may be declared.

While re-establishing synchronisation with a controller, a switch SHOULD maintain its connection state, deferring the decision about resetting the state until after synchronisation is re-established.

Once synchronisation is re-established the decision about resetting the connection state SHOULD be made on the following basis:

- If PFLAG = 1, then a new adjacency has been established and the state SHOULD be reset
- If PFLAG = 2, then adjacency has been re-established and the connection state SHOULD be retained. Verification that controller and connection state are the same is the responsibility of the controller.

11.5 Multiple Controllers per switch partition

Multiple switch controllers may jointly control a single switch partition. The controllers may control a switch partition either in a primary/standby fashion or as part of multiple controllers providing load-sharing for the same partition. It is the responsibility of the controllers to co-ordinate their interactions

with the switch partition. In order to assist the controllers in tracking multiple controller adjacencies to a single switch partition, the Adjacency Update message is used to inform a controller that there are other controllers interacting with the same partition. It should be noted that the GSMP does not include features that allow the switch to co-ordinate cache synchronization information among controllers. The switch partition will service each command it receives in turn as if it were interacting with a single controller. Controller implementations without controller entity synchronisation SHOULD NOT use multiple controllers with a single switch partition.

11.5.1 Multiple Controller Adjacency Process

The first adjacency for a specific partition is determined by the procedures described in section 11.2 and an Adjacency Update message will be sent. The next adjacencies to the partition are identified by a new partition request with the same Partition ID as the first one but with the different Sender Name. Upon establishing adjacency the Adjacency count will be increased and an Adjacency Update message will be sent.

When adjacency between one partition and a controller is lost, the adjacency count will be decremented and an Adjacency Update message will be sent.

Example:

A switch partition has never been used. When the first controller (A) achieves adjacency, an adjacency count will be initiated and (A) will get an Adjacency Update message about itself with Code field = 1. Since (A) receives an adjacency count of 1 this indicates that it is the only controller for that partition.

When a second adjacency (B), using the same Partition ID, achieves adjacency, the adjacency counter will be increased by 1. Both (A) and (B) will receive an Adjacency Update message indicating an adjacency count of 2 in the Code field. Since the count is greater than 1, this will indicate to both (A) and (B) that there is another controller interacting with the switch; identification of the other controller will not be provided by GSMP, but will be the responsibility of the controllers.

If (A) loses adjacency, the adjacency count will be decreased and an Adjacency Update message will be sent to (B) indicating an adjacency count of 1 in the Code field. If (B) leaves as well, the partition is regarded as idle and the adjacency count may be reset.

12. Failure Response Codes

12.1 Description of Failure and Warning Response Messages

A failure response message is formed by returning the request message that caused the failure with the Result field in the header indicating failure (Result = 4) and the Code field giving the failure code. The failure code specifies the reason for the switch being unable to satisfy the request message.

A warning response message is a success response (Result = 3) with the Code field specifying the warning code. The warning code specifies a warning that was generated during the successful operation.

If the switch issues a failure response in reply to a request message, no change should be made to the state of the switch as a result of the message causing the failure. (For request messages that contain multiple requests, such as the Delete Branches message, the failure response message will specify which requests were successful and which failed. The successful requests may result in a changed state.)

If the switch issues a failure response it MUST choose the most specific failure code according to the following precedence:

- Invalid Message
- General Message Failure
- Specific Message Failure A failure response specified in the text defining the message type.
- Connection Failures
- Virtual Path Connection Failures
- Multicast Failures
- QoS Failures
- General Failures
- Warnings

If multiple failures match in any of the following categories, the one that is listed first should be returned. The following failure response messages and failure and warning codes are defined:

Invalid Message

- 3: The specified request is not implemented on this switch.
The Message Type field specifies a message that is not implemented on the switch or contains a value that is not defined in the version of the protocol running in this session of GSMP.
- 4: One or more of the specified ports does not exist.
At least one of the ports specified in the message is invalid. A port is invalid if it does not exist or if it has been removed from the switch.
- 5: Invalid Port Session Number.
The value given in the Port Session Number field does not match the current Port Session Number for the specified port.
- 7: Invalid Partition ID
The value given in the Partition ID field is not legal for this partition.

General Message Failure

- 10: The meaning of this failure is dependent upon the particular message type and is specified in the text defining the message.

Specific Message Failure - A failure response that is only used by a specific message type

- Failure response messages used by the Label Range message
 - 40: Cannot support one or more requested label ranges.
 - 41: Cannot support disjoint label ranges.
 - 42: Specialised multipoint labels not supported.
- Failure response messages used by the Set Transmit Data Rate function of the Port Management message
 - 43: The transmit data rate of this output port cannot be changed.

44: Requested transmit data rate out of range for this output port.

The transmit data rate of the requested output port can be changed, but the value of the Transmit Data Rate field is beyond the range of acceptable values.

- Failure response message of the Port Management message

45: Connection Replace mechanism not supported on switch.

The R-flag SHOULD be reset in the Response Port Management message.

- Failure response message range reserved for the ARM extension

128-159: These failure response codes will be interpreted according to definitions provided by the model description.

Connection Failures

11: The specified connection does not exist.

An operation that expects a connection to be specified cannot locate the specified connection. A connection is specified by the input port and input label on which it originates. An ATM virtual path connection is specified by the input port and input VPI on which it originates.

12: The specified branch does not exist.

An operation that expects a branch of an existing connection to be specified cannot locate the specified branch. A branch of a connection is specified by the connection it belongs to and the output port and output label on which it departs. A branch of an ATM virtual path connection is specified by the virtual path connection it belongs to and the output port and output VPI on which it departs.

13: One or more of the specified Input Labels is invalid.

14: One or more of the specified Output Labels is invalid.

15: Point-to-point bi-directional connection already exists.

The connection specified by the Input Port and Input Label fields already exists, and the bi-directional Flag in the Flags field is set.

- 16: Invalid Service Selector field in a Connection Management message. The value of the Service Selector field is invalid.
- 17: Insufficient resources for QoS Profile.
The resources requested by the QoS Profile in the Service Selector field are not available.
- 18: Insufficient Resources.
Switch resources needed to establish a branch are not available.
- 20: Reservation ID out of Range
The numerical value of Reservation ID is greater than the value of Max Reservations (from the Switch Configuration message).
- 21: Mismatched reservation ports
The value of Input Port differs from the input port specified in the reservation or the value of Output Port differs from the output port specified in the reservation.
- 22: Reservation ID in use
The value of Reservation ID matches that of an extant Reservation.
- 23: Non-existent reservation ID
No reservation corresponding to Reservation ID exists.
- 36: Replace of connection is not activated on switch.
Only applicable for Add Branch messages. The Replace Connection mechanism has not been activated on port by the Port Management message.
- 37: Connection replacement mode cannot be combined with Bi-directional or Multicast mode. The R flag MUST NOT be used in conjunction with either the M flag or the B flag.

ATM Virtual Path Connections

- 24: ATM virtual path switching is not supported on this input port.

- 25: Point-to-multipoint ATM virtual path connections are not supported on either the requested input port or the requested output port.
One or both of the requested input and output ports is unable to support point-to-multipoint ATM virtual path connections.
- 26: Attempt to add an ATM virtual path connection branch to an existing virtual channel connection.
It is invalid to mix branches switched as virtual channel connections with branches switched as ATM virtual path connections on the same point-to-multipoint connection.
- 27: Attempt to add an ATM virtual channel connection branch to an existing ATM virtual path connection.
It is invalid to mix branches switched as virtual channel connections with branches switched as ATM virtual path connections on the same point-to-multipoint connection.
- 28: ATM Virtual path switching is not supported on non-ATM ports.
One or both of the requested input and output ports is not an ATM port. ATM virtual path switching is only supported on ATM ports.

Multicast Failures

- 29: A branch belonging to the specified point-to-multipoint connection is already established on the specified output port and the switch cannot support more than a single branch of any point-to-multipoint connection on the same output port.
- 30: The limit on the maximum number of multicast connections that the switch can support has been reached.
- 31: The limit on the maximum number of branches that the specified multicast connection can support has been reached.
- 32: Cannot label each output branch of a point-to-multipoint tree with a different label.
Some switch designs, require all output branches of a point-to-multipoint connection to use the same value of Label.
- 33: Cannot add multi-point branch to bi-directional connection.
It is an error to attempt to add an additional branch to an existing connection with the bi-directional flag set.

34: Unable to assign the requested Label value to the requested branch on the specified multicast connection. Although the requested Labels are valid, the switch is unable to support the request using the specified Label values for some reason not covered by the above failure responses. This message implies that a valid value of Labels exists that the switch could support. For example, some switch designs restrict the number of distinct Label values available to a multicast connection. (Most switch designs will not require this message.)

35: General problem related to the manner in which multicast is supported by the switch. Use this message if none of the more specific multicast failure messages apply. (Most switch designs will not require this message.)

QoS Failures

60-79: These failure response codes will be interpreted according to definitions provided by the model description.

80: Switch does not support different QoS parameters for different branches within a multipoint connection.

General Failures

2: Invalid request message.
There is an error in one of the fields of the message not covered by a more specific failure message.

6: One or more of the specified ports is down.
A port is down if its Port Status is Unavailable. Connection Management, Connection State, Port Management, and Configuration operations are permitted on a port that is Unavailable. Connection Activity and Statistics operations are not permitted on a port that is Unavailable and will generate this failure response. A Port Management message specifying a Take Down function on a port already in the Unavailable state will also generate this failure response.

19: Out of resources.
The switch has exhausted a resource not covered by a more specific failure message, for example, running out of memory.

- 1: Unspecified reason not covered by other failure codes.
The failure message of last resort.

Warnings

- 46: One or more labels are still used in the previous Label Range.

12.2 Summary of Failure Response Codes and Warnings

The following list gives a summary of the failure codes defined for failure response messages:

- 1: Unspecified reason not covered by other failure codes.
- 2: Invalid request message.
- 3: The specified request is not implemented on this switch.
- 4: One or more of the specified ports does not exist.
- 5: Invalid Port Session Number.
- 6: One or more of the specified ports is down.
- 7: Invalid Partition ID.
- 10: General message failure. (The meaning of this failure code depends upon the Message Type. It is defined within the description of any message that uses it.)
- 11: The specified connection does not exist.
- 12: The specified branch does not exist.
- 13: One or more of the specified Input Labels is invalid.
- 14: One or more of the specified Output Labels is invalid.
- 15: Point-to-point bi-directional connection already exists.
- 16: Invalid service selector field in a connection management message.
- 17: Insufficient resources for QoS profile.
- 18: Insufficient resources.
- 19: Out of resources (e.g., memory exhausted, etc.).
- 20: Reservation ID out of Range
- 21: Mismatched reservation ports
- 22: Reservation ID in use
- 23: Non-existent reservation ID
- 24: ATM virtual path switching is not supported on this input port.
- 25: Point-to-multipoint ATM virtual path connections are not supported on either the requested input port or the requested output port.
- 26: Attempt to add an ATM virtual path connection branch to an existing virtual channel connection.
- 27: Attempt to add an ATM virtual channel connection branch to an existing virtual path connection.
- 28: ATM Virtual Path switching is not supported on non-ATM ports.

- 29: A branch belonging to the specified point-to-multipoint connection is already established on the specified output port and the switch cannot support more than a single branch of any point-to-multipoint connection on the same output port.
- 30: The limit on the maximum number of point-to-multipoint connections that the switch can support has been reached.
- 31: The limit on the maximum number of branches that the specified point-to-multipoint connection can support has been reached.
- 32: Cannot label each output branch of a point-to-multipoint tree with a different label.
- 33: Cannot add multi-point branch to bi-directional connection.
- 34: Unable to assign the requested Label value to the requested branch on the specified point-to-multipoint connection.
- 35: General problem related to the manner in which point-to-multipoint is supported by the switch.
- 36: Replace of connection is not activated on switch.
- 37: Connection replacement mode cannot be combined with Bi-directional or Multicast mode.
- 40: Cannot support one or more requested label ranges.
- 41: Cannot support disjoint label ranges.
- 42: Specialised multipoint labels not supported.
- 43: The transmit data rate of this output port cannot be changed.
- 44: Requested transmit data rate out of range for this output port.
- 45: Connection Replace mechanism not supported on switch.
- 46: Labels are still used in the existing Label Range.
- 60-79: Reserved for QoS failures.
- 80: Switch does not support different QoS parameters for different branches within a multipoint connection.
- 128-159: Reserved for the ARM extensions.

13. Security Considerations

The security of GSMP's TCP/IP control channel has been addressed in [15]. 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 [15].

Appendix A Summary of Messages

Message Name	Message Number	Status
Connection Management Messages		
Add Branch	16	
ATM Specific - VPC.....	26	
Delete Tree.....	18	
Verify Tree.....	19	Obsoleted
Delete All Input.....	20	
Delete All Output.....	21	
Delete Branches.....	17	
Move Output Branch.....	22	
ATM Specific - VPC.....	27	
Move Input Branch.....	23	
ATM Specific - VPC.....	28	
Port Management Messages		
Port Management.....	32	
Label Range.....	33	
State and Statistics Messages		
Connection Activity.....	48	
Port Statistics.....	49	
Connection Statistics.....	50	
QoS Class Statistics.....	51	Reserved
Report Connection State.....	52	
Configuration Messages		
Switch Configuration.....	64	
Port Configuration.....	65	
All Ports Configuration.....	66	
Service Configuration.....	67	
Reservation Messages		
Reservation Request.....	70	
Delete Reservation.....	71	
Delete All Reservations.....	72	
Event Messages		
Port Up.....	80	
Port Down.....	81	
Invalid Label.....	82	
New Port.....	83	
Dead Port.....	84	

Abstract and Resource Model Extension Messages
 Reserved.....200-249

Adjacency Protocol.....10 Required

Appendix B IANA Considerations

Following the policies outlined in "Guidelines for Writing an IANA Considerations Section in RFCs" (RFC 2434 [19]), the following name spaces are defined in GSMPv3.

- Message Type Name Space [Appendix A]
- Label Type Name Space [3.1.3]
- Result Name Space [3.1.1]
- Failure Response Message Name Space [3.1.4],[11]
- Adaptation Type Name Space [4.1]
- Model Type Name Space [8.1]
- Port Type Name Space [8.2]
- Service ID Name Space [10.4]
- Traffic Control Name Space [8.4]
- Event Flag Name Space [6.1]

B.1. Message Type Name Space

GSMPv3 divides the name space for Message Types into four ranges. The following are the guidelines for managing these ranges.

- Message Types 0-99.
 Message Types in this range are part of the GSMPv3 base protocol. Message types in this range are allocated through an IETF consensus action [19].
- Message Types 100-199.
 Message Types in this range are Specification Required [19]. Message Types using this range must be documented in an RFC or other permanent and readily available references.

- Message Types 200-249.
Message Types in this range are Specification Required [19] and are intended for Abstract and Resource Model Extension Messages. Message Types using this range must be documented in an RFC or other permanent and readily available references.
- Message Types 250-255.
Message Types in this range are reserved for vendor private extensions and are the responsibility of individual vendors. IANA management of this range of the Message Type Name Space is unnecessary.

B.2. Label Type Name Space

GSMPv3 divides the name space for Label Types into three ranges. The following are the guidelines for managing these ranges.

- Label Types 0x000-0xAFF.
Label Types in this range are part of the GSMPv3 base protocol. Label Types in this range are allocated through an IETF consensus action [19].
- Label Types 0xB00-0xEFF.
Label Types in this range are Specification Required [19]. Label Types using this range must be documented in an RFC or other permanent and readily available reference.
- Label Types 0xF00-0xFFF.
Label Types in this range are reserved for vendor private extensions and are the responsibility of individual vendors. IANA management of this range of the Label Type Name Space is unnecessary.

B.3. Result Name Space

The following is the guideline for managing the Result Name Space:

- Result values 0-255.
Result values in this range need an expert review, i.e., approval by a Designated Expert is required [19].

B.4. Failure Response Name Space

GSMPv3 divides the name space for Failure Responses into three ranges. The following are the guidelines for managing these ranges:

- Failure Responses 0-59, 80-127, 160-255.
Failure responses in these ranges are part of the GSMPv3 base protocol. Failure Responses in these ranges are allocated through an IETF consensus action [19].
- Failure Responses 60-79, 128-159.
Failure responses in these ranges are reserved for vendor private extensions and are the responsibility of individual vendors. IANA management of these ranges of the Failure Response Name Space are unnecessary.

B.5. Adaptation Type Name Space

GSMPv3 divides the name space for Adaptation Types into two ranges. The following are the guidelines for managing these ranges:

- Adaptation Type 0x000-0x2FF.
Adaptation Types in this range are part of the GSMPv3 base protocol. Adaptation Types in this range are allocated through an IETF consensus action [19].
- Adaptation Type 0x300-0xFFFF.
Adaptation Types in this range are allocated by the first come first served principle [19].

B.6. Model Type Name Space

GSMPv3 divides the name space for Model Types into three ranges. The following are the guidelines for managing these ranges:

- Model Type 0.
Model Types in this range are part of the GSMPv3 base protocol. Model Types in this range are allocated through an IETF consensus action [19].
- Model Type 1-200.
Model Types in this range are Specification Required [19]. Message Types using this range must be documented in an RFC or other permanent and readily available references.
- Model Type 201-255.
Model Types in this range are reserved for vendor private extensions and are the responsibility of individual vendors. IANA management of these ranges of the Model Type Name Space are unnecessary.

B.7. Port Type Name Space

GSMPv3 divides the name space for Port Types into two ranges. The following are the guidelines for managing these ranges:

- Port Type 0-127.
Port Types in this range are part of the GSMPv3 base protocol. Port Types in this range are allocated through an IETF consensus action [19].
- Port Type 128-255.
Port Types in this range are Specification Required [19]. Port Types using this range must be documented in an RFC or other permanent and readily available references.

B.8. Service ID Name Space

GSMPv3 divides the name space for Service IDs into two ranges. The following are the guidelines for managing these ranges:

- Service ID 0-1023.
Service ID's in this range are part of the GSMPv3 base protocol. Service ID's in this range are allocated through an IETF consensus action [19].
- Service ID 1024-65535.
Service ID's in this range are Specification Required [19]. Service ID's using this range must be documented in an RFC or other permanent and readily available references.

B.9. Traffic Control Name Space

The following are the guidelines for managing Traffic Control Flags in GSMPv3:

- All Traffic Control Flags are allocated through an expert review, i.e., approval by a Designated Expert [19].

B.10. Event Flag Name Space

The following are the guidelines for managing Event Flags in GSMPv3:

- All Event Flags are allocated through an expert review, i.e., approval by a Designated Expert [19].

The TCP port for establishing GSMP connections has been defined as 6068.

References

- [1] "B-ISDN ATM Layer Specification", International Telecommunication Union, ITU-T Recommendation I.361, Feb. 1999.
- [2] "B-ISDN ATM Adaptation Layer (AAL) Specification", International Telecommunication Union, ITU-T Recommendation I.363, Mar. 1993.
- [3] "B-ISDN ATM Adaptation Layer specification: Type 5 AAL", International Telecommunication Union, ITU-T, Recommendation I.363.5, Aug. 1996.
- [4] Sjostrand, H., Buerkle, J. and B. Srinivasan, "Definitions of Managed Objects for the General Switch Management Protocol (GSMP)", RFC 3295, June 2002.
- [5] IANA Assigned Port Numbers, <http://www.iana.org>
- [6] Newman, P., Edwards, W., Hinden, R., Hoffman, E. Ching Liaw, F., Lyon, T. and G. Minshall, "Ipsilon's General Switch Management Protocol Specification Version 1.1", RFC 1987, August 1996.
- [7] Newman, P., Edwards, W., Hinden, R., Hoffman, E., Ching Liaw, F., Lyon, T. and G. Minshall, "Ipsilon's General Switch Management Protocol Specification Version 2.0", RFC 2297, March 1998.
- [8] ATM Forum Technical Committee, "Traffic Management Specification Version 4.1", af-tm-0121.000, 1999.
- [9] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997.
- [10] Jamoussi, B., Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heinanen, J., Kilty, T. and A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, January 2002.
- [11] ITU-T Recommendation I.233 Frame Mode Bearer Services, ISDN frame relaying bearer services and ISDN switching bearer service, Nov. 1991.
- [12] ITU-T Recommendation Q.933, Integrated Services Digital Network (ISDN) Digital Subscriber Signaling System No. 1 (DSS 1) Signaling Specifications For Frame Mode Switched And Permanent Virtual Connection Control And Status Monitoring, 1995.

- [13] ITU-T Recommendation Q.922, Integrated Services Digital Network (ISDN) Data Link Layer Specification For Frame Mode Bearer Services, 1992
- [14] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T. and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [15] Worster, T., Doria, A. and J. Buerkle, "General Switch Management Protocol (GSMP) Packet Encapsulations for Asynchronous Transfer Mode (ATM), Ethernet and Transmission Control Protocol (TCP)", RFC 3293, June 2002.
- [16] Doria, A. and K. Sundell, "General Switch Management Protocol Applicability", RFC 3294, June 2002.
- [17] IANAifType - MIB DEFINITIONS, <http://www.iana.org>, January 2001.
- [18] Anderson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [19] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [20] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [21] Conta, A., Doolan, P. and A. Malis, "Use of Label Switching on Frame Relay Networks Specification", RFC 3034, January 2001.

Authors' Addresses

Avri Doria
Div. of Computer Communications
Lulea University of Technology
S-971 87 Lulea
Sweden

Phone: +1 401 663 5024
EMail: avri@acm.org

Fiffi Hellstrand
Nortel Networks AB
S:t Eriksgatan 115 A
SE-113 85 Stockholm Sweden

EMail: fiffi@nortelnetworks.com

Kenneth Sundell
Nortel Networks AB
S:t Eriksgatan 115 A
SE-113 85 Stockholm Sweden

EMail: ksundell@nortelnetworks.com

Tom Worster

Phone: +1 617 247 2624
EMail: fsb@thefsb.org

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.

