

General Switch Management Protocol

<[draft-ietf-gsmp-03.txt](#)>

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

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.

Abstract

This memo provides the third draft of the standards track specification of GSMP. It is a revision of [draft-worster-gsmp-00](#) which itself was based on GSMP V2 [7].

Table of Contents

1. Introduction	4
2. GSMP Packet Encapsulation	7
3. Common Definitions and Procedures	8
3.1 GSMP Packet Format	9
3.2 Failure Response Messages	12
4. Connection Management Messages	18
4.1 General Message Definitions	18
4.2 Add Branch Message	27
4.3 Delete Tree Message	29
4.4 Verify Tree Message	29
4.5 Delete All Input Port Message	29
4.6 Delete All Output Port Message	30
4.7 Delete Branches Message	31
4.8 Move Branch Message	33
5. Port Management Messages	37
5.1 Port Management Message	37
5.2 Label Range Message	42
6. State and Statistics Messages	49
6.1 Connection Activity Message	49
6.2 Statistics Messages	52
6.2.1 Port Statistics Message	56
6.2.2 Connection Statistics Message	56
6.2.3 QoS Class Statistics Message	56
6.3 Report Connection State Message	56
7. Configuration Messages	63
7.1 Switch Configuration Message	63
7.1.1 Configuration Message Processing	65
7.2 Port Configuration Message	65
7.2.1 PortType Specific Data	68
7.3 All Ports Configuration Message	75
7.4 Service Configuration Message	77
8. Event Messages	82
8.1 Port Up Message	83
8.2 Port Down Message	83
8.3 Invalid Label Message	84
8.4 New Port Message	84
8.5 Dead Port Message	84
8.6 Adjacency Update Message	84
9. Service Model Definition	85

9.1	Overview	85
9.2	Service Model Definitions	85
9.2.1	Original Specifications	86
9.2.2	Service Definition, Traffic Parameters, QoS Parameters and Traffic Controls	86
9.2.3	Capability Sets	87
9.3	Service Model Procedures	87
9.4	Service Definitions	89
9.4.1	ATM Forum Service Categories	89
9.4.2	Integrated Services	94
9.4.3	MPLS CR-LDP	94
9.4.4	Frame Relay	95
9.4.5	Diff-Serv	95
9.5	Format and encoding of the Traffic Parameters Block in connection management messages	96
9.5.1	Traffic Parameters for ATM Forum Services	96
9.5.2	Traffic Parameters for the Int-Serv Controlled Load Service	96
9.5.3	Traffic Parameters for the CRLDP Service	97
9.5.4	Traffic Parameters for the Frame Relay Service	98
9.6	Traffic Controls (TC) Flags	99
10.	Reservation Management Messages	102
10.1	Reservation Request Message	102
10.2	Delete Reservation Message	104
10.3	Delete All Reservations Message	105
11.	Adjacency Protocol	106
11.1	Packet Format	106
11.2	Procedure	109
11.3	Partition Information State	112
11.4	Loss of Synchronisation	113
11.5	Multiple Controllers per switch partition	113
11.5.1	Multiple Controller Adjacency Process	114
12.	Summary of Failure Response Codes	115
13.	Summary of Message Set	117
14.	Security Considerations	119

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; 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 which 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, unpartitioned 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 is 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 which 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.

GSMP may be transported in three ways:

- GSMP may run across an ATM link connecting the controller to the switch, on a control connection (virtual channel) established at initialisation.
- GSMP operation across an Ethernet link is specified.
- GSMP operation across an IP network is specified.

A label switch is a frame or cell switch that supports 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 and MPLS Generic 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 arrive and the Labels 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 external to GSMP.

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 state 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 five classes of GSMP request-response message: Connection 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 does not acknowledge event messages. 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. Event messages have only a single format defined as they are not acknowledged by the controller.

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 subfields 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 state 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 does not currently have state synchronisation must be discarded.

The GSMP message type. GSMP messages fall into seven classes: Connection 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 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 Sate 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 were 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 repsonse, 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 results 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 requries acknowledgement, it will set the EventAck flag 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.

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. All messages in a segmented message, except for the last segment, will also have the More bit set.

Length

Message length including header.

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, "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.

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.

It is not an error for a GSMP message to contain additional data after the end of the Message Body. This is to support development and experimental purposes. However, the maximum transmission unit of the GSMP message, as defined by the data link layer encapsulation, must not be exceeded.

A success response message must not be sent until the requested operation has been successfully completed.

3.2 Failure Response Messages

[Editor's note: this section will be updated after the Feb 2 editing session. it is all out of whack right now. personally i think it ought to go to the end of the document.]

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.)

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

Invalid Message

Failure specific to the particular message type (failure code 16). (The meaning of this failure is dependent upon the particular message type and is specified in the text defining the message.)

A failure response specified in the text defining the message type.

Connection Failures

Virtual Path Connection Failures

Multicast Failures

QoS Failures (QoS failures are specified in [Section 9.7.](#))

General Failures

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 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.

- 5: 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.
- 4: 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.
- N1: Invalid Partition ID
The value given in the Partition ID field is not legal for this partition.

Connection Failures

- 8: 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 arrives. An ATM virtual path connection is specified by the input port and input VPI on which it arrives.
- 9: 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.
- 18: One or more of the specified input VPIs is invalid.
- 19: One or more of the specified Input Labels is invalid.
- 20: One or more of the specified output VPIs is invalid.
- 21: One or more of the specified Output Labels is invalid.
- 22: Invalid Service Selector field in a Connection Management message.
The value of the Service Selector field is invalid.
- 23: Insufficient resources for QoS Profile.
The resources requested by the QoS Profile in the Service Selector field are not available.

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 a 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 a 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.
- XX: 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

- 10: 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.
- 11: The limit on the maximum number of point-to-multipoint connections that the switch can support has been reached.
- 12: The limit on the maximum number of branches that the specified point-to-multipoint connection can support has been reached.

17: Cannot label each output branch of a point-to-multipoint tree with a different label.
Some early designs, and some low-cost switch designs, require all output branches of a multicast connection to use the same value of Label.

28: Only point-to-point bi-directional connections may be established.
It is an error to attempt to add an additional output branch to an existing connection with the bi-directional flag set.

13: Unable to assign the requested Label value to the requested branch on the specified point-to-multipoint 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 Label exists that the switch could support. For example, some switch designs restrict the number of distinct Label values available to a point-to-multipoint connection. (Most switch designs will not require this message.)

14: General problem related to the manner in which point-to-multipoint 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.)

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.

15: 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.

The following failure response messages are only used by the Label Range message. [Must come back and revise this --ed]

29: Cannot support requested label range.

The following failure response messages are only used by the Set Transmit Data Rate function of the Port Management message.

30: The transmit data rate of this output port cannot be changed.

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

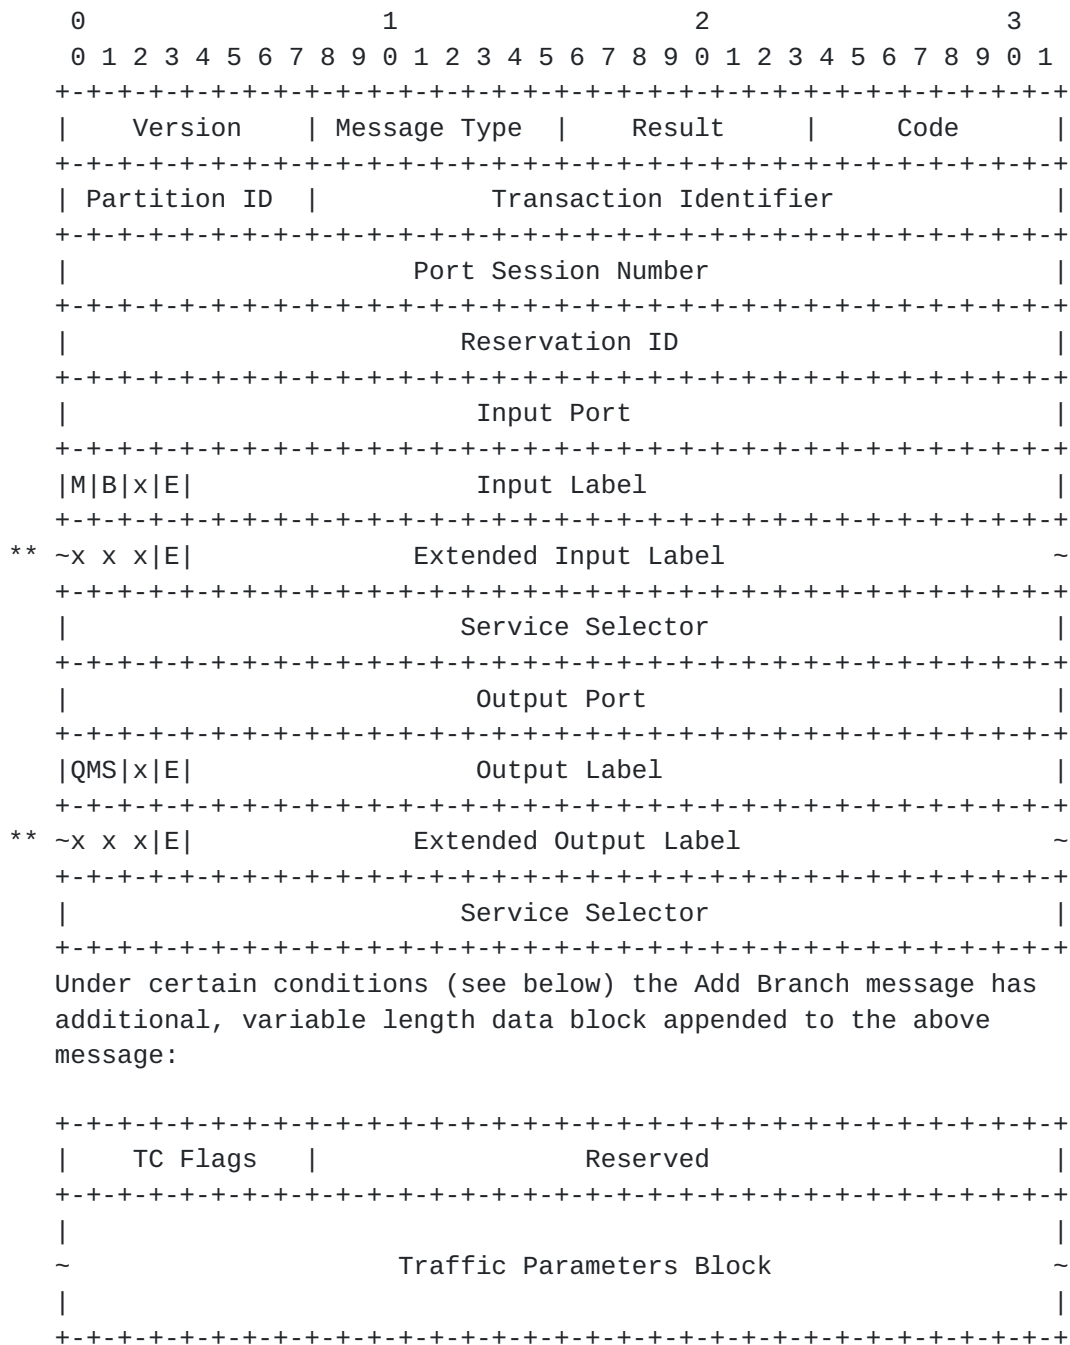
The following failure response message range is reserved for the ARM extension:

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

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:



** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

Input Port

Identifies a switch input port.

Flags

M: Multicast

Multicast flag is used as a hint for point-to-multipoint connections in the Add Branch message. It is not used in any other connection management messages and in these messages it should be set to zero. If set, 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.

The Multicast flag is 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 connection and on such connections it should be ignored by the receiver. (On receipt of the second and subsequent Add Branch messages the receiver knows that this is a point-to-multipoint connection.) If it is known that this is the first branch of a point-to-multipoint 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 this flag is not mandatory. It may be ignored by the switch. If unused the flag should be set to zero. Some switches use a different data structure for point-to-multipoint connections than for point-to-point connections. This flag avoids the switch setting up a point-to-point structure for the first branch of a point-to-multipoint connection which must immediately be deleted and reconfigured as point-to-multipoint when the second branch is established.

QMS: QoS Model Selector

The QoS Model Selector is used to specify a QoS Model for connection. The value of QMS indicates the value in the Service Selector should be interpreted as a priority, a QoS profile or a service specification as shown:

QMS	QoS Model	Service Selector
---	-----	-----
00	Simple Abstract Model	Priority
01	QoS Profile Model	QoS Profile
10	Service Model	Service Specification
11	Optional ARM	ARM Specification

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 arrives 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.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

x: Unused

Reservation ID

Identifies the reservation that must be deployed for the branch being added. Reservations are established using reservation management messages (see Chapter 10). A value of zero indicates that no Reservation is being deployed for the branch. If a reservation with a corresponding Reservation ID exists then the reserved

resources must be applied to the branch. If the numerical value of Reservation ID is greater than the value of Max Reservations (from the Switch Configuration message), a failure response is returned indicating "Reservation ID out of Range." If the value of Input Port differs from the input port specified in the reservation or if the value of Output Port differs from the output port specified in the reservation, a failure response must be returned indicating "Mismatched reservation ports." If no reservation corresponding to Reservation ID exists, a failure response must be returned indicating "Non-existent reservation."

If a valid Reservation ID is specified and the Service Model is used (i.e. QMS=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 implied.

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.

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

Service Selector

In the default QoS configuration, this field can contain either a Priority, a QoS Profile Identifier, or a Service Specification. When using an alternative QoS configuration, the format and semantics of data within the field are defined outside of GSMP.

In the default QoS configuration, if the QoS Model Selector is set to 0b00, the Service Selector field contains a Priority. If the QoS Model Selector is set to 0b01, the Service Selector field contains a QoS Profile. If the QoS Model Selector is set to 0b10, the Service

Selector field contains a Service Specification. If the QoS Model Selector is set to 0b11, the Service Selector field contains a service indicator which has its meaning defined by the optional ARM being used as indicated in the MType field of the configuration message. The Service Selector field is only used in the Add Branch and Move Branch messages.

Priority specifies the priority of the connection for Add Branch and Move Branch messages that choose not to use a QoS profile, or a service specification. The highest priority is numbered zero and the lowest priority is numbered "Q-1" where "Q" is the number of priorities that the output port can support. 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 is 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 output port can support is given in the Port Configuration message. In order to maintain backward compatibility with earlier versions of GSMP, the Priority octets will occupy the 2 right-most octets of the service selector.

A QoS Profile Identifier is an opaque 16-bit value. It is used to identify a QoS profile in the switch which specifies the Quality of Service required by the connection. QoS profiles are established by a mechanism external to GSMP.

A Service Specification is an alternative method of communicating the QoS requirements of a connection. The Service Specification is defined in Chapter 9.

TC Flags TC (Traffic Control) Flags are used in Add Branch messages for connections using the Service Model (i.e. when QMS=0b10). The TC Flags field is defined in [Section 9.6](#).

Traffic Parameters Block

This variable length field is used in Add Branch messages for connections using the Service Model (i.e. when QMS=0b10). Traffic Parameters Block is defined in [Section 9.5](#). 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 implied.

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 and the Number of Branches field indicating the number of branches on the connection after completion of the operation. 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 and the Number of Branches field indicating the number of branches on the connection.

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.

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. (An output branch is specified by an output port and output label.)

Label stacking is a technique used in MPLS that allows hierarchical labelling. MPLS label stacking is similar to but subtly different from the VPI/VCI hierarchy of labels in ATM. ...
[Must add blah --Ed]

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 state 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.


ATM Labels

If a port's attribute PortType=ATM then that port's labels must be interpreted as 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
+	-	-	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
										VPI																				VCI									
+	-	-	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+

ATM ports do not support Extension Labels so the VPI and VCI values always occupy the 28 bits following the flags in a connection management message. 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).



MPLS Label

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

4.2 Add Branch Message

The Add Branch message is a connection management message used to establish a connection or to add an additional branch to an existing connection. It may also be used to check the connection state stored in the switch. The connection is specified by the Input Port and Input Label fields. The output branch is specified by the Output Port and Output Label fields. The quality of service requirements of the connection are specified by the QoS Model Selector and Service Selector fields. To request a connection the Add Branch message is:

Message Type = 16

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, 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 previously established state. For messages that are reasserting 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 output branch specified by the Output Port and Output Label fields 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 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: "Only point-to-point bi-directional connections may be established."

It should be noted that different switches support multicast in different ways. There will be a limit to the total number of point- to-multipoint connections any switch can support, and possibly a limit on the maximum number of branches that a point-to-multipoint connection may specify. Some switches also impose a limit on the number of different Label values that may be assigned 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.

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, "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, "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, "Attempt to add a 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, "Attempt to add a 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 specified by the Input Port and 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. The Number of Branches field is not used in either the request or response messages of the Delete Tree message.

4.4 Verify Tree Message

The Verify Tree message has been removed from this version of GSMP. Its function has been replaced by the Number of Branches field in the success response to the Add Branch message which contains the number of branches on a connection after successful completion of an add branch operation.

Message Type = 19 is reserved.

If a request message is received with Message Type = 19 a failure response must be returned with the Code field indicating: "The specified request is not implemented in this version of the protocol."

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 arrive 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 Number of Branches field is not used in either the request or response messages of the Delete All Input Port message. 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.

The specified request is not implemented on this switch.

One or more of the specified ports does not exist.

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: "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 input 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 Number of Branches field is not used in either the request or response messages of the Delete All Output Port message. 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.

The specified request is not implemented on this switch.

One or more of the specified ports does not exist.

Invalid Port Session Number.

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

[4.7 Delete Branches Message](#)

The Delete Branches message is a connection management message used to request one or more delete branch operations. Each delete branch operation deletes a branch of a channel, or in the case of the last branch of a connection, it deletes the connection. The Delete Branches message is:

Message Type = 17

The 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																													
Reserved										Number of Elements																													
Delete Branch Elements																																							

Number of Elements

Specifies the number of Delete Branch Elements to follow in the message. The number of Delete Branch Elements in a Delete Branches message must not cause the packet

length to exceed the maximum transmission unit defined by the encapsulation.

Each Delete Branch Element specifies an output branch to be deleted and has the following structure:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Port Session Number                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Input Port                                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x x|I|0|                               |
+---+---+---+                               Input Label                         ~
~                                                                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Output Port                                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Error |                               |
+---+---+---+                               Output Label                       ~
~                                                                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

I: Input Extension Label

The Input Extension Label flag if zero indicates that the Input Label field is a 28 bit field. If the Input Extension Label flag is set then the Input Label field is a 60 bit field structured as a 28 bit first level label field followed by a 32 bit second level label field.

O: Output Extension Label

The Output Extension Label flag if zero indicates that the Output Label field is a 28 bit field. If the Output Extension Label flag is set then the Output Label field is a 60 bit field structured as a 28 bit first level label field followed by a 32 bit second level label field.

Error

Is used to return a failure code indicating the reason for the failure of a specific Delete Branch Element in a Delete Branches failure response message. The Error field is not used in the request message and must be set to zero. A value of zero is used to indicate that the delete operation specified by this Delete Branch Element was successful. Values for the other failure codes are specified in [Section 3.2](#), "Failure Response Messages."

All other fields of the Delete Branch Element have the same definition as specified for the other connection management messages.

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, "Failure specific to the particular message type," 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 Branch Message

The Move 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 Branch connection management message has the following format for both request and response messages:

The Move 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 Branch message is:

Message Type = 22

For the Move 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 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, "The specified branch does not exist."

For the Move 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.

ATM Specific Procedures:

The ATM VPC Move 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 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 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, "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, "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.

In 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.

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 arrive at the specified input port must be deleted and a new Port Session Number must be selected 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. The Port Status of the port afterwards 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. The Port Status of the port afterwards 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. The Port Status of the port afterwards 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. The Port Status of the port afterwards will be External Loopback.

Bothway Loopback:

Function = 5. Both internal and external loopback are performed. The Port Status of the port afterwards will be Bothway Loopback.

Reset Input Port:

Function = 6. All connections that arrive 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. The Port Status of the port afterwards 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 toggled; i.e. flow control for the corresponding event is turned off if 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: "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: "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 or frames/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 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.

The Event Flags field is only used in a request message with the Function field set to "Reset Event Flags." For all other values of the Function field, the Event Flags field is not used. In the success response message the Event Flags field must be set to the current value of the Event Flags for the port, after the completion of the operation specified by the request message, for all values of the Function field. Setting the Event Flags field to all zeros in a "Reset Event Flags" request message allows the controller to obtain the current state of the Event Flags and the current Event Sequence Number of the port without changing the state of the Event Flags.

The correspondence between the types of Event message and the bits of the Event Flags field is as follows:

```

          1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+--+--+--+--+--+--+--+--+--+--+
|U|D|I|N|Z|A|x x x x x x x x x x|
+-+--+--+--+--+--+--+--+--+--+--+

```

U: Port Up	Bit 0, (most significant bit)
D: Port Down	Bit 1,
I: Invalid Label	Bit 2,
N: New Port	Bit 3,
Z: Dead Port	Bit 4,
A: Adjacency Event	Bit 5,
x: Unused	Bits 6--15.

Flow Control Flags Field

This flags in this field are used to indicate whether

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																													
										Port																													
										Port Session Number																													
Q V x x										Min Label																													
x x x x										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.

V: Label

The Label flag use is port type specific.

x: Unused

The success response to a Label Range message requesting a change of label range is a copy of the request message with the Remaining Label Range fields 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 "Cannot support requested label range." In this failure response message the switch must use the Min Label and Max Label fields to suggest a label range that it would be 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, "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.

ATM Labels:

If PortType=ATM the label range fields have 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										
Q V x x										Min VPI																				Min VCI																			
x x x x										Max VPI																				Max VCI																			
										Remaining VPIS																				Remaining VCIS																			

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

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 VPIs

Remaining VCIs

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 VPIs and VCIs 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 VPIs and VCIs 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 VPIs and VCIs 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 VPIs 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 VCIs 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: "One or more of the specified input VPIs 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 VPIs to be allocated to the input port specified by the Port field. Whatever the range of VPIs previously allocated to this port it 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 VCIs to be allocated to each of the VPIs specified by the VPI range. Whatever the range of VCIs previously allocated to each of the VPIs 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 VPIs 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 "Cannot support requested label range." 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 VPIs 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 VPIs the VCI range that can be supported is often more constrained. Often the Min VCI must be 0 or 32. Typically all VCIs 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 VPIs and a range of VCIs within each VPI, the most likely use is to change either the VPI range or the range of VCIs within a single VPI. It is possible for a VPI to be valid but to be allocated no valid VCIs. 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 VCIs.

Frame Relay Labels:

If PortType=FR the label range fields have 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								
Q V x x x x Len										Min DLCI																													
x x x x x x x x										Max DLCI																													
										Remaining DLCIs																													

V: Label The Label flag is not used.

Len

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

Len	DLCI bits
0	10
1	17
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 DLCIs

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 DLCIs 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 DLCIs 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 DLCIs is available to every switch port.

MPLS Generic Labels:

The Label Range field for PortTypes using MPLS labels (e.g. Ethernet, SONET etc.) 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 x x x x x x x x x x x|           Min MPLS Label           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x x x x x x x x x x x x|           Max MPLS Label           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Remaining MPLS Labels          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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. 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.

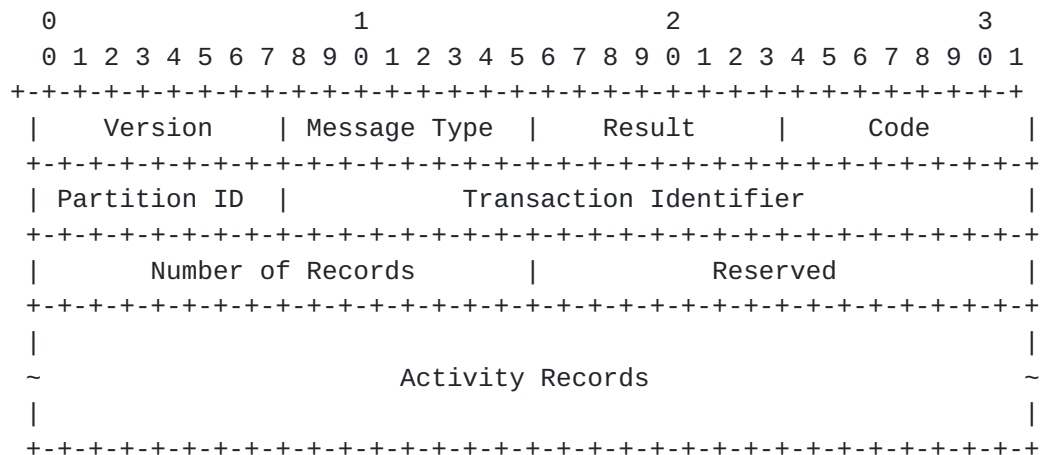
6.1 Connection Activity Message

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

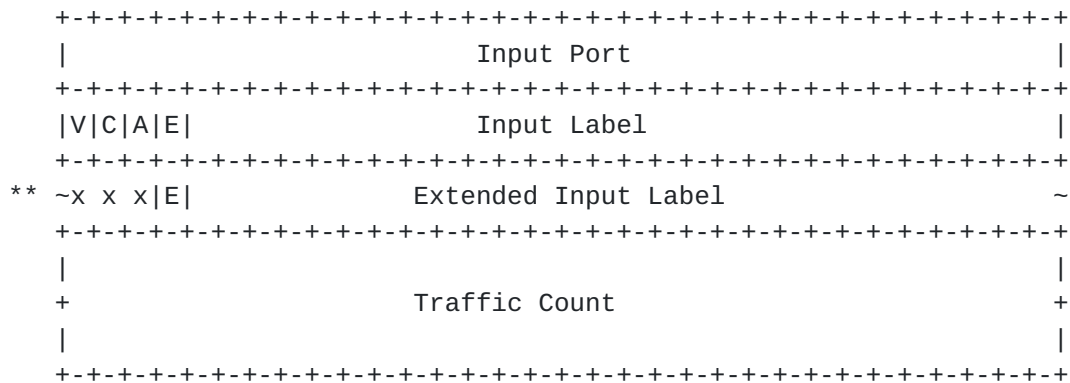
The Connection Activity request and success response messages have the following format:



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:



** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

Input Port

Identifies the port number of the input port on which the connection of interest arrives 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.

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.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

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

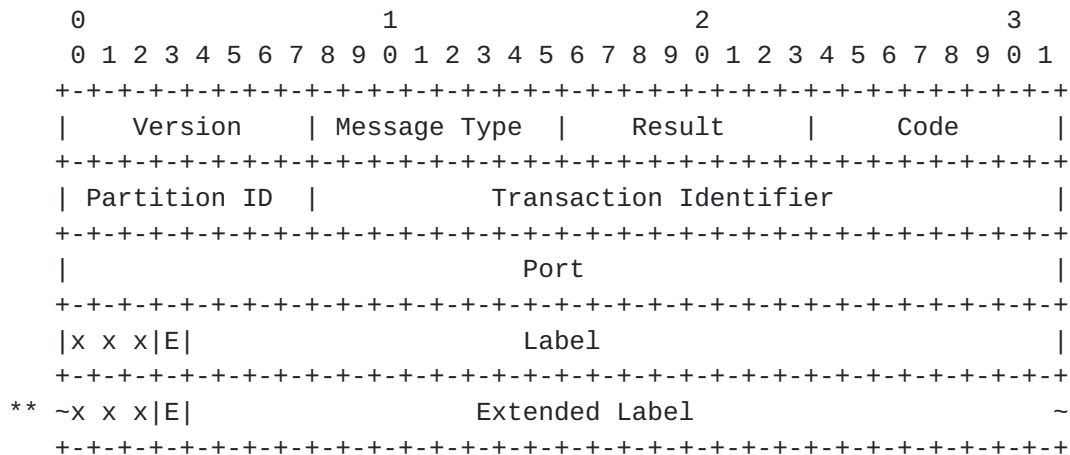
the value of a free running, connection specific, 64-bit traffic counter counting traffic flowing across the specified connection. The value of the traffic counter is not modified by reading it. If per connection traffic counting is supported, the switch must report the Connection Activity result using 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, "The specified request is not implemented on this switch."

6.2 Statistics Messages

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

The Statistics request messages have the following format:



** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

Label

The Label Field 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	2	3	4	5	6	7	8	9
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
Port																																							
x x x E										Label																													
~x x x E										Extended Label																													
Input Cell Count																																							
Input Frame Count																																							
Input Cell Discard Count																																							
Input Frame Discard Count																																							
ATM HEC Error Count																																							
Input Invalid Label Count																																							
Output Cell Count																																							
Output Frame Count																																							
Output Cell Discard Count																																							


```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
+           Output Frame Discard Count           +
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

E
Port
Label

Fields are the same as those of the request message.

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.

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.

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.

Input Frame Discard Count
Output Frame Discard Count

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

ATM HEC Error Count

Gives the value of a free running 64-bit counter counting ATM cells discarded due to header checksum errors on arrival at an input port.

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.

6.2.1 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

6.2.2 Connection Statistics Message

The Connection Statistics message requests the statistics for the connection specified in the Label field that arrives 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 ATM HEC 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

6.2.3 QoS Class Statistics Message

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

Message Type = 51 is reserved.

6.3 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																																															
Input Port																																																															
A V x E																Input Label																																															
** ~x x x E																Extended Input Label																~																															

** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

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 arrive 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 set and in which the input port specified by the Input Port field does not have PortType=ATM, the switch must return an Error Message "xxxxxx".

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 arrives 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 arrives 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.

x: Unused.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

Input Label

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

The Report Connection State 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																													
+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--									
										Input Port																													
+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--									
										Sequence Number																													
+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--									
~										Connection Records																				~									
+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--										+--+--+--+--+--+--+--+--+--									

Input Port

Is the same as the Input Port field in the request message. It identifies the port number of the input port for which the connection state is being reported.

Sequence Number

In the case that the requested connection state cannot be reported in a single success response message, each successive success response message in reply to the same request message must increment the Sequence Number. The Sequence Number of the first success response message, in response to a new request message, must be zero.

Connection Records

Each success response message must contain one or more Connection Records. Each Connection Record specifies a single point-to-point or point-to-multipoint connection. The number of Connection Records in a single Report Connection State success response must not cause the packet length to exceed the maximum transmission unit defined by the encapsulation. If the requested connection state cannot be reported in a single success response message, multiple success response messages must be sent. All success response messages that are sent in 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. The More flag of the last Connection Record in a success response message indicates whether the response to the request has been completed or whether one or more

further success response messages should be expected in response to the same request message.

Each Connection Record has the following format:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|A|V|P|M|                               Input Label                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Output Branch Records                    |
~                               ~                                         ~
|                               ~                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

[Editor's note: help, where do i put the extension label flag?]

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.

M: More

If the More flag is set, it indicates that another Connection Record, in response to the same request message, will follow either in the same success response message or in a successive success response message. If the More flag is zero it indicates that this is the last Connection record in this success response message and that no further success response messages will be sent in response to the current request message. It indicates that the response to the request message is now complete.

Input Label

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

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. The last Output Branch Record of each Connection Record is indicated by the Last Branch flag of the Output Branch Record. 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                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|x x|E|                               Output Label                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
** ~x x x|E|                               Extended Output Label           ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

Output Port

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

Flags

L: Last Branch

The Last Branch flag, if set, indicates that this is the last Output Branch Record of this Connection Record. If zero, it indicates that one or more further Output Branch Records are to follow. If this is the last Output Branch Record in the message and the Last Branch flag is zero, further output branches belonging to the same connection will be given in another Connection Record. This Connection Record will be the first Connection Record in the next success response message. This

Connection Record must have the same Input Label value as the current Connection Record.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

x: Unused.

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, "Failure specific to the particular message type." 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.

7. 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.

7.1 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																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																			
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	

MType

Represents an alternative QoS Configuration type. In the request message the requested MType is in the most significant (leftmost) MType octet; the other three MType octets 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 in the rightmost 3 octets of the reply message respectively, each of which indicates an available

alternative QoS Configurations. A switch that supports on the default QoS Configuration always returns MType=0 in all four MType fields. MType negotiation is discussed in [section 7.1.1](#).

0	i	Indicates use of the default GSMP model
1	i	Indicates use of IEEE qGSMP model
2 - 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.

[editor's note: some text may be added here with regard to the tcp/ip encapsulation since if tcp is used then the switch may adjust the receiver 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 10). A value of 0 indicates that the switch does not support Reservations.

7.1.1 Configuration Message Processing

After adjacency between a controller and a switch is first established the controller that opts to use a QoS Configuration 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.

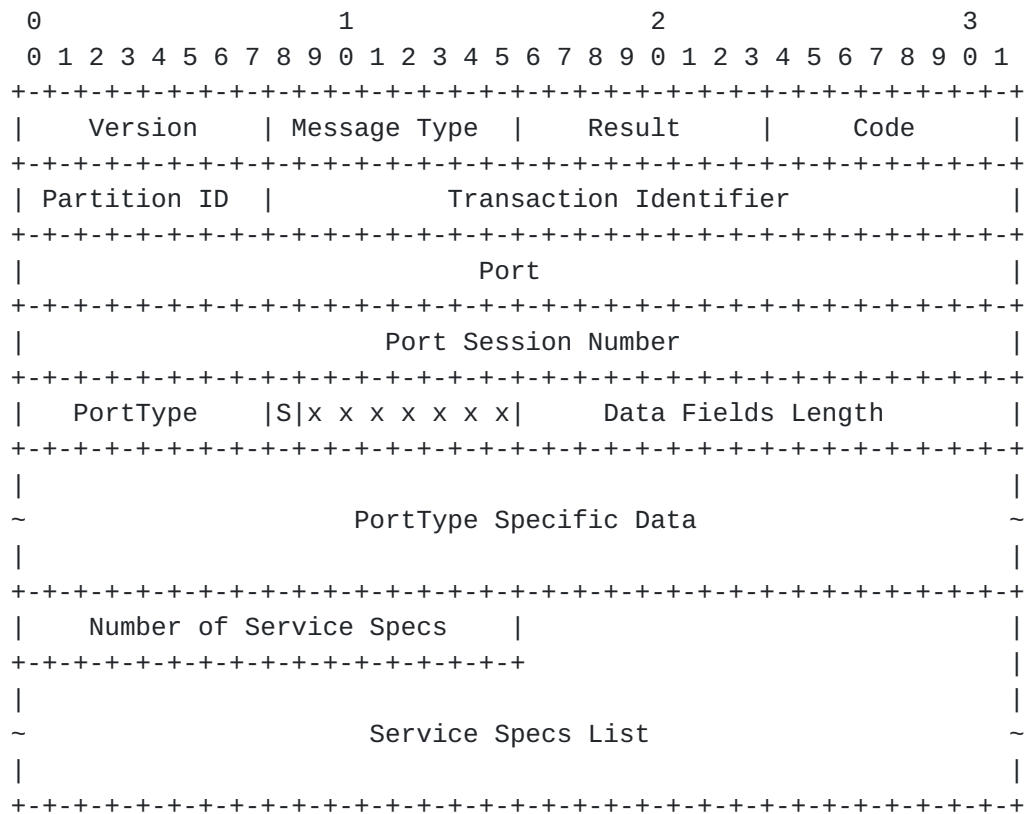
Once a MType has been established for the switch, it cannot be changed without full restart; that is the re-establishment of adjacency with the resetting of all connections.

7.2 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:



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.

PortType [Editor's note: words to be written. also, somewhere in chapter 1 or so we will need text that explains that certain protocol elements depend on the PortType value.]

PortType = 0d01 = ATM

PortType = 0d02 = FR

PortType = 0d03 = 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 depends also on the value of 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

Field contains a sequence of 1 or more Service Specs (defined below). If the Number of Service Specs is an even number then 16 bits of padding is inserted after the last Service Spec in order to justify the end of the Service Specs List at a 32bit word boundary.

Service Spec

The format of the Service Spec field is given below:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+-+--+--+--+--+--+--+--+--+--+--+--+
|  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 that share 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 8.6.

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.

x: Unused

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 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 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 [Note: what about other port types? Fiffi]

Only applicable for PortType=ATM. The type of physical transmission interface for this port. The values for this field are defined by the atmIfType object specified in the Ipsilon IP Switch MIB [4].

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 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.

7.2.1.1 PortType Specific data for PortType=ATM

If PortType=ATM the label range fields have 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|V|M|L|R|      Min VPI      |      Min VCI      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Q|x x x|      Max VPI      |      Max VCI      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

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 VPIs 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 VPIs 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 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 dynamically assigned

0 10


```

1    17
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.

7.2.1.3 PortType Specific data for PortType=MPLS

The Label Range field for PortTypes using MPLS labels (e.g. Ethernet, SONET etc.) 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|V|M|L|R|x x x x x x x x|           Min MPLS Label           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Q x x x|x x x x x x x x|           Max 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.

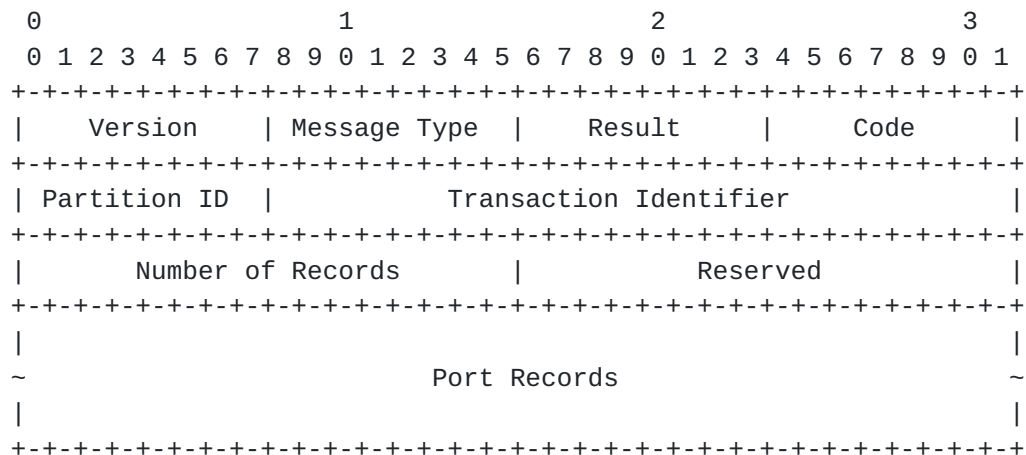
7.3 All Ports Configuration Message

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.

The All Ports Configuration success response message has the following format:



Number of Records

Field gives the total number of Port Records to be returned in response to the All Ports Configuration request message. The number of port records in a single All Ports Configuration success response must not cause the packet length to exceed the maximum transmission unit defined by the encapsulation. If a switch has more ports than can be sent in a single success response message it must send multiple success response messages. All success response messages that are sent in response to the same request message must have the same Transaction Identifier as the request message and the same value in the Number of Records field. All success response messages that are sent in response to the same request message, except for the last message, must have the result field set to "More." The last message, or a single success response message, must have the result field set to "Success." All Port records within a success response message must be complete, i.e. a single Port record must not be split across multiple success response messages.

Port Records

Follow in the remainder of the message. Each port 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								
+--+																																							

Number of Service Records

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

Service Records

A sequence of zero or more Service Records. The switch returns one Service Record for each Service that it supports 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 The Service ID Field identifies the Service supported by the port. The Services are defined with their Service ID values in Chapter 8.6.

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										Reserved										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 is 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 8.6. 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	6	7	8	9
+---+---+---+---+---+---+---+---+---+---+										+---+---+---+---+---+---+---+---+---+---+									
U D I E S V Res																			
+---+---+---+---+---+---+---+---+---+---+										+---+---+---+---+---+---+---+---+---+---+									

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.

Res: Reserved

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 8.6. 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 $10^{\exp(n)}$. 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 is given.

Frequency

The frequency field is coded as an 8 bit unsigned integer. Frequency applies to the MPLS CR-LDP Service (see [Section 9.4.3](#)). 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 binary integer.

CDV: Peak-to-peak Cell Delay Variation

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

8. 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:

[illegible]

**** Note:** There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

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.

E: Extension Label

The Extension Label Flag is used to extend the adjacent label field by inserting, after the adjacent label, an additional 32 bit word into the message. A 32 bit word

formatted according to the line marked ** in the message diagram follows the adjacent label field if and only if the E Flag is set.

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 port sends an Event message it must set the Event Flag on that port corresponding to the type of the event. The port is not permitted to send another Event message of the same type 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.

8.1 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

8.2 Port Down Message

The Port Down message informs the controller that the Line Status of a port has changed from the Up 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

8.3 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

8.4 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

8.5 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

8.6 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

9. Service Model Definition

9.1 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. 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 7.4](#)). 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, an eight bit identifier. Assigned numbers for the Service ID are given with the Service Definitions in [Section 9.4](#). The Capability Set is identified in GSMP messages by the Capability Set ID, an eight 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.

9.2 Service Model Definitions

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

9.2.1 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 [\[11\]](#).

9.2.2 Service Definition, Traffic Parameters, QoS Parameters and Traffic Controls

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.

9.2.3 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 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.

9.3 Service Model Procedures

A switch's Services and Capability Sets are reported to a controller in a Service Configuration messages. 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. Either the one stage or two stage connection set-up procedure may be used to change an extant connection's Traffic Parameters.

Move branch messages may be refused on the grounds of resource depletion. In the case of a one stage set-up the connection state does not change if a move branch request is thus refused.

A switch may support a bandwidth allocation function. If it does, a controller may choose to use it or not to use it. A controller indicates whether or not switch bandwidth allocation is requested using a bandwidth allocation (BA) flag in connection management messages. A switch indicates that it is honouring the bandwidth allocation request, and thus the QoS commitments indicated in the QoS of its Capability Sets, by responding with the BA flag set. If the switch does not have a bandwidth allocation function then it will always respond with the BA flag zeroed. If the controller ever sends a request with a zeroed BA flag, the switch is not

obliged to honour the QoS commitments for the requested connection, other extant connections or future connections. If the switch receives a request with the BA flag set it must not reject the connection based on a lack of bandwidth. If, after the controller has issued a request with the BA flag zeroed, the switch is still able to track whether or not it is honouring the QoS commitments then it may indicate that QoS commitments are honoured using the BA flag in its responses. The controller may poll the switch with connection refresh messages to determine if the switch is still honouring QoS.

9.4 Service Definitions

This section sets forth the definition of Services. 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.

9.4.1 ATM Forum Service Categories

9.4.1.1 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

9.4.1.2 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

9.4.1.3 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

9.4.1.4 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

9.4.1.5 ABR

ABR is not supported in this version of GSMP.

9.4.1.6 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

9.4.2 Integrated Services

9.4.2.1 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.

9.4.3 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.

9.4.4 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

9.4.5 Diff-Serv

For future study.

9.5 Format and encoding of the Traffic Parameters Block in connection management messages

Connection management messages that use the GSMP Service Model (i.e. those that have QMS=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 9.4](#). The format and encoding of these parameters are given below.

9.5.1 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			
+--+			
Reserved 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 9.4.1](#) in the order given there.

9.5.2 Traffic Parameters for the Int-Serv Controlled Load Service

The Traffic Parameters:

- Peak Data Rate,
- Peak Burst Size,
- Committed Data Rate,
- Committed Burst Size, and

- 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																																						
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
	Reserved																Weight																						
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								

9.5.4 Traffic Parameters for the 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 than 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								
+--+																																							

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

9.6 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 7.4](#)) 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.

10. 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 an 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 switch 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).

10.1 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 Port Management message has the following format for the request 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								
Version										Message Type										Result										Code									
Partition ID										Transaction Identifier																													
Port Session Number																																							
Reservation ID																																							
Input Port																																							
M B U E										Input Label																													
** ~X x x E										Extended Input Label																				~									
Service Selector																																							
Output Port																																							
QMS U E										Output Label																													
** ~X x x E										Extended Output Label																				~									
Service Selector																																							

When the value of QMS is set to 0b10 then the following Traffic Parameters Block is appended to the above message:

TC Flags										Reserved																			
Traffic Parameters Block																													

** Note: There can be zero or more 32 bit words containing Extended Labels (like those marked **) following an Input or Output Label field. A 32 bit word containing an Extended Label follows the previous label field if and only if the E Flag immediately preceding the previous label is set.

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

Flags

U: Unspecified Label

If Unspecified Label flag is set then the adjacent Label field is unused and the adjacent Extension Label flag must be set to zero. If Unspecified Label flag is zero then the adjacent Label field and Extension Label flag have the same meaning as in the Add Branch message.

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 "Insufficient Resources." Other failure responses are as defined for the Add Branch message.

10.2 Delete Reservation Message

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

Message Type = 71

The Delete Reservation message has the following format:

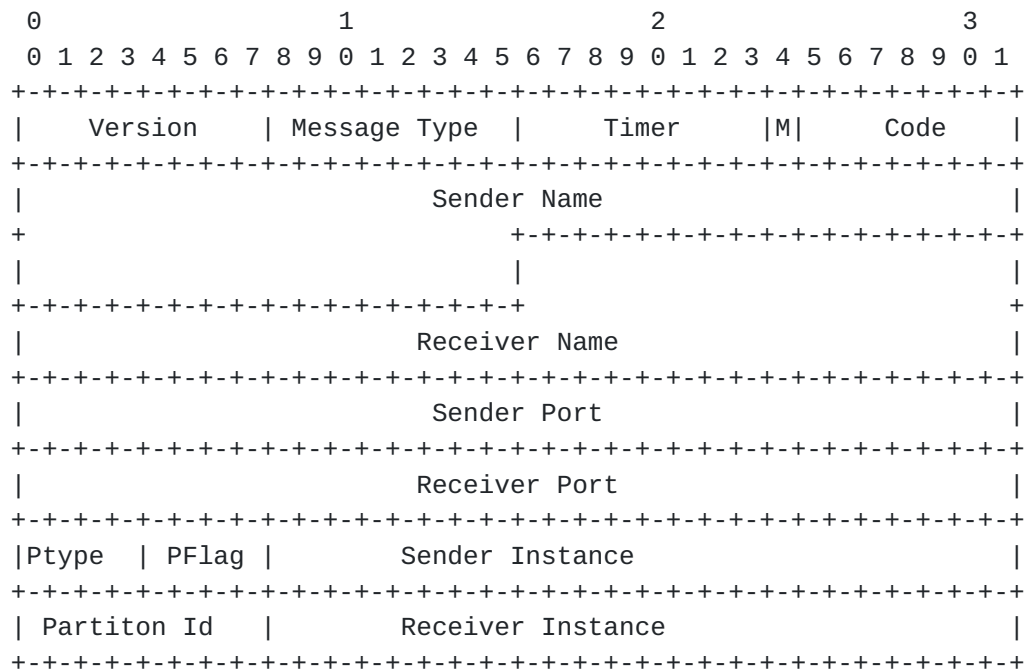
[illegible]

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. 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 = 2.

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

Type of partition being requested.

0 No Partition Request

1 Fixed Partition

PFLAG

Used to indicate 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 32-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, and Sender Name 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 Sender 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 Receiver Partition ID fields in the incoming message match the values of the Sender Instance, Sender Port, Sender Name and Sender 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.

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. Two scenarios are provided for:

- A controller can request a specific partition identifier with the switch having the option to either accept or to reject the request. In this case the adjacency message will include Partition ID != 0.
- The switch can assign partition identifiers to controllers based on its on pre-established mechanisms. In this case the adjacency message will include Partition ID = 0.

The assignment is determined by the following behaviour:

- An adjacency message from the controller with SYN || SYNACK && PTYPE is treated as a partition request
- An adjacency message from the switch with SYNACK || ACK && PTYPE is treated as a partition assignment
- An adjacency message from the switch with RSTACK && PTYPE is treated as partition unavailability.

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 state, deferring the decision about resetting the state until after synchronisation is re-established.

Once synchronisation is re-established the decision about resetting the 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 switch state should be retained. Verification that controller and switch 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 master/standby fashion or as part of multiple controllers providing load-sharing for the same partition. It is the

responsibility of the controller 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 the same partition. It should be noted that the GSMP will not co-ordinate cache synchronization information among controllers; i.e. the switch partition will service each command it receives in turn as if it were interacting with a single controller. Implementations without controller entity synchronisation are advised against using 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 chapter Partition Information State. The next adjacencies to the partition are identified by a new partition request with same Partition id as the first one but with different Sender Name. Adjacency loss is defined in the section on the Loss of Synchronization (Error! Reference source not found.).

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 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 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 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. Summary of Failure Response Codes

[Editor's note: this section is currently out of whack w.r.t. the rest of the spec and will be updated in a future revision of the draft.]

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: Invalid Port Session Number.
- N1: Invalid Partition ID
- 5: One or more of the specified ports does not exist.
- 6: One or more of the specified ports is down.
- 7: Unused. (This failure code has been replaced by failure codes 18--21.)
- 8: The specified connection does not exist.
- 9: The specified branch does not exist.
- 10: 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.
- 11: The limit on the maximum number of point-to-multipoint connections that the switch can support has been reached.
- 12: The limit on the maximum number of branches that the specified point-to-multipoint connection can support has been reached.
- 13: Unable to assign the requested Label value to the requested branch on the specified point-to-multipoint connection.
- 14: General problem related to the manner in which point-to-multipoint is supported by the switch.
- 15: Out of resources (e.g. memory exhausted, etc.).
- 16: Failure specific to the particular message type. (The meaning of this failure code depends upon the Message Type. It is defined within the description of any message that uses it.)
- 17: Cannot label each output branch of a point-to-multipoint tree with a different label.
- 18: One or more of the specified input VPIs is invalid.
- 19: One or more of the specified input VCIs is invalid.
- 20: One or more of the specified output VPIs is invalid.
- 21: One or more of the specified output VCIs is invalid.
- 22: Invalid Class of Service field in a Connection Management message.

- 23: Insufficient resources for QoS Profile.
- 24: Virtual path switching is not supported on this input port.
- 25: Point-to-multipoint virtual path connections are not supported on either the requested input port or the requested output port.
- 26: Attempt to add a virtual path connection branch to an existing virtual channel connection.
- 27: Attempt to add a virtual channel connection branch to an existing virtual path connection.
- 28: Only point-to-point bi-directional connections may be established.
- 29: Cannot support requested VPI range.
- 30: Cannot support requested VCI range on all requested VPIs.
- 31: The transmit cell rate of this output port cannot be changed.
- 32: Requested transmit cell rate out of range for this output port.
- 128: Weighted scheduling within this waiting room is unavailable.
- 129: This waiting room is unable to offer weighted sharing for a QoS class.
- 130: This waiting room is unable to offer weighted sharing for a connection.
- 131: Scheduler Identifier still in use.
- 132: QoS Class Identifier still in use.
- 133: Invalid QoS parameter.
- 134: Insufficient QoS resources.
- 135: Any point-to-multipoint connection arriving on this input port must use the same QoS parameters for all output branches.

13. Summary of Message Set

[Editor's note: this section is currently out of whack w.r.t. the rest of the spec and will be updated in a future revision of the draft.]

The following table gives a summary of the messages defined in this version of the specification. It also indicates which messages must be supported in a minimal implementation of the protocol. Those messages marked as "Required" must be supported by the switch for an implementation to be considered to conform to this specification. (While the controller should also implement those messages marked "Required," conformance cannot be tested for the controller due to the Master-Slave nature of the protocol.)

Message Name	Message Type	Status
Connection Management Messages		
Add Branch VCC.....	16	Required
VPC.....	26	
Delete Tree.....	18	Required
Delete All.....	20	
Delete Branches.....	17	
Move Branch VCC.....	22	
VPC.....	27	
Port Management Messages		
Port Management.....	32	Required
Label Range.....	33	
State and Statistics Messages		
Connection Activity.....	48	Required
Port Statistics.....	49	
Connection Statistics.....	50	
Report Connection State.....	52	
Configuration Messages		
Switch Configuration.....	64	Required
Port Configuration.....	65	Required
All Ports Configuration.....	66	Required
Service Configuration.....	N2	
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

14. Security Considerations

The security of GSMP's TCP/IP control channel has been addressed in [15]. Any potential remaining security considerations are not addressed in the current revision of this draft.

References

- [1] "B-ISDN ATM Layer Specification," International Telecommunication Union, ITU-T Recommendation I.361, Mar. 1993.
- [2] "B-ISDN ATM Adaptation Layer (AAL) Specification," International Telecommunication Union, ITU-T Recommendation I.363, Mar. 1993.
- [3] IEEE/WG 1520, Adam, C; Lazar, A; Nanadikesan, M; "Proposal for Standardizing the qGSMP protocol", P1520/TS/ATM-002, <http://comet.columbia.edu/pin-atm/docs/P1520-TS-ATM-002R1.pdf>, 19 Jan, 1999
- [4] Sjostrand, H., "Definitions of Managed Objects for the General Switch Management Protocol (GSMP)," Internet-Draft [draft-ietf-gsmp-mib-00](#), September 1999.
- [5] Reynolds, J., and J. Postel, "Assigned Numbers," STD 2, [RFC 1700](#), October 1994.
- [6] Newman, P, Edwards, W., Hinden, R., Hoffman, E. Ching Liaw, F., Lyon, T. and Minshall, G., "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 Minshall, G., "Ipsilon's General Switch Management Protocol Specification," Version 2.0, [RFC 2397](#), March 1998.
- [8] ATM Forum Technical Committee, "Traffic Management Specification Version 4.1," af-tm-0121.000, xxx 1999.
- [9] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," [RFC2211](#), Sep 1997.
- [10] B. Jamoussi, et. al. "Constraint-Based LSP Setup using LDP," Internet Draft [draft-ietf-mpls-cr-ldp-03.txt](#), Feb 1999.

- [11] ITU-T Recommendation I.233 Frame Mode Bearer Services 1992.
- [12] ITU-T Recommendation Q.933, Integrated Services Digital Network (ISDN) Digital Subscriber Signalling System No. 1 (DSS 1) ; Signalling 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] E. Rosen, et al, "MPLS Label Stack Encoding" Internet-Draft [draft-ietf-mpls-label-encaps-07.txt](#), Sep 1999.
- [15] T. Worster, "GSMP Packet Encapsulations for ATM, Ethernet and TCP," Internet-Draft [draft-ietf-gsmp-encaps-00](#), Jan 2000.

Authors' Addresses

Avri Doria
Nokia Telecommunications
5 Wayside Road
Burlington MA 01803
Phone: +1 781 993 4656
avri.doria@nokia.com

Fiffi Hellstrand
Nortel Networks AB
S:T Eriksgatan 115 A
P.O. Box 6701
SE-113 85 Stockholm Sweden
fiffi@nortelnetworks.com

Kenneth Sundell
Nortel Networks AB
S:T Eriksgatan 115 A
P.O. Box 6701
SE-113 85 Stockholm Sweden
ksundell@nortelnetworks.com

Tom Worster (Editor)
Ennovate Networks
60 Codman Hill Rd
Boxboro MA 01719 USA
Tel +1 978-263-2002
fsb@thefsb.org