Editors Note: The edits made herein are to TD-GEN-0622, the output from the January 2011 meeting, are only those which fix the remaining text which is now incorrect due to decisions made in the January meeting. These decisions included the elimination of GRS and VRS services, the elimination of sending data before a Response is received (thus including the elimination of static mode), the elimination of terms relating to IP protocol, addition of new definitions, some clarification, and the deletion of sections now inconsistent with the revised procedure.

Editor's notes to revise or create sections will be treated in another document.

Title

Signalling protocols and procedures relating to Flow State Aware QoS control in a bounded sub-network of a NGN

Summary

This Recommendation specifies the signalling format and procedures for Flow State Aware (FSA) Transfer Capability in a bounded sub-network of a Next Generation Network (NGN). This specification ensures (through [b-RFC-4775] procedures) that there is no adverse impact or modification to operational networks outside of the bounded sub-network. The FSA Transfer Capability provides QoS controls that operate on a per-flow basis, allowing flows within a bounded sub-network of a NGN to receive different treatment depending on signalled parameters. These parameters are requested using in-band signalling. The parameters contained in these signals are included in the "flow state" maintained on each flow (or each aggregate flow) at each FSA node.

Service options that may be selected include requested support of the highest available end-to-end (or FSA edge-to-edge) rate for data transfer, assuming some source-to-receiver forwarding paths are entirely within the bounded sub-network of a NGN. ~~Another option is immediate transmission, wherein a flow may start or assume a new rate immediately on the understanding that the network will provide a guaranteed rate as soon as possible. This will be provided when network resources permit. Yet another option is for a negotiated guaranteed rate.~~

The focus of this Recommendation is on broadband (including mobile) service access scenarios typically involving restricted bandwidth shared by many flows. In such circumstances, customer satisfaction, when there is temporary congestion, may be best handled by applying flow preferences and QoS differently for each customer and not simply on the type of media. This leads to the notion of customised QoS supported partly by signalling and partly through web-based tools.

These concepts may also be applied to flow aggregates, with associated signalling support acting at the aggregate level between Aggregation End-points, as defined in [ITU-T Y.2121] . Customisation of aggregates in this Recommendation is limited to the notions of preference priority and FSA Transfer Capability.

Keywords: Available Rate, Broadband service access, flow state, Flow State Aware (FSA), in-band signalling, NGN, QoS

Table of Contents

1. Scope

This document provides the recommendation for the signalling format and procedures for Flow State Aware (FSA) Transfer Capability in a bounded sub-network of a Next Generation Network (NGN). This specification ensures (through [b-RFC-4775] procedures) that there is no adverse impact or modification to operational networks outside of the bounded sub-network. The FSA Transfer Capability provides QoS controls that operate on a per-flow basis, allowing flows within a bounded sub-network of a NGN to receive different treatment depending on signalled parameters. These parameters are requested using in-band signalling. The parameters contained in these signals are included in the "flow state" maintained on each flow (or each aggregate flow) at each FSA node.

Out of scope are any FSA edge functions at network domain boundaries needed for interworking between two FSA networks, where one uses in-band signalling exclusively for all FSA service support, and the other uses out-of band signalling for resource reservation and clear-down coupled with in-band signalling to establish the agreed flow state. FSA signalling must not be used in the ~~IP/MPLS~~ core network as it impacts the scalability.

2. References

[ITU-T Y.1221] ITU-T Recommendation Y.1221 Amendment 2 (2005), *Traffic control and congestion control in IP-based networks*

[ITU-T Y.2111] ITU-T Recommendation Y.2111 (2008), *Resource and Admission Control Functions in Next Generation Networks*

[ITU-T Y.2121] ITU-T Recommendation Y.2121 (2008), *Requirements for the support of flow state aware transport technology in an NGN*

3. Definitions

3.1 Terms defined elsewhere:

This Recommendation uses the following terms defined elsewhere:

~~**3.1.1.Aggregation End-point [ITU-T Y.2121]:** An end-point within the network which attaches or deletes the common Flow Aggregate Identifiers to ensure commencement of/ cessation of common routing and QoS treatment of packets. This end-point also initiates/ terminates in-band signalling to control Flow State information retained for treatment of the flow aggregate.~~

~~**3.1.2.**~~3.1.1. **Available Rate Service (ARS) [ITU-T Y.2121]:** A FSA (Flow State Aware) transport service primarily for applications that can flexibly adapt to the current available capacity and can quickly adjust their sending rate as the available capacity changes.

~~**3.1.3.**~~3.1.2. **Flow [ITU-T Y.2121]:** A unidirectional sequence of packets with the property that, along any given network link, a Flow Identifier has the same value for every packet. A flow is a sequence of packets, all of which have the same flow identifier, where the flow identifier consists of ~~the source address, destination address, protocol, destination port and source port~~such fields within the packet that uniquely identify the flow. If such a flow is an encapsulation of other flows, the QoS controls are applied only to the external flow wrapper and the internal flows must be controlled to fit within the QoS specified for the external flow wrapper.

~~**3.1.4.**~~3.1.3. **Flow Aggregate [ITU-T Y.2121]:** A hierarchical flow construct that is associated with a group of flows. The carried flows may extend beyond the flow aggregate. Except for the end nodes, flow aggregate forwarders in general do not know that they are carrying flows within the flow aggregate. All packets belonging to a given flow aggregate are commonly routed between Aggregation End-points.

~~3.1.5.~~3.1.4. **Flow admission control [ITU-T Y.2121]:.** The determination, for authorised requests, of whether or not to accept a given flow.

~~3.1.6.~~3.1.5. **Flow Identifier [ITU-T Y.2121]:** A vector comprising the values of a number of elements taken from the ~~IP, TCP/UDP~~ header fields ~~, encapsulation header, and label fields~~of incoming packets ~~attached to a packet~~ which identify the flow. The Flow Identifier for a flow within a single FSA network is unique.

~~3.1.7.~~3.1.6. **Flow State [ITU-T Y.2121]:** A set of values stored per flow identifier at each Flow State Aware node. This set of values will determine controls applied on a per flow basis, dealing with forwarding rate, delay, and congestion recovery.

~~3.1.8.~~3.1.7. **Flow State Aware Node [ITU-T Y.2121]:** A network node that is capable of maintaining Flow State and applying per-flow QoS controls, based on recognising Flow Identifier and associated signals.

~~3.1.9.~~3.1.8. **Flow State Aware Signalling Edge Function [ITU-T Y.2121]:** A function that provides the origin and/ or termination of the Flow State Aware end-to-end signalling path, and participates in requests and responses on behalf of a user application or management action. It may be located, for example, in a user end-system or at a network edge node where it serves as the signalling end-point of multiple users and associated applications. ~~Alternatively, it may be located at an Aggregation End-point where it supports the signalling requirements of flow aggregates.~~

~~3.1.10.~~3.1.9. **In-band signalling [ITU-T Y.2121]:** A mode of signalling where the signalling messages follow a path that is tied to the data packets. Signalling messages are routed only through nodes that are in the data path. ~~[Editor's note: needs clarification showing difference between in-band and path-coupled]~~Alternatively, path-coupled or out-of-band signalling follows a different path.

~~3.1.11.~~3.1.10. **Maximum Rate Service (MRS) [ITU-T Y.2121]:** A FSA transport service for applications that want packet loss characteristics to be sufficient for streamed services as soon as possible ~~but are unwilling to wait or be rejected by network admission control, if network resource for this target QoS is not available immediately~~.

~~3.1.12.~~3.1.11. **Out-of-band signalling [ITU-T Y.2121]:** A mode of signalling where the signalling messages follow a different path to the data packets and are routed to one or more nodes that are not in the data path.

3.1.12. **Preference Priority [ITU-T Y.2121]:** A parameter used to determine whether to admit a flow in case of network overload. In network overload state, flow with the lower preference priority may be rejected while the one with higher preference priority level still admitted.

~~3.1.14.~~3.1.13. **QoS Structure [ITU-T Y.2121]:** The block of QoS signalling information in a signalling packet

3.2 Terms defined in this Recommendation

3.2.1 Bounded Subnet

A FSA bounded subnet is a network section where all network nodes support Q.Flowstatesig and is totally separated from any network nodes that do not support Q.Flowstatesig with "Flow State Aware Signalling Edge Function" protocol converters, otherwise called FSA Proxies. The FSA Proxy extracts all signalling packets as data leaves the bounded subnet and introduces signalling packets as data enters the bounded subnet. It is therefore required that FSA signalling packets never can leave the bounded subnet unless tunnelled inside another protocol.

3.2.2 Ethertype

An Ethertype is a 16 bit identifier assigned by the IEEE to designate a layer 3 protocol being carried over Ethernet. The Ethertype for FSA Signalling (Q.Flowstatesig) is 0x22EF.

3.2.3 Proxy

Proxy is another name for a Flow State Aware Signalling Edge Function, a process which encapsulates incoming traffic and inserts Q.Flowstatesig signalling in one direction and in the other direction deencapsulates and deletes Q.Flowstatesig signalling. This process can be located inside a users end system or exist as a separate system in the data path. ~~.~~

3.2.~~2~~ 4 RESPONCETIMEOUT:

The timer which checks if a signal packet of type request renegotiate, or close has not been responded to and the packet should be repeated a second and third time before the Q.Flowstatesig state is dropped.

3.2.45~~3~~ Signal Packet:

Signal packets are used to carry signalling information across the bounded subnet where FSA support is provided.

3.2.46: STATETIMEOUT:

STATETIMEOUT is the maximum period between packets before the network will drop a flow.

4. Abbreviations and acronyms

~~ACK~~	~~TCP response: acknowledgement number is valid~~
AR	Available Rate Parameter
ARS	Available Rate Service
~~ATM~~	~~Asynchronous Transfer Mode~~
BT	Burst Tolerance
~~C-mode~~	~~Connection mode (of GIST)~~
CD	Change Direction
~~D-mode~~	~~Datagram mode (of GIST)~~
~~Diffserv~~	~~Differentiated Services~~
~~DSL~~	~~Digital Subscriber Line~~
DoS	Denial of Service
DP	~~Delay Priority~~A blank field in the QoS Structure
~~DSCP~~	~~Diffserv Code Point~~
~~DSLAM~~	~~DSL Access Multiplexer~~
~~ECN~~	~~Explicit Congestion Notification~~
~~EDF~~	~~Edge Discovery and Feedback~~
~~EXP/LU~~	~~Experimental/ Local Use~~
FSA	Flow State Aware
FSD	Flow Sender Depth Parameter – Number of proxies entered but not exited
~~FTP~~	~~File Transfer Protocol~~
Gbps	Giga bits per second
~~GIST~~	~~Generic Internet Signalling Transport~~

GR	Guaranteed Rate Parameter
ID1	First part of packet ID in FSA Header
ID2	Second part of packet ID
ID3	Third part of packet ID - all 3 identify the flow
~~IHL~~	~~IP Header Length~~
~~IP~~	~~Internet Protocol~~
~~IPv4~~	~~Internet Protocol version 4~~
~~IPv6~~	~~Internet Protocol version 6~~
~~IPsec~~	~~Internet Protocol security~~
Kbps	Kilo bits per second
~~LAN~~	~~Local Area Network~~
M	Modified (marker)
~~MPLS~~	~~Multi-Protocol Label Switching~~
MR	Maximum Rate Service
~~MRI~~	~~Message Routing Information~~
~~MTU~~	~~Maximum Transmission Unit~~
~~NAT~~	~~Network Address Translator~~
NGN	Next Generation Network
~~NNI~~	~~Network-Network Interface~~
~~NSIS~~	~~Next Steps In Signalling~~
~~NSLP~~	~~NSIS Signalling Layer Protocol~~
~~OPWA~~	~~One Pass With Advertising~~
PP	Preference Priority
Ptype	Packet Type in FSA Header
QoS	Quality of Service
QoS Offset	Byte count to end of packet or start of QoS Header in FSA Header
~~RSTP~~	~~Rapid Spanning Tree Protocol~~
~~RSVP~~	~~Resource Reservation Protocol~~
~~RTT~~	~~Round Trip Time~~
~~SIP~~	~~Session Initiation Protocol~~
S bit	Bit indicating a signalling packet in FSA Header
~~SLA~~	~~Service Level Agreement~~
~~SP~~	~~Service Provider~~
~~SYN~~	~~TCP connection set-up request~~
~~TCP~~	~~Transmission Control Protocol~~
TP	Type (of FSA transmission mode)
Tbps	Tera bits per second
~~UDP~~	~~User Datagram Protocol~~
~~UNI~~	~~User-Network Interface~~
~~VLAN~~	~~Virtual LAN~~

5. Conventions

None

6. High-level description

Recommendation [ITU-T Y.2121] specifies the requirements for the support of the Flow State Aware (FSA) Transfer Capability in a Next Generation Network (NGN). This transfer capability provides Quality of Service (QoS) controls that operate on a per-flow basis, allowing flows to receive different treatment depending on stored parameters at FSA network nodes. ~~These~~

~~parameters may be either statically provisioned, assigned dynamically by a FSA network node based on policy (for example, in response to congestion conditions), or requested/ confirmed by an end-system using in-band signalling.~~ The set of values of all such parameters defines the "flow state" maintained on each flow at each FSA node.

~~A flow may be managed through either static or dynamic assignment of flow state parameters. Purely static assignment of all parameters may not require anything to be signalled by an end-system that desires to initiate a new flow (this is a special property of the Maximum Rate FSA service described in clause 6). More generally, signalling will be needed in terms of reservation of bandwidth and clear-down, and also to convey dynamically-assigned parameter values to new flows.~~

~~This Recommendation defines the protocol aspects of both parameter value assignment and modification, where such values are to be assigned dynamically rather than statically.~~

Dynamic provisioning allows the ability to request a traffic contract for an ~~IP~~ incoming flow (as defined in [ITU-T Y.1221]) from a specific source node to a destination node. In response to the request, the network determines if resources are available to satisfy the request and provision the required resources. Clause 6.5 of [ITU-T Y.1221] describes the case of dynamic provisioning within a Flow State Aware network, initiated by a Flow State Aware source node.

QoS requirements (as would be applied to services supported by Flow State Aware transfer) go beyond just the delay and loss that can occur in the transport of ~~IP~~ incoming packets. The requirements include:

– bandwidth/capacity needed by the application, and

– the priority that bandwidth is maintained during congestion and is restored after various failure events.

To achieve the required QoS for FSA transfer, networks must incorporate the following functions:

1) Functions supporting the FSA Packet forwarding behaviours that are applied per flow.

2) Flow admission control recognising and processing requests for associated FSA transport services.

3) Functions supporting the signalling for allocating necessary resources for each flow.

~~4) In addition it may be noted that, if either:~~

- ~~aggregate bandwidth between the source FSA QoS Manager and receiver FSA QoS Manager is not reserved, or if~~
- ~~that aggregate bandwidth does not, at least, have a very high probability of being available when needed (especially because the core is significantly under-utilised)~~

~~then packet discards may occur within the core network that are not detected by the FSA QoS Managers at either end. So as not to exacerbate congestion conditions in the core network and in order to improve the QoS of all users, end-system elastic applications should reduce their rate of sending when experiencing significant packet discards. Furthermore, although low bit-rate real time applications (especially voice) may continue to send without reducing their rate when congestion conditions exist in the core, it would be inadvisable to allow much higher bit-rate inelastic applications to continue in such circumstances. Therefore aggregate bandwidth management across the core should be included for all service scenarios that involve high bit-rate inelastic applications.~~

Figure 6-1 shows the main functions which are involved in establishing and ceasing FSA transfer and ensuring the correct provisioning of resources to meet QoS objectives.

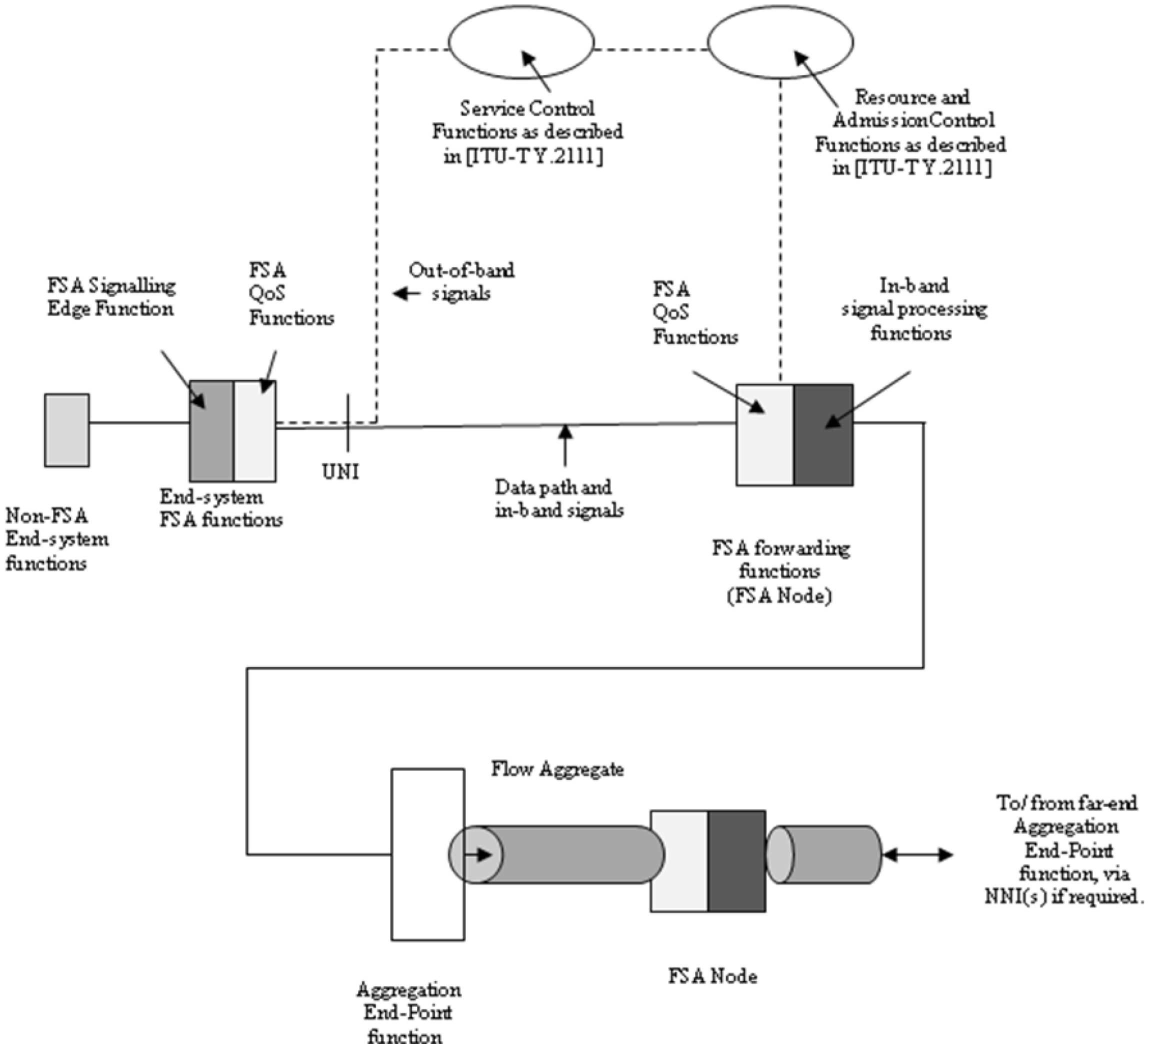

Figure 6-1 Overview of FSA functions

It may be noted that Figure 6-1 shows Aggregation End-points which are defined in [ITU-T Y.2121].

Two alternative methods of signalling are shown in Figure 6-1. One method uses in-band signalling exclusively, the other method uses out-of-band signalling for resource reservation and clear-down coupled with in-band signals to establish agreed flow state in each FSA node. Definitions of the terms “in-band signalling” and “out-of-band signalling” are included in Recommendation [ITU-T Y.2121] to describe the meaning of such terms when linked to the concept of a flow rather than pre-established channels designated for signalling or content transport. It is a Network Operator or Service Provider choice on what options are supported, but any FSA network shall be capable of at least passing every type of FSA in-band signal transparently to allow interoperation with networks that exclusively use such signals.

For the case of in-band signalling, the signalling messages are incorporated into the user data packets themselves, allowing the QoS requirements to be setup during the initial network traversal from sender to receiver (and back if needed). Each FSA node in the path examines the in-band signal and agrees to or adjusts parameter values it can support.

It is assumed that signalling messages originate and terminate at Flow State Aware Signalling Edge Functions as defined in [ITU-T Y.2121]

As specified in [ITU-T Y.2121] there are four general types of FSA transport services of which two will be used.   The first is a ~~fully guaranteed rate flow, which implies no oversubscription of network resources. The second is a~~ maximum rate flow, which allows some oversubscription but very low delay. ~~The third is a variable rate flow, where available rate is combined with a minimum rate guarantee.~~ The ~~fourth~~ second is an available rate flow, one that can be used to determine the highest rate the network can immediately support, eliminating slow-start problems.

The focus of this Recommendation is FSA signalling protocols in support of the QoS requirements of broadband and mobile service access control. The focus may be further sub-divided into three main cases, as shown in Figures 6-2, 6-3 and 6-4.

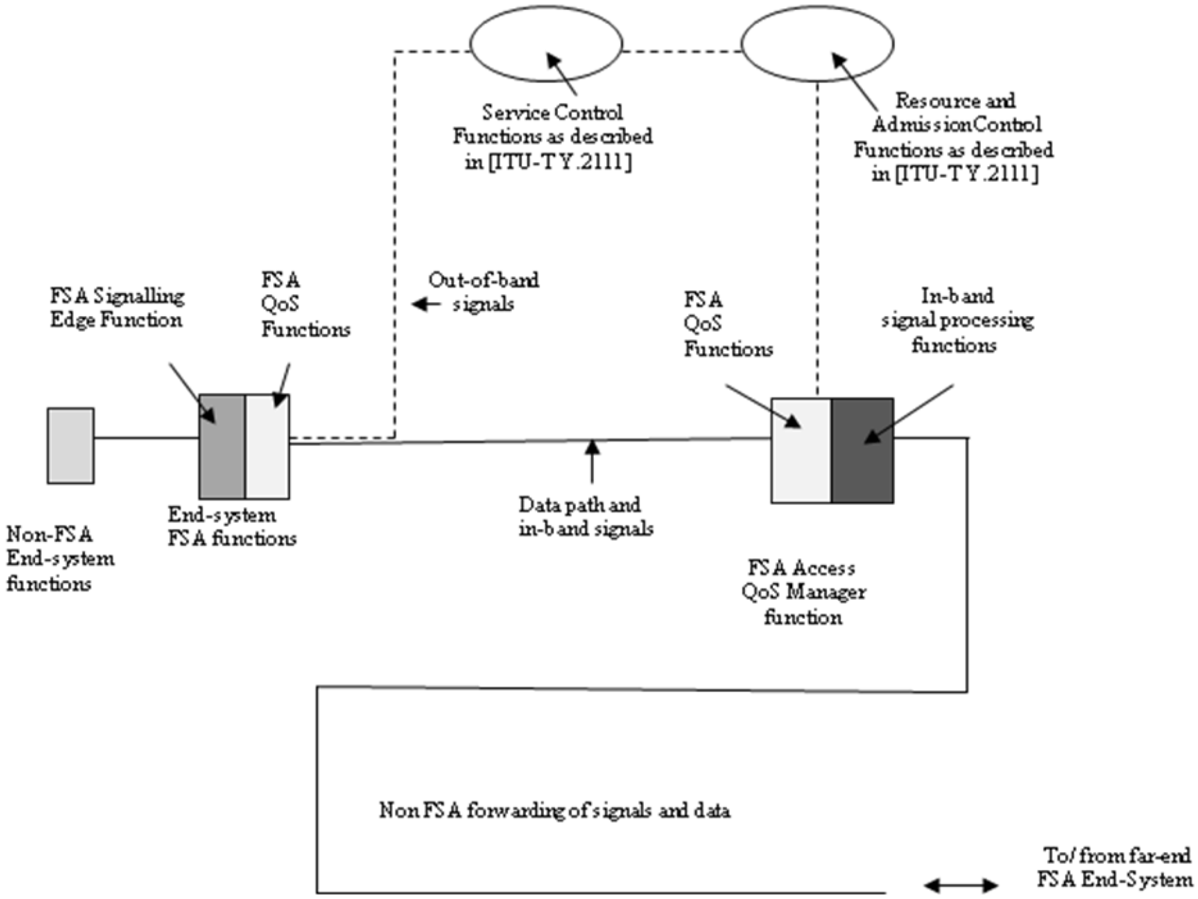

Figure 6-2 Case 1: Asymmetric access control with FSA-aware End-systems

Figure 6-2 shows Case 1 which assumes that End-systems have FSA functions to initiate signals as appropriate for each flow. The FSA Access QoS Manager is a network edge node that manages:

- the QoS of flows passing in either direction between itself and the FSA End-system.
- Optionally, the QoS down to the Non-FSA End-system functions, where the bandwidth limitations of this are known.

This Recommendation covers protocol aspects involving both near-end and far-end FSA End-systems as well as the FSA Access QoS Manager. For out-of-band signals, this Recommendation also includes Service Control and Resource Reservation functions as described in [ITU-T Y.2111].

The inclusion of far-end FSA End-systems is only related to their optimal management with respect to the limited access bandwidth. There is no inclusion of other FSA network nodes. This implies that there is an asymmetry wherein the far-end systems are not strongly constrained by access bandwidth and do not require a FSA Access QoS Manager. Case 1 is therefore a special derivative of Case 2, described next.

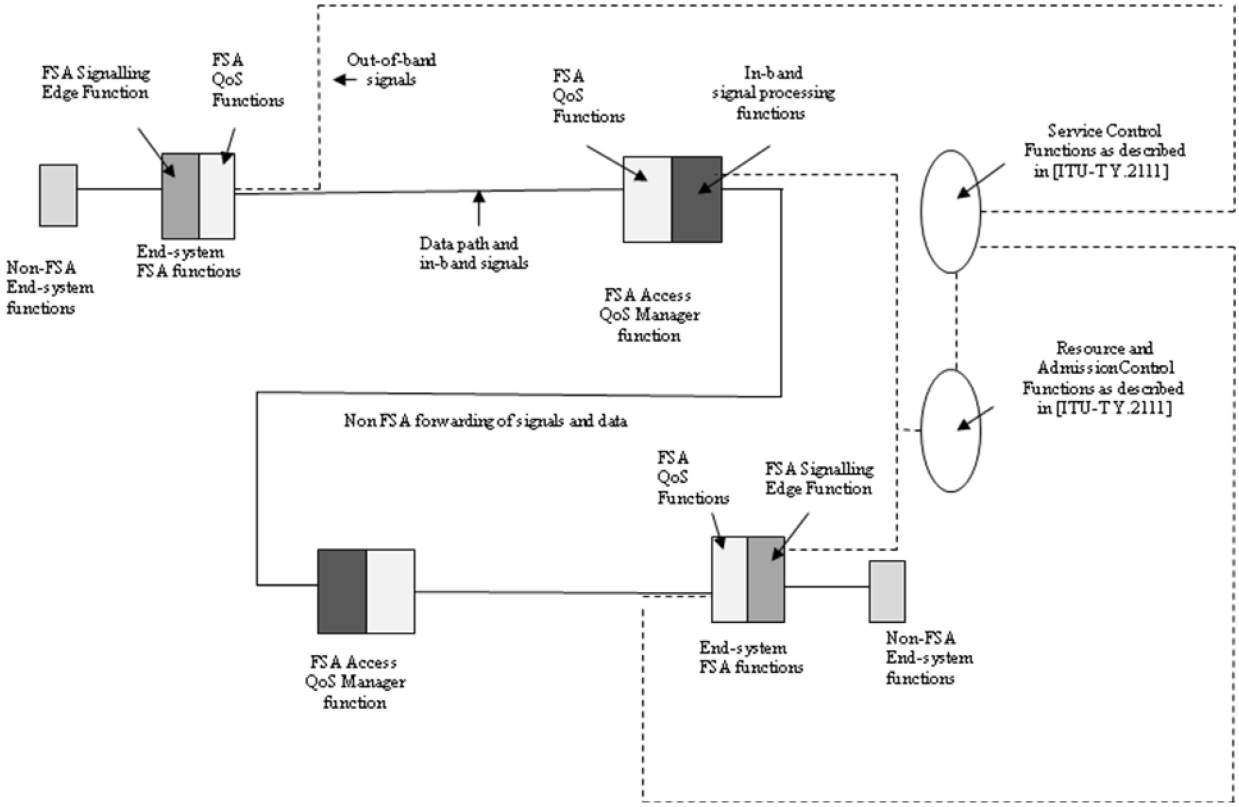
FSA Signalling Edge Function
FSA QoS Functions
Out-of-band signals
FSA QoS Functions
In-band signal processing functions
Service Control Functions as described in [ITU-T Y.2111]
Non-FSA End-system functions
End-system FSA functions
Data path and in-band signals
FSA Access QoS Manager function
Resource and Admission Control Functions as described in [ITU-T Y.2111]
Non FSA forwarding of signals and data
FSA QoS Functions
FSA Signalling Edge Function
End-system FSA functions
Non-FSA End-system functions
FSA Access QoS Manager function

Figure 6-3 Case 2: Symmetric access control with FSA-aware End-systems

In Case 2, the communicating End-systems are each managed by a FSA Access QoS Manager. Furthermore, signalling enables each Access QoS Manager to inform flow state information to an upstream Access QoS Manager.

Note that there are no routing requirements to consider in any of the cases under consideration in this Recommendation. The upstream forwarding path towards any Access QoS Manager may be realised using any appropriate technology that supports the desired QoS.

The Access QoS Manager is limited in its functions for supporting QoS on the path between an upstream and a downstream Access QoS Manager. It may manage the aggregate traffic forwarded along this path so that it is limited to a given ceiling value, where this given value is provided by network management. A further option is addressed in this Recommendation that allows the signalling of aggregate traffic requests or modifications between Access QoS Managers. Such signals can be passed in-band through the addition of Aggregation End-Point functions in the FSA Access QoS Manager (see Figure1). The requirements relating to Aggregation End-Point functions are captured in [ITU-T Y.2121].

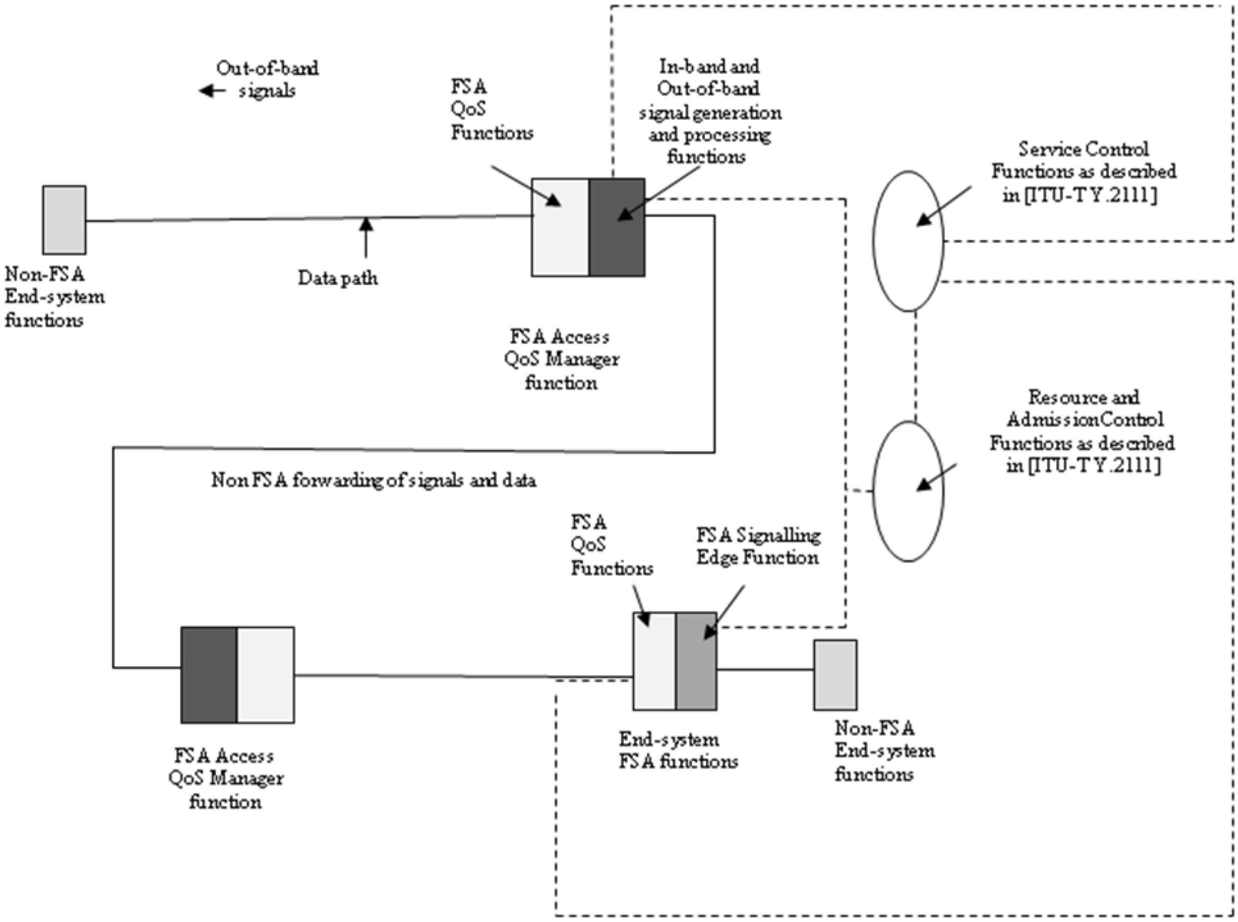
Out-of-band signals
FSA QoS Functions
In-band and Out-of-band signal generation and processing functions
Service Control Functions as described in [ITU-T Y.2111]
Non-FSA End-system functions
Data path
FSA Access QoS Manager function
Resource and Admission Control Functions as described in [ITU-T Y.2111]
Non FSA forwarding of signals and data
FSA QoS Functions
FSA Signalling Edge Function
FSA Access QoS Manager function
End-system FSA functions
Non-FSA End-system functions

Figure 6-4 Case 3: Non FSA-aware End-system

In Case 3, the FSA Access QoS Manager provides “intelligence” functions to support non-FSA end systems. In this case, the FSA Access QoS manager will detect new flows and can apply policies relating to such new flows. These policies can be based partly on initial, or default QoS parameters that may be associated with a particular end user. Policies may also allow modification of flow state values based on, for example, measurement of flow rate.

To manage upstream FSA end systems, in-band signals are inserted into the forwarding path towards the upstream source. Out-of-band signals may also be generated that may result in reservations being granted. This, in turn, may result in a change of status and in-band modification of flow state that takes account of granted resources.

Appendix I describes service scenario examples that may be associated with all three cases described above.

Editors Note – Redo this prior section and condense it.

7. Protocol description

Clause 6 details the main models for access QoS management. In this clause we will first formulate the parameters that can be requested or modified relating to the QoS treatment of an individual flow or flow aggregate.

7.1 QoS parameters related to flow-based QoS management

7.1.1 Overview of the QoS Structure

The critical part of the in-band QoS Signalling Protocol consists of a QoS Structure that represents a set of fields containing values and indications on the requested flow treatment or on network responses to this request.

The uses of the QoS Structure are as follows:

7.1.1.1 Travelling in the first packet, the QoS Structure is examined by each FSA node to determine if the QoS request can be supported. If it can be supported the packet proceeds to the next FSA node without change. If the FSA node cannot provide the rates or delay requested, it reduces the request in the QoS Structure to what it can support~~. It then forwards the packet. Be sure to recompute any checksum which includes any modified field~~.

~~**7.1.1.2** The QoS Structure will be provided in the first packet. It should not be encrypted even when the remaining data information is encrypted. In this way the FSA Nodes can process flows quickly and efficiently with the correct QoS.~~

7.1.1.~~3~~ 2 If any FSA node finds it cannot continue to accept the rates it has approved, it may selectively discard packets from that flow. This loss of packets is the indication that the receiver may use to create a response towards the source that includes a QoS Structure indicating a reduced new rate if appropriate.

Note that the FSA Edge QoS Manager will need to adjust the Available Rate of long duration flows as load builds up.

7.1.2 Overview of QoS Structure parameters

~~The IPv4~~ Request/Response QoS Structure is 13 fields as follows:

- **AR**: Available Rate – floating point rate for network assigned rates.
- **GR**: Guaranteed Rate –floating point rate for requested guaranteed rate.
- **PP**: Preference Priority – indicates the override or relative rate priority of the flow.
- **CD:** Change/Direction field – indicating an action type, such as request or response.
- **TP:** Type of FSA transport, such as available rate, or maximum rate.
- **Second QoS structure attached:** Indicates two QoS structures are included in the same packet, for example so that one relate to a response to a forward request and one relates to a separate request for QoS in the reverse direction.
- Security Structure attached. Indicates that a Security Structure follows the QoS structure(s).
- **QoS Version:** QoS Structure Version Field – Set to 2
- **M:** Modified marker. Set to 0 by sender on request. Set to 1 by FSA nodes if any field changed during a request or renegotiate. Set to 0 and not changed on Response.
- **FSD:** Flow Sender Depth – Number of proxies entered but not exited

7.2 Message sequences related to FSA Transport service set up and clear down.

Editors Note: Rework section to it shows bounded subnet and show proxies and Internet Core

Add separate figure with core & subnet.

First a figure on the global subnet and then these figures on the internal subnet.

7.2.1 FSA Maximum Rate Service (MRS), sender requests set-up.

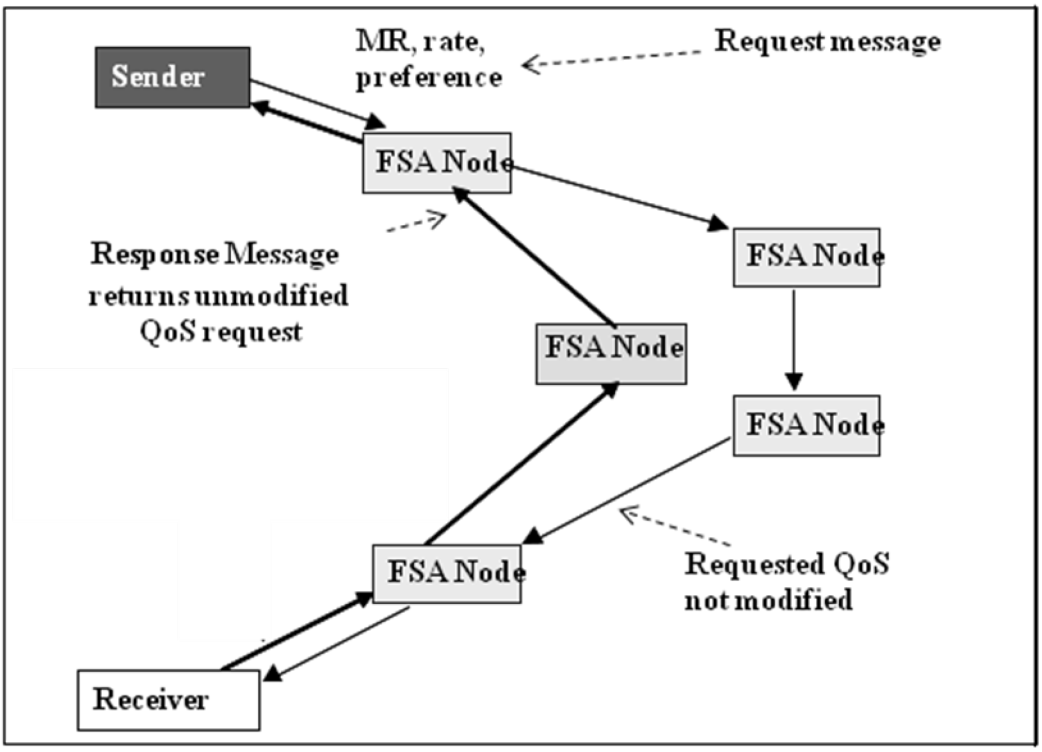

Figure 7-1 Maximum Rate message sequence

The "Sender" and "Receiver" in Figure 7-1 and 7-2 are the parts of a proxy that manage the Q.Flowstatesig signalling. The figures show the signalling process within the Q.Flowstatesig bounded subnet.

In Figure 7-1, the source (Sender) forwards to the Receiver a request packet with the requested QoS. This packet will be read by intermediate FSA nodes as well as the Receiver and includes the following QoS parameters:

- Requested FSA Transport Service = MRS
- Requested Maximum Rate
- Requested preference priority

After sending the request, if no response is received within RESPONSETIMEOUT then repeat the request for up to three times. If that fails, the FSA signalling should be terminated by sending a close (CD=7) packet and proceeding with the data flow without any FSA. If a valid response is received, then the data packets may start to be sent at the rate received in the response. If the response has zero as the rate, no data may be sent. The default value of RESPONSETIMEOUT is one second.

7.2.2. MR clear-down

Flow state times out and is deleted if data packets corresponding to that flow are not seen at an FSA node for a period equal or greater than STATETIMEOUT. The default value for STATETIMEOUT is 2 seconds.

7.2.3 FSA Available Rate (AR), sender requests set-up.

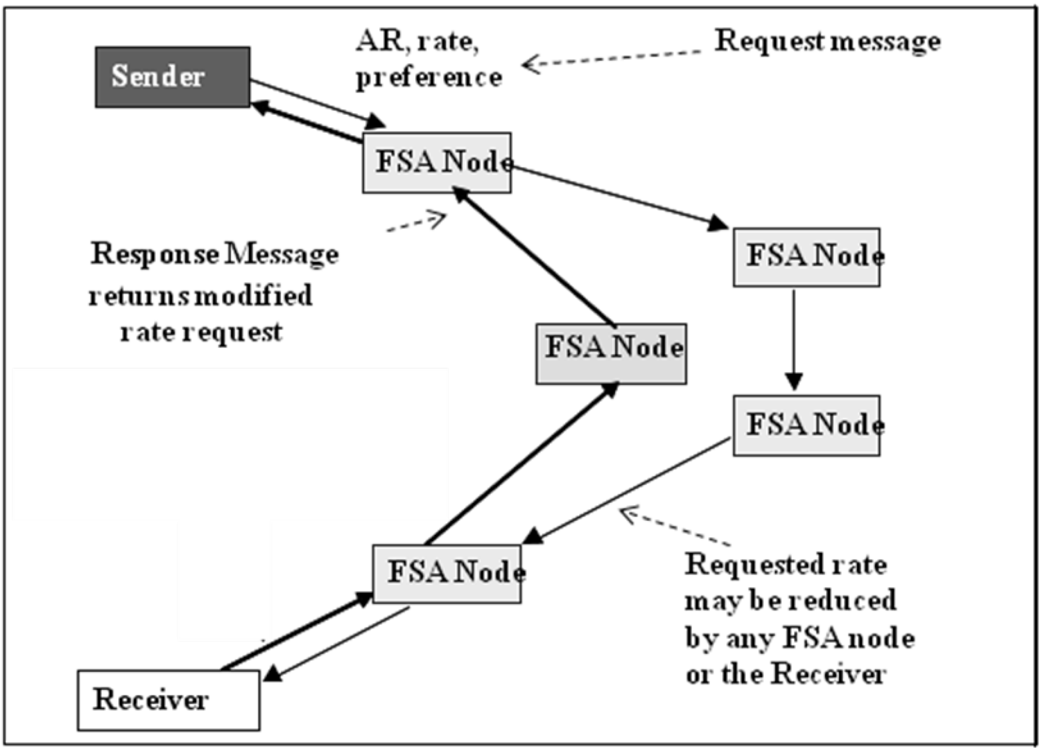
AR, rate, preference
Request message
Sender
FSA Node
Response Message returns modified rate request
FSA Node
FSA Node
FSA Node
FSA Node
Requested rate may be reduced by any FSA node or the Receiver
Receiver

Figure 7-2 Available Rate message sequence

In Figure 7-2, the Sender forwards to the Receiver a sequence of request packets. All such packets, after the first, are treated as re-negotiate packets, with CD=5 (see also sections 7.1.3 and 7.2.2). It is a requirement that the Sender forwards a new AR request packet after every 128 data packets or after 1 second since it sent the last AR request packet (whichever is the sooner). Each request packet will be read by intermediate FSA nodes as well as the Receiver and includes the following QoS parameters:

- Requested FSA Transport Service = ARS
- Requested Available Rate = AR
- Requested preference priority

Following each request, the sender may maintain its previous rate until a response message is received and, thereafter must shape its forwarding to the new rate received in the response packet.

ARS has no initial rate and so the sender must wait for the first response packet before forwarding data packets.

If no response packet is returned to the sender, as determined following the RESPONSETIMEOUT at the sender, then the request packet must be retried 3 times and then if no response packet is received the FSA signalling should be terminated by sending a close (CD=7) packet and proceeding with the data flow without any FSA signalling.

7.2.4 AR clear down.

Flow state times out and is deleted if data packets corresponding to that flow are not seen at an FSA node for a period equal or greater than STATETIMEOUT.

7.3 Procedures and Format of Layer 2 FSA Ethertype packets

7.3.1 FSA Ethertype Data Packets

At the sending FSA Signalling Edge Function, if the Ethertype of arriving packets is not the FSA Ethertype, then a 4 byte FSA header should be inserted between the Ethernet header and the packet body. If the Ethertype of the incoming packet was already the FSA Ethertype, then it will already have a FSA header and no restructure is required. The FSA header has parameters of Ptype (Packet type), ID1 (first packet ID), and QoS Offset (Bytes to the end of the packet or to the start of the QoS header). These FSA Parameters must then be computed (see Section 7.3.3) and stored into the FSA header. Then the Ethertype in the Ethernet Header must be set to the FSA Ethertype. For data packets the S bit (bit 0 of the FSA header) must be set to zero. The original packet is thereby encapsulated in a FSA header. The format of a FSA data packet is:

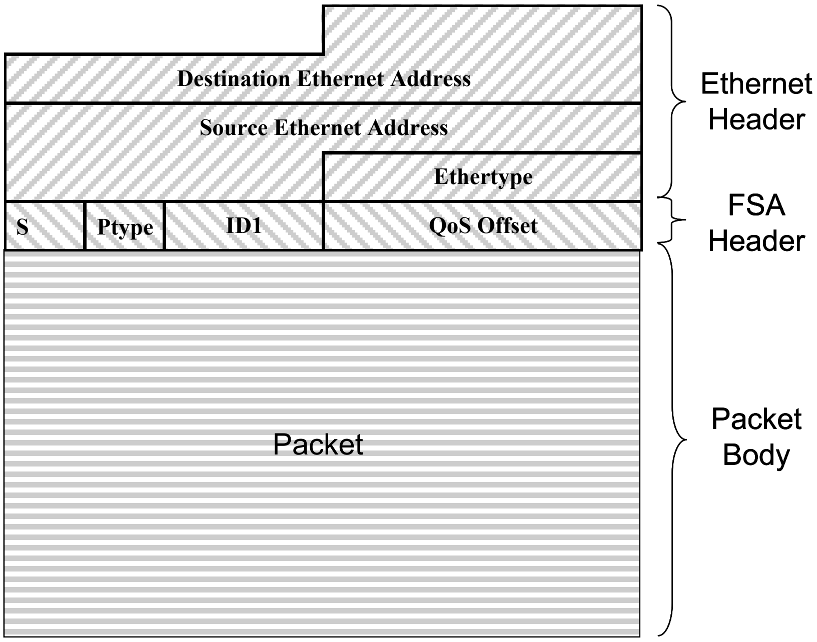

Figure 7-3: Data Packet Format

7.3.2 Signalling Packets

Signalling packets are indicated by the S bit, bit 0 of the FSA header, being set to one. Also in a signalling packet the Ptype, ID1 and QoS Offset need to be set (see section 7.3.3). Following the FSA header, signalling packets need to have a section copied from one of the flow's data packets which includes whatever information is needed identify the flow. The length of the packet section is the number of 32 bit words specified by the QoS Offset. Right after this portion of the packet follows the 1st QoS Structure. The Flow Type in the FSA header can be found in the packet. The signalling packets structure is:

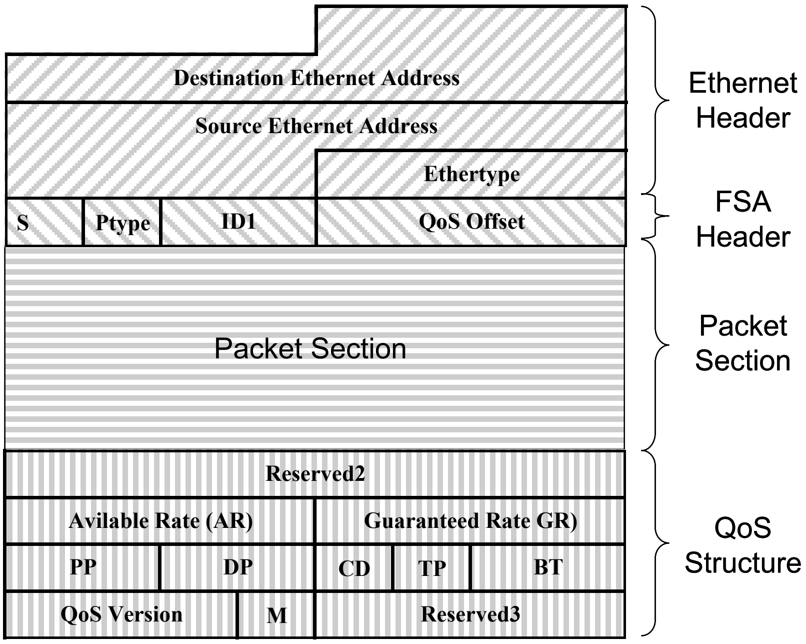

Figure 7-4: Signalling packet structure

The signalling packet components are:

- **FSA Header:**
  - **S:** Signalling Bit (4 bits) S=0x8 for a signalling packet, S=0x0 for a data packet
  - **Ptype:** (4 bits) Packet type code
  - **ID1:** (8 bits) Flow identification, first part
  - **QoS Offset:** Number of bytes between the Ethernet header and the QoS Structure
- **QoS Structure:**
  - **Reserved2:** (32 bits) Reserved and set to zero
  - **Available Rate (AR):** (16 bits) Floating point rate assigned by the network. Floating point format is covered in Annex A
  - **Guaranteed Rate (GR):** (16 bits) Floating point rate assigned by the network. Floating point format is covered in Annex A
  - **PP:** (8 bits) Preference Priority – Indicates a relative rate priority for ARS ~~and VRS~~ and an acceptance priority for MRS ~~and GRS~~ flows; 64 levels in the high order 6 bits (bits 0-5), 0=lowest, 63=highest. The two low order bits are reserved and should be set to zero
  - **DP:** (8 bits) Reserved and set to zero
  - **CD:** (4 bits) Change/Direction field – Bit 0: set to zero, Bits 1-3: 0= No action required, 1=Request at the start of a flow to negotiate the QoS parameters, 2= Response returning agreed parameters to the Sender, 3=Reserved, 4=Reserved, 5=Renegotiate, a sender request to renegotiate rates on the continuation of a flow, 6=Reserved, 7=Close, sent by Sender to close out FSA Network Nodes state

- **TP:** (4 bits) Bits 0-1: Reserved and set to zero; Bits 2-3: Type of flow – 0=Available Rate Service (ARS, 2=Maximum Rate Service (MRS)
- **CH:** (4 bits) Bit 2: Second QoS Attached, 0 = single QoS Structure, 1 = Second QoS Structure follows; Bit 3: Security Structure attached, 0 = No Security Structure follows, 1 = Security Structure follows
- **QoS Version:** (12 bits) QoS Protocol Version Field – set to 2
- **M**: (4 bits) Bit 3: Modified marker. Set to 0 by Sender on Request or Renegotiate. Set to 1 by FSA Network Nodes if any field changed during a request or renegotiate; Bits 0-2: Flow Sender Depth (FSD) – Set to 1 by Sender, increased by one entering a proxy and decremented by one when leaving a proxy
- **Reserved3:** (16 bits) Reserved and set to 0

7.3.3 Procedure to Determine Parameters to store in a FSA Header

The sender (the sending FSA signalling edge function) must determine the value of the parameters which will go into the FSA Header. The following procedure uses Table 7-1 and Table 7-2 to help compute QoS Offset, ID1, and Ptype. The function C(n) in these computations means the contents of byte n. C(n-m) means the 16 or 32 bit number stored in bytes n through m. The function Cquad(n) means the contents of quad n (4 bits). All these count from the end of the Ethernet Header.

The first step is to lookup in table 7-1 the Ethertype that was in the arriving packet and compute Ptype, ID1, and Q based on the formulas in that row.

Ethertype	Ptype=	ID1=	Q=
800	1	C(14)	4*CQuad(10)+4
86DD	2	C(11)	44
22EF	CQuad(2)	C(2)	C(3-4)

Table 7-1: Lookups based on Ethertype

Once Ptype, ID1, and the initial value of Q are determined, the next step is to finish evaluating ID1 and Q. Normally this is one step but on occasion the process needs to be repeated several times to determine the correct values of ID1 and Q. This is called a cycle and means that once ID1 and Q are updated the lookup should repeat.

If ID1=	new ID1=	new Q=	Next Action
0	C(Q)	Q=Q+C(Q+1)*	cycle
1	C(Q)	Q=Q+C(Q+1)	cycle
4	C(Q+10)	Q=Q+20	cycle
6	ID1	Q=Q+4	done
17	ID1	Q=Q+4	done
41	C(Q+7)	Q=Q+40	cycle
43	C(Q)	Q=Q+C(Q+1)	cycle
44	C(Q)	Q=Q+C(Q+1)	cycle
50	ID1	Q=Q+8	done
51	ID1	Q=Q+4	done
60	C(Q)	Q=Q+C(Q+1)	cycle
other	ID1	Q	done

* If not first Table 7-2 lookup then Q=Q+1

Table 7-2: Lookups based on ID1

Once the lookups are done, Ptype, ID1, and Q should be saved into the FSA header. Q is saved into QoS Offset. To extend Q.Flowstatesig to other packet types, it is only necessary to augment these tables.

7.3.4 Procedure at the Sending FSA Signalling Edge Function

When a new flow is received at the sending FSA Edge, a Request type signalling packet needs to be sent before the packet. The signalling packet is created as specified in section 7.3.2 using a copy of the first packet. A 16 byte QoS Structure should then be affixed to the end of the packet section. The service type, rates, and QoS parameters should be set and the CD field set to 1 for Request. The data packet should be held and the signalling packet sent.

After sending the request signalling packet, the sender should wait for the receiver to return a response packet. If a response is not received within the RESPONSETIMEOUT period the sender should resend the request a second and then a third time. If still there is no response, the sender should abort using the FSA protocol and send the data packets of the flow without change.

When a response packet is received, the sender should set the sending rate and QoS as returned in the response packet and proceed to encapsulate arriving data packets with the FSA Ethertype and the FSA header as in 7.3.1 and send them forward.

7.3.5 Procedure at FSA Network Nodes

For each incoming packet which has a FSA Ethertype, a FSA Network Node should compute the full packet ID to determine the relevant flow. The full ID consists of three segments; ID1, ID2, and ID3. The FSA header of the packet has Ptype, ID1, and Q=QoS Offset. ID2 and ID3 are found as follows:

If Ptype=	ID2=	ID3=
1	C(13-20)	C((Q-4)-(Q-1))
2	C(9-40)	C((Q-4)-(Q-1))

Table 7.4: ID2 and ID3 Lookup

The full ID completely defines a specific individual flow. Next, if the FSA header has S=1 then the QoS Structure found at byte Q should be examined and if a request or renegotiate the rates and QoS requested should be adjusted downward to the best available. If this is different than is in the QoS Structure, the new values should be inserted and the M bit set to a one indicating that a change has been made. The packet should then be forwarded.

7.3.6 Procedure at the Receiving FSA Signalling Edge Function~~ for TCP~~

When the Receiving FSA Edge receives a packet with the FSA Ethertype it first needs to determine if it is nested. Upon the receipt of a Request Signalling packet it should examine the value of FSD in the M field in the QoS Structure. If FSD>1 then the packets of this flow should be passed forward intact. A flag should be set to note this for the data packets in the flow. If FSD=1 then the receiving FSA Edge should create a response packet when it receives a Request packet. The Response packet is the Request packet with the Ethernet addresses reversed and the CD field in the QoS Structure changed to Response (2). The other fields of the QoS Structure should remain as received except that the rate can be decreased to the maximum the receiver can support. The same process of returning Response packets should be used when receiving Renegotiate packets. The signalling packets should no be forwarded. When data packets are received, the Ethertype related to the Ptype should be saved into the Ethernet header, the FSA header removed and the packet forwarded. The Ethertype is found by a lookup into Table 7-3.

If Ptype=	Ethertype=
1	800
2	86DD

Table 7-3: Ethertype Lookup

7.3.7 Renegotiate Signalling Packets (All Services)

After the first16 packets or 1 second, whichever is sooner, the sender must insert a Renegotiate packet to allow the Network FSA Nodes to adjust the rate. After that, the sender must send a Renegotiate packet every 1 second or every 128 packets, whichever is sooner. The receiving FSA Edge should return a Response to each Renegotiate. When a Response is received the sender must adjust to the new rate (or less).

8. Security Considerations and Requirements

8.1 Authentication

Flow State Aware user authentication requirements are addressed in [ITU-T Y.2121]. FSA nodes within a domain may authenticate to peer FSA nodes within the domain. FSA nodes communicating as peers across a domain boundary should authenticate with each other. Authentication security requires an additional security data structure which, if the CH bit 3 it set to a one, would follow the QoS structure(s). This additional security data structure is for further study.

~~8. Security Considerations and Requirements~~

~~8.1Authentication~~

~~Flow State Aware user authentication requirements are addressed in [ITU-T Y.2121]. FSA nodes within a domain may authenticate to peer FSA nodes within the domain. FSA nodes communicating as peers across a domain boundary should authenticate with each other. Authentication security requires an additional security data structure which, if the CH bit 3 it set to a one, would follow the QoS structure(s). This additional security data structure is for further study.~~

8.2 Authorisation

Flow State Aware authorisation requirements are addressed in [ITU-T Y.2121].

8.3 Data Confidentiality

In the case where user flows with data confidentiality requirements also invoke the ARS, and MRS , the parameters describing the in-band service request shall not be encrypted.

8.4 Data Integrity

Flow State Aware parameters may be protected against unauthorized modification while in transit. Flow State Aware parameter requests may be protected against replay attacks, in conjunction with data integrity protection binding a set of Flow State Aware parameters to a specific flow.

8.5 Accountability

It is recommended that Flow State Aware service invocations are logged, including the identity of the entity requesting the service, the actual service request, and actual service granted.

8.6 Availability/accessibility

Flow State Aware services shall respect the priority preference of each authenticated entity in making admission decisions.

8.7 Privacy

It is recommended that Flow State Aware services ensure the privacy of user specific policy profiles defining QoS parameter limits and privileges.

8.8 Protection against network attacks, from within or outside

It is recommended that Flow State Aware services include mechanisms to protect against malformed service invocations and to mitigate Denial of Service (DoS) attacks.

9 State Diagrams for FSA Signalling Edge Functions at origination and destination

FSA Signalling Edge Functions (also called FSA Proxies) separate the standard ~~IP~~ network from the FSA supported bounded subnet. They receive ~~standard IP on the IP~~packets of any incoming protocol on the input side, add the FSA signalling packets and send FSA signalled traffic on FSA side. They also receive FSA signalled traffic from the FSA side, delete the signalling and send ~~standard IP on the IP~~the original packets on the other side. Each such system contains two processes, a Signalling Origination process and a Signalling Termination process. These two processes are shown in section 9.1 and 9.2 using a state machine presentation. One difference between a FSA Edge Function and where the FSA process is installed into the sending computer is that the edge function does not have as much information as to the intent of the user. Thus the FSA Edge Function should classify ~~TCP~~ available rate flows as ARS and ~~UDP~~ fixed rate flows as MRS. The duration of the flows is not an issue. The rate of MRS flows can be often be determined from the ~~DSCP code which is~~ information in the incoming ~~IP~~ packet ~~from the IP network~~, but if not, the peak rate may be set to that of a high rate video and since the network do not reserve capacity for MRS flows, there is little harm is specifying a rate higher than needed. The only impact will be that a MRS request may be rejected if it is higher than the remaining capacity in a network trunk. For typical high capacity trunks, this is no problem.

9.1 Origination State Machine

FSA signalling is originated by the source of the data flow. The FSA signalling component will maintain a separate FSA signalling origination for each outbound connection. Figure 9.1-1 shows the four states for the sender.

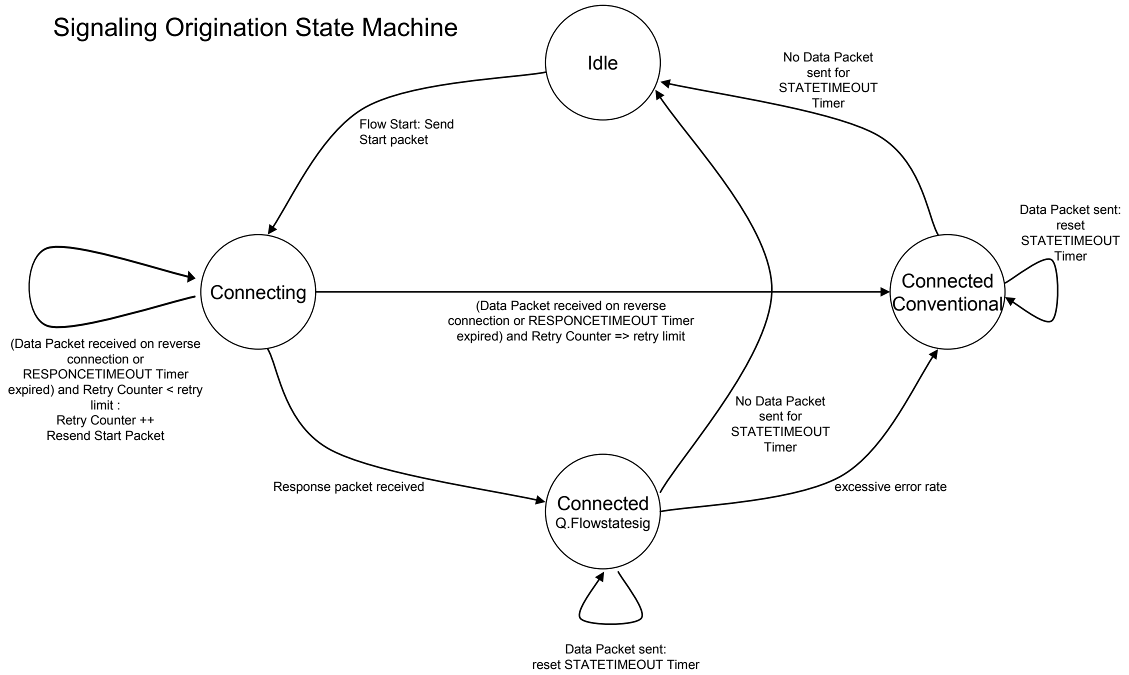
Signaling Origination State Machine
Idle
No Data Packet sent for STATETIMEOUT Timer
Flow Start: Send Start packet
Data Packet sent: reset STATETIMEOUT Timer
Connecting
Connected Conventional
(Data Packet received on reverse connection or RESPONCETIMEOUT Timer expired) and Retry Counter => retry limit
(Data Packet received on reverse connection or RESPONCETIMEOUT Timer expired) and Retry Counter < retry limit : Retry Counter ++ Resend Start Packet
No Data Packet sent for STATETIMEOUT Timer
Response packet received
excessive error rate
Connected Q.Flowstatesig
Data Packet sent: reset STATETIMEOUT Timer

Figure 9.1-1: Signal Origination State Machine

Whenever a packet is received ~~from the TCP/IP stack~~at the input a ~~five-tuple~~full ID1-3 based lookup is used to determine if there is a connection state entry for the flow the packet belongs to. If there is no originating state machine state record then one will be created.

The FSA signalling origination state machine will send a start packet for the new connection and then transition to the connecting state. The FSA signalling origination state machine will remain in the connecting state and resend start packets each time the Q.Flowststesig RESPONCETIMEOUT timer expires until the retry limit is exceeded or a Q.Flowststesig response packet is received.

9.2 Termination State Machine

The transition out of the idle state on the FSA signalling termination state machine is triggered by the reception of a data or start packet for a new flow. When a data packet is received it is not known whether the packet belongs to a Q.Flowstatesig flow or a conventional flow. The start packet could have been lost of potentially delayed in the network. The connection state is used to wait for a clear indication that the flow is either a Q.Flowstatesig flow or a conventional flow. Since all flows ~~in IP~~ are expected to be uni-directional, the origination and termination processes both exist on each system and run independently. Figure 9.2-1 shows the termination state machine.

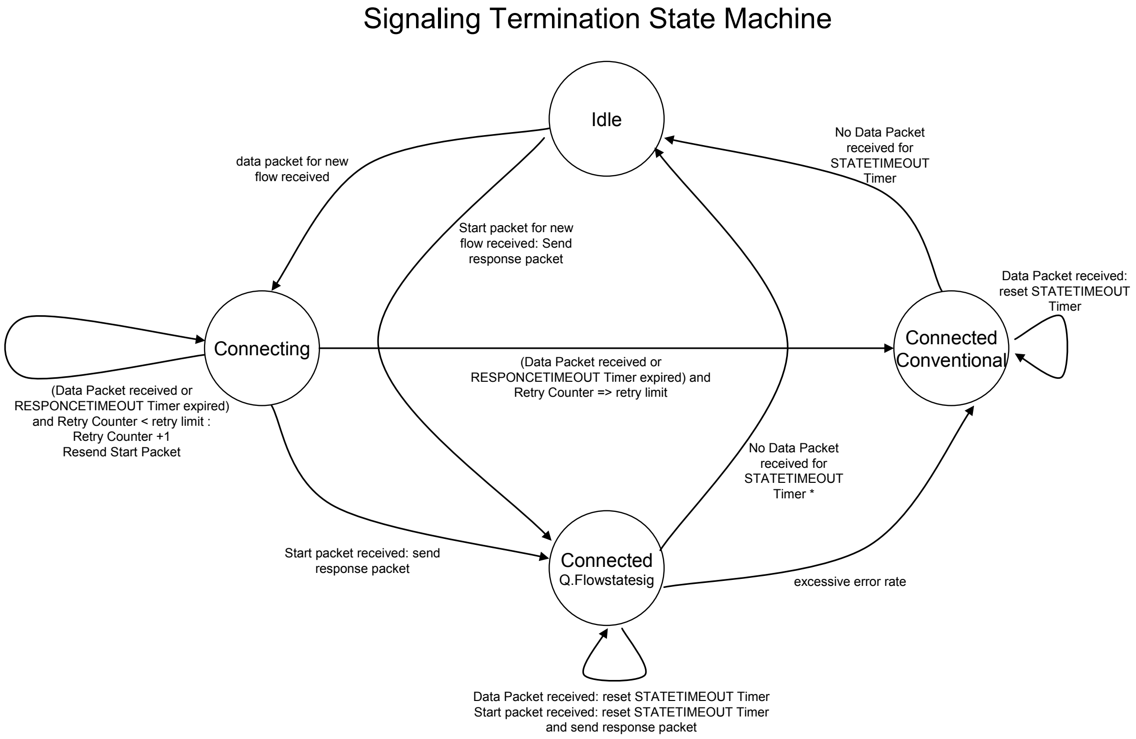
Signaling Termination State Machine
Idle
data packet for new flow received
Start packet for new flow received: Send response packet
No Data Packet received for STATETIMEOUT Timer
Data Packet received: reset STATETIMEOUT Timer
Connecting
Connected Conventional
(Data Packet received or RESPONCETIMEOUT Timer expired) and Retry Counter => retry limit
(Data Packet received or RESPONCETIMEOUT Timer expired) and Retry Counter < retry limit : Retry Counter +1 Resend Start Packet
No Data Packet received for STATETIMEOUT Timer *
Start packet received: send response packet
Connected Q.Flowstatesig
excessive error rate
Data Packet received: reset STATETIMEOUT Timer Start packet received: reset STATETIMEOUT Timer and send response packet

Figure 9.1-2: Termination State Machine

The connected state for the FSA signalling termination state machine must also deal with reception of additional Start packets to deal with the potential loss of response packets. When a Start packet is received in the connected Q.Flowstatesig state a response packet is sent.

9.3 Signalling Renegotiation

To keep the state diagrams readable renegotiation was not included in either diagram.

Renegotiation is handled in the connected Q.Flowstatesig states. The FSA signalling origination state machine maintains a renegotiation packet counter and renegotiation timer. The FSA signalling origination state machine sends a renegotiation packet when the renegotiation count is exceeded or the renegotiation timer expires. The default value for the renegotiation packet count is 128 and the default value for the renegotiation timer is one second. The FSA signalling termination state machine sends a response packet every time a renegotiation packet is received.

10 FSA QoS Manager

As the Signalling packets flow through the bounded subnet, each network node needs to add a FSA Flow Manager capability which will check the rates and ~~QoS~~ priority requests and if necessary, downward adjusts the requests to the rates ~~and QoS~~ which can be supported on the next link. This capability is best shown as a flow chart which acts on each packet passing through. This is shown in Figure 10-1

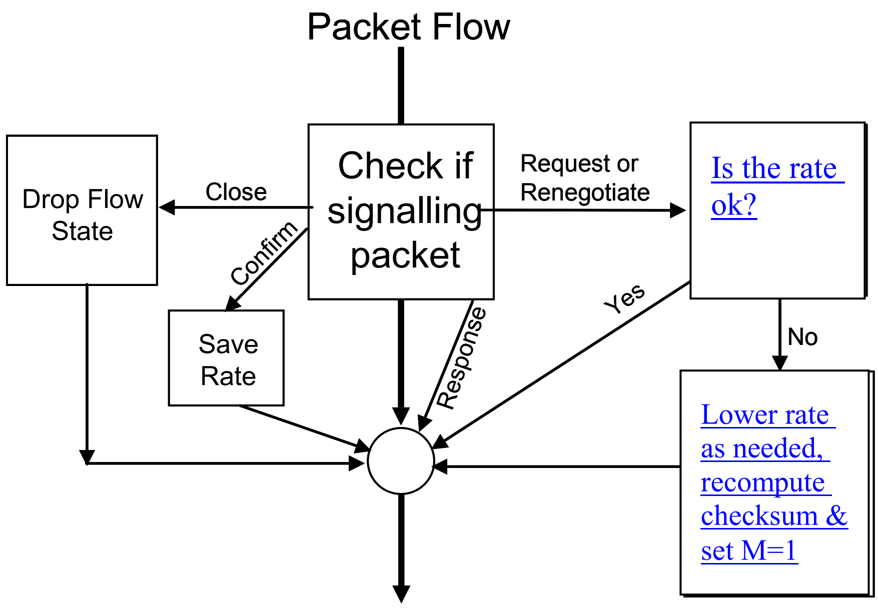

Figure 10-1 FSA QoS Manager Flow Chart

Annex A: Rules for encoding floating point rates

(This annex forms an integral part of this Recommendation)

A.1 AR and GR – Encoding Floating Point

AR and GR fields are encoded according to the following rules:

- The most significant bit of AR and GR is zero and reserved.
- The next bit is nz where nz=1 indicates if the number non-zero and nz=0 means the number is zero.
- The next 5 bits of AR and GR are the exponent E,
- The next 9 bits are the mantissa M.
- The rates AR or GR= (1+M/512)*2^E kilobits per second.
- All zero is interpreted as zero.
- Since E can be as large as 31 and M 511, the maximum rate is 4.291 Tbps. The lowest positive rate would be 1 Kbps since 0 is zero.

This is the same type of floating point number used in ATM except that in ATM its units are cells per second. Since there are no cells in ~~IP~~Q.Flowstatesig, the units are kilobits per second. A benefit of using kilobits/sec as the units is that 64 kilobits per second is represented exactly, as are 2^n multiples of 64 Kbps.

- ~~Rates are to be measured for TCP over the round trip time (RTT) if possible so as to determine the current TCP average rate. Since TCP sends a burst and then waits for the ACK the RTT can be deduced from the longest inter-packet gap. If the RTT cannot be determined and for UDP where there is no gap, the default rate measurement time is the value of a exponential decay filter with a half life of 500 ms (satellite RTT).~~
- The rate is measured using all the bytes in the ~~IP~~ packet. The rate on a given trunk would include in some cases the Ethernet ~~or MPLS~~ header~~s~~. This overhead must be accounted for at each interface so as to compute the trunk load correctly. However, the ~~IP~~ packet byte count is a usually constant so the user will know what is being asked for and supported. One exception is when fragmentation occurs and in this case those FSA nodes before the fragmentation will have set their rates based on the whole packet, the same as the sender. FSA nodes after fragmentation is done may see increased traffic due to the fragmentation and thus lower the rate. The lower rate will be enforced. A second case exists if the sender uses header compression. If so, it is the compressed packet size that determines the rate since this is the trunk load incurred across the network. Another case is when header compression is used mid-path to reduce the load on a particular link. Here, the rate marked in the QoS Structure should be computed based on the uncompressed header, but the link rate required may be less. Since the average packet size is important to compute the ratio, the first estimate should be based on a best guess assuming large packets, and as the flow proceeds if the packet size is smaller, the extra link capacity can be assigned to other flows.

Appendix 1

Mapping FSA services to Y.1221 Transfer Capabilities

Maximum Rate (MRS) Service

Key attributes of this service that determine the underlying transfer capability required are;

Q.Flowstatesig Characteristic	**Implication on Transfer capability**
After a Responce packet, the sender may transmit data packets at any speed up to the requested maximum rate.	Until resources are confirmed as available with a response packet for the flow no QoS commitments can be met.
Only a maximum rate is requested	Once resources have been confirmed then this become a single token bucket transfer capability (Rp= Max Rate, Bp= the default Service Provider value). From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming.
If a higher rate is required it must send a new request. Having sent the request the sender may send at the new rate immediately	If all packets up to the new maximum rate are considered conforming, then QoS commitments cannot be met until resources at the new rate have been confirmed.

The initial request would signal commencement of a flow in the “discard first” state with Rp= max rate.

Default values for Bp and the maximum packet size, M, are pre-defined by the Service Provider rather than requiring to be signalled.

MRS Conclusion

The CDBW transfer function of Y1221 6.5 provides the correct transfer functionality for this service.

Available Rate (ARS) Service

Key attributes of this service that determine the underlying transfer capability required are;

Q.Flowstatesig Characteristic	**Implication on Transfer capability**
The sender forwards to the receiver a sequence of request packets. All such request packets after the first are labeled as re-negotiate packets.	Each flow request could be regarded as a stand alone transfer capability request.
AR has no initial rate and so the sender must wait for the first response packet before forwarding data packets	Packets are not conforming until a Response is received from the first request packet that resources are available.

QoS Parameters • Requested available rate • Highest Available Rate (per Preference Priority)	A single bucket GBRA(Rp,Bp), with Rp = Highest Available Rate and Bp set to the default Service Provider value. From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming. Y1221 6.2 supports QoS commitment for conforming packets.
It is a requirement the sender forwards a new AR request after every 128 data packets or 1 sec since it sent the last AR request packet. Following each request, the sender may maintain its previous rate until a response message is received and, thereafter must shape its forwarding to the new rate in the response packet	Each flow request can be regarded as a stand alone transfer capability request from one response to the next. From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming.

Default values for the Highest Available Rate per Preference Priority, Bp and the maximum packet size, M, are pre-defined by the Service Provider rather than requiring to be signalled.

ARS Conclusion

The SBW transfer function of Y.1221 6.2 provides the correct transfer functionality for this service.