International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.2121

(01/2008)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

Next Generation Networks – Quality of Service and performance

Requirements for the support of flow-state-aware transport technology in NGN

Recommendation ITU-T Y.2121

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	**Y.2100–Y.2199**
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T Y.2121

Requirements for the support of flow-state-aware transport technology in NGN

Summary

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). The FSA transfer capability provides QoS controls that operate on a per-flow basis, allowing flows 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 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. Another option is immediate transmission, wherein a flow may start or assume a new rate immediately on the understanding that the network is required to provide a guaranteed rate as soon as possible. This is required to be provided when network resources permit. Yet another option is for a negotiated guaranteed rate. These services are targeted at access scenarios where media flows may result in temporary congestion and where best effort would not act selectively on the last few flows that had contributed to the onset of congestion. These services may also be applied to flow aggregates, providing the possibility of highest available rate between the aggregation end-points or the option of supporting immediate aggregate rate changes that act in conjunction with per-flow controls.

Source

Recommendation ITU-T Y.2121 was approved on 25 January 2008 by ITU-T Study Group 13 (2005-2008) under the Recommendation ITU-T A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at http://www.itu.int/ITU-T/ipr/.

CONTENTS

Recommendation ITU-T Y.2121

Requirements for the support of flow-state-aware transport technology in NGN

1 Scope

This Recommendation provides flow-state-aware requirements in support of per-flow service options providing for edge-to-edge QoS and transport resource control (including resource reservation and admission control) in next generation networks (NGNs). The pertinent protocol specifications and measurement requirements will be described in separate Recommendations. Note that network management functionality is outside the scope of this Recommendation. Also out of scope are any edge functions at network domain boundaries needed for interworking QoS support between FSA and non-FSA networks. Similarly, out of scope are any 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 in the involved FSA nodes. Administrations may require operators and service providers to take into account national regulatory and national policy requirements in implementing this Recommendation.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T Y.1221] Recommendation ITU-T Y.1221 (2002), *Traffic control and congestion control in IP-based networks*.

[ITU-T Y.2012] Recommendation ITU-T Y.2012 (2006), *Functional requirements and architecture of the NGN release 1*.

[ITU-T Y.2111] Recommendation ITU-T Y.2111 (2006), *Resource and admission control functions in Next Generation Networks*.

3 Definitions

This Recommendation defines the following terms:

3.1 aggregation end-point: An end-point within the network which attaches or deletes the common flow aggregate identifiers to ensure commencement/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.

3.2 available rate service (ARS): The name of one of the flow-state-aware (FSA) transport services. ARS is primarily for applications that can flexibly adapt to the current available capacity and can quickly adjust their sending rate as the available capacity changes.

3.3 flow: A unidirectional sequence of packets with the property that, along any given network link, a flow identifier has the same value for every packet.

3.4 flow aggregate: 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.

3.5 flow admission control: The determination, for authorized requests, of whether or not to accept a given flow.

3.6 flow identifier: A vector comprising the values of a number of elements taken from the IP, TCP/UDP header fields, encapsulation header, and label fields attached to a packet. The flow identifier for a flow within a single FSA network is unique. Clause II.10 describes some examples of suitable identifiers purely for information and as an aid to understanding.

3.7 flow state: A set of values stored per flow identifier at each flow-state-aware node. This set of values determines controls applied on a per-flow basis, dealing with forwarding rate, delay and congestion recovery.

3.8 flow-state-aware node: A network node that is capable of maintaining flow state and applying per-flow QoS controls, based on recognizing flow identifier and associated signals.

3.9 flow-state-aware signalling edge function: 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-end-system (UES) application or management action. It may be located, for example, in the UES 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. Out of scope is a network domain edge function required for interworking QoS support between FSA and non-FSA networks. Similarly, out of scope is a network domain edge function that may be required when two FSA networks interwork, 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 in the involved FSA nodes.

3.10 guaranteed rate service (GRS): The name of one of the flow-state-aware (FSA) transport services. GRS is for applications that require guaranteed bandwidth for the duration of the flow.

3.11 in-band signalling: A mode of signalling where the signalling messages are within the flow of the data packets, and follow a path that is tied to the data packets. Signalling messages are routed only through nodes that are in the data path.

3.12 maximum rate service (MRS): The name of one of the flow-state-aware (FSA) transport services. MRS is 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.

3.13 out-of-band signalling: A mode of signalling where the signalling messages are not in the same flow of the data packets, and may follow a different path to the data packets and are routed to one or more nodes that are not in the data path.

3.14 preference indicator: A parameter used to determine whether to admit a flow in case of network overload. In a network overload state, the flow with the lower preference indicator value may be rejected while the one with higher preference indicator value may still be admitted.

3.15 QoS structure: The block of QoS signalling information in a signalling packet.

3.16 variable rate service (VRS): The name of one of the flow-state-aware (FSA) transport services. VRS is for applications that want an ARS and can flexibly take advantage of additional available capacity, but with a minimum capacity consistent with MRS transport characteristics, allowing the option of immediate transmission at or above this minimum rate.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

AAA	Authentication, Authorization and Accounting
ACK	Acknowledge (TCP response)
ANF	Access Node Function
ABR	Available Bit Rate
ARS	Available Rate Service
ATM	Asynchronous Transfer Mode
BA	Behaviour Aggregate
BRAS	Broadband Remote Access Server
Diffserv	Differentiated Services
DSL	Digital Subscriber Line
DoS	Denial of Service
eFSA	egress FSA
FR	Fixed Rate
FSA	Flow-State-Aware
GRS	Guaranteed Rate Service
GRE	Generic Routing Encapsulation
iFSA	ingress FSA
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
L2TP	Layer 2 Tunnelling Protocol
LAC	Layer 2 Tunnelling Protocol Access Concentrator
LNS	Layer 2 Tunnelling Protocol Network Server
LSP	Label Switched Path
MPLS	Multi-Protocol Label Switching
MRS	Maximum Rate Service
NACF	Network Attachment Control Function
NAT	Network Address Translator
NGN	Next Generation Network
NNI	Network-Network Interface
NR	Network Rate
PDA	Personal Digital Assistant
PD-FE	Policy Decision Functional Entity
PE-FE	Policy Enforcement Functional Entity
PPP	Point-to-Point Protocol
PT	Payload Type

QoS Quality of Service

PHB Per-Hop Behaviour

RACF Resource and Admission Control Function

RM Resource Management

RTP Real-time Transport Protocol

SCF Service Control Function

SIP Session Initiation Protocol

SYN Initiation Flag of Transmission Control Protocol connection

TCP Transmission Control Protocol

UDP User Datagram Protocol

UES User-End-System

UNI User-Network Interface

VP Virtual Path

VRS Variable Rate Service

5 Overview

To meet specific network performance requirements, a network operator needs to implement capabilities such as those specified in [ITU-T Y.1221], which now includes clause 6.5 describing a conditionally dedicated bandwidth transfer capability.

To implement this transfer capability as defined in [ITU-T Y.1221], a network needs to provide specific user plane functionality at the UNI and NNI. This Recommendation specifies the flow-state-aware (FSA) requirements for such functionality.

Appendix V provides example implementation principles associated with FSA nodes purely for information and as an aid to understanding.

In terms of QoS aspects of flow-state-aware transfer, a network may be provisioned to meet performance requirements either statically or dynamically. The provisioning would be applied on a per-flow basis. This Recommendation defines the requirements for such provisioning.

Static network provisioning is typically performed by a network management system. Static provisioning takes into account both overall network performance requirements and performance requirements for individual customers based on traffic contracts between the customer and the network operator.

Dynamic network provisioning at a UNI and/or NNI node allows the ability to dynamically request a traffic contract for an IP flow (as defined in [ITU-T Y.1221]) from a specific source node to one or more destination nodes. In response to the request, the network determines if resources are available to satisfy the request and provision the network. 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 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 recognizing and processing requests for associated FSA transport services.

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

Figure 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.

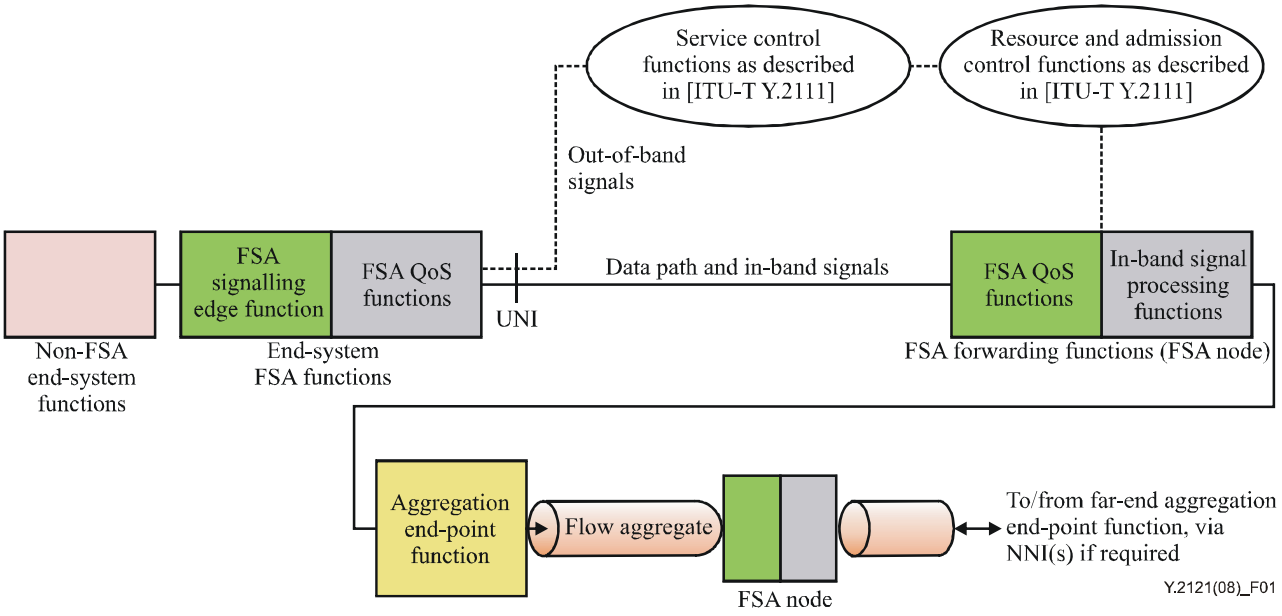

Figure 1 – Overview of FSA functions

Two alternative methods of signalling are shown in Figure 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 this Recommendation 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.

The requirements in this Recommendation cover the following areas:

- Requirements associated with flow-state-aware dynamic provisioning requests initiated by a flow-state-aware signalling edge function, for example located in a UES. See clause 6.1.
- Resource modification. See clause 6.1.
- Response to flow-state-aware requests. See clause 6.2.
- iFSA function answer to network response to flow-state-aware requests. See clause 6.2.
- Form of a verifiable flow-state-aware request. See clause 6.3
- Request performance requirements (e.g., negotiation delays). See clause 6.3.
- Release of resources no longer required. See clause 6.3.
- Error reporting. See clause 6.3.

- Preference indicator resolution. See clause 6.4.
- Parameters and values for transport connections. See clause 6.5.
- Security considerations and requirements. See clause 6.6.

6 Requirements

All of the requirements are numbered consecutively across all subordinate clauses using the format R-#.

6.1 Dynamic provisioning requests from an FSA signalling edge function

6.1.1 Flow identifier

R-1) Signalling from the FSA signalling edge function is required to always carry the flow identifier.

R-2) The flow identifier is required to be used at each flow-state-aware node along the data path to recognize which data packets belong to which flow.

R-3) Aggregation end-points are required to be able to create a flow aggregate identifier and notify the next FSA signalling nodes about the flow aggregate identifier.

- Aggregation end-points are required to be able to aggregate selected flows into fewer flow aggregates, based on some criteria such as the FSA transport services parameters (see clauses 6.1.4, 6.1.6 and 6.1.7) and the path in the network.

R-4) Aggregation end-points are required to be able to change the flow aggregate identifier to which a flow belongs.

- For example, if aggregation is based on a common preference indicator value (see clause 6.1.4) and a flow violates the contract, it causes an automatic reallocation of a flow to a different aggregate.

Further information on the flow identifier is provided in clause I.1. This is purely for information and as an aid to understanding.

6.1.2 Signalling negotiations

R-5) For applications on UESs with an FSA signalling edge function and wanting service to be supported via an FSA transport service, the FSA signalling edge function is required to forward an in-band signalling request across the user-network interface whenever a new flow is starting.

- All other per-flow negotiations involving signals to or from an FSA signalling edge function, such as response, confirmation and close, may be conducted using either in-band signalling or out-of-band signalling. In-band signalling is recommended as the default for all negotiations between an FSA signalling edge function and one or several FSA nodes in the end-to-end data path (see also clause 5).

R-6) For UES's not capable of, or not willing to, support flow-state-aware signalling, an ingress edge node may provide the option that it generates flow-state-aware signalling requests on behalf of the UES. In this case, the edge node performs the FSA signalling edge function.

R-7) If a UES without the FSA signalling capability wants this option to be actively supported, it is required to register for the FSA transport services with information of its device type if such information has not been pre-configured. The method for registration is for further study.

R-8) The ingress edge nodes may provide the option of alternative choices of pre-defined parameters to be used in the request signal and may determine whether a flow should belong to a specific choice within such alternatives.

R-9) The ingress edge nodes may determine the FSA parameters by mapping from the IP level flow descriptor, which is distributed by a central admission entity such as the RACF through the authorization and admission procedure.

R-10) In order to perform signalling requests on behalf of an end-system, an edge node should maintain the registered users for future active identification, by means suitable for each network.

R-11) The edge node may aggregate the flows into a flow-aggregate if appropriate.

R-12) Edge nodes should be able to aggregate signalling messages of end-to-end flows into aggregated signalling messages, and de-aggregate the aggregated signalling messages into signalling messages of the end-to-end flows. In order to do this, edge nodes should be able to:

- set an indication in every signalling packet that downstream FSA nodes may ignore such packets but transparently forward them;
- remove this indication at any downstream FSA node where end-to-end signalling packets are to be restored.

See also clauses A.1 and I.2. The information added in this appendix is purely for information and as an aid to understanding.

6.1.3 Requested rate

R-13) The new flow request is required to specify the requested rate.

6.1.4 Preference indicator

Preference indicator values may be used by different network operators for the development of different service propositions. The policy associated with a preference indicator value depends on the service context (see also clause 6.2.6).

R-14) The new flow request is required to specify a preference indicator value.

R-15) A well-known value is required to be used in a request signal to signify to downstream FSA nodes that the flow is a flow aggregate containing multiple flows of different preference indicator values.

See also clause I.3. The information added in this appendix is purely for information and as an aid to understanding.

6.1.5 Authorization and enforcement

R-16) The new flow request may contain authorization information.

R-17) It is required that QoS-related transport resources are not provided for new flow-state-aware requests unless they can be associated with an authorization.

R-18) If the request is admitted, admission decisions are enforced as follows:

- The ingress edge node, which performs signalling on behalf of registered UESs, is required to map the traffic descriptors distributed by a central admission entity to the FSA parameters, prior to sending the request.
- Policy enforcement is required to occur for all resource requests related to a guaranteed rate (GRS) or maximum rate (MRS) flow request (see clause 6.1.6).
- Policy enforcement is required to occur for all resource requests related to the MRS component of a variable rate (VRS) flow, i.e., for resource requests related to the minimum rate of a VRS flow (see clause 6.1.6).

See also clause A.2.

6.1.6 FSA transport service requirements

R-19) Available rate (ARS)

The FSA signalling edge function is required to initiate and then frequently repeat a request, consistent with the requirements listed in clauses 6.1.1, 6.1.2 and 6.1.3, such that:

- each such request is required to include the requested rate whose value is pre-configured. For FSA signalling edge functions in the UES, the configured value for the requested rate is recommended to take account of the highest rate that the source application can sustain, as well as entitlement captured in service profile information;
- the rate of sending data is required to be controlled by the source to be always less than the latest offered rate (as provided to the source, utilizing the response requirement listed in clause 6.2.4). If the sending rate is sustained at a value above the offered rate, packet discards may be applied;
- the signalling edge function is required to always update the source with the latest offered rate.

See also clause A.3.1.

R-20) Guaranteed rate (GRS)

- In this context, the required response to the signal for a new flow with a requested rate, as described further in the requirements listed in clauses 6.2.3 and 6.2.4 below, indicates either: rejection of the flow or acceptance of the flow with an assignment of the flow to the "discard last" packet discard priority (see clause 6.2.2).
- If the sending rate is sustained at a value above the guaranteed rate (the requested rate), packet discards may be applied.

See also clause A.3.2.

R-21) Maximum rate (MRS)

- In this service context, the required response to the requested rate indicates either rejection of the flow or acceptance of the flow with conditional guarantee of the requested rate in accordance with the principles of clause 6.5 of [ITU-T Y.1221].
- In the case where the in-band signalling response is used, MRS is required to support user applications transmitting immediately after sending an in-band request signal. This is termed the "immediate transmission" option. However, this is subject to authorization and also subject to the flow rejection requirements discussed in clause 6.2.3.
- Where the immediate transmission option is supported, applications can transmit at any rate up to the requested rate specified in the in-band new flow request signal. If the sending rate is sustained at a level above this rate, packet discards may be applied.
- Where the immediate transmission option is not supported, applications are recommended to wait for an in-band response signal (see clause 6.3) prior to entering the data transmission phase.

See also clause A.3.3.

R-22) Variable rate (VRS)

- The initial requested rate (consistent with the requirements listed in clauses 6.1.1, 6.1.2 and 6.1.3) is required to be interpreted as the minimum rate. In other words, policing is required to never reject packets sent at or below this rate, subject to authorization and subject to the rejection requirements in clause 6.2.3.
- Subsequent request signals that follow the initial request signal are required to be interpreted as requests for bandwidth over and above this minimum rate. The additional

bandwidth is required to be assigned to a flow through the available rate transport service.

- The policed rate is recommended to be updated following each ARS request/response but is required to never be set to a rate less than the rate specified in the minimum rate requested.

See also clause A.3.4.

6.1.7 Resource modification

R-23) Resource modification request is recommended to be done through in-band signalling as a default. It is required to further contain the requested rate, preference indicator value and service context.

R-24) For the ARS and VRS context, the source is required to wait for the network response to the requested rate before changing its sending rate upwards.

6.2 Network and FSA signalling edge function responses to flow-state-aware requests

6.2.1 Requiring the flow identifier to reference the per-node flow state

R-25) Each accepted flow identifier is required to be assigned a flow state, including a packet discard priority value, at each FSA node along the data path.

6.2.2 Packet discard priority assignment requirements

R-26) The packet discard priority is not included in signalling information.

R-27) The packet discard priority is required to allow for the possibility of distinguishing between at least two values, namely "discard first" and "discard last".

6.2.3 Flow rejection response

R-28) The flow rejection response may be conveyed through in-band or out-of band signalling. In-band signalling is recommended to be the default.

R-29) The rejection response is recommended to be used to reject any flow according to network operator policies.

6.2.4 Flow acceptance response to requested rate

R-30) The acceptance response is required to confirm that all FSA nodes along the edge-to-edge route have either accepted or modified the requested rate from the source (or otherwise the rejection response is required to be sent).

R-31) The acceptance response may be an in-band signal or an out-of-band signal. In-band signalling is recommended to be the default.

6.2.5 Flow acceptance response to preference indicator

R-32) The acceptance response is required to confirm that all FSA nodes along the edge-to-edge route have either accepted or modified the requested value for preference indicator from the source (or otherwise the rejection response is required to be sent).

6.2.6 Priority of packet discard, including service context use of the preference indicator

R-33) Flows in the "discard last" state are required to always have priority with respect to the available buffer capacity.

R-34) When congestion conditions are such that congestion persists even after the removal of packets of all flows with the "discard first" priority, the network should re-mark additional flows as "discard first", starting in order from the lowest preference indicator value.

R-35) For any service contexts, it may be possible (subject to network operator conditions) for a UES to establish a pre-assigned requested rate for any preference indicator value.

See also clause I.5. The information added in this appendix is purely for information and as an aid to understanding.

6.2.7 Congestion notification

R-36) An FSA node is required to send a congestion notification signal to the iFSA function (from any congested network node) as determined by:

- Whether congestion conditions require that the ARS flow rate be immediately reduced to avoid packet discards.
- Whether congestion conditions are so severe that it may be desirable to suppress the sending of congestion notifications if they would exacerbate the problem. The mechanisms for achieving this are for further study.

R-37) The congestion notification signal is required to be sent to the eFSA function if the congested network node can only communicate control signals in the direction of the flow forwarding path. After receiving the congestion notification, the eFSA function is required to forward the message to the iFSA function, using either in-band or out-of-band signalling.

R-38) This congestion notification signal may be sent using either in-band or out-of-band signalling. In-band signalling is recommended to be the default.

R-39) It is required that a congestion notification signal is not generated within a congested FSA node for the purposes of forwarding to and notifying an end-system of packet losses. Detection of packet losses in end-systems shall be provided for by the applications themselves and/or protocols operating at the end-to-end level.

See also clause I.6. The information added in this appendix is purely for information and as an aid to understanding.

6.3 Signalling requirements

6.3.1 Form of flow-state-aware signalling packets

R-40) Signalling packets are required to be uniquely marked so that the FSA signalling edge functions and the FSA nodes can easily recognize them.

R-41) Data packets are required to be uniquely marked so that they are recognizable as data packets under FSA forwarding treatment.

See also clauses II.1 and II.2. The information added in this appendix is purely for information and as an aid to understanding.

6.3.2 Performance requirements for requests and responses

R-42) Signalling exchanges are required to be completed fast enough so as to meet the needs of frequent ARS rate adjustments during the lifetime of a flow and to meet the needs of the MRS immediate transmission option.

See also clause II.3. The information added in this appendix is purely for information and as an aid to understanding.

6.3.3 Release of resources no longer required for GRS

R-43) For obtaining the exact start time for the utilization of resources, an in-band confirmation signal may be sent from the FSA signalling edge function in the iFSA (i.e., the edge function that sent the initial request), so that all nodes in the path know the final capacity that is agreed upon.

R-44) For obtaining the exact ending time for the utilization of resources, an in-band "close" signal may be sent from the FSA signalling edge function in the iFSA to ensure all nodes know that the capacity can be released.

See also clause II.4. The information added in this appendix is purely for information and as an aid to understanding.

6.3.4 Protection against lost signalling packets

Requirements R-45 and R-46 follow from the observation that, for all services, it is important that the sender know if the initial signalling packet was lost.

R-45) If a response is required and is not received within a predetermined retransmission time-out period, the iFSA function may send a second request.

R-46) For GRS, to ensure against lost in-band signals (if sent) conveying either a "confirmation" or "close" indication, the FSA signalling edge function in the eFSA (which may be in the UES) is required to resend the response if data arrives before a confirmation is received and is required to send an in-band confirmation signal to a received in-band "close" signal.

6.3.5 Service contexts

R-47) The service context information is recommended to be included in all request, response and confirmation signals.

R-48) There are four service contexts: GRS, MRS, VRS and ARS. Each service context is required to have a code.

GRS

R-49) When the destination receives a GRS request, it is required to send a response back to the iFSA function with the requested rate received or, if desired, a lower rate.

R-50) The initiating FSA signalling edge function is required to send a confirmation after receiving the response telling all FSA nodes the final agreement.

R-51) The eFSA function is required to confirm the close to ensure that there are no lost close packets.

R-52) The iFSA function is required to send an additional close signal (see clause 6.3.8) after a short timeout has elapsed indicating either the close signal or its confirmation has been lost.

MRS

R-53) The sender need not wait for a response and can send at the requested rate immediately after the request.

R-54) The eFSA function is required to send a response reflecting the rate received, or a lower rate if desired.

R-55) No confirmation is required.

VRS

R-56) The initial minimum rate request and response is required to be treated like an MRS request.

R-57) The eFSA function is required to send a response with a rate no higher than received.

R-58) No confirmation is required.

ARS

R-59) The iFSA function is recommended to send a request with a requested rate set to the maximum rate that can be supported by the end equipment or application.

R-60) The FSA nodes are required to forward the request with the rate reduced to a rate they can support.

R-61) The destination is required to return the rate value received or some lower rate to the iFSA function in a response packet.

R-62) The sender is required to conform to the offered rate received in the response.

R-63) Frequently the iFSA function may send a new request to see if a higher rate is available.

R-64) The network FSA nodes may send to the iFSA function signalling packets specifying a new lower rate.

R-65) No confirmation is required.

See also clause II.6. The information added in this appendix is purely for information and as an aid to understanding.

6.3.6 Preference indicator

R-66) This parameter is recommended to be included in all requests, responses and confirmation signals.

See also clause II.7. The information added in this appendix is purely for information and as an aid to understanding.

6.3.7 Delay priority

R-67) This parameter is required to be included in all requests, responses and confirmations.

See also clause II.8. The information added in this appendix is purely for information and as an aid to understanding.

6.3.8 Signalling type

R-68) Every QoS signalling packet is required to specify the signalling type.

R-69) There are at least five types of in-band signalling packets as follows:

- **Request**: The start of the signalling process is a request packet. A signalling code point is required to indicate a request.
- **Response**: When the eFSA function receives a request packet which may have been modified by the network, it is recommended that it returns a response packet containing all the requested parameters and how they have been modified. A signalling code point is required to indicate a response.
- **Confirm**: For GRS there needs to be a confirm packet. A signalling code point is required to indicate a confirm. A confirm packet is required to be used to confirm the GRS response packet including setting the fixed rate (FR) indication, as further illustrated for information purpose in clause II.5.2. This ensures all FSA nodes know the final rate. A confirm packet is also required to be used after the GRS close packet to ensure that all FSA nodes actually received the close and remove any guarantees.
- **Renegotiate**: A signalling code point is required to indicate a renegotiate. Depending upon the application, the iFSA function may wish to try for a different or higher rate, preference, delay priority or burst tolerance. To do this and minimize confusion for the FSA nodes and the eFSA function, it is important to use a different code rather than "request". Thus renegotiate requests a new QoS for the ongoing flow. It should be noted that increasing the delay priority may cause out of order packets due to the queue change, and thus this type of change would be at the sender's risk.
- **Close**: For GRS, all the FSA nodes in the path must be informed that the reserved bandwidth can be released. A signalling code point is required to indicate a close. For MRS, ARS and VRS, the FSA nodes shall time out the flow if no packets are seen for a given period.

In conjunction with a central admission entity such as the RACF, response, confirmation, and close may be omitted. See also Appendix VI. The information added in this appendix is purely for information and as an aid to understanding.

6.3.9 Signalling aggregation indication

R-70) **Ignore indication.** An ignore indication is required to be included within every signalling packet. This indication enables aggregated signalling, such that the end-to-end signalling messages are hidden from (ignored by) the core nodes. The edge nodes of an aggregation region are required to set this indication to show whether to ignore signalling by downstream nodes.

6.3.10 Charging direction indication

R-71) In order to allow services where the flow is paid for by the receiver, it is required that there be a signalling code point to indicate the charging direction. This indication could be set by the iFSA function. If configuration or policy associated with the eFSA function conflicts with the indicated charging direction, the response from the eFSA function may clear the indication and such a conflict between requests and responses may be treated as appropriate to network policy (e.g., revert to sender is charged). This means that all flows are required to be identified as to the paying party.

6.3.11 QoS structure extension

R-72) It is required that any QoS signalling packet carries an indication that enables an outgoing request for a flow in one direction to be included in the same signalling packet that carries a response in the opposite direction, where these both relate to flows between the same two end-points.

See also clause B.1.

6.3.12 Authorization information attached

R-73) It is required that there is an indication within a signalling packet that authorization information is present.

See also clause B.2.

6.3.13 Flow aggregation request

R-74) It is required that a request signal carries an indication allowing FSA nodes to recognize that the request relates to a flow aggregation.

See also clause B.3.

6.3.14 Burst tolerance

R-75) It is required that there is some tolerance to rates that exceed the requested rate for a short duration. A parameter defining such a tolerance is required to be included in all requests, responses and confirmations.

See also clause II.9. The information added in this appendix is purely for information and as an aid to understanding.

6.3.15 QoS approval indication

R-76) An FSA node is required to set the QoS approval indication if it has approved the request from a flow.

R-77) An FSA node may clear this indication to inform the iFSA and/or eFSA functions that the request is not approved.

6.3.16 FSA signalling edge function

R-78) The eFSA function is required to create and send a response to a request to show the acceptance or rejection of a flow request, or modify the request so that it can be accepted.

- When the iFSA function receives a response that indicates a modified request, it may notify the application.

6.3.17 FSA node requirements

6.3.17.1 FSA node operation

R-79) The ingress edge FSA node is required to be responsible for checking an attached authorization indication and also for determining if the preference indicator value requested is authorized for this user.

R-80) All FSA nodes receiving a flow request are required to check their available capacity and resources and adjust downward the requested rates, delay priority, preference indicator value, and burst tolerance requested to what they believe they can support with reasonable assurance.

R-81) For available rate (ARS) and variable rate (VRS) services, any FSA node along the path may send a message indicating that the rate needs to be lowered.

- It is recommended for this message to go to the iFSA function, but if this is difficult, it is required to be sent to the terminating FSA signalling edge function which is required to then send a response to the iFSA function.

See also clause B.4.

6.4 Admission decision

This clause defines minimum requirements for managing the admission decision process.

6.4.1 Admission decision for maximum rate (MRS) flows

R-82) The decision process is required to be fast enough to support the immediate transmission capability, taking account of preference indicator value and available capacity.

See also clause III.1. The information added in this appendix is purely for information and as an aid to understanding.

6.4.2 Admission decision for available rate (ARS) flows

R-83) The decision process is required to be capable of supporting frequent rate update requests on each ARS flow.

See also clause III.2. The information added in this appendix is purely for information and as an aid to understanding.

6.4.3 Admission decision for variable rate (VRS) flows

R-84) VRS flows are expected to be treated like ARS flows except for their minimum guaranteed capacity which is required to be treated like MRS. Thus, the admission decision for VRS flows is a combination of the one for the MRS and the one for the ARS.

6.4.4 Admission decision for guaranteed rate (GRS) flows

R-85) Total capacity assigned for GRS flows on a port is required to be checked before the admission decision.

6.5 General architectural requirements on the management of transport connections carrying flow-state-aware traffic and other traffic

R-86) Flow-state-aware QoS controls are required to be agnostic to the underlying transport technology.

R-87) A given network link (between two network nodes) may be configured so that it is not dedicated to carrying flow-state-aware traffic only.

See also clause IV.1. The information added in this appendix is purely for information and as an aid to understanding.

6.6 Security considerations and requirements

R-88) **Authentication**: User authentication is addressed in clause 6.1.5. 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.

R-89) **Authorization**: Authorization is addressed in clause 6.1.5.

R-90) **Data confidentiality**: It is required that no additional data confidentiality requirements are imposed through the use of flow-state-aware transport technology. It is required that, in the case where user flows have data confidentiality requirements and invoke ARS, MRS, VRS or GRS, the parameters describing the in-band request signal are not encrypted.

R-91) **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.

R-92) **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.

R-93) **Availability/accessibility**: Flow-state-aware services are required to respect the priority preference of each authenticated entity in making admission decisions.

R-94) **Privacy**: It is recommended that flow-state-aware services ensure the privacy of user-specific policy profiles defining QoS parameter limits and privileges.

R-95) **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.

Annex A

Dynamic provisioning requests from an FSA signalling edge function

(This annex forms an integral part of this Recommendation)

A.1 Negotiations for a user-end-system without signalling capability

This clause references requirements in clause 6.1.2.

With respect to requirement R-8, as stated there, the edge nodes may have pre-defined parameters. Each new flow may be monitored, and if the flow under observation reaches some threshold that determines which parameter choice to use, the node generates an in-band signalling request. This starts negotiation for the flow.

Thus, if a data packet from a registered UES is received, the edge node records the flow identifier as an 'unidentified flow identifier' and monitors for the unidentified flows. If conditions are met (parameters are set *a priori*), the edge node puts the unidentified flow into the 'identified flow' list with proper request items a UES would request (requested rate, preference indicator value, and service context), and then performs signalling as a UES with flow-state-aware capability would do. The signalling functions further required are performed by the edge node. The edge node may aggregate the flow into a flow-aggregate if appropriate.

A.2 Authorization

According to the principles of [ITU-T Y.2111], the following requirements may be derived for the authorization of flow-state-aware in-band request signals while recognizing the fast processing nature of services created by FSA nodes. It is based on a "two-phase scheme" as described in [ITU-T Y.2111]: authorization is performed in one step (see Figure A.1), followed by reservation and commitment in another step (see Figure A.2).

Before the UES forwards the in-band signal associated with a new flow request, it first initiates the QoS authorization procedure. This procedure (see Figure A.1) may be triggered by a service-establishment signalling message (e.g., a SIP invite message) that in turn instigates an SCF request to the PD-FE. The PD-FE may, as a result, generate an authorization token as optional for a given service and send it to the SCF. The SCF may forward the authorization token in the service signalling message to the UES (see Figure A.1).

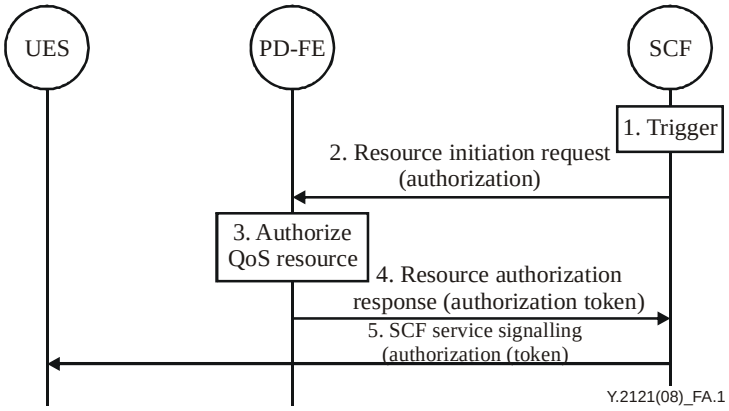

Figure A.1 – SCF-requested QoS initial authorization procedure

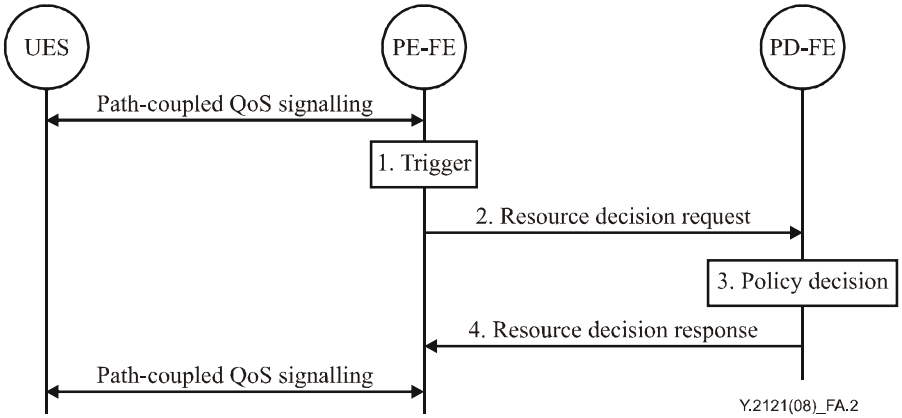

Figure A.2 – Resource reservation procedure following a UES request

A.3 Service context

This clause references requirements in clause 6.1.6.

Figure A.3 summarizes the grouping of new flow requests and responses into service contexts. This greatly increases the utility of the flow-state-aware capability as supported through the underlying requests and responses related to [ITU-T Y.1221].

Requests and responses used for ARS	Requests and responses used for GRS	Requests and responses used for MRS	Requests and responses used for VRS
Basic requests and responses of clause 6.5 of [ITU-T Y.1221], dealing with the association of a flow state with a flow identity, together with a requested maximum rate and preference indicator.			

Figure A.3 – Service options

A.3.1 Available rate (ARS)

Utilizing the service context rules of clause 6.1.6, the transfer capability described in clause 6.5 of [ITU-T Y.1221] provides support for an available rate service (ARS), similar to that described in clause 6.2.4 of [b-ITU-T I.371]. In [b-ITU-T I.371] it is stated that, in ABR, the UES regularly polls the network for the currently available bandwidth by sending resource management (RM) cells conveying a requested rate to the network. RM cells are a form of in-band signal.

A.3.2 Guaranteed rate (GRS)

This service context (termed the guaranteed rate or GRS context) may be offered for applications requiring guaranteed bandwidth for the entire duration of the flow. Additional requirements apply to this service to ensure the reservation is aborted if not all nodes can support the request and to ensure that clear down of reservations is completed successfully at the end of the flow.

A.3.3 Maximum rate (MRS)

In effect the service offer is that the requested flow may proceed, with this "as soon as possible" condition on meeting the loss component of the requirement. We term this the maximum rate service (MRS). As stated in clause 6.5.2 of [ITU-T Y.1221], while in "discard first" status, a flow may be subject to a higher probability of packet discards during moments of congestion. It is up to the application to choose to terminate the flow or continue and compensate for any bursts or losses as and when they occur.

A.3.4 Variable rate (VRS)

The service context rules of clause 6.1.6 may be offered as a service that supports applications that require a minimum rate and have some degree of elasticity above this minimum.

Annex B

Signalling requirements

(This annex forms an integral part of this Recommendation)

B.1 Second QoS structure attached

This clause references requirement R-72.

When the eFSA function (which may be in the end-system) receives a QoS signalling request and desires to establish a return connection, it may proceed as follows. This FSA signalling edge function may send both a response to the request and also a QoS signalling request for the opposite direction in the same packet. This fits ideally into the normal TCP start-up handshake where the first packet is a SYNC and would include the forward request, the next packet is the SYNC/ACK which would contain the forward response and the reverse request, and the third packet is an ACK which would carry the reverse response.

B.2 Authorization information attached

This clause references requirement R-73.

It is important to allow a cryptographic authorization header to be attached to each flow request so that the network can verify the user and the user's privileges.

B.3 Flow aggregation request

This clause references requirement R-74.

There is a strong need in all networks for flow aggregation which groups and maps some of the flows into a smaller set of flow aggregates. With the full capabilities of the signalling available to specify the flow aggregate, a flow aggregate may be one of GRS, MRS, VRS or ARS, with any rate, preference indicator, and delay priority desired. It may utilize encapsulating IP header (IPv4 or IPv6), MPLS label, VLAN ID, or the likes to specify the destination such that, with this header, the flow aggregate could be routed across any transport network to its destination. At the egress FSA node, a flow aggregate would be terminated and all the traffic inside the flow aggregate directed to the specified address. A signalling edge function may request flow aggregation.

B.4 FSA node operation

This clause references requirements in clause 6.3.17.1.

With reference to these requirements, when a compliant FSA node receives an in-band signalling packet, it is required to inspect the QoS structure to determine what action it needs to take, if any. If this is a new flow, it is required to check the capacity of the output port and determine what rate it can accept.

Note that, with respect to requirement R-80, it is not necessary for FSA nodes to guarantee support under all possible conditions, only that they have high confidence that they can support the resulting request under normal operating conditions.

Appendix I

Supplementary information on information exchanges via requests from an FSA signalling edge function and associated responses

(This appendix does not form an integral part of this Recommendation)

I.1 Flow identifier

[ITU-T Y.1221] describes the combination of parameters "source IP address", "destination IP address", "source and destination port numbers", "protocol" and "experimental/Diffserv value" as a basis for a flow identifier.

Requirement 1 is derived from clause 6.5.1 of [ITU-T Y.1221], where it can be inferred that signalling from the FSA signalling end function always carries flow identifier information that is the same as that carried in the header of each data packet of the flow. This information about the flow identifier is required to be used at each flow-state-aware node along the data path to recognize which data packets belong to which flow.

Ingress edge nodes may support the option to aggregate selected flows into fewer flow aggregates, based on some criteria such as the service context (see clause 6.1.6), the preference indicator value (see clause 6.1.4), and the path in the network.

An alternative method for carrying the flow identifier, or especially for carrying a flow aggregate identifier, may be through the multi-protocol label switching (MPLS) label. [b-RFC 3270] describes two standard methods to map the Diffserv behaviour aggregates (BAs) onto MPLS label switched paths (LSPs).

For networks without MPLS functions, an encapsulating protocol (e.g., IP in IP encapsulation [b-RFC 2003] or generic routing encapsulation (GRE) [b-RFC 2748]) header with IPv4 address may be used as the identifier for flow aggregates within a single administrative network domain.

I.2 In-band signalling negotiations

Requirement R-5, while not fully explained in clause 6.5.1 of [ITU-T Y.1221], supports the signalling performance requirement R-42. It also facilitates service options such as available rate as discussed further in clause 6.1.6.

Again referring to requirement R-5, there is a reference to the same terminology in clause 6.5.1 of [ITU-T Y.1221], i.e., that there is an in-band signal.

Requirement R-7 envisages an edge node that is used for flow negotiations on behalf of a UES. To perform this function, the edge node may utilize information about the UESs that was obtained during registration. The following steps describe this in more detail:

1) UES registers itself for flow identification with its IP address to an edge node. Or an access node (e.g., access node function (ANF) in access transport) registers its attached UESs for the flow identification with their IP addresses and end nodes information to an edge node.
2) The UES information may be terminal type (if the terminal is a user terminal), or application type (if the terminal is an application server).
3) The edge node stores the IP address and end node information.
4) When this edge node receives a packet with an unknown identifier from a registered UES, it retrieves the UES registration information. It may monitor the flow for additional information to determine the flow characteristics and perform signalling as necessary.

The following are a few examples of possible end nodes which would benefit from active flow identification.

- VoIP terminals.
- Mobile handheld devices (including cellular phones, PDAs, etc.): Traffic from such terminals may be either voice, streaming video, or message (text or multimedia). The voice and streaming video are with distinct encoding rates. If an end node registers as a mobile handheld device, the edge node may utilize this information together with information obtained by monitoring a few packets to determine the characteristics of the flow.
- Media streaming servers: If such a server is the source of consecutive packets with a relatively regular interval, the edge node may infer that the flow is a media stream. The encoding rate information may be retrieved from information stored during registration.
- Emergency services-related equipments: Light-weight emergency or military equipments may not have the capability to support flow-state-aware signalling. Registration information may enable a node to assign such equipment the pertinent service contexts, rates and priorities.

One of the possible ways to expedite the identification is to look at the upper layer header (e.g., RTP payload type (PT) field in the header), only when the end node of the packet is configured to be registered. The PT field provides detailed information on whether the flow contains audio or video, the type of encoding, and the encoding rate.

I.3 Preference indicator request

The higher preference indicator values may be reserved for such purposes as emergency services and military commands, etc.

I.4 Authentication

The following service scenarios illustrate authentication aspects of flow-state-aware transport and service.

Scenario 1: Mobile user access

With reference to Figure I.1, a mobile user access scenario is shown. For example, a terminal (end-A) could connect to a public WLAN hotspot (node 1 of Figure I.1). It may establish an IPsec tunnel to an IPsec gateway (node 2) to access enterprise services. However, assuming IPv4 is used, flow-state-aware QoS controls cannot be applied to the QoS experience obtained along the IPsec tunnel. If IPsec is used, there is no way for the FSA nodes to recognize the flow since the ports and protocol are hidden. Thus with IPsec and IPv4, the protocol cannot be used and only Diffserv is available. IPsec can be used with IPv6 and FSA protocols.

With IPv4, to make use of flow-state-aware QoS, end-A may make use of split IPsec tunnelling, whereby:

- non-FSA traffic destined to the service provider domain is sent via IPsec,
- flow-state-aware traffic bypasses IPsec.

Thus, in Figure I.1, the FSA IPv4 end-A is shown gaining access to an FSA content services platform (e.g., a content distribution and delivery platform) that is accessible from a flow-state-aware node 1.

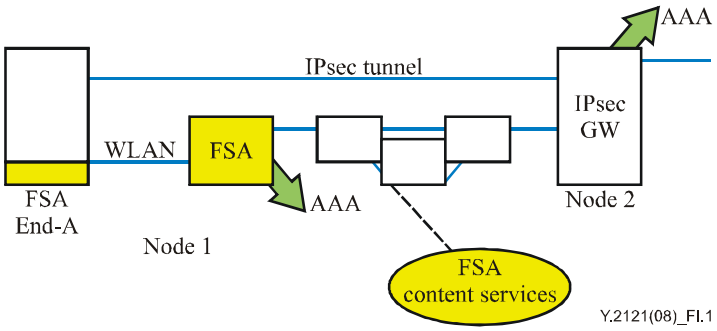

Figure I.1 – Scenario 1: Mobile user access

Scenario 2: Broadband access

Figure I.2 shows end-A connected over DSL to a BRAS. This BRAS acts as an LAC and forwards the traffic by using L2TP to a second BRAS acting as an LNS. Node 1 may issue a RADIUS request to obtain attributes for the tunnel to be established (e.g., [b-RFC 2868]). The second BRAS terminates the PPP state machine and may issue a RADIUS request to perform user authentication. Policy enforcement is generally only done in node 2. Part of that policy enforcement could be a flow aggregate with associated control that manages downstream traffic (towards end-A) at the point where it is aggregated and manages upstream requested rate (GRS, MRS or VRS) requests and ARS-controlled back-pressure.

An alternative arrangement is shown in Figure I.3. Here, the L2TP tunnel is replaced by a flow-state-aware flow aggregate. The objective is to handle any packet discards, either upstream or downstream at the flow aggregate entrance. For example, the sudden insertion of an upstream emergency call may cause this.

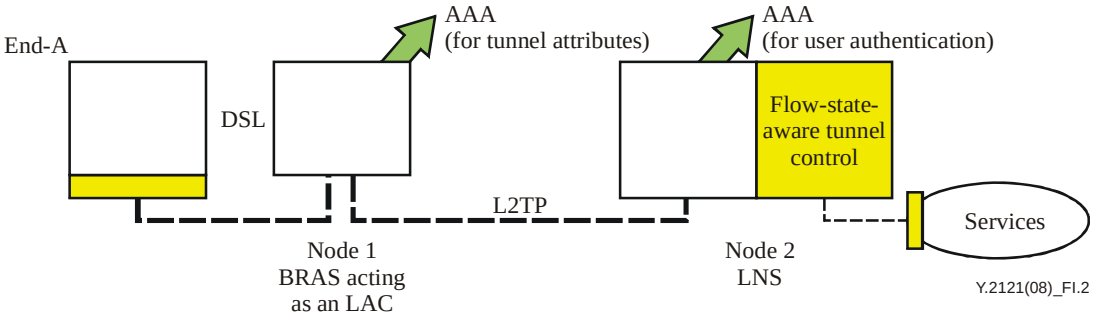

Figure I.2 – Broadband services via an L2TP tunnel

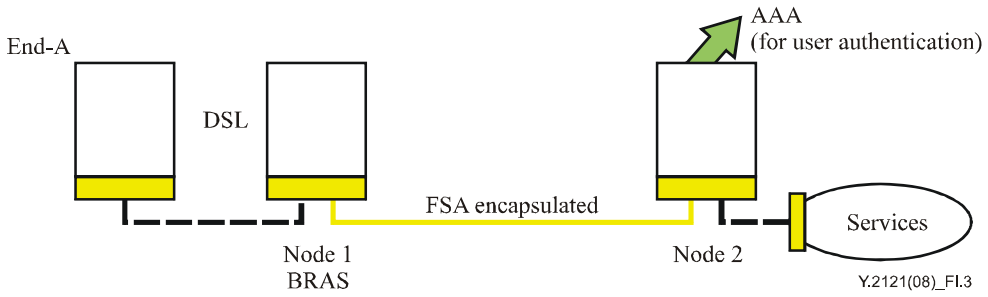

Figure I.3 – Flow aggregate with FSA outer header established by in-band signals between nodes 1 and 2

Network management commands are used to trigger nodes 1 and 2 to exchange in-band flow-state-aware signals. These establish, say, a flow aggregate with GRS service context including a defined capacity (requested rate). All IP packets from end-A or from "services" are encapsulated with an outer packet header bearing the flow identifier of the aggregate flow. Flow-state-aware packets sent by end-A are QoS-managed along the aggregate flow at node 1. Similarly, flow-state-aware "services" packets are managed along the aggregate flow at node 2.

I.5 Priority of packet discard, including service context use of preference indicator values

This clause references requirements in clause 6.2.6.

Typically, a flow with a high preference indicator value may be established even when network capacity is already fully loaded with flows of lower preference priorities.

For example, an MRS flow with a high preference indicator value could be established for emergency services, where transmission could begin immediately, without even waiting for a network response, provided the authorization was accepted. This would be one possible service context. Another context may allow parental configuration of content so that certain types of content or certain user identities determine the preference indicator value.

Therefore, two parameters (preference indicator and packet discard priority) determine the probability of packet discard. The relationship between the nodal behaviours and these parameters in flow-state-aware network nodes are as follows:

- Admission decisions are not governed by packet discard priority. But newly admitted flows are usually marked as "discard first". Depending upon policy, some newly admitted flows (for example, those with a high preference indicator value) may be immediately marked "discard last".
- Packet discard (buffer management) decisions are governed by packet discard priority, but this may be influenced by preference indicator as follows:
  - when packet discard is necessary, if there are packets with "discard first" priority, drop those packets;
  - if buffer congestion persists, remark additional flows with the lowest preference indicator values as "discard first".

I.6 Congestion notification

This clause references requirements in 6.2.7.

It is desirable that the implementation has queue buffers that are large enough to absorb the remaining round-trip delay prior to the rate reduction.

Appendix II

Supplementary information to signalling requirements

(This appendix does not form an integral part of this Recommendation)

II.1 Recognition of QoS signalling packets

This clause references requirement R-40.

There is always a need to identify some in-band signalling packets. However, in-band signalling need only be done at the start of a flow.

The exact method for identification of data packets is for further study.

II.2 Form of QoS information

This clause references requirement R-40.

The requirement R-40 follows from the fact that data packets, by chance, could include information that looks similar to the information in a signalling packet. Therefore, a unique marking for signalling packets is required.

FSA nodes should be able to read and modify a QoS structure within a signalling request and, if used, a response, confirmation or close in real time at any port speed. Port speeds are never faster than the logic speeds currently possible for simple operations such as reading and/or writing packets into memory, checking sum checks, operating linked lists, reading and interpreting fixed length fields, etc. Usually, the serial input stream is first converted to a parallel word stream so that standard CMOS logic can keep up. Thus, so long as the QoS structure follows some simple rules, any port speed currently feasible for high speed FSA nodes could be processed.

The following text illustrates an example implementation, not the only implementation, facilitating fast processing.

II.2.1 All information fields are in a fixed location in the QoS structure.

II.2.2 All numerical values which need to be treated as a number are structured as an integral number of bytes and are byte aligned.

- If the number is 2 bytes, it is dual byte aligned and if 4 bytes it is quad byte aligned. This is because the parallel input stream in IP is quad byte aligned and it would require more steps to shift and mask to find a value.

II.2.3 IP protocols also works with 8-byte boundaries and thus the total QoS structure is a multiple of 8 bytes.

- Also, the number of 8-byte groups required is as small as possible to reduce overhead and processing.

II.2.4 Since overhead is of great importance, options within the QoS structure may be packed to a byte.

II.3 Performance requirements for requests and responses

This clause references requirement R-42.

The immediate transmission option relies on the fast establishment of flow state at every FSA node such that, in the extreme, the request packet arrives at an FSA node and is immediately followed by the first data packet of the flow.

II.4 Release of resources no longer required

This clause references requirements in clause 6.3.3.

An FSA node may need to reject new traffic when too many absolute capacity guarantees are no longer required but are not yet released.

II.5 QoS signalling parameters

Clauses II.5.1 and II.5.2 provide information for illustrative purposes relating to an example implementation, not the only implementation.

II.5.1 IPv6 header

The QoS structure in IPv6 may be chosen to be a hop-by-hop option.

II.5.2 Rates

II.5.2.1 As an example implementation, two rates may be utilized, one that is network-selected (network rate or NR) and one that is an application requirement (fixed rate or FR). NR would be used where the application requires to send buffered data where the rate may vary (typically TCP). FR would be used where a fixed rate must be available at all times.

II.5.2.2 These may have multiple uses and may be used together so as to support the four types of service (GRS, MRS, ARS and VRS). As an example:

II.5.2.2.1 **GRS**: The GRS could use the FR field to specify the requested rate with the NR field set to zero.

II.5.2.2.2 **MRS**: The MRS could use the FR field to specify the requested rate with the NR field set to zero.

II.5.2.2.3 **ARS**: The ARS could use the NR rate to specify the requested rate, setting the NR rate to the maximum rate that the application or computer can support. The FR field could be set to zero.

II.5.2.2.4 **VRS**: The VRS could use the NR plus FR rates to specify the requested rate. The FR portion could carry the minimum required rate (lowest value) for the flow. The NR rate could be set as in the ARS service (see clause II.5.2.2.3).

II.5.2.3 The lowest rate may be low enough that the FSA node would not find any significant value in reserving or managing a flow to a lower rate.

- The lowest rate could be 1 kbit/s.

II.5.2.4 Zero may be represented since NR or FR may be zero.

II.6 Service contexts

This clause references requirements in clause 6.3.5.

GRS

Guaranteed rate has the most stringent requirement and is for the equivalent of a leased line where the rate is permanent until closed.

GRS requires commitments from FSA nodes even in the absence of traffic.

MRS

The second type of flow is maximum rate where a network timeout after a period of no data is sufficient. MRS may be used for video, voice or other streaming media whenever the QoS requirement is for low loss and low delay variance, but where such requirements allow for a network-conditional guarantee that facilitates immediate transmission. The conditional guarantee is

that the network is required to support the loss/delay QoS targets as soon as possible, while allowing the immediate transmission to continue.

ARS

This service is typically used for TCP flows but can be used with any protocol. It offers an available rate that can be immediately supported across the network.

VRS

This service offers a minimum rate with the conditional guarantee and immediate transmission characteristics of MRS. Additionally, it allows an application to exploit the latest available rate so that it may send at a higher rate but need never send at a lower rate than the MRS value. For example, a stock trade may be required to be transacted or reported in some known and acceptable time but faster is better.

II.7 Preference indicator

This clause references requirement R-66.

The preference indicator is a parameter used to determine which flows should be admitted in the case of a network overload. The overload could be on an access line that has too many video requests, or it could be due to trunk or network equipment failures. This type of capability is necessary as the network moves toward streaming media flows (GRS and MRS) and away from available rate flows. It has always been implemented in telephone and military networks. However, since the Internet was mostly TCP to start, it has not been required or standardized before. Now it is required. The number of preference priorities needs to be sufficient for the multiple military and civilian emergency systems, corporate priorities, and home priorities. Assigning the preference codes is beyond the scope of these requirements but the number of levels needs to be set.

II.8 Delay priority

This clause references requirement R-67.

Although absolute delay is not controllable in a network due to the speed of light, delay variance or jitter can be controlled and may have different requirements for different services. Typically, video and voice require lower delay variance than file transfer but there may be many other services with many different requirements.

II.9 Burst tolerance

This clause references requirements R-75.

The rates negotiated in GRS, MRS, VRS and ARS are maximum rates and the UES is free to send at lower rates. But if the transmission rate momentarily exceeds the agreed rate, it is typical in packet networks to include some burst tolerance.

II.10 Flow identifier fields

This clause references requirements in clause 6.1.1.

IPv6 flow identifier field

The following is an example implementation for illustrative purposes. In IPv6, the 20-bit sender flow label may be used as part of the vector of parameters that identify the correct flow from a response packet. The triplet source address, destination address and the flow label could together identify a flow. The response has the same source and destination address (reversed) but the sender's flow label is still required to identify the matching request. Thus, in this example, there would be a 20-bit source flow label field in the response packet. The eFSA function fills this in from the request.

IPv4 flow identifier fields

The following is an example implementation for illustrative purposes. In IPv4, if unencrypted, the source address, destination address, source port, destination port and protocol could define the flow. When the eFSA function sends a response, these fields are all in the packet although source and destination are reversed. Thus, the sender needs no extra information to identify the flow. It may be noted that most senders are behind NAT devices and the source address and source port have been changed, perhaps several times, by the time the eFSA function gets the request. However, they are all restored by the time the response gets to the iFSA function. Thus, no extra issues are raised by NAT.

However, if the iFSA function moves (mobile user), the eFSA function will not recognize the new source address and source port. Nor does the iFSA function know what the source address and port were that the eFSA function recorded. Therefore, in order to support mobility of the source, a simple solution is for the eFSA function to copy the source address and source port it receives in the request into the response. The iFSA function could save this information as the identifier of the flow. Then when the iFSA function moves, it makes a new QoS request and includes the saved original source address and port in the request instead of zero in these fields. The network treats this as a new request and times out the old one. However, the eFSA function can see that these fields are not zero and match up the original source address and port with ongoing flows and thus determine that this is a continuation of the flow.

An example implementation could utilize two fields for IPv4, the "original source address" (32 bits) and the "original source port" (16 bits). Since IPv6 requires 32 bits of header information, this space can be used for the IPv4 original source address. Then the IPv4 original source port can be placed in the same 20-bit space that IPv6 requires for the source flow label.

Appendix III

Illustrative QoS support for different preference indicator values

(This appendix does not form an integral part of this Recommendation)

III.1 Preference resolution for maximum rate (MRS) flows

This clause references requirement R-82.

The preference indicator value of a newly admitted MRS flow may alter the QoS support of MRS flows which have lower preference indicator values. If necessary, when there is insufficient capacity to support all MRS flows, some or all of the flows with lower preference indicator values may continue to transmit as in the case of the immediate transmission option. For such flows, the network is required to revert to the conditional guarantee that loss/delay targets are supported as soon as possible.

III.2 Preference resolution for available rate (ARS) flows

This clause references requirement R-83.

With ARS traffic (typically TCP), the network decides on the rate it could support for each flow. There may be different classes of ARS traffic where some flows are permitted more capacity, but within a class it assumed that the goal would be for rate equality or fairness. When a new ARS flow is received, the network equipment would examine the available capacity for the class and assign what it believed was a fair rate to the flow. Typically, the process would allow the total ARS traffic in the class to utilize as much of the class capacity as is presumed to be safe and thus most of the time there would be a preference decision to be considered.

Requirement R-83 does not specify how much more or less capacity is given to flows of different preference indicator values, and that may well be a function that the network operator may wish to control. This requirement however leads to a very simple process, the preference indicator value can be converted to a weight by a function or a table lookup so that a weight is determined for each flow. Then each new flow could be assigned a weighted fraction of the total capacity. One measurement is required, the sum of the flow weights in process. The new flow would then receive its weight times the total safe capacity divided by the sum of the active flow weights. This is one solution that illustrates how a large number of separate preference indicator values could be supported without the computation and memory requirements increasing with the number of such values.

Appendix IV

Supplementary information relating to requirements on the management of transport connections carrying flow-state-aware traffic and other traffic

(This appendix does not form an integral part of this Recommendation)

IV.1 General architectural assumptions

This clause references requirements in clause 6.5

The main requirement is that flow-state-aware QoS controls are agnostic to the underlying transport technology. This can be Ethernet, ATM or any other choice of the network provider, and may be a mixture of different transport technologies at different locations.

Additional assumptions

A given network link (between two network nodes) may not be dedicated to carrying flow-state-aware traffic only.

Where a given network link is carrying a mixture of flow-state-aware traffic and other traffic, the FSA node may assume that part of the capacity of that link is guaranteed to be available for flow-state-aware traffic. The role of FSA QoS control should be to control the utilization of that capacity among the competing flows taking account of FSA services (ARS, MRS, VRS, GRS). The network may achieve capacity guarantees in various ways, including:

- IP-layer scheduling functions to manage and limit the capacity available to the non-FSA traffic, together with FSA QoS controls to limit the capacity available to the FSA traffic. Note that this method may allow for the dynamic "borrowing" of unused capacity by either FSA or non-FSA traffic as discussed in the next bullet.
- FSA management of link capacity does not imply a dedicated FSA management function per link. For example, consider the broadband access scenario shown in Figures IV.1 and IV.2. These illustrate an arrangement where a single FSA access QoS functional entity is capable of managing both DSL and Ethernet link capacity limits. Clearly the scale capabilities of this arrangement will be vendor-specific. In the case of Figure IV.2, the FSA access QoS management functional entity is shown as being combined with an Ethernet VLAN switch. Of course, these two functions could also be split. The main difference with Figure IV.1 is that all traffic, FSA and non-FSA, passes through FSA QoS control management. The non-FSA traffic is either passed transparently or managed, e.g., throttled (as necessary), if the network provider chooses to offer such a service.
- Figures IV.1 and IV.2 also illustrate the requirement that FSA management of link capacity is required to take into account the capacity reserved on that link for non-FSA traffic on both the DSL and Ethernet links, i.e.,:
  - the non-FSA traffic within (in the scenario of Figures IV.1 and IV.2) an Ethernet VLAN carrying a mix of FSA and non-FSA traffic, or
  - the additional non-FSA traffic on the physical link towards an access node, in dedicated non-FSA VLANs.

  Note that, within a single VLAN, dynamic "borrowing" of capacity between FSA and non-FSA traffics may be performed by the FSA access QoS functional entity.

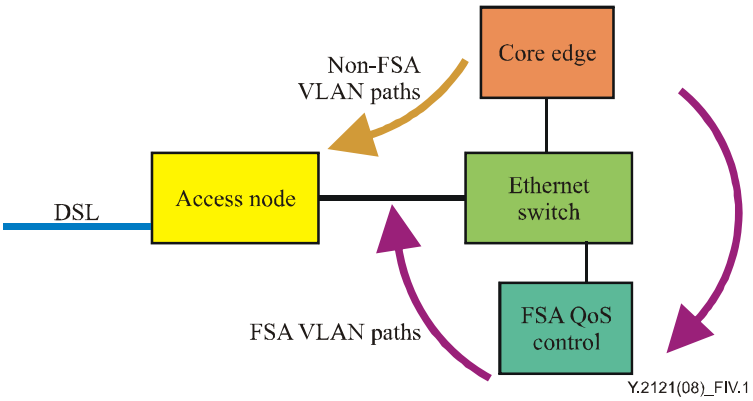

Figure IV.1 – FSA control of some of the access VLANs

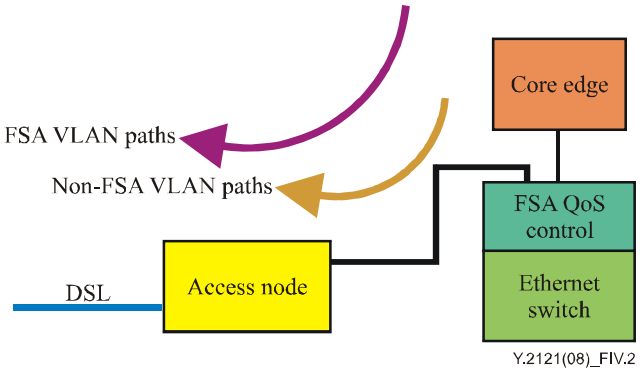

Figure IV.2 – FSA IP-layer inspection and management of all access links

IV.2 General issues on the management of access links shared by FSA and non-FSA traffic

With reference to Figure IV.3 below, an access link is shown with downstream traffic being forwarded on to it from an edge FSA node. The downstream traffic consists of both FSA and non-FSA components.

Figure IV.3 shows a grooming process that consists of two stages. Stage 1 is the separate grooming of FSA flows (with the non-FSA traffic by-passing this stage). Stage 2 is the grooming of the FSA traffic with the non-FSA traffic.

It will be appreciated that this description of a two-stage process could be realized without actually implementing two physically separate stages. The description has been chosen for clarification of QoS management, and not to suggest any particular implementation.

Similarly, although Figure IV.3 shows both FSA and non-FSA capabilities in an edge FSA node, this does not suggest an actual implementation. The stage 1 function could be implemented within the edge FSA node or external to it.

In more detail, stage 1 may groom traffic for one or several Ethernet VLANs or ATM VPs according to different policies, including:

- Provider-specific requirements on preference indicator handling.
- Aggregate maximum rate limit per VLAN or VP and UES maximum rate limit for downstream traffic.
- A multiple-provider shared access with an overall maximum rate limit and with equal or preferential policies on ARS rate allocation or MRS discard first.

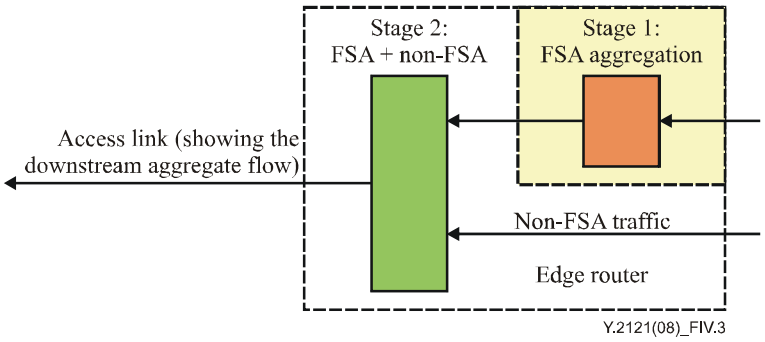

Figure IV.3 – Conceptual 2-stage grooming of FSA + non-FSA for a shared access link

As shown in Figure IV.4 below, stage 1 grooms the FSA traffic into virtual access links corresponding to one or several different access products. In addition, several of these virtual access links may share a single physical access link. The method of sharing may allow:

- strict adherence to the maximum rate on each of these virtual access links;
- a virtual access link borrows unused physical link capacity (in terms of what is pre-set as the notional physical link capacity assigned for stage 1). It returns to its guaranteed maximum rate as other virtual links demand more.

In the simplest of these cases (strict adherence to each assigned maximum rate), the stage 1 grooming process:

- manages ARS rate shares within any virtual link on the basis of the available maximum rate of that virtual link, subtracting any GRS or MRS discard last rates from the total available;
- manages MRS packet discards to ensure strict adherence to the maximum rate.

In the alternative case, MRS packet discards may be reduced if there is some unused physical capacity. An alternative, and possibly more complex, arrangement would allow ARS rates to exploit the unused capacity and rapidly relinquish it as demands from other virtual links increase.

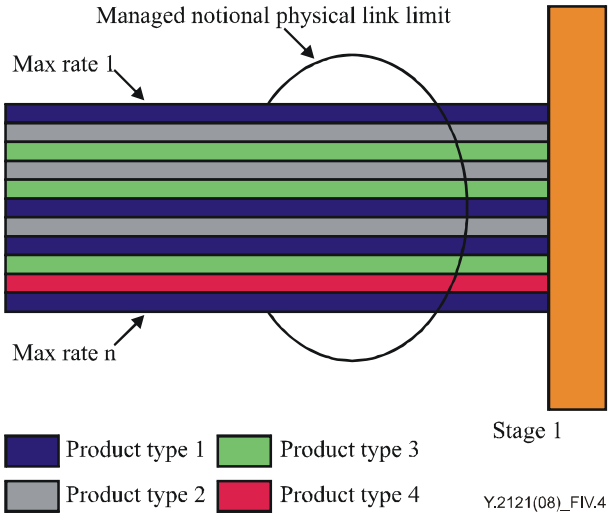

Figure IV.4 – FSA traffic groomed into different virtual access links at stage 1

The stage 2 grooming process forwards these virtual access links onto a number of physical access links, combining this traffic with the non-FSA traffic. Prior to stage 2, the non-FSA traffic may also

be groomed into one or several virtual access links per physical access link (e.g., associated with one or several Ethernet VLANs).

If the FSA traffic is groomed into Ethernet VLANs at stage 1 and if the non-FSA traffic has also been groomed into Ethernet VLANs, then stage 2 grooming could be performed by an Ethernet switch. In this case, for QoS reasons, there should be strict adherence to the maximum rate of each virtual access link. Where several virtual access links share the same physical access link, the sum of their maximum rates should be managed to be less than the physical link capacity.

On the other hand, it was noted that one type of access product may borrow unused capacity and may temporarily exceed its guaranteed maximum rate. To perform stage 2 grooming in this case, it is necessary to take account of congestion. One option is to perform stage 2 grooming in the edge FSA node. Another possibility is to perform this function in an external function that may or may not be FSA-capable (if the latter, then this external function may also simultaneously perform stage 1 – see, for example, Figure IV.2).

IV.2.1 Case 1: FSA-aware stage 2

Consider Figure IV.5, which shows the case of an FSA-aware stage 2. It is assumed that any reduction in bandwidth available to a virtual access link at the output of stage 2 may not immediately cause a reduction in the output load from stage 1. Therefore, on an FSA virtual access link that is allowed to borrow unused bandwidth above its guaranteed maximum, the simplest procedure is:

- stage 2 allows unused capacity to be used for forwarding packets of this virtual access link, but reduces the aggregate forwarding rate to the guaranteed maximum rate when necessary;
- MRS "discard first" packets may be discarded if there are too many packets waiting to be forwarded at stage 2;
- the sum of the rates of all GRS flows added to the sum of MRS flows in "discard last" state (including the MRS component of VRS) should not exceed the guaranteed maximum rate of the virtual access link;
- stage 2 should reduce ARS rates as appropriate when it is necessary to reduce the aggregate rate of the virtual access link to its guaranteed maximum rate.

Therefore, stage 2 supports all guaranteed bandwidth flows (whether in FSA or non-FSA virtual access links). This is provided that the call acceptance process does not allow any guaranteed bandwidth flow to be accepted if it should cause a virtual access maximum rate to be exceeded (and provided the sum of the set of maximum rates does not exceed a physical link rate).

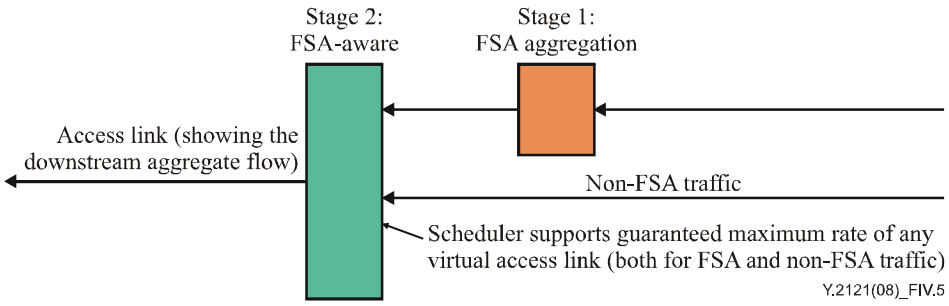

Figure IV.5 – Case 1 – FSA-aware stage

IV.2.2 Case 2: Non-FSA-aware stage 2

This case requires more complexity in the marking of packets so that stage 2 correctly discards excess packets when it has to limit the virtual access link to its guaranteed maximum rate. One possibility is that MRS "discard first" packets are marked differently to MRS "discard last" packets

(requiring the marking to be changed on a flow if the "discard first" state changes to "discard last"). Then stage 2 operates as follows:

- stage 2 allows unused capacity to be used for forwarding packets of this virtual access link, but reduces the aggregate forwarding rate to the guaranteed maximum rate when necessary;
- packets appropriately marked for discard are required to be discarded if there are too many packets waiting to be forwarded at stage 2;
- the sum of the rates of all GRS + ARS flows added to the sum of VRS + MRS flows in "discard last" state should not exceed the guaranteed maximum rate of the virtual access link.

This implies that only MRS "discard first" packets are exploiting the excess unused capacity on the physical link. With these restrictions, the same statement may be made that a non FSA-aware stage 2 supports all guaranteed bandwidth flows (whether in FSA or non-FSA virtual access links). Again, this is provided that the call acceptance process does not allow any guaranteed bandwidth flow to be accepted if it should cause a virtual access maximum rate to be exceeded (and provided the sum of the set of maximum rates does not exceed a physical link rate).

IV.3 Combined flow-level and aggregate flow-level FSA controls

It may be anticipated that network providers may want to combine FSA controls at the flow level and aggregate flow level. This creates the opportunity of providing variable-rate aggregation products that are adapted to the needs of the service providers and their customers.

The following text describes examples of variable rate aggregation products for the purposes of clarifying requirements only:

- products that connect between (see Figure IV.6):
  - an access product service edge point, carrying traffic to/from a service provider point of presence;
  - an access product service edge point, carrying traffic to/from a given set of user-end-systems;
- connecting between the service edge points of access products that carry traffic to/from enterprise sites (see Figure IV.7);
- assigning capacity to the aggregate based on the latest available rate, where the flow level consists of only a set of ARS flows that share this capacity on a basis that recognizes preference indicator values (see Figure IV.8);
- assigning capacity to the aggregate based on a minimum (MRS-requested) rate plus an available rate top-up. Here the set of flows can be MRS, VRS and ARS and, again, the capacity sharing of the top-up portion recognizes preference indicator values (see Figure IV.9);
- assigning capacity (perhaps to cover sudden needs) based on a new MRS-requested aggregate rate while recognizing that this new rate may not be instantly available. However it is required to be made available as soon as possible and, meanwhile, the policed rate is required to be adjusted to the new need. Here the flow level consists of a set of MRS flows only, and discard control again recognizes preference indicator values (see Figure IV.10);
- assigning capacity based on a new GRS-requested aggregate rate, with the possibility that this request may be rejected. Accepted requests may remove some of the capacity available to other services and trigger flow-level controls to operate and modify, for example, ARS rates (again taking preference indicator values into account). (see Figure IV.11).

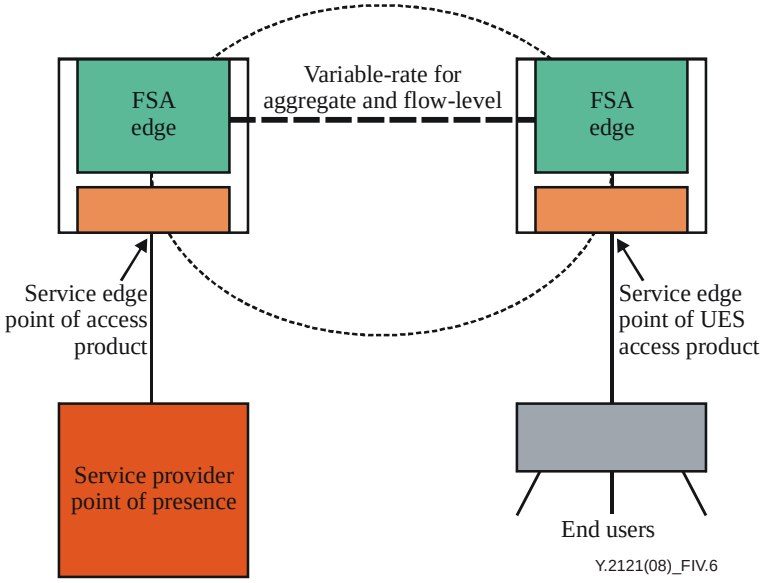

Figure IV.6 – FSA variable capacity controls supporting the connection between a service provider and specified user-end-systems

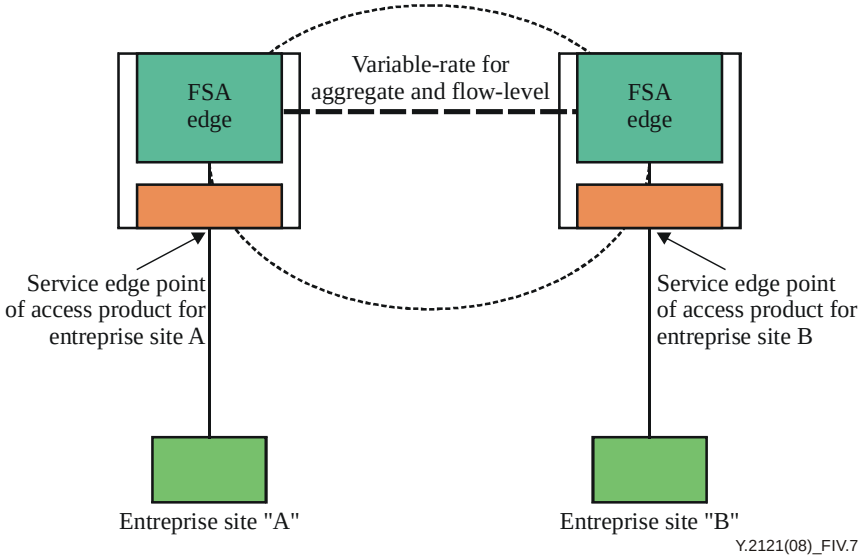

Figure IV.7 – Enterprise site-to-site FSA controlled variable capacity

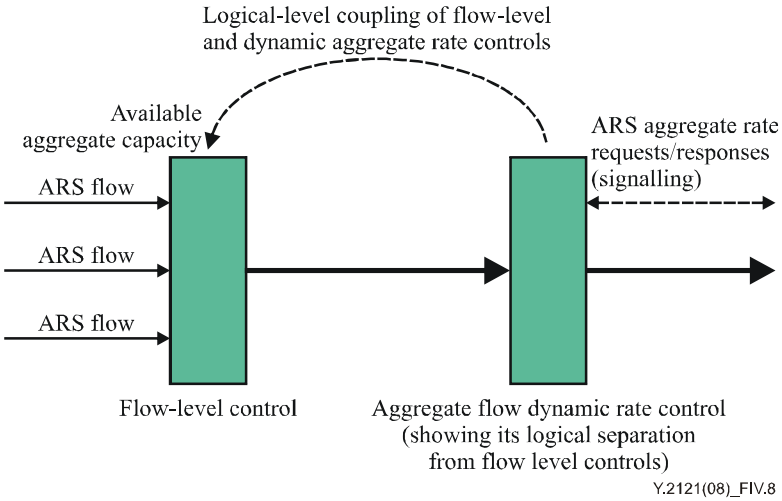

Figure IV.8 – ARS flows within an ARS aggregate access

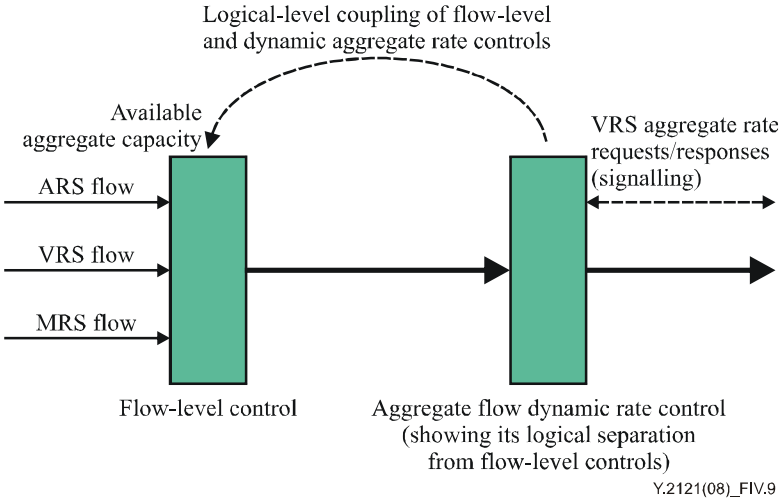

Figure IV.9 – ARS, VRS and MRS flows within a VRS aggregate access

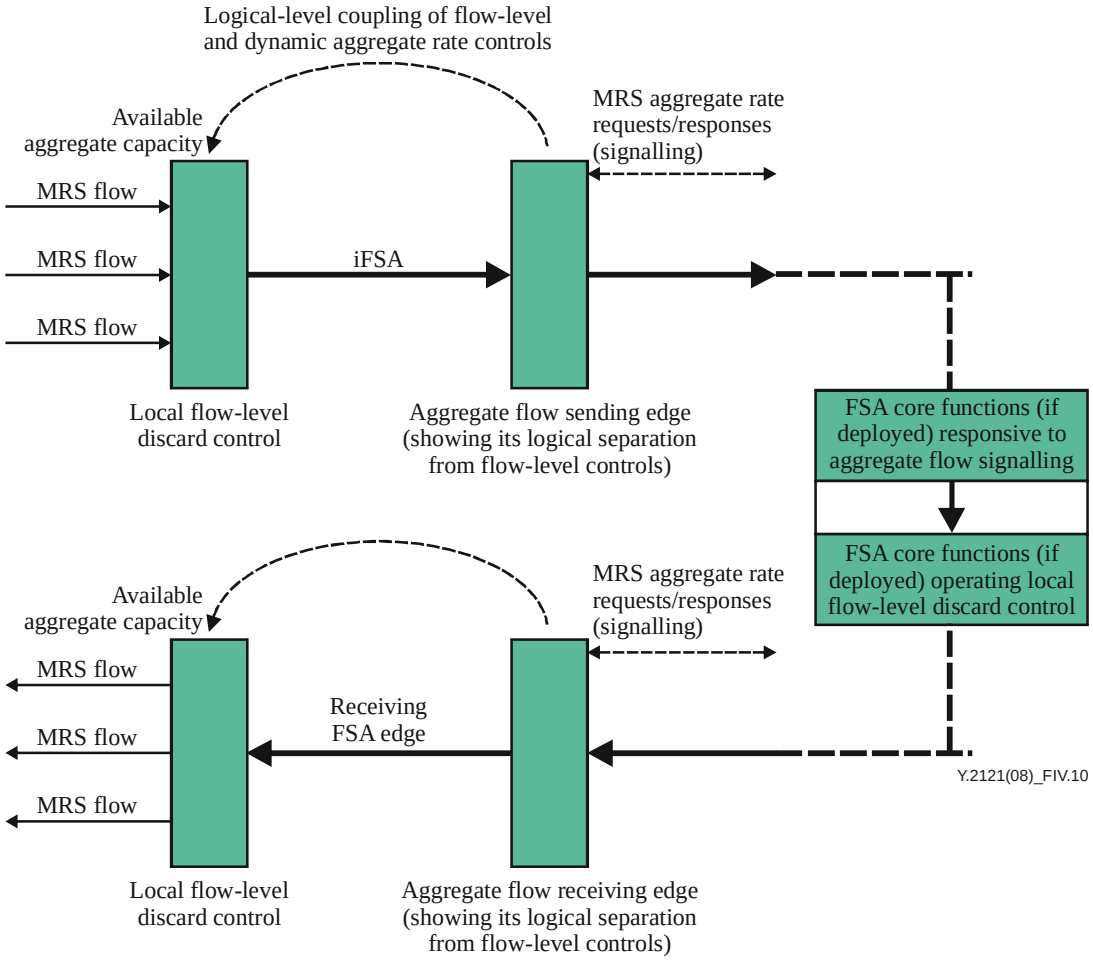

Figure IV.10 – MRS flows within an MRS aggregate

It may be noted in Figure IV.10 that each FSA node is making a local decision about the rate that could be supported on reception of an MRS aggregate rate increase request. Unlike ARS, the sending end may immediately send at the new requested aggregate rate. However, network providers may choose to prevent any sending node from increasing a rate beyond some pre-determined maximum rate or beyond some pre-determined maximum allowable percentage increase above a current rate. Other policy options may be applied.

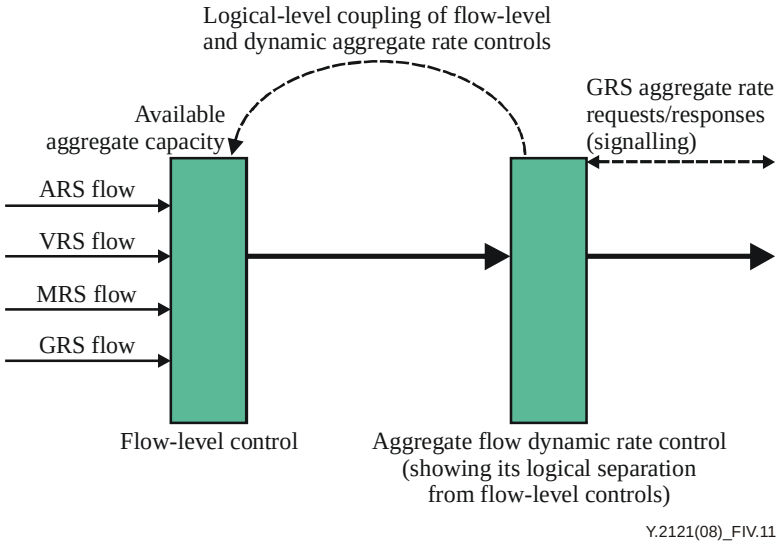

Figure IV.11 – GRS negotiated aggregate rate changes

Appendix V

Example implementation principles associated with FSA nodes

(This appendix does not form an integral part of this Recommendation)

The principles below describe an example implementation, not the only possible implementation of an FSA node:

1) The FSA node determines the identifier of each QoS signalled flow.

2) Having identified a flow, some memory of the flow is required. This memory may be as little as a saved hash entry and a bit to indicate if the flow is "discard last" or "discard first". It may also contain the QoS information for the flow that was in the initial in-band signalling packet.

3) The FSA node also needs to be monitoring the loading of the port on which the flow is exiting the FSA node. From this load information, it must determine what capacity is available for a new flow. There are many techniques for doing this and deciding what rate an ARS flow may have or if a MRS flow should be "discard first", "discard last", or must be denied altogether.

4) The FSA node only needs to route this first packet of the flow to determine what route it should take, what QoS it should have, and if it is subject to denial of service (DoS).

5) Once a flow has been accepted at a rate, if the rate is saved in the state information, then the flow may be policed to that rate. If no rate information is kept, the policing could be based on the total port load and, for ARS, this plus the number of flows may be sufficient. For GRS the rate is required since the total commitment must be controlled. For MRS and the MRS portion of VRS, the rate would need to be saved if per-flow policing is required, but if only flow acceptance is needed then the "discard first" bit would be sufficient.

6) The flow entries need to be cleared out if no packets arrive for a period, thus there also needs to be a time stamp for the last packet arrival.

Figure V.1 shows an example implementation of an FSA node where the data flow follows the path through the switching logic and the first packet is routed with normal L3 FSA node logic.

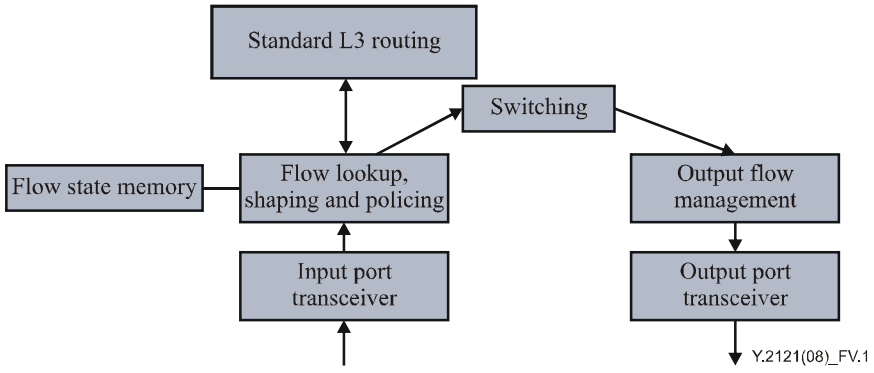

Figure V.1 – Example implementation illustrating an FSA node

Once the packet is identified with a flow record which includes the agreed rate, the flow can be policed and thus the total output load on a port controlled to any desired loading. New flows can be accepted by examining the remaining capacity and assigning a rate to the flow. MRS flows may be controlled not to exceed some load limit, but if too many arrive, they could be accepted as "discard first" or rejected if necessary. Thus, the ability to support the envisioned QoS signalling protocol only requires a modest memory and logic capability which need not add any significant cost to the FSA node, and in fact saves cost by reducing the routing logic.

Appendix VI

Out-of-band signalling with a central admission entity

(This appendix does not form an integral part of this Recommendation)

Under assumption that a central admission entity such as RACF performs the admission function accurately, the messaging overhead for the response and the confirmation is not required. This we call the simplified FSA signalling.

The simplified FSA signalling can further benefit by combining with the proxy signalling. In this case, any registered UES for FSA treatment, without having the FSA signalling capability, can simply initiate a call. Then through the call-level authorization and admission process, the IP-level traffic parameters (e.g., RACF traffic descriptor) are notified to the RACF. The RACF then distributes the traffic descriptor to the appropriate ingress edge node, which is the signalling proxy. The proxy maps the IP-level traffic descriptor distributed by the RACF to the FSA parameters, prior to sending the start packet. The UESs then do not have to register the individual FSA parameters to the proxy. The simplified signalling procedure is depicted in Figure VI.1.

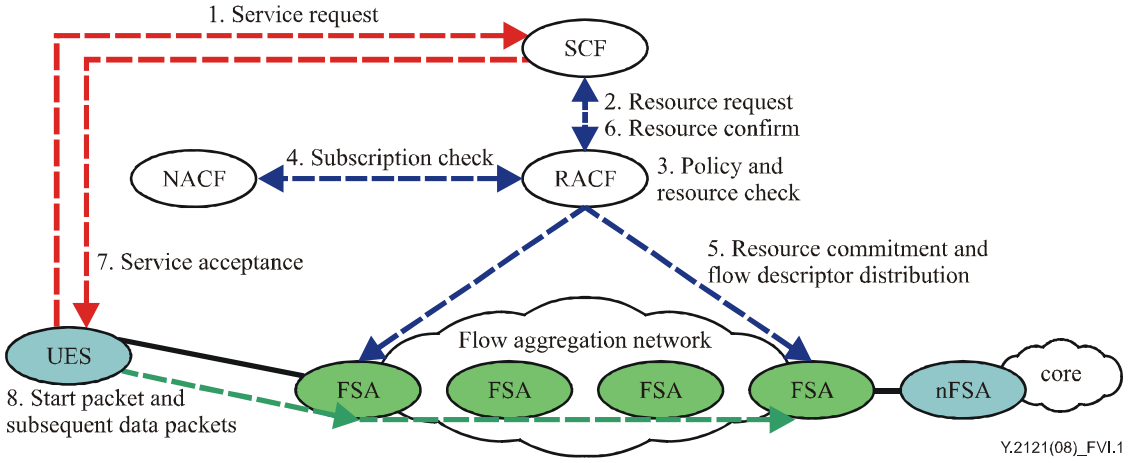

Figure VI.1 – Simplified signalling procedure

1) A UES requests a service for a call.
2) The SCF requests for available resource to the RACF.
3) The RACF checks the applicable policy to the call, thus to the flow, then checks for the available resources in the network.
4) The RACF also checks for the subscription status.
5) Upon passing the checks, the RACF decides to admit and commit resource to the call. The RACF also distributes the IP-level flow descriptor. If any of the ingress edge nodes work as a signalling proxy, then such an edge node maps the IP level flow descriptor into the FSA QoS parameters such as service context, burst tolerance and delay priority. Those FSA parameters are stored in the ingress edge node.
6) The RACF then confirms the available resource.
7) The SCF notifies the UES of the service acceptance.

8) The UES generates the in-band FSA request signal. Every FSA in the path recognizes it, and stores the FSA parameters for the flow. If the ingress edge node performs as a proxy, the UES just transmits the ordinary data packet. The ingress edge node realizes the first data packet sent from the UES and generates the start packet using the FSA parameter mapped from the IP-level flow descriptor.

It is likely that the central admission entity examines only the requested data rate of a flow, which corresponds to the requested rate in FSA. Other parameters such as preference indicator, service context, burst tolerance and delay priority can be inferred from the IP-level flow descriptor received from the central admission entity.

In summary, the negotiation procedure can be through the complete in-band signalling (including the response and the confirmation) or the in-band signalling plus the authorization signalling to the central admission entity. In the latter case, the start packet can be generated at the ingress edge node by mapping the central admission entity IP level descriptor (such as RACF traffic descriptor) to FSA parameters.

Bibliography

[b-ITU-T I.371] ITU-T Recommendation I.371 (2004), *Traffic control and congestion control in B-ISDN*.

[b-IETF RFC 2003] IETF RFC 2003 (1996), *IP encapsulation within IP*.
<http://www.ietf.org/rfc/rfc2003.txt>

[b-IETF RFC 2748] IETF RFC 2748 (2000), *The COPS (Common Open Policy Service) Protocol*.
<http://www.ietf.org/rfc/rfc2748.txt>

[b-IETF RFC 2868] IETF RFC 2868 (2000), *RADIUS Attributes for Tunnel Protocol Support*.
<http://www.ietf.org/rfc/rfc2868.txt>

[b-IETF RFC 3270] IETF RFC 3270 (2002), *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
<http://www.ietf.org/rfc/rfc3270.txt>

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	**Global information infrastructure, Internet protocol aspects and next-generation networks**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2008