

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.

113

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/>.

© ITU 2008

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Page

1	Scope	1
2	References.....	1
3	Definitions	1
4	Abbreviations and acronyms	3
5	Overview	4
6	Requirements	6
6.1	Dynamic provisioning requests from an FSA signalling edge function.....	6
6.2	Network and FSA signalling edge function responses to flow-state-aware requests	9
6.3	Signalling requirements.....	10
6.4	Admission decision	14
6.5	General architectural requirements on the management of transport connections carrying flow-state-aware traffic and other traffic	14
6.6	Security considerations and requirements	15
Annex A –	Dynamic provisioning requests from an FSA signalling edge function.....	16
A.1	Negotiations for a user-end-system without signalling capability	16
A.2	Authorization.....	16
A.3	Service context	17
Annex B –	Signalling requirements	19
B.1	Second QoS structure attached.....	19
B.2	Authorization information attached.....	19
B.3	Flow aggregation request	19
B.4	FSA node operation.....	19
Appendix I –	Supplementary information on information exchanges via requests from an FSA signalling edge function and associated responses.....	20
I.1	Flow identifier	20
I.2	In-band signalling negotiations	20
I.3	Preference indicator request	21
I.4	Authentication	21
I.5	Priority of packet discard, including service context use of preference indicator values.....	23
I.6	Congestion notification	23
Appendix II –	Supplementary information to signalling requirements	24
II.1	Recognition of QoS signalling packets	24
II.2	Form of QoS information	24
II.3	Performance requirements for requests and responses.....	24
II.4	Release of resources no longer required.....	25
II.5	QoS signalling parameters.....	25

	Page
II.6 Service contexts.....	25
II.7 Preference indicator.....	26
II.8 Delay priority.....	26
II.9 Burst tolerance.....	26
II.10 Flow identifier fields	26
Appendix III – Illustrative QoS support for different preference indicator values.....	28
III.1 Preference resolution for maximum rate (MRS) flows	28
III.2 Preference resolution for available rate (ARS) flows.....	28
Appendix IV – Supplementary information relating to requirements on the management of transport connections carrying flow-state-aware traffic and other traffic	29
IV.1 General architectural assumptions.....	29
IV.2 General issues on the management of access links shared by FSA and non-FSA traffic.....	30
IV.3 Combined flow-level and aggregate flow-level FSA controls	33
Appendix V – Example implementation principles associated with FSA nodes.....	38
Appendix VI – Out-of-band signalling with a central admission entity	40
Bibliography.....	42

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.

Copyright International Telecommunication Union (ITU) 2008

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.