International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1223

(07/2008)

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

Internet protocol aspects – Architecture, access, network capabilities and resource management

Interworking guidelines for transporting assured IP flows

Recommendation ITU-T Y.1223

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.1223

Interworking guidelines for transporting assured IP flows

Summary

In order to transport IP flows with assured end-to-end quality and reliability in a multi-provider environment, coordinated decisions must be made regarding the admission, policing, and assignment of resources to particular offered flows. To do this, a uniform way of characterizing such IP flows is needed, and some shared decision rules for handling them. Recommendation ITU-T Y.1223 defines a set of IP flow specifications that could offer a basis for such cooperation.

Source

Recommendation ITU-T Y.1223 was approved on 14 July 2008 by ITU-T Study Group 12 (2005-2008) under Recommendation ITU-T A.8 procedure.

Keywords

IP, IP flows, IP priority, IP QoS, IP traffic, Pspec, QoS, Qspec, Tspec.

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 not 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

Introduction

A major challenge facing network providers in transporting IP flows with assured quality and reliability levels is the fact that several independently-operated networks, using different QoS control and resource management technologies, will typically share the responsibility for providing the end-to-end services. This Recommendation addresses this challenge by extracting, from existing standards, a set of flow specifications that network operators need to share in making coordinated admission, policing, and resource assignment decisions for offered IP flows. This Recommendation also suggests a set of basic decision rules that cooperating network providers could follow, in interpreting and acting on such shared information, to ensure that the end-to-end quality and reliability requirements of admitted flows are consistently met.

Recommendation ITU-T Y.1540 provides the parameters and definitions for IP QoS; Recommendation ITU-T Y.1541 gives a set of end-to-end IP QoS classes with numerical objectives for these parameters; and Recommendation ITU-T Y.1542 provides a framework of approaches for achieving end-to-end IP QoS. This Recommendation is the logical next step in this sequence – a uniform way of dealing with IP flows, so that multiple providers have a basis for coordinating to deliver end-to-end IP QoS with assured quality and reliability.

Recommendation ITU-T Y.1223

Interworking guidelines for transporting assured IP flows

1 Scope

This Recommendation provides a set of specifications for characterizing IP flows requiring assured quality and reliability levels, and identifies possible decision rules for interpreting and processing that information to ensure that the requested quality and reliability requirements of admitted flows are consistently met. Extracted from existing standards, a set of flow specifications is provided that network operators need to share in making coordinated admission, policing, and resource assignment decisions for offered IP flows. A set of basic decision rules are suggested that cooperating network providers could follow, in interpreting and acting on such shared information, to ensure that the end-to-end quality and reliability requirements of admitted flows are consistently met. This Recommendation makes no assumptions about how the shared information is exchanged among cooperating network providers. The focus here is on the *semantics* and *processing* of the shared information, rather than on how it is exchanged.

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.1540] Recommendation ITU-T Y.1540 (2007), *Internet protocol data communication service – IP packet transfer and availability performance parameters*.

[ITU-T Y.1541] Recommendation ITU-T Y.1541 (2006), *Network performance objectives for IP-based services*.

[ITU-T Y.1542] Recommendation ITU-T Y.1542 (2006), *Framework for achieving end-to-end IP performance objectives*.

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

[ITU-T Y.2171] Recommendation ITU-T Y.2171 (2006), *Admission control priority levels in Next Generation Networks*.

[ITU-T Y.2172] Recommendation ITU-T Y.2172 (2007), *Service restoration priority levels in Next Generation Networks*.

3 Definitions

This Recommendation defines the following terms:

3.1 **Qspec**: A set of QoS-related IP flow parameters (extracted from existing standards).

3.2 **Pspec**: A set of Priority-related IP flow parameters (extracted from existing standards).

3.3 **Tspec**: A set of Traffic-related IP flow parameters (extracted from existing standards).

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

AS	Autonomous System
ASBR	Autonomous System Border Router
DiffServ	Differentiated Services
EXP bits	Experimental bits (used to support DiffServ) in the MPLS header
FIFO	First In, First Out
IP	Internet Protocol
IPDV	IP Packet Delay Variation
IPLR	IP Packet Loss Ratio
IPTD	IP Packet Transfer Delay
ISP	Internet Service Provider
MPLS	MultiProtocol Label Switching
OAM	Operations, Administration and Maintenance
QoS	Quality of Service
SLA	Service Level Agreement
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network

5 Conventions

There are no conventions specific to this Recommendation.

6 Interworking guidelines for transporting assured IP flows

End-to-end, IP-based service delivery is typically shared by several independently-operated networks, often using different QoS control and resource management technologies. Previous studies [b-T1A1/2003-133] and [b-PRQC] have identified the general elements of a standards-based approach to this problem, and a number of standards (e.g., [ITU-T Y.1221], [ITU-T Y.1541], [ITU-T Y.2111] and [b-IETF RFC 4804]) have addressed particular solution elements. The next step is to produce a set of standards-based flow specifications that network operators need to share in making coordinated admission, policing, and resource assignment decisions for offered IP flows. Also needed is a set of basic decision rules that cooperating network providers could follow, in interpreting and acting on shared information, to ensure that end-to-end quality and reliability requirements of admitted flows are consistently met.

No assumptions are made here regarding how the shared information is exchanged among cooperating network providers. At one extreme of complexity (and flexibility), the information could be explicitly encoded in signalling messages and communicated among providers on a per-flow basis, e.g., using procedures like those defined in [b-ITU-T Q-Sup.51] or [ITU-T Y.2111]. At the other extreme, the shared information could be captured in static network management databases derived from inter-provider service level agreements (SLAs) or known network attributes, and no explicit signalling of IP flow specifications might be needed. In intermediate situations, the shared information could be communicated indirectly through other signalled information (e.g.,

Ethernet traffic types, DiffServ code points, MPLS EXP bits). In some cases, the shared information could be inferred from routing information, VPN addresses, service classes, or topological cues, such as the ingress or egress port at which a particular flow appears. The focus here is on the *semantics* and *processing* of the shared information, rather than how on it is exchanged.

To simplify the problem, we focus on the general interworking model shown in Figure 1. The figure depicts two independently-operated networks (Ni and Nj), which represent separate domains or autonomous systems (ASs) connected to each other (and potentially, to other networks or to user equipment) by a number of links between autonomous system border routers (ASBRs). Ni and Nj are independently administered and are viewed as separate entities from a modelling point of view. They may be comprised internally of multiple ASs. In general, the two networks will have different internal topologies, point-to-point transmission capacities, and point-to-point QoS characteristics. Their instantaneous offered and carried traffic levels will also differ. They may use different internal transport technologies, signalling protocols, and routing protocols. Their internal admission control, policing, and resource assignment policies and mechanisms may also differ. Despite these differences, the network operators will need to find ways to cooperate based on certain shared information, in transporting IP flows. Their mutual goal is to accept and successfully transport a high proportion of offered IP flows, requiring a wide range of assured quality and reliability levels, under widely varying network conditions including, *inter alia*, variable offered and carried traffic levels, variable network topologies and routing states, and variable impairment levels on particular transmission paths.

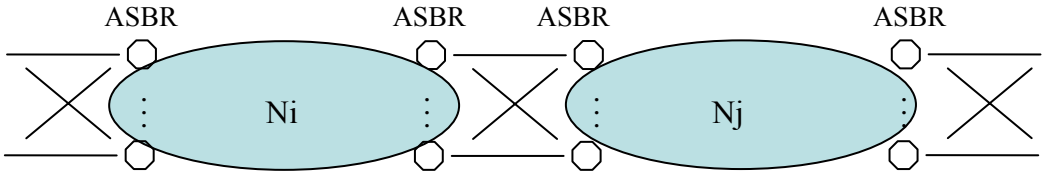

Figure 1 – General interworking model

This rest of this clause addresses the information about offered IP flows (and the ability of a network to support them) to be shared among concatenated networks (e.g., Ni, Nj) so that networks cooperate successfully in transporting IP flows with assured quality and reliability levels.

6.1 IP flow specifications

It is suggested that IP flows requiring assured quality and reliability levels can be described, from a technical (non-administrative) point of view, in terms of three types of IP flow specifications: a traffic specification (Tspec), a priority specification (Pspec), and a quality specification (Qspec). The proposed specifications are presented and described, using information abstracted from existing ITU-T Recommendations, in the following subclauses.

6.1.1 Traffic specification (Tspec)

The most likely IP flow attribute that will need to be specified (and shared among network providers supporting end-to-end services with assured quality and reliability levels) is *flow rate*. The flow rate of an IP stream will largely determine the bandwidth and buffer space that must be allocated to its transport, as well as the associated transport "cost". A continuous flow can be specified with one or more upper limiting values (e.g., peak rate in bytes per second), and some important NGN flows (e.g., flows supporting circuit emulation service) will be continuous. However, many NGN IP flows will have variable rates, and it will not be cost effective to

dimension networks to support a "worst case" peak rate for each. To allocate resources to variable rate IP flows economically, cooperating networks will need to specify and share statistics that describe and limit the flow variability – and "police" input flows to ensure that they respect the established "traffic contracts".

A practical and generally accepted method of specifying variable rate flows is defined in [ITU-T Y.1221]. Four flow parameters are defined in terms of a flow description algorithm called a token bucket, as summarized in Annex A. Briefly, a token bucket algorithm characterizes a flow in terms of two parameters:

- a transfer rate *R*, expressed in bytes/second;
- a bucket size *B*, expressed in bytes.

[ITU-T Y.1221] specifies IP flows in terms of two token buckets: a peak token bucket, with parameters peak rate *Rp* and peak bucket size *Bp*, and a sustainable token bucket, with parameters sustainable rate *Rs* and sustainable bucket size *Bs*.

[ITU-T Y.1221] defines one additional IP flow parameter: the maximum allowable packet size, *M*. Thus a total of five parameters should be specified in a complete IP flow Tspec.[1]

6.1.2 Priority specification (Pspec)

In many countries, NGNs will be required to afford authorized "critical users" priority access to, and restoral of, telecommunication services under specified emergency conditions. [ITU-T Y.2171] and [ITU-T Y.2172] specify, respectively, admission control and restoral priority levels for NGN emergency services. The Y.2171 admission priority levels make it possible for networks to recognize and accept call requests (and other IP flows) from designated critical users in preference to those of non-critical users under congested network conditions. The Y.2172 restoral priority levels make it possible for networks to restore facilities and services supporting critical users prior to those supporting only non-critical users when network outages or other service disruptions occur. Since IP flows will typically traverse multiple networks, it will be important to share this priority information among cooperating networks.

Three admission control priority levels are specified and described in [ITU-T Y.2171]:

Priority level 1: Traffic with this priority level receives the highest assurance for admission to the network. This level is reserved for emergency telecommunications over NGN.

Priority level 2: Examples include real-time services (VoIP, video), VPN, and data services. The selection of this priority level is expected to be determined by the appropriate service level agreements (SLAs) between network operators and customers for the desired service.

Priority level 3: Traffic with this priority level receives the least assurance for admission to the network. Examples include "traditional" Internet service provider (ISP) services (e-mail, web surfing). The selection of this priority level is expected to be determined by appropriate SLA agreements between network operators and customers for the desired service.

Three corresponding restoral priority levels are specified and described in [ITU-T Y.2172]:

Priority level 1: Traffic with this priority receives the highest assurance of restoration. This class must include control services crucial to the operation of a network and emergency telecommunications. Other services may be included depending on availability of restoration capacity and service level agreements (SLAs) between network operators and customers for the desired service.

---

[1] Fewer parameters may be specified in describing constrained flows, e.g., flows supporting circuit emulation or best-effort services. See [ITU-T Y.1221].

Priority level 2: Traffic with this priority will receive lower assurance than priority level 1 traffic but will receive higher assurance than priority level 3 traffic for restoration. Examples include real-time services (VoIP, video), VPN, and data services. The selection of this priority class is expected to be determined by appropriate SLA agreements between network operators and customers for the desired service.

Priority level 3: Traffic with this priority receives the least assurance for restoration. Examples include "traditional" Internet service provider (ISP) services (e-mail, web surfing). The selection of this priority class is expected to be determined by appropriate SLA agreements between network operators and customers for the desired service.

Although the priority levels defined in these two Recommendations are conceptually similar, they differ operationally and the admission and restoration priority levels assigned to particular critical users may differ. Accordingly, the admission and restoration priority levels should be specified and shared among cooperating networks as independent requirements in a complete Pspec. This raises the issue of the rules and policies to be used by individual operators for interpreting and processing the priorities of IP flows, to avoid transporting malicious and non-committed traffic. This is a subject for further study.

6.1.3 QoS specification (Qspec)

A number of international, regional, and technology-based standards organizations have developed specifications for use in describing and controlling the QoS of IP flows in a multi-provider environment. Some of the most detailed and generally-applicable results are presented in [ITU-T Y.1540] and [ITU-T Y.1541]. [ITU-T Y.1540] defines a set of performance parameters to be used in specifying and assessing the speed, accuracy, dependability, and availability of packet transfer in IP-based networks. The parameters apply to end-to-end, point-to-point IP services and to network portions that provide, or contribute to the provision of, such services. They are defined on the basis of packet layer reference events that may be observed at IP network boundaries. [ITU-T Y.1541] specifies end-to-end objectives for the Y.1540 parameters and defines six QoS classes (plus two new "provisional" classes), each of which captures the performance requirements of a group of related IP applications in a corresponding set of end-to-end performance parameter values. The QoS classes are intended to be communicated from end users to network providers and among network providers as a basis for coordinating QoS control decisions and ensuring that the specified end-to-end QoS objectives are met.

Three of the Y.1540 parameters are particularly important, and provide the basis for specifying the Y.1541 QoS classes. Abbreviated definitions for these parameters are provided below.

IP packet transfer delay (IPTD) – The time, $(t_2 - t_1)$ between the occurrence of two corresponding IP packet reference events: an ingress event at time t_1 and an egress event at time t_2, where $(t_2 > t_1)$ and $(t_2 - t_1) \leq T_{max.}$

IP packet delay variation (IPDV) – The difference between the IPTD of an observed IP packet and that of a defined reference packet communicated between the same ingress and egress points. The reference IPTD is the shortest delay observed between the ingress and egress points for the population of interest.[2]

IP packet loss ratio (IPLR) – The ratio of total lost IP packet outcomes to total transmitted IP packets in a population of interest. A lost IP packet outcome occurs when a packet input at an ingress point does not appear at a permissible egress point within the specified maximum packet transfer time, $T_{max.}$[3]

---

2 This can be estimated using the delay experienced by the first IP packet transferred between the relevant points during a measurement.

3 Some or all of the lost packet contents may be misdirected to an impermissible egress point.

A minimal description of IP flow QoS requires that values for these three performance parameters be specified. The simplest and most useful way to do that is by specifying one of the Y.1541 QoS classes, since each class defines a value for each of the three parameters.[4]

An important further study topic, not addressed in this Recommendation, is whether (and if so, how) timing synchronization requirements should be specified (and coordinated among networks) in an NGN context.

6.2 Using IP flow specifications in network planning and operation

Clause 6.1 recommended that IP flows requiring assured quality and reliability levels be described in terms of a traffic specification (Tspec), a priority specification (Pspec), and a QoS specification (Qspec), and defined a particular set of parameters, abstracted from existing ITU-T Recommendations, that should be included in each of these specifications. In this clause, we assume that two or more cooperating networks (e.g., Ni and Nj in Figure 1) have access to the three specifications characterizing a particular IP flow, and consider how the networks could interpret and act on this information to ensure that the quality and reliability requirements of IP flows traversing them are met. A very basic set of decision rules the cooperating network providers could follow is proposed.

6.2.1 Using traffic specifications

It is recommended that the traffic characteristics of IP flows be specified using a Tspec comprising five parameters: a peak rate *Rp* and peak bucket size *Bp*, a sustainable token rate *Rs* and sustainable bucket size *Bs*, and a maximum allowable packet size, *M,* as defined in [ITU-T Y.1221].

When two (or more) networks are interconnected in tandem to support an offered IP flow, the "bottleneck" principle applies; i.e., for a flow's Tspec to be supported end to end, each cooperating network must commit sufficient resources to support that flow's Tspec between its ingress and egress points. Any network that cannot support a flow's Tspec should reject the flow.[5] A network could of course assign capacity to an offered flow in excess of the Tspec, gaining "margin" or other benefits at the cost of less efficient resource utilization. Similarly, a network could support a maximum packet size larger than that specified in a flow's Tspec. In general, a network can support a flow requesting "statistical bandwidth" capability (indicated by the inclusion of sustainable rate and bucket size parameters in its Tspec, as described in [ITU-T Y.1221]) using a corresponding "dedicated bandwidth" capability (indicated by omission of those two parameters), but the converse is not true.

As long as each cooperating network knows an offered flow's Tspec, there should be no need for the networks to "negotiate" or otherwise interact in deciding if, and how, to support the flow. Each network must simply meet (or exceed) the flow's Tspec, or not accept the flow at the requested level.

6.2.2 Using priority specifications

It is recommended that the admission and restoral priority requirements of IP flows be specified using a Pspec comprising two parameters: an admission priority level, as defined in [ITU-T Y.2171], and a restoral priority level, as defined in [ITU-T Y.2172].

---

[4] Each Y.1541 QoS class also defines a value for IP packet error ratio. A somewhat different set of QoS classes has been defined by 3GPP for UMTS; relationships and a possible mapping among the ITU-T and 3GPP QoS classes are described in [b-T1A1/2003-075].

[5] Some networks may offer to support a less stringent flow specification rather than simply rejecting the flow. In no case should a network change an IP flow's original Tspec.

Although the Pspec levels are not quantitative, something similar to the Tspec "bottleneck principle" applies to their end-to-end fulfilment in a multi-provider environment. For a flow's Pspec to be supported end to end, each network that admits the flow must commit sufficient resources to support the Pspec between its ingress and egress points. A network could accept a flow it would otherwise not be able to support by internally giving the flow a higher admission or restoral priority level than that specified in the Pspec (at the risk of disadvantaging later flow requests with the higher priority level), or (more conservatively) could inform the IP end user that the flow could be accepted if it were requested with a higher Pspec level.

There may be cases where a network could accept an offered flow, but would not be able to support the flow's specified restoral priority. In such cases, the network should inform the requesting user (and intervening networks, as appropriate) of the situation so that alternatives can be considered. There may also be cases where a network could admit a flow with an admission priority lower than that specified in the Pspec. If a flow can be admitted, its admission priority level within a network would seem to be immaterial unless it could affect the network's support of the flow in some other way, e.g., eligibility for restoration. In the latter case, the network should inform "upstream" entities of the situation as described above. In no case should a network change an IP flow's original Pspec levels.

6.2.3 Using QoS specifications

It is recommended that the QoS requirements of IP flows be specified in terms of a Qspec designating one of the eight QoS classes defined in [ITU-T Y.1541]. As discussed in clause 6.1.3, the Y.1541 QoS classes are based primarily on specified values for three Y.1540 parameters: IPTD, IPDV, and IPLR. Some simple decision rules for interpreting and acting on the specified values for each of these parameters are introduced below.[6]

6.2.3.1 IP packet transfer delay

Meeting an end-to-end delay objective for an IP flow involves limiting the *total delay* introduced by all of the networks the flow traverses. In general, this requires that each network share an estimate of the delay the flow will experience, between that network's ingress and egress points, with one or more interworking networks (or with a "third party" acting as a decision entity). As noted earlier, such exchanges may be accomplished through signalling or various other means. The method of exchange is not addressed here.

The IPTD objectives specified in [ITU-T Y.1541] are mean delays. Such statistics have a desirable property of additivity, i.e., the means delays for individual concatenated networks can be summed to produce an unbiased estimate of the total end-to-end delay. Such calculation is a type of *accumulation*. Examples are given in [ITU-T Y.1542]. It is also possible, in cooperative network planning, to assign each of several concatenated networks a portion of a specified end-to-end IPTV objective, to be achieved (for example) through routing constraints. Such assignment has been called apportionment or *allocation*. Accumulation and allocation are the principal alternatives for relating end-to-end performance objectives with the objectives for individual concatenated networks.

To establish an end-to-end IP path with an assured IPTD value, the cooperating IP networks (or the deciding "third party") should:

1) estimate the end-to-end IPTD value the flow will experience in transiting the proposed path (e.g., by adding the IPTD values for the individual concatenated networks); and

---

[6] One important topic not addressed here is the possible need to coordinate timing synchronization requirements for particular IP flows among cooperating networks.

2) compare this end-to-end IPTD estimate with the specified IPTD for the requested Y.1541 QoS class.[7]

If the estimate exceeds the specified IPTD, a decision must be made by individual operators as to whether to still admit the flow, reject it, or to seek an alternate path capable of meeting the end-to-end IPTD objective. The mechanisms used in the latter case may involve "crankback", and a different chain of networks may ultimately be included in the established end-to-end path.

There are IP network QoS specifications that describe IP packet transfer delay using statistics other than mean IPTD. One example is [b-3GPP TS 23.107], which specifies a maximum (rather than a mean) IP packet transfer delay. As discussed in [b-T1A1/2003-075], adding the maximum delays for concatenated networks can produce unrealistically high end-to-end delay estimates. However, such estimates may be the best obtainable in some situations, and will be conservative from a performance assurance point of view. The overestimates will have less impact if the delays in question are small, as may be the case in describing the performance of access networks. This is a subject for further study.

6.2.3.2 IP packet delay variation

Delay variation must be specified and controlled in addition to absolute delay in IP networks to limit packet loss, and the sizes of "jitter buffers" required to prevent it. IPDV is a difficult IP flow characteristic to control in a multi-provider environment because it is specified in terms of a distribution range (rather than a simpler statistic such as a mean), and because its observed values depend strongly on traffic level, transmission link capacity, and packet size. Appendix IV of [ITU-T Y.1541] provides guidelines for combining IPDV values for individual router hops to establish an upper bound on the IPDV experienced by an end-to-end IP flow. The following factors are identified in that specification as the most significant contributors to IPDV for variation-sensitive flows:[8]

- Packet-to-packet differences in the processing delay for packet forwarding decisions (routing table look-up)
- Queuing of variation-sensitive packets behind other variation-sensitive packets
- The need to finish servicing (i.e., transmitting) a variation-insensitive packet already in service.

[ITU-T Y.1541] provides a procedure for calculating the aggregate IPDV for an end-to-end path as a function of the number of router hops in tandem and the relevant traffic levels and transmission link capacities for an assumed packet size, taking each of these factors into account. The procedure assumes the delays introduced in successive router hops are independent, but it is conservative in several other respects. Briefly, the procedure calculates delay variation limits for each of the three listed factors and uses convolution or simple addition to calculate an overall end-to-end delay distribution from the per-hop distributions. IPDV can then be calculated as the difference between the defined distribution quantiles.

This procedure can be used to estimate delay variation for relevant ingress-to-egress paths within a service provider network. Hypothetical reference paths like those defined in Appendix III of [ITU-T Y.1541] could be used to model internal network topologies, or (more directly) a network provider could determine the topologies of selected paths from routing tables or network

---

[7] The IPTD values characterizing particular networks may be pre-specified (or measured) and stored in routing tables or network management databases, or may be estimated or measured during the flow establishment process.

[8] [ITU-T Y.1541] assumes that variation-sensitive and variation-insensitive flows are handled separately; that packets of variation-sensitive flows are scheduled with non-pre-emptive priority over packets from variation-insensitive flows; and that the scheduling within each of these two categories is FIFO.

management information. A similar procedure could be used to estimate end-to-end IPDV from the IPDV values for a number of tandem networks, analogous to router hops. The entity estimating end-to-end IPDV would need to know 1) an IPDV value for each tandem network and 2) the number of networks that would be interconnected in supporting a proposed end-to-end path. Delay distribution information might also need to be obtained or assumed. Ultimately, the decision entity responsible for admitting or rejecting a requested flow would compare the estimated end-to-end IPDV value with the requested IPDV value. If the requested value was exceeded, the flow would be rejected. Alternative routings could be explored through "crankback" as described earlier.

As in the case of IPTD, there are published specifications that define IP packet delay variation differently than it is defined in [ITU-T Y.1541], and such differences can affect the process of coordinating end-to-end delay variation objectives among networks.

6.2.3.3 IP packet loss ratio

IP packet loss ratio is affected by the same factors that affect IPDV (traffic level, transmission link capacity, and packet size), but also by transmission errors and, in some cases, by queuing disciplines. In limiting situations, IPLR values may also be affected by delay variation (since jitter buffer overflow causes packet loss) and by long IP packet transfer delays (since excessively delayed packets may be counted as lost). However, IPLR is relatively easy to measure in an individual network (e.g., using OAM), and estimating end-to-end IPLR from tandem network values is straightforward if losses in the tandem networks can be assumed to be independent. In this case (and assuming the individual network IPLR values are relatively low), the end-to-end IPLR for a path can be estimated by simply adding those values together. As in the case of IPDV, the responsible entity would need to decide whether to admit or reject a requested flow by comparing the estimated and requested end-to-end IPLR values.

6.3 Further study items

This Recommendation provides a set of specifications for characterizing IP flows requiring assured quality and reliability levels, and identifies some possible decision rules for interpreting and processing that information to ensure that the requested quality and reliability requirements of admitted flows are consistently met.

Further study could include specifying more completely the proposed rules for interpreting and processing shared Tspec, Pspec, and Qspec information, and using shared IP flow specifications in making coordinated decisions about the admission, policing, and assignment of resources to particular offered IP flows. For example, the IPDV aggregation procedure defined in Appendix IV of [ITU-T Y.1541] could be adapted for use in estimating end-to-end IPDV values.

Further study could also consider adding one or more parameters to specify any timing synchronization requirements of IP flows.

Annex A

Definition of IP flow parameters using the continuous-state token bucket algorithm

(This annex forms an integral part of this Recommendation)

Clause A.2 of [ITU-T Y.1221] defines IP flow parameters in terms of a "continuous-state token bucket" algorithm, summarized below.[9]

The continuous-state token bucket has two fixed parameters per IP flow:

– The token bucket rate R (in bytes per second) for the flow.

– The token bucket size B (in bytes) for the flow.

The continuous-state token bucket uses the following variables:

– The token count Tc (in bytes) of the flow.

– LCT is the last conformance time of the flow (in seconds).

Initially (at time ta of the arrival of the first packet of the flow):

– Tc = B

– LCT = ta

At arrival of a packet with size N (bytes) at time ta:

Tc' = Tc + R*(ta – LCT)

If Tc' < N

Then packet is not conforming

Else packet is conforming

Tc = min (Tc', B) – N

The variables Tc and LCT are only modified at packet arrival.

[9] An equivalent "generic byte rate algorithm" is also defined. The continuous-state token bucket has the advantage that it is easily implemented.

Bibliography

[b-ITU-T Q-Sup.51] ITU-T Q-series Recommendations – Supplement 51 (2004), *Signalling Requirements for IP-QoS*.

[b-T1A1/2003-133] T1A1/2003-133[10] (May 2004), *Survey of IP Network QoS Architecture and Protocol Standardization Activities*.

[b-PRQC] PRQC-2005-180, *Towards an NGN QoS Standards Solution: New Information and Related Discussion Issues*, October 2005.

[b-IETF RFC 4804] IETF RFC 4804 (2007), *Aggregation of Resource ReSerVation Protocol (RSVP) Reservations over MPLS TE/DS-TE Tunnels*. <http://www.ietf.org/rfc/rfc4804.txt?number=4804>

[b-3GPP TS 23.107] 3GPP TS 23.107, V7.0.0, June 2007, *Quality of Service (QoS) concept and architecture*. <http://www.3gpp.org/ftp/specs/html-info/23107.htm>

[b-T1A1/2003-075] T1A1/2003-075[10] (February 2004), *Mapping between ITU-T (Y.1541/Y.1221) and 3GPP (TS 23-107) QoS Classes and Traffic Descriptors*. <>

[10] T1 standards are maintained since November 2003 by ATIS.

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, 2009