International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1542

(06/2010)

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

Internet protocol aspects – Quality of service and network performance

---

Framework for achieving end-to-end IP performance objectives

Recommendation ITU-T Y.1542

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
IPTV over NGN	Y.1900–Y.1999
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
Future networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999

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

Recommendation ITU-T Y.1542

Framework for achieving end-to-end IP performance objectives

Summary

Recommendation ITU-T Y.1542 considers various approaches toward achieving end-to-end (UNI-UNI) IP network performance objectives. Detailed examples are provided as to how some approaches might work in practice, including how service providers might handle cases where the aggregated impairments exceed those specified in a requested QoS class (such as those of Recommendation ITU-T Y.1541). The advantages and disadvantages of each approach are summarized.

History

Edition	Recommendation	Approval	Study Group
1.0	ITU-T Y.1542	2006-07-14	12
1.1	ITU-T Y.1542 (2006) Amend.1	2009-03-19	12
2.0	ITU-T Y.1542	2010-06-29	12

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/.

© ITU 2010

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

CONTENTS

Introduction

Compared to networks and systems that are circuit-based, those based on IP pose distinctly different challenges for planning and achieving the end-to-end performance levels necessary to adequately support the wide array of user applications (voice, data, fax, video, etc.). The fundamental quality requirements for these applications are well understood and have not changed as perceived by the user; what has changed is the technology (and associated impairments) in the layers below these applications. The very nature of IP-based routers and terminals, with their queuing methods and de-jitter buffers, respectively, makes realizing good end-to-end performance across multiple network operators a very major challenge for applications with stringent performance requirements.

Fortunately, Recommendations ITU-T Y.1540 and Y.1541 together provide the parameters needed to capture the performance of IP networks, and specify a set of "network QoS" classes with end-to-end objectives specified. It is widely accepted (i.e., beyond the ITU-T) that the network QoS classes of Recommendation ITU-T Y.1541 should be supported by next generation networks, and thus by networks evolving into NGNs.

Thus, while there is general agreement that the IP network QoS classes of ITU-T Y.1541 are what should be achieved, what is missing is the methodology for satisfying the end-to-end objectives over paths involving multiple network operators, and in some cases, unusual topologies and distances. The guidance provided here is intended to accelerate the planning, deployment and management of networks and systems that can interoperate with a clear goal of supporting the end-to-end performance objectives detailed in Recommendation ITU-T Y.1541.

Regardless of the approach, there is no guarantee that the end-to-end objectives can be met for a highly congested path through a complex network topology and/or over extremely long distances. However, the guidance provided in this Recommendation should facilitate network design and operation capable of nearly always meeting the desired levels of performance.

Recommendation ITU-T Y.1542

Framework for achieving end-to-end IP performance objectives

1 Scope

This Recommendation includes a broad consideration of approaches toward achieving end-to-end IP performance objectives on as many UNI-UNI paths as possible, including some detailed examples of how some approaches might work in practice. Examples include how service providers might handle cases where the aggregated impairments exceed those specified in a requested QoS class (such as those of [ITU-T Y.1541]).

The advantages and disadvantages of each approach are evaluated to the extent currently possible.

For the purposes of this Recommendation, dynamic inter-AS routing using border gateway protocol (BGP) is assumed as per current practices.

Other approaches for achieving end-to-end IP performance objectives have been mentioned, such as a "Costed Bids Method" and "Bid Discovery Using a Global Registry". As these methods differ fundamentally from those treated here, with significantly different implications for deployment, these and other possible approaches are for further study.

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 G.826]	Recommendation ITU-T G.826 (2002), *End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections*.
[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*.

3 Terms and definitions

This Recommendation defines the following terms:

3.1 access segment: The network segment from the customer interface (UNI) to the interface on the customer side of the first gateway router. It is understood that the access segment may be managed by different entities. In the case where network providers deploy and manage their network devices in user segments (e.g., routers, enterprise LAN, home network), these network devices should be included in "UNI-to-UNI".

3.2 allocation: Formulaic division or assignment of a performance impairment objective among segments.

3.3 apportionment: Method of portioning a performance impairment objective among segments.

3.4 total transit segment: The segment between gateway routers, including the gateway routers themselves. The network segment may include interior routers with various roles.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations:

AS	Autonomous System
BGP	Border Gateway Protocol
DV	Delay Variation
ER	Edge Router
GW	GateWay router
IPDV	IP Packet Delay Variation
IPLR	IP Packet Loss Ratio
IPTD	IP Packet Transfer Delay
LAN	Local Area Network
NSIS	Next Step In Signalling
RSVP	Resource ReserVation Protocol
UNI	User-Network-Interface

5 Problem statement and consideration of approaches

How can QoS classes (e.g., network performance according to [ITU-T Y.1541]) be assured for users? The fundamental challenges to achieving end-to-end QoS are present when:

- multiple network providers are necessary to complete the path;
- the number of networks in the path will vary request by request;
- distances between users is generally unknown;
- the impairment level of any given network segment is highly variable;
- it is desirable to estimate the actual performance levels achieved on a path;
- the operator must be able to say if the requested performance can be met or not; and
- the process must eventually be automated.

Also, solving the problem of delivering UNI-UNI IP QoS through the standards process will require development and agreement on many new tools and capabilities, and the extent of new work required should be assessed for each candidate solution. These too are challenges.

There are two basic approaches to solve this problem. One involves *allocating* performance to a limited number of network segments, which allows operators to contribute known levels of impairments per segment, but restricts the number of operators that can participate in the path. (If a given segment does not need to use all of its allocation, the balance is wasted.) The other approach is *impairment accumulation*, which allows any number of operators to participate in a path. On the surface this may appear too relaxed, but assuming operators in a competitive environment will actively manage and improve performance, the likelihood of the concatenated segments satisfying the end-to-end objectives can be predictably good.

Figure 1 provides the basis for expressing the problem and how various solutions might work.

Generally, the approaches that could be taken in allocating total impairment targets among network segments can be characterized by the amount of information shared among segments. Each approach has advantages and disadvantages. We describe them here with simple examples. (Detailed examples of the various approaches are appended to this Recommendation.)

For all allocation-based approaches, a "top-down" or "bottom-up" method could be applied. That is, percentages of the aggregated target (top-down) or fixed/negotiated values for impairments (bottom-up) may be allocated for each segment. A hybrid of these methods, with percentages for some segments and fixed/negotiated values for others, could also be used.

Brief examples for some approaches as applied to Figure 1 below are provided. Note that the provider which sends traffic over a peering link is assumed to be responsible for that link's performance and its impairments must be included in the segment total.

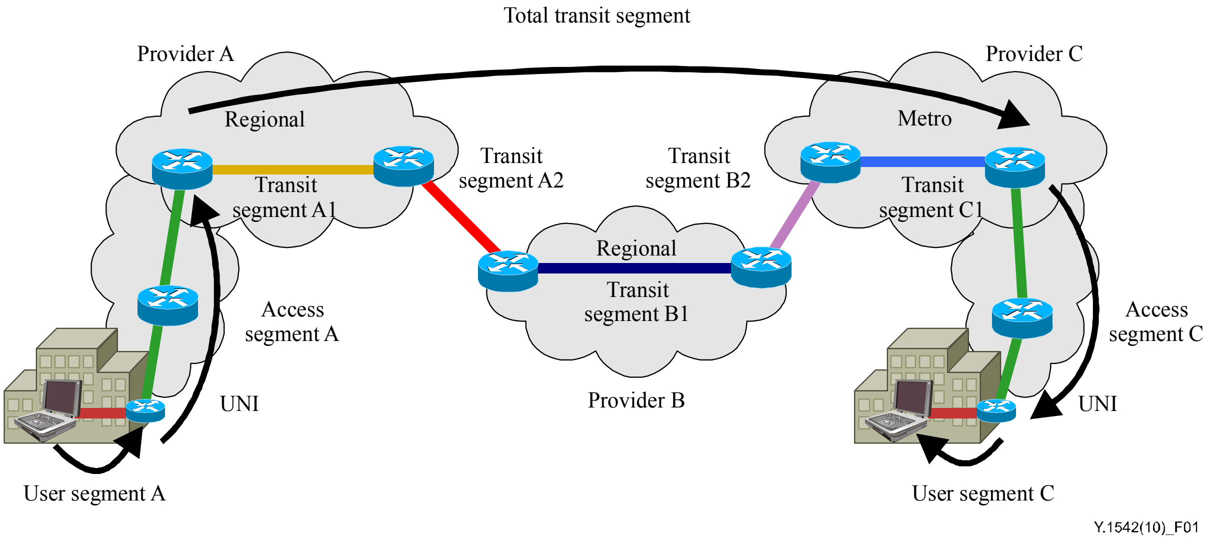

Figure 1 – Example topology for impairment allocation

For some approaches, transit segment distances are required to estimate distance dependence metrics such as mean delay. Ground level distance between any two (user) points may be readily estimated despite the traffic's signal being carried over varying altitude, the non-spherical shape of the earth, etc. Distance-inefficient routing over multiple segments may result in traffic travelling over a significantly longer distance than expected between two user points. The approaches to accounting for these inefficiencies can also be characterized by the amount of information shared among segments.

Regardless of the approach, there is no guarantee that the desired end-to-end objectives will be met. Any approach can fail to achieve a specific set of objectives on a highly congested path through a complex network topology and/or over extremely long distances. Therefore, a key attribute is how each approach accommodates such failures and whether the failure response is acceptable to users.

NOTE – The phrase "end-to-end" has a different meaning in Recommendations concerning user QoS classes, where end-to-end means, for example, from mouth to ear in voice quality Recommendations. Within the context of this Recommendation, end-to-end is to be understood as from "UNI-to-UNI", as described in [ITU-T Y.1541].

5.1 Static approaches

5.1.1 Static divisor approach

This approach "divides" the UNI-to-UNI path into a fixed number of segments and budgets the impairments such that the total objective is met in principle. It requires that individual segments have knowledge of the distance and traffic characteristics between the edges of their domains, as these properties of the segment affect the resulting allocations. For example, the delay budget allocated to a network segment depends on whether it is access or transit, and whether the transit distance is metro or regional. Similarly, packet loss and delay variation will have to be allocated according to whether the segment is access or transit, as the traffic aspects can differ significantly.

Appendix I provides examples of this approach.

An important aspect of the static allocation is its dependence on the number of providers, as the allocation has to be done accordingly. This can result in undershooting or overshooting the objective because any actual path may traverse a different number of network segments from what was assumed to be the case in the allocation scheme.

Service providers may reallocate impairment targets among the segments under their control.

5.1.2 Static reference allocation approach

This approach requires that individual segments have knowledge of the distance between the edges of their domains. In this approach, Appendix III of [ITU-T Y.1541] example router delay values and the ITU-T G.826 air-to-route distance conversion are used, which accounts for major delay contributions of each provider. This approach calculates the delay margin and allocates a proportion of that margin to each provider, as follows:

Step 1: Calculate propagation delay for each provider attributable to distance.

Step 2: Calculate the processing and queuing delays of each provider using Appendix III of [ITU-T Y.1541] example values.

Step 3: Calculate the delay margin by subtracting the sum of the providers' propagation delays from the ITU-T Y.1541 network QoS class objectives.

Step 4: Calculate the prorated fraction of processing and queuing delays of each provider to the sum of the processing and queuing delays of all providers.

Step 5: For each provider, the allocated delay is equal to its propagation delay plus that provider's prorated fraction of the delay margin.

See Appendix II for static reference allocation model, values and detailed examples.

Note that the scope of this approach extends between UNIs, and excludes the user segments.

5.1.3 Weighted segment approach

This approach allocates a significant proportion of the impairment budget to each access segment, with each core segment having a lesser fixed budget. This approach also allocates a fixed budget for core network segments, irrespective of the number of core network segments in any resulting services. This core network segment budget can be concatenated within bounds to create end-to-end services that have a high probability of still being within the overall end-to-end class targets.

An additional allowance for propagation delay for long network segments is also possible. In this case, core segments must have knowledge of the distance between their edges when the total distance between the edges of any core network segment exceeds an air path distance of a given distance, e.g., 1200 km.

As this approach runs the risk of confounding IPTD and IPDV (because the weighted proportions may not match), it is for further study.

5.2 Pseudo-static approach

In a "pseudo-static" approach, each provider would have knowledge of how many providers are present in the traffic path and allocate among each other without wasting part of the impairment budget. Service providers may reallocate their impairment target among the segments under their control. This approach also requires further study.

5.3 Signalled approaches

Given the flexibility of a signalled approach, multiple examples are given to investigate its flexibility. For a signalled approach, the use of resource management and signalling for the purposes of impairment apportionment is assumed.

5.3.1 Negotiated allocation approach

In some situations, for the static and pseudo-static approaches, certain segments will not meet their formulaic targets, while others will, and thus have an "impairment budget" excess.

Access providers which require less than the normal allocation of impairments may be able to have the unneeded part of their allocation allocated instead to a transit or user link. They may reallocate their impairment allocation within their control or negotiate the unneeded part to other segments.

A transit provider may negotiate to use the unneeded part, or to make its unneeded part negotiable for other segments.

Similarly, in a managed user segment, the user may require a greater or lesser impairment allocation based upon access sub-type, e.g., by broad category (enterprise, home, wireless) or specific capability ([b-IEEE 802.11g], 100 Mbit/s Ethernet) and negotiate with their access provider.

Starting with initial segment impairments targets, based possibly upon the static and pseudo-static allocations in this Recommendation; the distributed use of negotiation among providers allows the opportunity to negotiate for any "impairment budget" excesses, and to advertise to multiple interested parties if they can provide a network service that is within their collective impairment budget.

First, assume that an extension to BGP can provide for multiple advertisements to a prefix, depending upon whether particular network classes are supported along a path. Then starting with the provider closest to the destination, the advertisement is conditionally transitive depending upon whether a collaborative impairment target for the network class is met.

Referring to Figure 1, provider C advertises a real-time network class to provider B indicating that provider C can meet their impairment budget for that class. If provider B can meet their impairment budget, then they will advertise the path to provider A.

However, if provider B cannot meet the impairment target that has been set for them, they may negotiate with provider C for the right to use any excess available impairment of provider C. Similarly, provider A may in a cascade fashion negotiate with provider B.

Pairwise negotiations between segment owners may occur either by signalling or manually, and are assumed to change infrequently.

This approach appears to support multiple connections among providers, where provider's BGP advertising policies and aggregation would influence the solution.

5.3.2 Ranged allocation approach

For the purpose of meeting the overall impairment target as well as optimizing resource utilization, this clause provides another signalled approach example, namely ranged allocation.

In this approach, the range between the minimum and maximum of the allocated impairment budget for every segment along the data path is negotiated and calculated out by the use of resource management and signalling among the segments. Any value within each segment impairment budget range, when added with those of other segments, can meet the total impairment budget target for the whole data path. Thus, every segment itself can choose an appropriate value within its allocated budget range under the consideration of optimizing its resource utilization.

The key points of ranged allocation are that, firstly, the minimum of the allocated impairment budget for every segment along the data path is negotiated out; secondly, the remnant impairment along the whole data path equals the total impairment budget target minus the amount of the allocated segment impairment budget minimums; thirdly, the ratio of the range minimum to maximum equals the total minimum allocated impairment divided by the total impairment budget target along the whole data path, and finally, the maximum of the allocated impairment budget for every segment along the data path is calculated out.

In this example, three network providers are interconnected (providers A, B and C) as shown in Figure 1. Provider A and provider C have access network segments. The following steps outline the process:

1) The user determines the desired UNI-UNI performance objectives, and requests provider A for the total impairment target (e.g., IPTD).
2) Provider A:
    a) calculates the remnant impairments by subtracting its own minimum impairments for the packets from the UNI-UNI target, and inserts its AS number and its own minimum impairments to the request message;
    b) sends the request message containing the remnant impairments, the UNI-UNI target and previous segment's minimum impairments list to the next provider B along the data path.
3) Provider B calculates the new remnant impairment as provider A does, and then sends the new request to its next provider along the data path in turn.
4) At last, provider C, which is the last provider along the data path:
    a) calculates the new remnant impairment;
    b) calculates the total allocated impairment by subtracting the new remnant impairment from the total impairment target;
    c) calculates the allocated impairment proportion by dividing the total allocated impairment by the UNI-UNI target;
    d) calculates its max impairment by dividing its min impairment by the allocated impairment proportion;

       Max impairment = min impairment/allocated impairment proportion

       and chooses the appropriate impairment between the min and the max;
    e) sends the allocated proportion back to the previous provider B.
5) Provider B chooses the appropriate impairment as provider C and then sends the allocated proportion to its previous provider A.
6) At last, provider A can also choose its appropriate impairment, and then sends the success message to the user.

7) If a provider (e.g., provider C) along the data path detects that the path does not meet the requested objectives, because its min impairment is greater than the remnant impairment received from its previous provider, it sends the failure message back to its previous provider (e.g., provider B). The previous provider sends the failure message back to its previous provider in turn. At last, the first provider (e.g., provider A) negotiates with the user about an alternative service class or relaxed objectives to be offered. The alternative opportunity for negotiation is the path negotiation that an alternative path might be sought, requiring a routing change based upon these minimum impairments offered by these providers.

5.4 Impairment accumulation approach

Accumulation approaches are defined here as those that include requests of what performance level each provider can offer, followed by decisions based on the calculated estimate of UNI-UNI performance. The requesting provider may be the customer-facing provider only (hub and spoke) or include all the providers along a path (cascade). The responding provider may be one or more providers in the path or their proxy.

In this approach:

1) The requesting provider:
   a) Determines the path that packets will follow (e.g., based on inter-domain routing information);
   b) Requests from each provider the performance level that they will commit to for each segment of the path for packets identified by source/destination pair, possibly using an on-path QoS signalling protocol.
2) A responding provider makes a commitment (which might have been derived by one of the other approaches described in this Recommendation) which is good for the session (unless modified).
3) The requesting provider:
   a) Combines the segment performance levels (according to rules that are defined in [ITU-T Y.1541]); and
   b) Compares the estimated performance with the desired UNI-to-UNI QoS Class/Objectives.

If the path does not meet the requested objectives, there are two opportunities for negotiation:

1) Path negotiation: An alternative path might be sought, requiring a routing change based upon parallel or subsequent request of other providers.
2) User negotiation: An alternative service class or relaxed objectives could be offered to the user. (Note that, in many cases, the estimation process will result in a total that is slightly beyond a particular class's objectives but considerably better than the target performance level of a different service class.)

Advantages of this approach are:

- No formulaic impairment allocation agreements are required to use this approach.
- No explicit knowledge of distance is required.
- It is completely consistent with the vision of achieving UNI-UNI performance objectives (ITU-T Y.1541 network QoS classes) with signalling protocols that automate the process of reserving bandwidth and accumulating impairment levels. [b-ITU-T Q-Sup.51] (on IP QoS signalling) codifies one set of requirements for this task, but clear parallels may be found in Integrated Services/RSVP and in the Next Steps in Signalling (NSIS) Qspec template.

- As no allocations are performed, the fact that it is not known how to decompose certain parameters (especially IPDV) is not a problem.

Disadvantages of this approach are:

- Users' segment impairments are not taken into account.
- If the initial process fails, multiple passes of request/estimation cycle may be required.
- Requires customer or customer proxy (rules driven agent or equivalent) involvement.
- Commitments for each network segment must be pre-calculated taking distance into account.
- Commitments for "all time" may need to be over-conservative for low-utilization circumstances.

A detailed example of the impairment accumulation approach is given in Appendix III.

Network operators who implement the impairment accumulation approach usually derive performance design incentives from general guidance, rather than the numerical design objectives that are part of other approaches. Appendix IV gives detailed guidance for both the design and daily operation phases of a provider's network life-cycle.

6 Advantages and disadvantages of the approaches considered

Tables 1 and 2 give the advantages and disadvantages of the allocation and the accumulation approaches considered, respectively.

Table 1 – Summary of performance impairment apportionment approaches

Approach	Description	Information required at each segment	Advantages	Disadvantages
Static (simplest/least flexible) – No information is required to be shared among segments.	A fixed number of segments is assumed. Allocation is formulaic among user, access, and transit segments.	Information required is: a) type of link; b) traffic service class; and c) transit distance.	No information is required to be shared among segments. Providers may reallocate among their user, access and transit segments.	Over-engineered when number of segments is less than number assumed. Paths with more than assumed number of segments are not covered. No negotiation. Works best with static routing, which is not common anymore.
Pseudo-static – Some information is required to be shared among segments.	The exact number of transit providers is determined. Impairment alloc. is formulaic among user, access, and transit segments.	Information required is: a) type of link; b) traffic service class; and c) transit distance; d) destination address; e) BGP tables.	Impairment allocation may be efficient and scalable.	Signalling among providers required to determine the number of transit providers in each traffic path, e.g., from BGP, number of ASs. Negotiation not supported. Works best with static routing.

Table 1 – Summary of performance impairment apportionment approaches

Approach	Description	Information required at each segment	Advantages	Disadvantages
Signalled (least simple but most flexible) – Some info. is required to be shared among segments and possibly with users.	Exact number and sub-type of all segments may be known, e.g., if user segment is wireless or wireline. Impairment apportionment may be negotiated among segments and with users.	Information required is: a) type of link; b) traffic service class; c) destination address; d) BGP tables, or other means to determine path(s) at operator-level; e) network edge-edge performance information. Additional information required may include: f) transit distance.	Negotiation is supported allowing highly flexible apportionment among segments. Transit distance may not be required. Able to address cases where objective cannot be met by user for relaxed objective. Consistent with proposed direction of methods automated by QoS signalling (e.g., RSVP/NSIS).	Signalling required to apportion amounts to each segment, and to negotiate with user when the requested objective cannot be met. Performance and routing info. must be signalled to determine identities of transit providers in each path (e.g., from BGP, number of ASs) and their performance. However, there are alternative ways to determine path, and some providers give performance info in real time.
NOTE – All allocation approaches suffer from not being able to decompose IP delay variation according to agreed-upon methods (the technique for combining IP delay variation was agreed to only in 2005).				

Table 2 – Approach to impairment apportionment based on accumulation

Approach	Description	Information required at each segment	Advantages	Disadvantages
Impairment accumulation, where some information is required to be shared among segments.	The path through various network operator domains is determined. Impairment levels and other parameters may be requested for various network segments or their proxy, combined and compared with desired objectives. If not met, then path or user negotiation takes place, or the request is rejected.	Information required is: a) traffic service class; b) destination address (always known); c) BGP tables, or other means to determine path at the operator-level; d) network edge-edge performance.	No allocations required, so no process to achieve agreements. Impairment accumulation is simple and scalable. No distance and route-to-air conversion factors required. Negotiation is supported. Consistent with methods automated by QoS signalling (RSVP/NSIS). Agreement on how to decompose IPDV not needed.	Performance and routing information must be exchanged among providers to determine identities of providers in each traffic path (e.g., from BGP, number of ASs) and their performance. However, there are alternative ways to determine path, and many providers publish performance info in real time. Cannot guarantee that objectives will be met (true for all approaches).

7 Summary of approaches according to the problem statement challenges

The problem statement in clause 5 lists the challenges to delivering UNI-UNI QoS and recognizes further challenges for solutions in the standards development process. With the various approaches now described, it is possible to compare them according to these challenges. Table 3 provides a comparison for the QoS delivery challenges. (As noted in the Scope, dynamic inter-AS routing using BGP is assumed.)

Table 3 – Summary of approaches according to problem statement challenges

	Multiple networks	Variable number of nets	Works w/ unknown distances	Variable impairment levels	Actual perform. estimated	Response to request	Automated
Static	Yes	No	No	No	No	No	May be
Pseudo-static	Yes	Somewhat	Somewhat	May be	No	No	May be
Signalled apportion.	Yes	Yes	Somewhat	Somewhat	May be	Yes	Yes
Impairm. accumul.	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Each approach will require development and agreement on new tools and capabilities, representing challenges to the standards process. Table 4 summarizes the various aspects of new development required for each approach.

Table 4 – Comparison of approaches in terms of standards development challenges

	Allocation of UNI-UNI objectives required?	**Decomposition methods required?**	**Segment weighting factors required?**	**Signalling protocol needed?**	**Segment measurement collection support?**	**Composition methods required?**
Static	Yes	Yes	No	No	Yes	No
Pseudo-static	Yes	Yes	Yes	No	Yes	No
Signalled apportionment	Yes	Yes	No	Yes	Yes	May be. If yes, already developed
Impairment accumulation	No	No	No	Yes, but optional at small scale	Yes	Yes, already developed

Appendix I

Detailed example of a static divisor approach

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

To gain an appreciation of how a static allocation scheme may look, assume a maximum of three transit providers in a path interconnecting user segments.

User segment impairment budgets are dependent upon the nature and size of the enterprise, home, etc., but as a simplifying approximation, a 1% static allocation is made to user segments for packet loss and delay variation (for delay, 2 ms is allocated to the user segment).

NOTE – User segment is not to be included in the portion where allocation of IP network performance objectives is implemented.

The following impairment allocations apply to the user, access and transit segments (independent of application). The percentages are of total site-to-site impairment targets for each service class.

Table I.1 – Allocations for user, access and transit portions

Parameter	User segments (each)	Access segments (each)	Transit segment (total)
Packet loss	1%	47.5%	5%
Delay variation	1%	40%	40%
Mean delay	2 ms	30 ms	Distance-dependent (see text below)

A budget for each parameter needs to be assigned for each of the three provider networks which could comprise the total transit segment. For packet loss, it is 33% and for delay variation it is 40%. For delay, the budget for each transit provider is based on geographic distance. Each is allowed up to 33% of the appropriate transit delay listed in Table I.2, depending on the category of the transit segment.

Table I.2 – Total transit delay by distance

Categories	Distance (km)	Shortest path propagation delay (ms)	Total transit delay (ms)
Metro	< 100	0.56	5
Regional	< 1000	5.6	15
Continental	< 5000	27.8	45
International	< 20 000	111.2	140
NOTE – Total Transit Delay = Shortest path propagation delay + Allowance for inefficient topology + Allowance for queuing delays.			

The route length calculation used here is based on [ITU-T G.826], only for the distances listed.

Appendix II

Detailed example of the static reference allocation approach

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

The static reference allocation approach uses the following steps for determining IP delay time.

i) Establish interconnection network section model (e.g., ITU-T Y.1541 UNI-to-UNI reference path).

ii) Establish network element model for each network section (see Figure II.1).

iii) Calculate propagation delay of each network section distance (use [ITU-T G.826] route-to-air mileage scaling factors).

iv) Calculate each network section's processing and queuing delay using network element models and per-element delay times. Table III.1 of [ITU-T Y.1541] gives this calculation.

v) Subtract the sum of propagation delays (step iii) above from the ITU-T Y.1541 delay objective. This value is the delay margin.

vi) Divide the processing and queuing delay of each network section (step iv) by the sum of all the section's processing and queuing delays. This gives the prorated fraction of the processing and queuing delay that is assigned to each section. Multiply this fraction by the total delay margin (step v) to get the prorated delay margin for each section.

vii) The allocated delay time of each network section is the sum of its propagation delay (step iii) and its prorated fraction of the delay margin (step vi).

Figure II.1 is an example of each network element model, and Table II.1 provides typical delay contribution by router role. These models and values should be consistent with [ITU-T Y.1541].

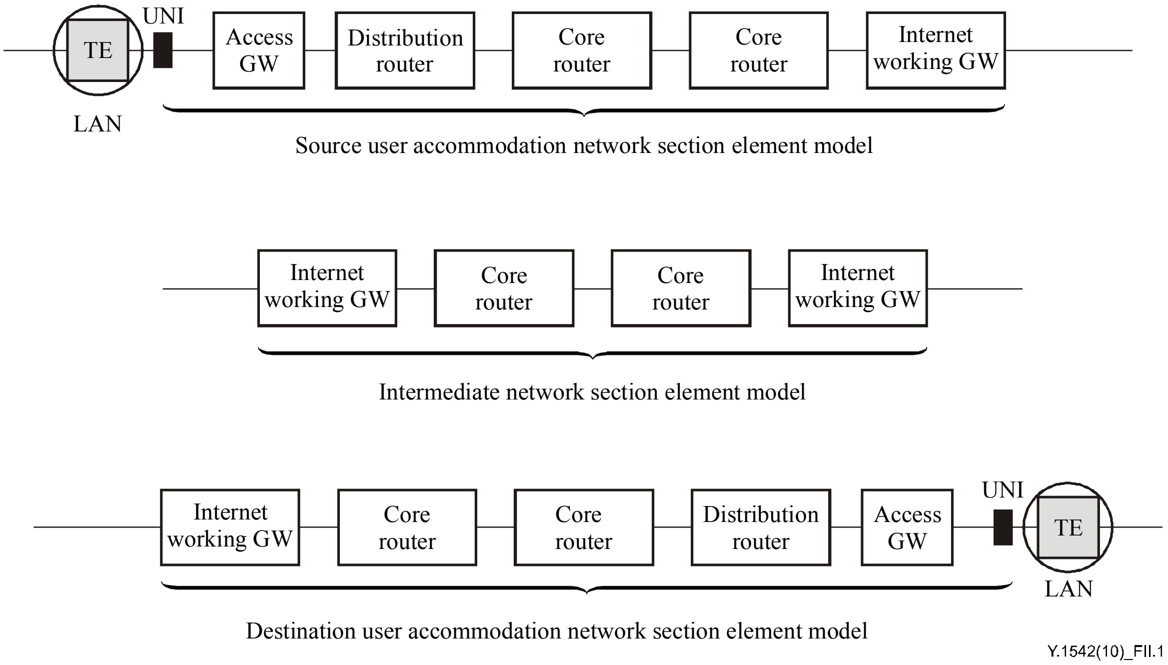

Figure II.1 – Example of network element model for each network section

Table II.1 – Example of typical delay contribution by router role (Table III.1 of [ITU-T Y.1541])

Role	Average total delay (sum of queuing and processing)	Delay variation
Access gateway	10 ms	16 ms
Internetworking gateway	3 ms	3 ms
Distribution	3 ms	3 ms
Core	2 ms	3 ms

Detailed Example

In this example, three network providers interconnect (Figure II.2) with these assumptions:

a) Three network providers are connected (providers A, B and C).

b) Network providers A and C have access network which admit user direct.

c) Air route distance across provider A's network is 1500 km; 4000 km across provider B; and 900 km across provider C.

d) Non-IP networks are not needed in UNI to UNI.

e) UNI-to-UNI delay time-limit is 100 ms (class 0, 1 delay time of [ITU-T Y.1541]).

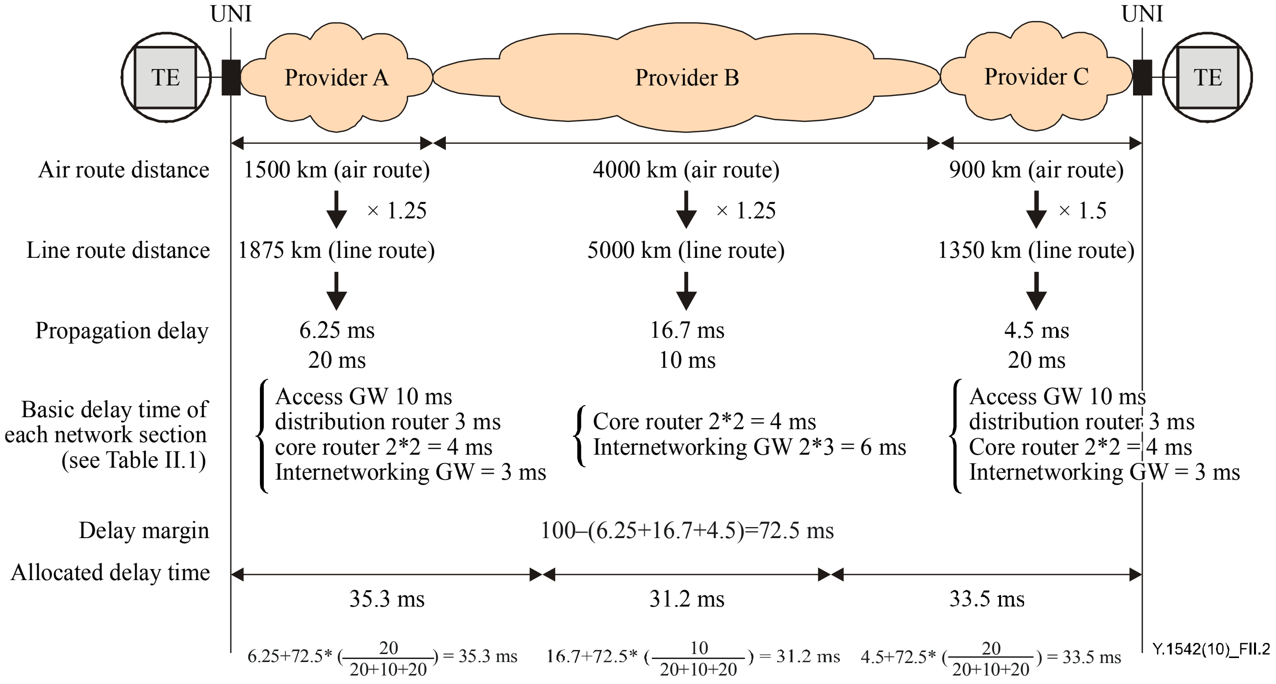

Figure II.2 – Static reference allocation example

Appendix III

Detailed example of impairment accumulation approach

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

This appendix describes a process to accumulate network performance levels along an end-to-end path and compare the combined performance estimate with specified objectives, consistent with procedures envisioned with quality of service signalling protocols such as those that meet the requirements of [b-ITU-T Q-Sup.51] on IP QoS. We do not address capacity reservation aspects here, or subscription, authorization, and accounting, though they are critical aspects of a premium service offering as well.

The following steps outline the process at high level:

1) Determine the desired UNI-UNI performance objectives and any acceptable alternatives (e.g., the desired ITU-T Y.1541 network QoS class).
2) Determine the user-network interfaces (UNI) and network-network interfaces (NNI) that appear in the end-to-end path.
3) Determine the performance of each segment of the path (each operator domain from UNI to NNI, NNI to NNI, etc.) for each parameter with an end-to-end objective. If there is uncertainty which NNI will be traversed from among several possibilities, then separate calculations can take each one into account (although such instances should be minimized, especially where the performance differences are significant).
4) Combine the segment performance levels according to composition relationships.
5) Determine if the combined performance estimate meets the desired objectives.
6) If the objectives were not achieved, then take one or more of the following actions:
    a) user negotiation: An alternative QoS class or modified objectives could be offered to the user;
    b) path negotiation: An alternative path may be assessed based upon parallel or subsequent request of other providers, and possibly requiring a routing change.

There are only three pieces of information exchanged among the provider networks:

- End-to-end objectives.
- Path UNI and NNI list, including operator identifications.
- Performance of each path segment between specific edge interfaces.

Assuming that this process will be automated (with on-path signalling), then the ingress edge router at each UNI/NNI may play the primary role for each autonomous system (AS) on the source-destination path (step 3 above). When a QoS signalling request enters an AS, the following operations might take place:

1) The edge router identifies the packet as one requiring exception processing (possibly after inspecting the protocol number in the IP header), and sends the packet to the central processor (the packet has not been processed previously in this AS).
2) The router processor inspects the destination address and determines the BGP next hop (or other equivalent egress point) for this AS. This provides the local loopback addresses of the ingress and egress edge routers and NI.

3) The AS ingress and egress points can be mapped to a matrix of performance measurements (likely stored elsewhere on a server known to the router, so the router might encapsulate the signalling packet with the ingress/egress points into one packet and forward it to the measurement server). The performance matrix would be updated frequently as new loss, delay and delay variation measurements become available, and the most recent valid measurements are always used.

4) The signalling packet is augmented with the AS number and the edge-edge performance measurements (again, the measurement server might perform this function, and it may encapsulate the signalling packet in an IP header to send it back to the edge router).

5) The edge router (extracts and) forwards the augmented signalling packet along the normal path.

6) Interior routers in the same AS would inspect the packet, find that their AS is already listed, and take no action on the performance fields.

Note that this is a process using operator domain (AS) performance as the building blocks. Other processes use network elements and the links between them as the building blocks, such as those envisioned for integrated services supported by RSVP signalling. It may be possible to perform capacity/traffic management on an element-by-element basis, while managing performance aspects on a domain basis as long as sufficient capacity is available on the path through the domain.

Example calculations

Figure III.1 gives an example path with three network segments.

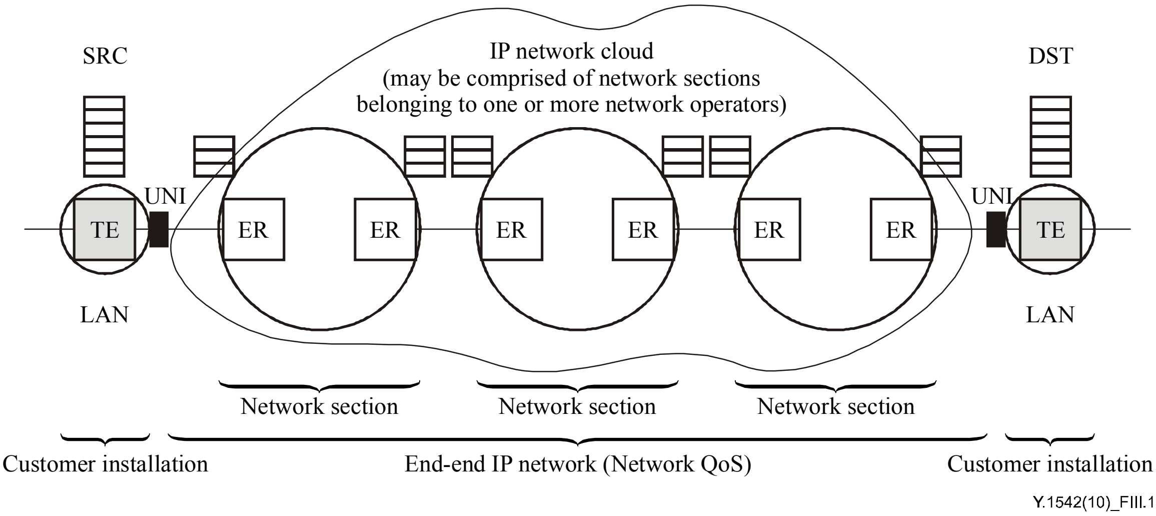

Figure III.1 – Impairment accumulation example UNI-UNI path

In this example, the user has requested QoS class 0, and the network sections above are determined to compose the UNI-UNI path. The next step is to request the performance (impairment) levels from each segment of the path. The results are shown in Table III.1:

Table III.1 – Example of accumulating and estimating UNI-UNI performance

	Requested	**Network 1**	**Network 2**	**Network 3**	**Estimated UNI-UNI**
QoS Class	Class 0				Class 0
Mean transfer delay (IPTD)	100 ms	22.4 ms	10.6 ms	32.4 ms	65.4 ms
99.9% – min delay var. (IPDV)	50 ms	25 ms	2 ms	25 ms	47.5 ms
Minimum transfer delay	–	10 ms	10 ms	20 ms	–
Variance of transfer delay	–	52.4 ms	0.23 ms	55.1 ms	–
Loss (IPLR)	10^{-3}	10^{-4}	10^{-4}	10^{-4}	3×10^{-4}
Errored packets (IPER)	10^{-4}	3×10^{-5}	3×10^{-5}	3×10^{-5}	9×10^{-5}

The performance levels from Networks 1 through 3 are combined according to the composition relationships in clause 8 of [ITU-T Y.1541] to produce the estimated UNI-UNI performance.

In this first example, the Class 0 performance objectives will be achieved on the path, so the response to the user confirms the request for Class 0 and may optionally report the estimated UNI-UNI values for this path.

We illustrate the steps when the path does not meet the desired objectives in a second example below. Again, the user has requested QoS Class 0, and the three network sections above are determined to compose the UNI-UNI path. The next step is to solicit the performance (impairment) levels from each segment of the path. The results are shown in Table III.2.

Table III.2 – Example of accumulating and estimating UNI-UNI performance

	Requested	**Network 1**	**Network 2**	**Network 3**	**Estimated UNI-UNI**
QoS Class	Class 0				Class 1
Mean transfer delay (IPTD)	100 ms	42.4 ms	20.6 ms	42.4 ms	105.4 ms
99.9% – min delay var. (IPDV)	50 ms	25 ms	2 ms	25 ms	47.5 ms
Minimum transfer delay	–	30 ms	20 ms	30 ms	–
Variance of transfer delay	–	52.4 ms	0.23 ms	55.1 ms	–
Loss (IPLR)	10^{-3}	10^{-4}	10^{-4}	10^{-4}	3×10^{-4}
Errored packets (IPER)	10^{-4}	3×10^{-5}	3×10^{-5}	3×10^{-5}	9×10^{-5}

In this example, the estimated delay exceeds the limit for Class 0. The process allows two alternatives when a failure occurs.

User negotiation requires the response to reject the Class 0 request, but it may offer Class 1 with a commitment to meet 105.4 ms IPTD, making the Class 1 response much more palatable.

Path negotiation requires the soliciting network operator to seek alternative paths between the UNI source and destination. The process returns to solicitations of performance levels for the new path segments and repeating the calculations to estimate UNI-UNI performance.

Appendix IV

Performance guidance for providers

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

IV.1 Qualitative guidance statements

Composition of the end-to-end objectives highlights the performance areas to be emphasized. When working to achieve ITU-T Y.1541 classes with unspecified delay variation, different techniques are used to achieve those objectives than those that might be used for Class 0 or 1 (with delay variation limits).

Performance guidance need not be quantitative (e.g., X ms/km of delay allowed) to be useful. General guidance like: "Minimize delay by keeping the route-to-air distance ratio as small as feasible" should achieve nearly the same result. Economic factors cannot be ignored in this exercise. These factors usually set the point of diminishing returns when seeking to improve performance in any area.

Other simple, but meaningful, statements of performance guidance are:

"Minimize delay by providing sufficient link capacity to keep queue occupation low."

"Minimize delay variation by giving queue scheduler priority to traffic that is sensitive to variation, or by grooming or shaping such traffic."

"Minimize packet loss by planning sufficient link capacity to avoid queue tail-drops."

It is expected that additional guidance statements will be developed, thus this set is just the start.

IV.2 Circumstances when guidance is useful

There are several phases in the life of a network, such as when new construction or expansion is in progress. A stable phase would be where the network's geographic assets are fixed, and customers are connected to the closest existing node. Capacity may be added in any phase. Adding links from network locations to reach remote customer sites is simply the expected growth under normal/stable operation, unless new network nodes (points-of presence or concentration) are constructed. During construction or expansion, Table IV.1 indicates how guidance may influence network design.

Table IV.1 – Areas for action given qualitative design guidance

Performance enhancement area	Design aspects		
Delay	Location of nodes	Capacity (avoid queuing)	
Delay variation	Capacity (avoid queuing)	QoS mechanism provisioning	
Loss ratio	Failure protection Restoration time	Capacity (avoid queue overflow = drops)	Transport facility types (bit errors cause loss)

During stable operation, these same three forms of guidance translate into:

- monitoring and maintaining the network according to design levels plus some tolerance;
- managing load to avoid bottlenecks or congestion;
- adding capacity when necessary.

In a competitive environment, network operators are under pressure to follow these guidelines.

Appendix V

Additional considerations for achieving end-to-end performance objectives in an NGN environment

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

V.1 Introduction

This appendix provides some considerations that need to be addressed when the NGN environment is used. Even though NGN is based on the same IP technology as that of non-NGN IP networks, there are some technological differences compared with generic IP networks. In particular, with respect to QoS aspects of the NGN, each segment is appropriately managed, and the capability of providing QoS services is supplied as one of the key features of NGN. With regard to services, session-based services such as IP telephony are particularly important in the NGN because they are often considered to be fundamental services to be provided in a public network. Thus, some countries or regions have special interest in specifying the QoS performance of the NGN, [b-JDTEL A35], so that speech services with PSTN-like quality can be supported. As there are end-to-end QoS requirements to the NGN [b-ITU-T Y.2001], the QoS assignment for each NGN segment of interconnected NGNs needs to be considered (Figure V.1). This appendix describes an overview of the relationship between this Recommendation and other QoS-related Recommendations, and provides some information about the applicability of this Recommendation with respect to end-to-end QoS achievement in the interconnected NGN environment.

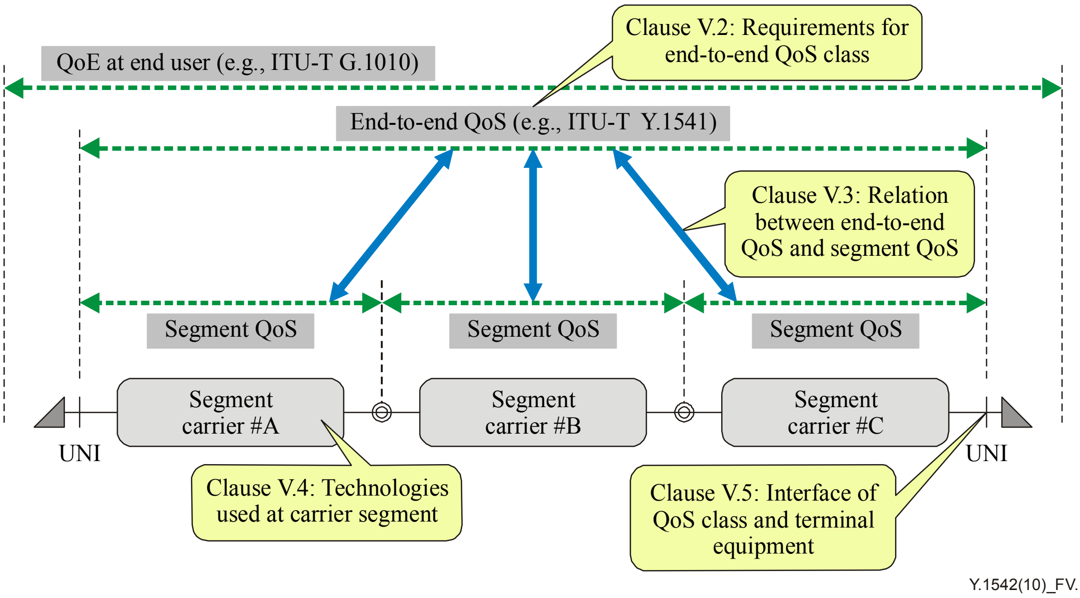

Figure V.1 – End-to-end QoS and NGN segments

V.2 Requirements for end-to-end QoS class

End-to-end QoS classes are defined in [ITU-T Y.1541]. A QoS class consists of a set of network performance parameters and each QoS class is characterized by the objectives of those network performance parameters to fulfil various performance requirements. There are many performance requirements depending on various applications, terminal equipment and users, but negotiation and determination of all parameters and their values respectively are too complicated and not so efficient. Showing too much detail about QoS class in a protocol to end users would be useless since most end users do not know the exact meaning of these parameters. Thus, aggregating parameters into a few service classes is a good way of managing QoS services and the concept of QoS class has been introduced.

One thing worth mentioning here is that the scope of ITU-T Y.1541 QoS classes is end-to-end, not a segment of end-to-end. Therefore, specifying a segment QoS class by referring to the same objectives of the QoS classes in [ITU-T Y.1541] is meaningless, confusing, and should be avoided.

The concept of QoS class can be applied to a part of the end-to-end network in terms of efficiency of QoS-related information handling. It is obvious that end-to-end characteristics have some relation with those of the segments. A segment QoS class should be mapped from an end-to-end QoS class. The mapping should vary depending on relationships among segment carriers that provide each segment. The QoS class specification of each segment should be determined with consideration of other segments along the end-to-end path.

V.3 End-to-end QoS performance objectives and multi-domain network

In a multi-domain network, such as interconnected NGNs, the method of guaranteeing an end-to-end QoS is not clear. The relationship between QoS class and performance requirements between segment carrier and end-to-end path varies depending on issues such as the selection of a neighbouring carrier and the location of the source and the destination. Several models related to these issues are described in this Recommendation. There are two major approaches: one is a top-down approach (such as "static divisor" approach), which allocates a performance budget that is given in advance, and the other is a bottom-up approach (such as the "impairment accumulation" approach), which first considers actual network performance and then compares that performance to performance objectives (Figure V.2).

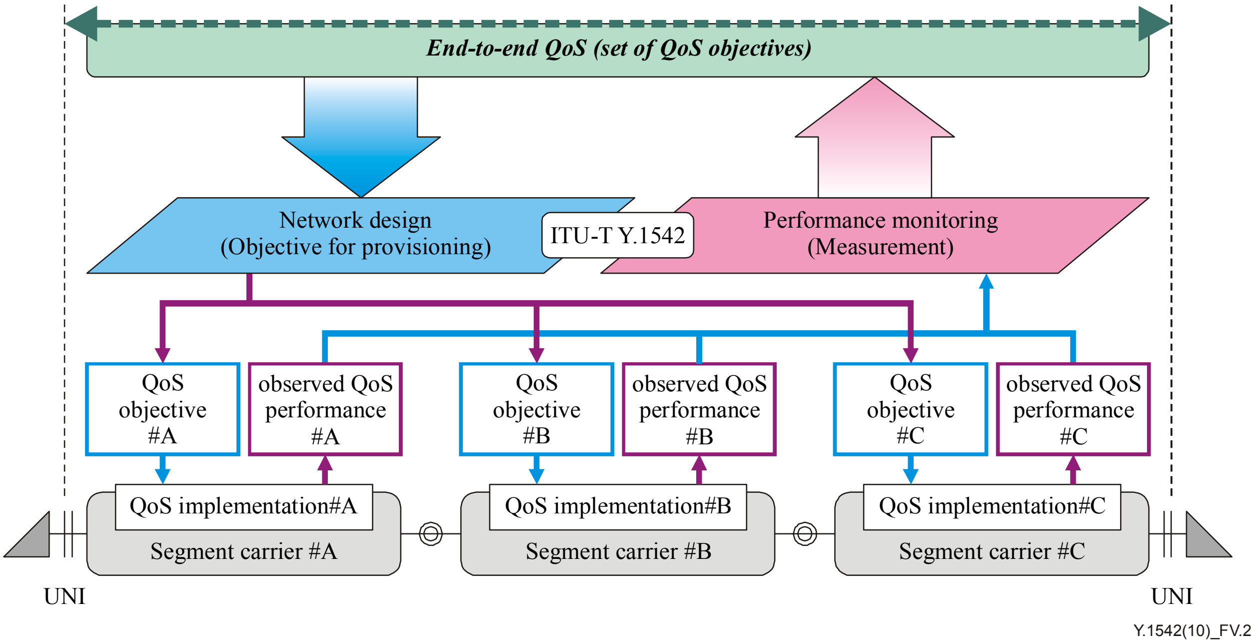

Figure V.2 – End-to-end QoS performance objectives and multi-domain network

Each approach is developed from different background requirements. In situations like NGN interconnections under well-organized regional arrangement within a small or medium size area, the top-down approach with all carriers conforming to the required performance would be suitable. On the other hand, a top-down approach might not work well in a dynamically routed network, in which it is hard to predict the number of networks involved and their impairments. In the case of worldwide interconnections, a static allocation would not be applicable because the performance varies too much. These are cases for which the bottom-up approach is well suited. If QoS assignment is based only on on-the-fly measurements with predetermined performance objectives, naturally, the bottom-up approach would be selected.

There is often the common misunderstanding that these approaches are mutually exclusive, which is not correct. One can also use these approaches in combination. Here is one example scenario: At the beginning, the end-to-end performance of multiple network domains is designed (top-down) to meet a fixed set of regional objectives. However, the pre-configured pattern-based apportionment design may not cover all the situations. For those situations, the bottom-up type of approach may also be used alternatively to achieve end-to-end QoS.

V.4 Technologies used at carrier segment

End-to-end QoS in a multi-domain environment can only be achieved when all segment carriers along the path are well managed to satisfy a given quality requirement. Since ITU-T Y.1541 QoS classes are only a set of end-to-end performance values, the performance requirement for each segment would also be a set of figures. There seems to be some attempts to seek the unique technology solution for all network segments. The implementation used to achieve the required QoS performance, however, does not matter as long as the required performance is being fulfilled. There are a number of technologies to ensure QoS: DiffServ and MPLS, for example. Each technology has its advantages and disadvantages regarding the type of service, terminal and network structure, network scale, and regional regulation, for example. More time is needed for the convergence to a single universal solution out of multiple candidate technologies. As a practical solution, each segment carrier should be free to choose any technologies as far as it meets the required performance.

Carriers need to negotiate minimum QoS-related parameters, which are neutral with respect to any specific technology. [b-ITU-T Y.1223] describes an adequate set of these parameters. Thus, it could be used as a specification for this set of exchange parameters.

V.5 Interface of QoS class and terminal equipment

As discussed regarding inter-carrier negotiation, there should also be a negotiation mechanism to share the common view on the QoS class at the user network interface (UNI) between an end-user terminal and the network. There are currently few methods that are standardized and well deployed. One way that has been used in some cases is to use session initiation protocol and session description protocol (SIP/SDP) to deduce the required QoS class, which are specified in [b-3GPP TS 29.213]. When using SIP to establish a session, SDP gives a simple description of the flow. Parameters specified in SDP such as "m" meaning media, "b" meaning bandwidth, and "a" meaning attribute are used to decide an appropriate QoS class. Currently, this method offers the only practical solution.

V.6 Interface of QoS class and enterprise networks

In addition to considering the interaction between the network QoS class and the terminal, it will be necessary to consider the interaction between the network QoS class and the QoS scheme of interconnected enterprise networks. Such enterprise networks may be regional, or even global, in reach and may have a QoS scheme that is as complex as that of the public carrier networks. This may be based on the top-down or bottom-up approaches described above and it may be necessary to treat the enterprise network as another network cloud. Detailed discussion of this interaction remains for further study.

It should be borne in mind that enterprise networks may add significant amounts of delay, IPDV, etc., above those recommended by [ITU-T Y.1541].

Bibliography

[b-ITU-T Y.1223]	Recommendation ITU-T Y.1223 (2008), *Interworking guidelines for transporting assured IP flows*.
[b-ITU-T Y.2001]	Recommendation ITU-T Y.2001 (2004), *General overview of NGN*.
[b-ITU-T Q-Sup.51]	ITU-T Q-series Recommendations – Supplement 51 (2004), *Signalling requirements for IP.QoS*.
[b-3GPP TS 29.213]	3GPP TS 29.213 V8.0.0 (2008), *Policy and charging control signalling flows and Quality of Service (QoS) parameter mapping*.
[b-IEEE 802.11g]	IEEE 802.11g-2003, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*.
[b-JDTEL A35]	Japanese Directive: Directive for Commercial Telecommunication Network Equipments (rules that Japanese telecom carriers need to obey) Article 35 (2008.4) http://law.e-gov.go.jp/htmldata/S60/S60F04001000030.html (Japanese only)

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	Terminals and subjective and objective assessment methods
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, 2010