

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

Page

1	Scope	1
2	References.....	1
3	Terms and definitions	1
4	Abbreviations and acronyms	2
5	Problem statement and consideration of approaches.....	2
5.1	Static approaches	4
5.2	Pseudo-static approach	5
5.3	Signalled approaches	5
5.4	Impairment accumulation approach	7
6	Advantages and disadvantages of the approaches considered.....	8
7	Summary of approaches according to the problem statement challenges	10
	Appendix I – Detailed example of a static divisor approach	12
	Appendix II – Detailed example of the static reference allocation approach	13
	Appendix III – Detailed example of impairment accumulation approach.....	15
	Appendix IV – Performance guidance for providers	18
	IV.1 Qualitative guidance statements	18
	IV.2 Circumstances when guidance is useful	18
	Appendix V – Additional considerations for achieving end-to-end performance objectives in an NGN environment.....	19
	V.1 Introduction	19
	V.2 Requirements for end-to-end QoS class	19
	V.3 End-to-end QoS performance objectives and multi-domain network.....	20
	V.4 Technologies used at carrier segment.....	21
	V.5 Interface of QoS class and terminal equipment.....	21
	V.6 Interface of QoS class and enterprise networks	21
	Bibliography.....	22

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.

Copyright © 2010 ITU-T. All rights reserved.

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].