COMMUNICATIONS ALLIANCE LTD

INDUSTRY GUIDELINE

G633:2012

Quality of Service for networks using the Internet Protocol – Test Methods

G633:2012 Quality of Service for networks using the Internet Protocol – Test Methods Industry Guideline

First published as G633:2007
Second edition as G633:2012

Communications Alliance Ltd was formed in 2006 to provide a unified voice for the Australian communications industry and to lead it into the next generation of converging networks, technologies and services.

Disclaimers

1. Notwithstanding anything contained in this Industry Guideline:
   - (a) Communications Alliance disclaims responsibility (including where Communications Alliance or any of its officers, employees, agents or contractors has been negligent) for any direct or indirect loss, damage, claim, or liability any person may incur as a result of any:
     - (i) reliance on or compliance with this Industry Guideline;
     - (ii) inaccuracy or inappropriateness of this Industry Code; or
     - (iii) inconsistency of this Industry Code with any law; and
   - (b) Communications Alliance disclaims responsibility (including where Communications Alliance or any of its officers, employees, agents or contractors has been negligent) for ensuring compliance by any person with this Industry Guideline.
2. The above disclaimers will not apply to the extent they are inconsistent with any relevant legislation.

Copyright

© Communications Alliance Limited 2012

This document is copyright and must not be used except as permitted below or under the Copyright Act 1968. You may reproduce and publish this document in whole or in part for your or your organisation's own personal or internal compliance, educational or non-commercial purposes. You must not alter or amend this document in any way. You must not reproduce or publish this document for commercial gain without the prior written consent of Communications Alliance. Organisations wishing to reproduce or publish this document for commercial gain (i.e. for distribution to subscribers to an information service) may apply to subscribe to the Communications Alliance Publications Subscription Service by contacting the Communications Alliance Business Manager at info@commsalliance.com.au. If you publish any part of this document for any purpose, you must also publish this copyright notice as part of that publication.

EXPLANATORY STATEMENT

This Explanatory Statement is to be read in conjunction with the G633:2012 ***Quality of Service for networks using the Internet Protocol – Test Methods*** Industry Guideline.

This Explanatory Statement outlines the purpose of this Industry Guideline (the Guideline) and the factors that have been taken into account in its development.

Context of this Guideline

This is a companion Guideline to G632:2012 ***Quality of Service for networks using the Internet Protocol*** Industry Guideline, which defines the IP Network QoS Classes that this Guideline refers to. This Explanatory Statement should be read in conjunction with the Explanatory Statement in G632.

The Objective of this Guideline

This Guideline recommends methods for testing and monitoring IP traffic to verify that IP traffic is achieving the required network performance in accordance with performance objectives for IP Network QoS Classes in G632.

It also provides guidance on reporting occurrences when QoS performance objectives may not have been met, noting that such reporting is a matter for bilateral network interconnection agreements between Carriage Service Providers (CSPs).

Anticipated Benefits to Subscribers

Subscribers will benefit from improved service performance guarantees that may be offered once CSPs can be assured that QoS-enabled traffic is achieving the expected performance levels while crossing other CSP networks, and that interconnected CSPs are testing and monitoring IP traffic performance in a consistent manner.

Anticipated Benefits to Industry

Industry will benefit from reduced time, cost and complexity to establish bilateral network interconnection agreements between CSPs, by defining agreed methods for performance measurement and verification.

Anticipated Cost to Industry

Costs to CSPs will vary depending on the size and scale of network deployments, on the test and measurement equipment procured or developed, and on associated IT systems. These costs may be offset by some degree through increased revenue from enhanced services.

Other Public Interest Benefits or Considerations

Many end-user services, such as telephony and interactive client/server databases, make stringent demands on packet-based network performance. Enhanced reliability and guaranteed end-to-end network performance is required before end-users will have the confidence to transition their services from legacy networks to NGN-style services that are built on IP packet network infrastructure, while maintaining 'any-to-any' connectivity.

Agreed and consistent test methods are an important enabler for building the required level of confidence.

Paul Brooks
Chairman
IP Network Quality of Service Working Committee

COPYRIGHT

TABLE OF CONTENTS

COPYRIGHT

1 GENERAL

1.1 Introduction

1.1.1 The development of this Guideline has been facilitated by the Communications Alliance through a Working Committee comprised of representatives from the telecommunications industry and Government regulatory agencies.

1.1.2 The Guideline should be read in the context of other relevant Codes, Guidelines and documents, including G632:2012 ***Quality of Service parameters for networks using the Internet Protocol*** and other documents listed in Section 8 REFERENCES.

1.1.3 Statements in boxed text are a guide to interpretation.

1.2 Scope

1.2.1 This Guideline applies to CSPs providing IP transport services within Australia in accordance with G632.

1.2.2 This Guideline recommends test methods for the IP Network QoS Classes defined in G632.

1.3 Objectives

1.3.1 This Guideline recommends procedures for testing by a CSP to ensure IP packets for a particular IP Network QoS Class achieve the required performance in all required QoS parameters (IPTD, IPDV, IPLR) for that IP Network QoS Class while being transported across that CSP's network. This testing is regarded as intrusive testing, meaning it may be detrimental if performed in the presence of live end-user traffic.

1.3.2 This Guideline also recommends procedures for continuous or periodic monitoring of network performance in a non-intrusive way that enables a CSP to be assured that its network continues to transport QoS-enabled IP traffic within the indicated performance objectives.

1.3.3 With ongoing monitoring, this Guideline recommends principles for reporting to interconnected CSPs instances when the CSP's network has been unable to meet the performance objectives, noting that such reporting is a matter for bilateral agreement between interconnected CSPs.

1.4 Guideline review

Review of the Guideline will be conducted within five years of publication, or when G632 is reviewed.

COPYRIGHT

2 ACRONYMS, DEFINITIONS AND INTERPRETATIONS

2.1 Acronyms

For the purposes of the Guideline:

CA

means Communications Alliance

CE

means Customer Equipment

CSP

means Carriage Service Provider

GW

means Gateway Router

IETF

means Internet Engineering Task Force

IP

means Internet Protocol

IPDV

means IP packet Delay Variation

IPLR

means IP packet Loss Ratio

IPPM

means IP Performance Metrics

IPTD

means IP packet Transfer Delay

ITU-T

means International Telecommunication Union – Telecommunications Standardization sector

MP

means Measurement Point

MP-NNI

means Measurement Point at an NNI

MP-UNI

means Measurement Point at a UNI

MTU

means Maximum Transmission Unit

COPYRIGHT

OWAMP

means One Way Active Measurement Protocol

NGN

means Next Generation Network

NNI

means Network-to-Network Interface

QoS

means Quality of Service

RFC

means Request For Comment

TCP

means Transmission Control Protocol

UDP

means User Datagram Protocol

UNI

means User-to-Network Interface

UTC(AUS)

means Coordinated Universal Time (Australia).

2.1 Definitions

For the purposes of the Guideline, the following definitions apply:

Act

means the Telecommunications Act 1997.

Carriage Service Provider (CSP)

has the meaning given by section 87 of the Act.

Carrier

has the meaning given by section 7 of the Act.

Customer Equipment

has the meaning given by section 21 of the Act.

Internet Protocol

means the protocol defined in IETF RFC 791.

IP packet Transfer Delay (IPTD)

has the meaning given by G632.

IP packet Delay Variation (IPDV)

has the meaning given by G632.

COPYRIGHT

IP packet Loss Ratio (IPLR)

has the meaning given by G632.

IP Network QoS Class

has the meaning given by G632.

Measurement Point

has the meaning given by G632.

Measurement Point at a UNI

means the location defined as a Measurement Point in G632.

Measurement Point at a NNI

means the interface of a CSPs network equipment that either forms the NNI, or is closest to the NNI.

> *NOTE: Refer to Appendix A for more information on the MP-NNI.*

Network-to-Network Interface (NNI)

has the meaning given by G632.

User-to-Network Interface (UNI)

has the meaning given by G632.

> *NOTE: Refer to G632 and its Appendix C for more information on the UNI.*

COPYRIGHT

3 OVERVIEW

3.1 Performance Parameters

The three performance parameters that define an IP Network QoS Class in a network and which can be measured for a packet population of interest are:

(a) IPTD;

(b) IPDV; and

(c) IPLR.

3.2 IP Network QoS Classes

G632 defines a number of IP Network QoS Classes, including performance objectives and methods of marking packets to signal the desired IP Network QoS Class as the packets are passed across network boundaries.

3.3 Related International Activities

3.3.1 This Guideline references a number of international standards and recommendations (see Section 8 for details).

3.3.2 The QoS Working Group within the Communications Futures Program of MIT released a whitepaper titled "*Interprovider Quality of Service*" in 2006, which was used in the creation of ITU-T Y.1543.

3.3.3 ITU-T Rec. Y.1541 includes some guidance on aspects of performance verification including recommendations of packet sizes and measurement intervals.

3.3.4 ITU-T Rec. Y.1542 provides a framework for achieving end-to-end IP performance objectives.

3.3.5 ITU-T Rec. Y.1543 includes guidance on measurements in IP networks for inter-domain performance assessment.

3.3.6 The IETF IPPM Working Group has developed a series of RFCs on IP traffic performance measurement, based on a framework described in RFC2330. Other documents are listed in Section 8.

NOTES:

1. In several of the reference documents listed in section 8 the ITU-T and the IETF discuss a number of differing definitions of IPDV and calculation of IPDV statistics. Care should be taken to understand the precise form of IPDV calculation used to determine a reported IPDV value when comparing performance measurements or specifications.

2. It is highly desirable that a CSP is able to reliably compare performance specifications and reported achieved values for

 COPYRIGHT

network sections and to be able to reliably estimate likely end to end service outcomes when multiple interconnected CSPs are used to provide the end to end path. It is also highly desirable to minimise the volume and complexity of data that may need to be exchanged between interconnected CSPs to achieve that. Statistical values for IPDV cannot be reliably converted from one calculation method to another unless the distribution of singleton measured values is provided. It is therefore important that all interconnected CSPs, as a minimum, report IPDV statistical values using the same IPDV definition, calculation method, and maximum measurement interval.

3. The definition of IPDV used in G632 and in this Guideline is from paragraph 6.2.4 of ITU-T Rec. Y.1540.

COPYRIGHT

4 REFERENCE ARCHITECTURE

4.1 Architecture

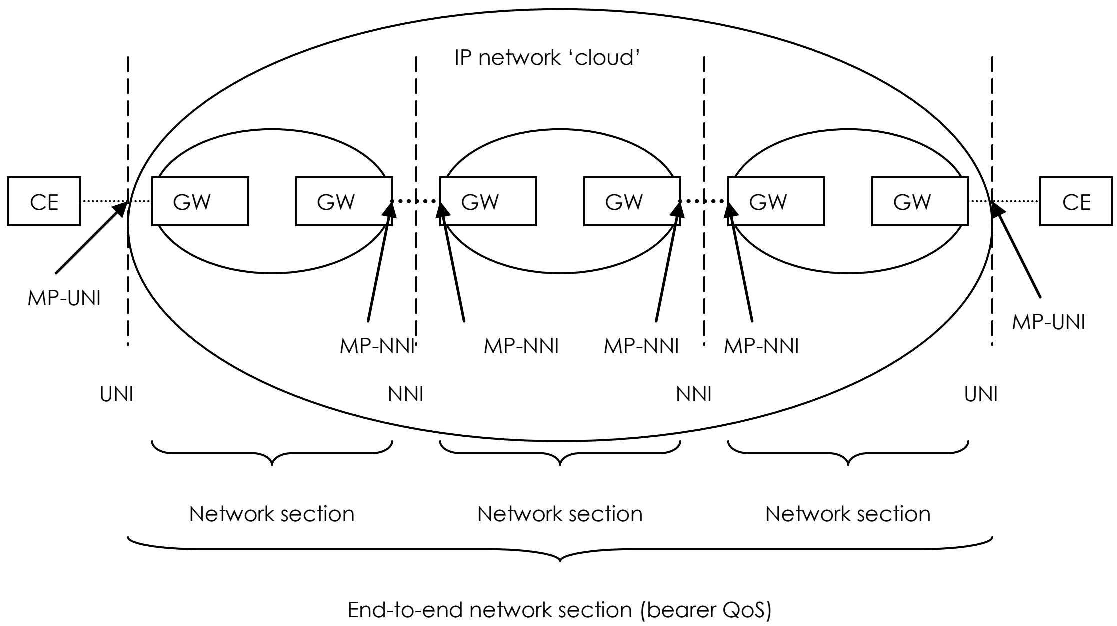

Figure 1

Reference Architecture for IP Network QoS Class Performance

4.1.1 CSPs are limited in their ability to measure the UNI-to-UNI performance directly, because they do not necessarily participate in the end-user's applications.

> *NOTE: The end-user may measure the end-to-end performance of the service directly using information carried in their application packets or measurement data passed between end-user applications.*

4.1.2 The IP network performance may be measured by a CSP over one or more packets that cross its own network section, as shown diagrammatically in Figure 1 and Figure 2 (in Appendix A). A CSP is expected to only test and/or monitor its own network section (UNI-to-NNI or NNI-to-NNI), and rely on reports from other interconnected CSPs to derive the UNI-to-UNI performance likely to be experienced by an end-user's traffic.

4.2 Measurement Points

4.2.1 The measurement point for a UNI (MP-UNI) is the closest point to the end user side of the UNI which can transmit and/or receive a packet over an agreed standards based protocol.

4.2.2 For each CSP, the ideal measurement point for an NNI between two network sections (MP-NNI) is the interface of that CSP's

COPYRIGHT

network equipment that either forms the NNI, or is closest to the NNI.

4.2.3 In many cases an NNI will be carrying operational traffic for other end-users, and cannot be disconnected to attach to test and measurement equipment. In this case an alternative MP-NNI should be used, being another physical interface of the same type and on the same gateway router as the NNI, for the purposes of the performance testing (see Appendix A). In this case, any impairments due to other traffic on the real NNI will not be included in measurements.

5 PERFORMANCE VERIFICATION INTRODUCTION

5.1 General

5.1.1 Traffic performance is verified by measuring the performance parameters of a population of packets as they cross one or more network sections.

5.1.2 A CSP should measure all the packet performance parameters that are relevant for each contracted Network QoS Class.

5.1.3 It is recommended that prior to testing, the network transmission equipment is verified to be working within specification. This is to ensure that there are no lower layer impairments, such as excessive bit errors, which may be degrading service quality.

5.1.4 This Guideline distinguishes two broad categories of measurement:

(a) Testing, using test traffic; and

(b) Monitoring, using end-user payload traffic, or test traffic designed to be unobtrusive to the end-user.

5.1.5 Measurement IP packets should traverse as much as possible the same path as user packets having the same IP Network QoS Class. The same QoS treatment (including network design, routing/forwarding decisions, and queuing strategies) should be provided to test traffic as to end-user traffic along the path. In particular, the packet marking of test traffic should be the same as for user traffic for the IP Network QoS Class to be measured.

5.2 One-way and Two-way measurement

5.2.1 The performance of traffic in each direction of a bidirectional path may differ significantly, due to factors such as asymmetric routing, asymmetric link capacity (especially in access network links), and unequal background traffic levels.

5.2.2 UNI-to-UNI performance objectives in G632 are stated as one-way measures, and performance should be measured for traffic across each network section in each direction separately (MP-UNI to MP-NNI and then MP-NNI to MP-UNI as separate

COPYRIGHT

measurements, or one MP-NNI to the other MP-NNI, then back again, as separate measurements).

5.2.3 For operational reasons it may sometimes be necessary to measure both directions simultaneously using two-way or round-trip traffic, where test packets are transmitted at one MP, pass through the network section to the other MP and back again, to be received and measured at the original MP.

> *NOTE: If supported by the network architecture and technology, a hardware or software loop-back may be placed on one MP and then test traffic may be both transmitted and received at the other MP. The loopback function should ensure the source and destination IP addresses in the IP packet headers are swapped when the packet swaps direction.*

5.2.4 If the two-way test results indicate the combined path does not achieve the performance objectives, then each direction should be separately measured to determine the direction that causes excessive impairment in one or more performance measures.

5.3 Clocks and Time References

5.3.1 A common time reference for each interconnected CSP is required to enable test reports from multiple network segments to be combined and compared.

5.3.2 It is recommended that UTC(AUS) should be the common time reference for all measurements – however this does not preclude interconnected CSPs from using alternate time references as per their bilateral agreements.

5.3.3 Where statistical values are reported by the interconnected CSPs that are calculated from a population of measurements over a measurement interval, the time at the beginning of the measurement interval, and the duration of the measurement interval should be reported by the interconnected CSPs.

5.3.4 Precise time comparisons for one-way measurements between two measurement devices requires:

(a) both measurement devices to be synchronized to a common time reference; or

(b) the clocks of each measurement device to be synchronized with each other.

5.3.5 Measurements recorded by the interconnected CSPs should include an indication of the maximum level of uncertainty in the synchronization of the clocks compared to their agreed common time reference.

COPYRIGHT

NOTES:

1. Interconnected CSPs should periodically compare and check their reference clocks, and measurement devices should be actively synchronized to their agreed common time reference.

2. It is up to each interconnected CSP to develop its own method for verifying their time reference.

3. For more detail on UTC (AUS), refer to the National Measurement Institute at www.measurement.gov.au/Services/Pages/TimeandFrequencyDisseminationService.aspx.

5.4 Format of Test Packets

5.4.1 When verifying performance of an end-user's service, test packets should be generated only for the IP Network QoS Classes contracted to be carried for the end-user by the CSP.

5.4.2 Test packets for each IP Network QoS Class should be evenly distributed throughout the test packet population.

5.4.3 If possible, the IP packet payload should be formatted as an OWAMP-test packet (refer to RFC 4656). If the packet generator does not support the OWAMP protocols, the packet should be formatted as a UDP packet with the payload filled with a pseudo-random sequence of byte values, to remove the effects of any link compression technologies. The UDP source port and destination port values should be higher than 10,000, and should remain the same for all packets in the test stream.

5.4.4 Any individual stream of test packets should not be sent at a higher data rate than is committed for an individual stream of end-user traffic. The end-user traffic offered to a network is likely to consist of multiple streams of UDP and TCP packets, and test traffic should reflect this.

COPYRIGHT

6 TESTING

6.1 General

6.1.1 The purpose of performance testing is to ensure the performance objectives are achieved when:

(a) one contracted IP Network QoS Class; or

(b) multiple contracted IP Network QoS Classes of a multi class service instance

is/are offered valid traffic at the maximum data rate committed by the CSP. Where feasible two such tests are recommended – one with packets at or near the maximum MTU size (smallest packet rate) and one with small packet size, maximizing packet rate.

6.1.2 These tests are designed to verify that when traffic is offered to the UNI or NNI on one or more contracted IP Network QoS Class(es) at the maximum data rate committed by the CSP, that the performance specification for each contracted IP Network QoS Class is still met across the network section being tested.

6.1.3 When testing performance, measurements are made on test traffic introduced directly within the end-user's traffic path. In this way all parameters can be measured directly, however an existing end-user's service may be degraded. Testing should therefore only be done either prior to service handover, or in an intermittent fashion such as during fault investigation.

6.2 Test Traffic

6.2.1 Test traffic should be transmitted with uniform packet length.

NOTE: As IPTD is defined as the time between when the first bit of the packet enters the network, and the time the last bit of the packet leaves the network, test packets must be of uniform length to ensure variations in packet length do not cause variations in IPTD measurements.

6.2.2 Test traffic may be transmitted with uniform inter-packet spacing, or the inter-packet spacing may have a random distribution. If a uniform inter-packet spacing is used then the test procedure should conform to RFC 3432.

NOTES:

1. As stated in RFC 3432, the start times of successive test intervals should not be correlated or at regular intervals, and the test duration should be short enough to minimize the chance of other traffic flows that share part of the test traffic path from becoming synchronized.

COPYRIGHT

> *2. CSPs are recommended to migrate towards random inter-packet spacing in future as international standards for IP traffic measurement are developed and compatible measuring equipment becomes available cost-effectively in Australia.*
>
> *3. RFC 2330 and RFC 3432 provide further guidance on inter-packet spacing distributions.*

6.2.3 Subject to clause 6.2.5, the average data rate for each IP Network QoS Class should be the maximum data rate committed by the CSP within the service contract.

6.2.4 Each test should consist of two sub-tests with different packet lengths. The packet lengths should be:

(a) Less than or equal to 200 bytes (including IP header); and

> *NOTE: This is to test the service under maximum packet rate conditions, so the chosen packet size should match the smallest size of packets in the traffic class that the CSP has committed to support.*

(b) Either:

i. the MTU committed to be supported within the service contract; or

ii. 1500 bytes (including IP header) if no MTU is specified.

> *NOTE: The values of 200 bytes and 1500 bytes are derived from ITU-T Rec. Y.1541 clause 5.3.3.*

6.2.5 When the sum of the committed data rates for all contracted IP Network QoS Classes exceeds the maximum committed capacity of the service then the service is overbooked. In such cases the test data rate for each IP Network QoS Class with the lesser priority should be reduced in accordance with the service contract such that the service is not overbooked.

6.2.6 When the service is offered over data transmission links of varying capacity, the testing should be conducted on a link of the lowest capacity offered.

6.3 Test Duration

6.3.1 The testing should be consistent with ITU-T Rec. Y.1541. This suggests:

(a) The measurement interval duration should be recorded with the corresponding measured metrics.

(b) A measurement interval of 1 minute (refer to section 5.3 of ITU-T Rec. Y.1541 for more information).

(c) Not less than 1500 singleton measurements in each measurement interval.

COPYRIGHT

NOTE: This enables the estimation of the 1 - 10^{-3} quantile (i.e. 99.9%) IPDV performance objective of Table 1 of G632.

(d) The measurement of performance metrics for multiple measurement intervals to increase the measured quantile.

NOTE: The combination of (c) and (d) enables the estimation of a more precise value for one network in order to assess an end-to-end performance objective from Table 1 of G632 across multiple networks.

6.3.2 At the end of a test, measuring equipment receiving the test packets should continue to listen and count packets for at least 3 seconds after the traffic generator has ceased transmitting the test packet stream.

COPYRIGHT

7 MONITORING

7.1 General

7.1.1 The purpose of performance monitoring and analysis is to confirm that ongoing delivery of the IP transport service is achieved within the performance objectives for the relevant IP Network QoS Classes. Monitoring is intended to enable ongoing service assurance while the end-user's service is in operation.

7.1.2 The measurement approach should be non-intrusive – it should not significantly impact end-user traffic, either through excessive link load from the addition of test traffic in the traffic path or as the result of load placed on gateway routers by the measurement processes such as generating and responding to test traffic.

7.1.3 When monitoring, performance of a connection may be inferred by:

(a) observing the end-user's traffic; or

(b) by introducing test traffic into the same physical path taken by the end-user's traffic.

The method adopted will depend on the network architecture, amongst other factors.

7.2 Monitoring Traffic

7.2.1 Each CSP should measure the three IP packet performance parameters for each Network QoS Class supported on its network.

7.2.2 Packet performance should be measured and recorded in each monitoring period over a population-of-interest of at least 1500 packets, spread approximately uniformly throughout the monitoring period.

7.2.3 It is the responsibility of each CSP to establish the duration and frequency of its monitoring period in accordance with its operational policies & procedures.

7.2.4 Measurements should:

(a) not have an impact on end-user traffic;

(b) consist of not less than 1500 singleton measurements in a measurement interval. This is to achieve a statistically valid population-of-interest, as in clause 6.3.1;

(c) be obtained over a measurement interval of 1 minute (refer to section 5.3 of ITU-T Rec. Y.1541 for more information); and

COPYRIGHT

NOTE: Monitoring a low speed access link may require an increase in the measurement interval in order not to have an impact on end-user traffic.

(d) include a repetition of measurement intervals to assist ongoing monitoring. Measurements should be repeated every 5 minutes, unless to do so would impact on end-user traffic in which case the measurements should be repeated every 15 minutes.

NOTE: The time between measurement intervals should be randomized – refer to RFC 3432 for more information.

7.2.5 Test traffic may be used for monitoring purposes, provided it follows the guidance for test traffic in section 6.2 and in section 7.1.2.

7.3 Reporting

7.3.1 When a performance measurement of a network section exceeds an agreed threshold value for a parameter in Table 1, the CSP should report the incident to its interconnected CSPs.

7.3.2 To minimize the incidence of reporting performance values under normal conditions, parameter measurements should be compared by interconnected CSPs to an agreed threshold value for the relevant parameter.

7.3.3 Industry agreed example threshold values for the network parameters are in Table 1. Refer to Appendix B for a structure to define industry agreed threshold values for reporting on access and core networks.

7.3.4 The agreed threshold values should be set with a margin such that normally there is no need to report, but are not so large that the end-to-end performance objective might be exceeded when all network sections are below the agreed threshold.

7.3.5 CSPs should agree bilaterally on the reporting methods and frequency.

NOTE: Matters that might be included in bilateral agreements include:

(a)content and format of such reports;

(b) agreed processes for the exchange of hard copies of the measurement results;

(c) methods for the electronic exchange of measurement reports;

(d) time frames for delivery of the report after the identification of exceeding an agreed threshold value; and

COPYRIGHT

(e) One or more agreed threshold value(s) against each parameter with indications of severity.

7.3.6 CSPs should agree on the reporting methods. At the minimum there should be agreed processes for the exchange of hard copies of the performance results, including the content and format of such reports. It is highly desirable that CSPs agree on methods for the electronic exchange of measurement reports. Such an agreement would include both the content of the reports, a protocol for exchange of the reports, and a time limit for delivery of the report after the threshold breach is identified.

7.3.7 It is recommended that reports when an agreed threshold value is exceeded should, as a minimum, contain the following information:

(a) Date;

(b) Time (at the commencement of the measurement interval);

(c) Location of end points;

(d) Measurement/report period;

(e) Measurement type;

(f) Measurement statistics;

(g) Brief Reason; and

(h) Optionally, the report may also include service identifiers for services that might be affected by the issue.

NOTE: For some purposes, reporting intervals on a similar timescale to billing activities (say monthly) might be sufficient for matching performance issues with fault reports, and be relatively easily implemented.

At the other extreme, a CSP's monitoring system upon detecting an agreed threshold value breach might automatically raise an alarm, which might be disseminated to an interconnected CSP's alarm management system. This would permit service issues to be diagnosed rapidly, but would be operationally complex and may be expensive to implement.

7.4 Considerations for services spanning multiple networks

7.4.1 When a service spans the networks of multiple CSPs, there may be additional challenges in providing accurate end-to-end measurements for a given end-user's service. For example, it may be difficult for any one CSP to determine the precise path that is taken by a particular end-user's traffic. Even if the path is known, it may be difficult to conduct measurements along that exact path, e.g. due to a lack of devices to respond to measurement probes at various points on the path.

7.4.2 The goals of the monitoring techniques described above, therefore, are more modest than the delivery of precise performance data to a particular end-user. Instead, the primary goal is to allow a CSP to make certain QoS assurances to a end-user, knowing that:

(a) the impairments that can be expected from other CSPs in the path will enable those assurances to be met if all interconnected CSPs meet their impairment targets; and

(b) the reported measurements of each CSP should indicate when an interconnected CSP has failed to meet the targets.

COPYRIGHT

8 REFERENCES

Publication	Title
IETF RFCs	
RFC 2330	Framework for IP Performance Metrics (IPPM) http://tools.ietf.org/html/rfc2330
RFC 2679	A One-way Delay Metric for IPPM http://tools.ietf.org/html/rfc2679
RFC 2680	A One-way Packet Loss Metric for IPPM http://tools.ietf.org/html/rfc2680
RFC 3393	IP Packet Delay Variation Metric for IPPM http://tools.ietf.org/html/rfc3393
RFC 3432	Network performance measurement with periodic streams http://tools.ietf.org/html/rfc3432
RFC 4656	One Way Active Measurement Protocol (OWAMP) http://tools.ietf.org/html/rfc4656
ITU-T Recommendations	
Y.1540 (12/02)	IP data communication service – IP packet transfer and availability performance parameters http://www.itu.int/rec/T-REC-Y.1540-201103-I
Y.1541 (02/06)	Network performance objectives for IP-based services http://www.itu.int/rec/T-REC-Y.1541-201112-I
Y.1542	Framework for achieving end-to-end IP performance objectives
Y.1543	Measurements in IP networks for inter-domain performance assessment
Industry Guidelines	
G632:2012	Quality of Service parameters for networks using the Internet Protocol http://commsalliance.com.au/Documents/all/guidelines/g632
Other	
MIT Communications Futures Program	Inter-provider Quality of Service – White Paper http://cfp.mit.edu/publications/index.shtml

COPYRIGHT

APPENDIX A – THE NNI MEASUREMENT POINTS

A.1.1 Figure 2 below is an example of reference architecture for a possible combination of interconnecting the IP networks of CSPs. It expands on Figure 1 to examine the structure of the inter-provider link in more detail, and highlights the NNI for QoS purposes.

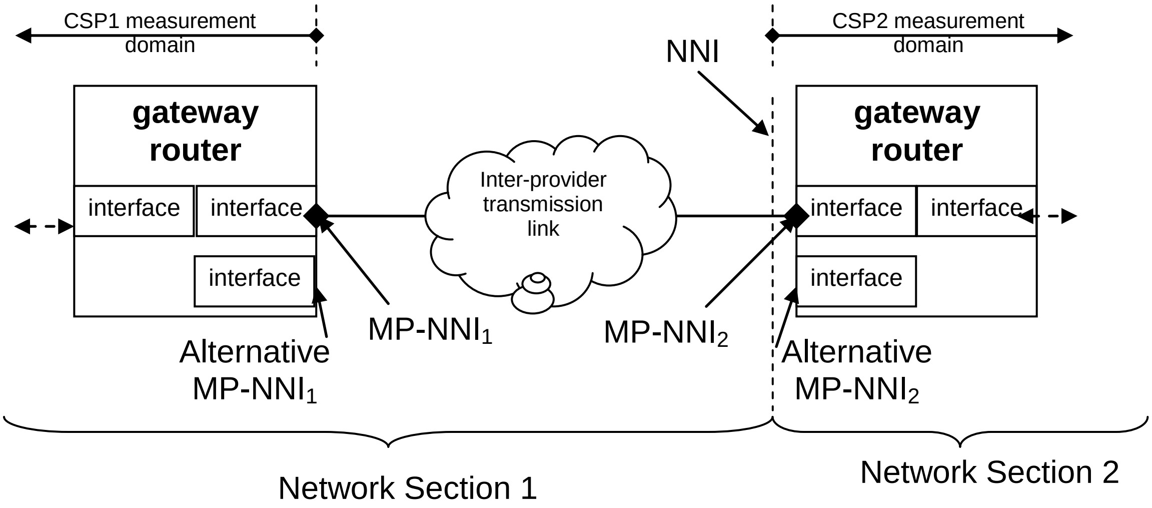

Figure 2
Reference Architecture for NNI

A.1.2 Where the link between two CSPs' gateway routers is provided by one of the CSPs, the NNI and MP-NNI is taken as the interface of the gateway router of the CSP not providing the link.

A.1.3 In some cases the NNI is assumed to be at the midpoint of the link between two CSPs. The MP-NNI for each CSP is the gateway router interface of their network equipment that connects to the inter-provider link.

COPYRIGHT

APPENDIX B – THRESHOLD VALUES FOR REPORTING AMONG CSPS

B.1.1 Table 1 is a recommended structure to document agreed threshold values for reporting between CSPs on access and core networks when parameter(s) are out of specification. However CSPs may choose to adopt alternate reporting structures.

B.1.2 The example threshold values in Table 1 are derived from data in Appendix III to Y.1541, as reflected in G632 and the MIT communications future paper.

B.1.3 A threshold value for an access network parameter is likely to be different from the corresponding threshold value for a core network parameter.

Network Type	**Parameter**	**Threshold**			
		QoS Class 0	**QoS Class 2**	**QoS Class 6**	**QoS Class 5**
Core	Mean IPTD	Largest of: (a) 15 ms; or (b) 10ms + (Airpath distance in km) x 1.4 x0.005ms			-
Core	99% IPDV	2ms	Not Applicable	2ms	-
Core	IPLR	1×10^{-5}	1×10^{-5}	1×10^{-7}	-
Access	Mean IPTD	25ms			-
Access	99% IPDV (Note 2)	16ms	-	16ms	-
Access	IPLR	4×10^{-4}	4×10^{-4}	4×10^{-6}	-

Table 1

Example Threshold values for reporting among CSPs.

NOTES:

1. For the purpose of comparing against a threshold, IPDV is measured to the 99% level for QoS classes 0, 2, 5 and 6, rather than the default level (e.g. 99.9% or 99.99%). This is for reasons of measurement practicality with a relatively low packet rate.

2. A CSP should add 300 ms IPTD for segments containing a geostationary satellite hop.

 COPYRIGHT

PARTICIPANTS

The Working Committee that developed the Guideline consisted of the following organisations and their representatives:

Organisation	Membership	Representative
Individual	Voting	Allan Horsley
ACMA	Non-voting	Cuong Nguyen
ACMA	Non-voting	Chris Wong
Alcatel-Lucent	Voting	Evan Stanbury
Internet Society of Australia	Voting	Paul Brooks
NBN Co	Voting	Rod Ferguson
SingTel Optus	Voting	Krishna Ramalingam
SingTel Optus	Non-voting	Sam Mangar
Telstra	Voting	Hee Juan Ho
Telstra	Non-voting	Daniel Wetuschat

This Working Committee was chaired by Paul Brooks.

James Duck of Communications Alliance provided project management support.

COPYRIGHT

Communications Alliance was formed in 2006 to provide a unified voice for the Australian communications industry and to lead it into the next generation of converging networks, technologies and services.

In pursuing its goals, Communications Alliance offers a forum for the industry to make coherent and constructive contributions to policy development and debate.

Communications Alliance seeks to facilitate open, effective and ethical competition between service providers while ensuring efficient, safe operation of networks, the provision of innovative services and the enhancement of consumer outcomes.

It is committed to the achievement of the policy objective of the *Telecommunications Act 1997* - the greatest practicable use of industry self-regulation without imposing undue financial and administrative burdens on industry.

Published by:
COMMUNICATIONS ALLIANCE LTD

Level 9
32 Walker Street
North Sydney
NSW 2060 Australia

Correspondence
PO Box 444
Milsons Point
NSW 1565

T 61 2 9959 9111
F 61 2 9954 6136
TTY 61 2 9923 1911
E info@commsalliance.com.au
www.commsalliance.com.au
ABN 56 078 026 507

Care should be taken to ensure the material used is from the current version of the Guideline and that it is updated whenever the Guideline is amended or revised. The number and date of the Guideline should therefore be clearly identified. If in doubt please contact Communications Alliance.