ETSI EG 202 765-3 V1.1.2 (2010-07)

ETSI Guide

**Speech and multimedia Transmission Quality (STQ);
QoS and network performance metrics and measurement methods;
Part 3: Network performance metrics and measurement methods in IP networks**

Reference

REG/STQ-00174-3

Keywords

performance, QoS

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This ETSI Guide (EG) has been produced by ETSI Technical Committee Speech and multimedia Transmission Quality (STQ).

The present document is part 3 of a multi-part deliverable covering the QoS and network performance metrics and measurement methods, as identified below:

EG 202 765-1: "General considerations";

ES 202 765-2: "Transmission Quality Indicator combining Voice Quality Metrics";

EG 202 765-3: "Network performance metrics and measurement methods in IP networks";

ES 202 765-4: "Indicators for supervision of Multiplay services".

Introduction

The need to define Internet performance metrics and measurement methodologies stems from the need to compare different measurements and to measure performance with a reproducible and unambiguous methodology, independent from transmission technology and implementation details. Both the ITU-T Study Group 12 and the IETF IPPM Working Group have produced such definitions (see table 1), although each with a different emphasis closely linked to the historical background of both organizations. The ITU has its origins in telephony, while the IETF has a data networking background. Whereas the ITU emphasizes the evaluation of a service and its quality, the IETF measures the network and wants to provide the IT-community with an accurate, common understanding and measurement of the performance and reliability the Internet [i.2].

In most cases this results in different terminology rather than in incompatibilities; most differences in approach and emphasis serve the different intended use of each metric, but have no operational significance. In some cases the terminology used by each organisations can be mapped to the other, while in some others there is only approximate equivalence (e.g. ITU network section versus an IPPM cloud; one focuses on corresponding events while the other measures the fate of a single packet). Other terms have no correspondence. For example, ITU-T Recommendation I.380 [i.28] has a notion of an IP packet transfer reference event while IPPM defines "wire time".

Other differences between IETF and ITU-T metrics result from their intended application. ITU-T metrics seek to provide a common language for providers to communicate about performance, so the ITU-T metrics do not concentrate on performance within a single network, while the IETF focuses on performance measurement protocols and implementation. ITU-T seeks to evaluate service and to exclude unfair use, while the IETF seeks to measure network quantities and avoid biased measurement results. Due to their respective backgrounds, the ITU generally produces statistical metrics geared towards a quantitative representation of the complete end-to-end user experience while the IETF IPPM working group mainly focuses more on statistical metrics which provide a detailed technical view of different aspects of transmission quality along the network path.

Table 1: Overview of Relevant Standards

	IETF RFCs	ITU-T Recommendations
Framework	RFC 2330 [i.2]	Y.1540 [i.1], sections 1 through 5
Loss	RFC 2680 [i.5]	Y.1540 [i.1], section 5.5.6 G.1020 [i.17]
Delay	RFC 2679 [i.4] (One-way) RFC 2681 [i.6] (Round Trip)	Y.1540 [i.1], section 6.2 G.1020 [i.17] G.114 [i.16] (One-way)
Delay Variation	RFC 3393 [i.9]	Y.1540 [i.1], section 6.2.2 G.1020 [i.17]
Connectivity / Availability	RFC 2678 [i.3]	Y.1540 [i.1], section 7
Loss Patterns	RFC 3357 [i.8]	G.1020 [i.17]
Packet Reordering **Packet Duplication**	RFC 4737 [i.11]	Y.1540 [i.1], sections 5.5.8.1 and 6.6 Y.1540 [i.1], sections 5.5.8.3, 5.5.8.4, 6.8, and 6.9
Link/Path Bandwidth Capacity, Link Utilization, Available Capacity	RFC 5136 [i.21]	
Bulk Transport Capacity	RFC 3148 [i.7], RFC 5136 [i.21]	

The goal of the present document is to define network performance metrics for applications sensitive to quality of service such as Voice over IP, referring to the existing work produced by both IETF and ITU-T. The present document highlights the differences between the two standards and provides guidelines on resolving these differences, when they are due to addressing different goals.

The scope of the present document is limited to IP performance metrics relevant for data transmission over IP-based networks for use in QoS sensitive applications. For each addressed metric, the document recommends one or more measurement methods. The document only focuses on intrinsic network QoS metrics; perceived QoS metrics applicable for voice transmission are out of scope of the present document.

The remainder of the present document is organised as follows: Clause 4 describes the definitions of the most important performance metrics as defined by the standard bodies and methods for measuring them, and discusses the applicability of the definitions and the differences between them. Clause 5 discusses other metrics applicable to QoS. Finally, clause 6 gives an overview of relevant QoS measurement standards, which can be used in end to end performance evaluation.

1 Scope

The present document provides an overview of the common metric definitions and measurement method specifications upon which the interoperability of network performance measurement (also called QoS measurement) is based. Two different standardisation bodies, the Internet Engineering Task Force (IETF) and the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T), have addressed this issue. The present document addresses the following points:

- Survey the existing network performance related IETF standards and how these standards can be applied to end-to-end network performance measurements. The scope of this work is also to discuss the relationship of those standards to those of ITU-T and ETSI.
- Discuss and compare definitions of metrics used to specify and assess performance in IP networks. The metrics addressed in the present document are those defined by the IETF IPPM working group and ITU-T Study Group 12. Besides comparing the different definitions, the present document gives applicability guidelines on which metric is more appropriate for a particular application, configuration or scenario.
- Define measurement methods for selected performance metrics in IP networks, addressing both active and passive methods. Clarifying guidelines are given.

NOTE: All text sections in the remainder of the present document which are enclosed in quotation marks (") and *formatted in italic style* denote citations taken verbatim from referenced documents.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are necessary for the application of the present document.

Not applicable.

2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ITU-T Recommendation Y.1540: "Internet protocol data communication service - IP packet transfer and availability performance parameters".

[i.2] IETF RFC 2330: "Framework for IP Performance Metrics". V. Paxson, G. Almes, J. Mahdavi, M. Mathis. May 1998.

[i.3] IETF RFC 2678: "IPPM Metrics for Measuring Connectivity". J. Mahdavi, V. Paxson. September 1999.

[i.4] IETF RFC 2679: "A One-way Delay Metric for IPPM". G. Almes, S. Kalidindi, M. Zekauskas. September 1999.

[i.5] IETF RFC 2680: "A One-way Packet Loss Metric for IPPM". G. Almes, S. Kalidindi, M. Zekauskas. September 1999.

[i.6] IETF RFC 2681: "A Round-trip Delay Metric for IPPM". G. Almes, S. Kalidindi, M. Zekauskas. September 1999.

[i.7] IETF RFC 3148: "A Framework for Defining Empirical Bulk Transfer Capacity Metrics". M. Mathis, M. Allman. July 2001.

[i.8] IETF RFC 3357: "One-way Loss Pattern Sample Metrics". R. Koodli, R. Ravikanth. August 2002.

[i.9] IETF RFC 3393: "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)". C. Demichelis, P. Chimento. November 2002.

[i.10] IETF RFC 4656: "A One-way Active Measurement Protocol (OWAMP)". S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, M. Zekauskas. September 2006.

[i.11] IETF RFC 4737: "Packet Reordering Metrics". A. Morton, L. Ciavattone, G. Ramachandran, S. Shalunov, J. Perser. November 2006.

[i.12] IETF RFC 5101: "Specification of the IPFIX Protocol for the Exchange of IP Traffic Flow Information". B. Claise, S. Bryant, S. Leinen, T. Deitz, B.Trammell. January 2008.

[i.13] Internet-Draft, work in progress: "IPFIX Architecture". N.Brownlee et Al.

[i.14] IETF RFC 5102: "IPFIX Information Model". J. Quittek et Al. January 2008.

[i.15] Internet-Draft, work in progress: "IPFIX Applicability Statement". T. Zseby, E. Boschi, N.Brownlee, B. Claise.

[i.16] ITU-T Recommendation G.114 (05/03): "One-way transmission time".

[i.17] ITU-T Recommendation G.1020 (07/06): "Performance parameter definitions for quality of speech and other voiceband applications utilizing IP networks".

[i.18] IETF RFC 3917: "Requirements for IP Flow Information Export". J. Quittek, T. Zseby, B. Claise, S. Zander. October 2004.

[i.19] draft-morton-ippm-reporting-metrics-02 work in progress: "Reporting Metrics: Different Points of View", A. Morton, G. Ramachandran, G. Maguluri.

NOTE: http://tools.ietf.org/html/draft-morton-ippm-reporting-metrics-02, and the derived presentation "Reporting Metrics: Different Points of View" presented by Al Morton on IETF66 July 2006, http://www3.ietf.org/proceedings/06jul/slides/ippm-2.pdf.

[i.20] IETF RFC 3611: "RTP Control Protocol Extended Reports (RTCP XR)", T. Friedman, R. Caceres, A. Clark. November 2003.

[i.21] IETF RFC 5136: "Defining Network Capacity", P. Chimento, J. Ishac. February 2008.

[i.22] IETF RFC 2581: "TCP Congestion Control", M. Allman, V. Paxson, W. Stevens. April 1999.

[i.23] IETF RFC 5357: "A Two-Way Active Measurement Protocol (TWAMP)", K. Hedayat, R. Krzanowski, A. Morton, K. Yum, J. Babiarz. October 2008.

[i.24] IETF RFC 1122: "Requirements for Internet Hosts - Communication Layers", R. Braden ed. October 1989.

[i.25] IETF RFC 3550: "User Accounts for UCSB On-Line System".

[i.26] IETF RFC 1633: "Integrated Services in the Internet Architecture: an Overview".

[i.27] IETF RFC 2216: "Network Element Service Specification Template".

[i.28] ITU-T Recommendation I.380: "Internet protocol data communication service - IP packet transfer and availability performance parameters".

3 Definitions, symbols and abbreviations

3.1 Definitions

For the purposes of the present document, the terms and definitions given in RFC 2330 [i.2], ITU-T Recommendation G.1020 [i.17] and RFC 2680 [i.5] apply.

3.2 Symbols

For the purposes of the present document, the following symbols apply:

T, t	Time
T_{max}	Time threshold
dT	Time difference

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ASON	Automatically Switched Optical Network
ATM	Asynchronous Transfer Mode
BTC	Bulk transport Capacity
DNS	Domain Name System
ESD	End System Delay
FTP	File Transfer Protocol
HTTP	HyperText Transfer Protocol
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IPDV	IP Packet Dealy Variation
IPFIX	IP Flow Information eXport
IPLR	IP Packet Loss Ratio
IPPM	IP Performance Metrics
IPTD	IP Packet Transfer Delay
ITU-T	International Telecommunication Union - Telecommunication standardisation sector
MIB	Management Information Base
NSE	Network Section Ensemble
OP	Observation Point
OWAMP	One Way Active Measurement Protocol
OWD	One Way Delay
PDV	Packet Delay Variation
PIA	Percent IP service Availability
PON	Passive Optical Network
PSAMP	Packet SAMPling
QoS	Quality of Service
RFC	Request For Comments
RTCP	Real Time Control Protocol
RTD	Round Trip Delay
RTP	Real-Time Transport Protocol
RTT	Round Trip Time
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
TCP	Transmission Control Protocol
TWAMP	Two-Way Active Measurement Protocol
UTC	Coordinated Universal Time
VoIP	Voice over IP

4 Performance Metrics Definitions and Measurement Methods

This clause provides common definitions for network performance metrics. These definitions are based, whenever possible, on existing definitions proposed by other relevant standard bodies such as IETF or ITU-T. Note that the different definitions of similar metrics are in most cases compatible, that is, semantically equivalent or easily convertible into one another.

For each metric, passive and active measurement methods are defined. Note that we chose to focus on commonly used measurement methods rather than on standards; when a standard exists, a reference is provided as well. Note also that throughout this text we refer for each metric to active and passive measurements in the following way:

- **Active measurements**

  Active measurement methods inject traffic into the network and compute traffic metrics based on monitoring the injected traffic or the response to the injected traffic. Active test traffic may perturb other traffic already present on the network; therefore its scheduling and volume should be carefully configured. One can distinguish active monitoring systems based on the position of sender and receiver and the observed traffic; this is specified in detail for the considered metrics in the following text.

- **Passive measurements**

  Passive measurements provide information about traffic in the observed network by capturing all or a selected subset of the IP packets traversing a monitoring point. Since no test traffic is generated, passive measurements can only be applied when the traffic of interest is already present on the network. The physical deployment of monitoring probes in the network can be realised in different ways, depending on the metrics of interest, but also on the network technology, e.g. via a physical line splitter, via a normal client connection in broadcast networks, or via a dedicated monitoring port on a switch or router.

4.1 One Way Delay vs. IP Packet Transfer Delay

Delay is used to measure the expected time for an IP packet to traverse the network from one host to another. Delay is applicable to QoS for latency-sensitive protocols. The IETF and ITU-T metrics for measuring delay are essentially compatible, though there are minor differences; the details of these metrics are given in this clause.

4.1.1 IETF Definition

RFC 2679 [i.4] distinguishes between a "singleton analytic metric", called Type-P-One-way-Delay, and a "sample", called Type-P-One-way-Delay-Poisson-Stream. The singleton is introduced to measure a single observation of one-way delay, while the sample is used to measure a sequence of singleton delays measured at times taken from a Poisson process. Based on these samples, several statistics are defined, such as Type-P-One-way-Delay-Percentile, Type-P-One-way-Delay-Median, Type-P-One-way-Delay-Minimum, and Type-P-One-way-Delay-Inverse-Percentile.

Since the value of many of these metrics depends on the type of the IP packet used to perform the measurements, IPPM metrics definitions include the generic notion of "a packet of type P", which should be further specified when making actual measurements.

RFC 2679 [i.4] defines:

"For a real number dT, >>the *Type-P-One-way-Delay* from Src to Dst at T is dT<< means that Src sent the first bit of a Type-P packet to Dst at wire-time* T and that Dst received the last bit of that packet at wire-time T+dT."

The notion of wire time is introduced in RFC 2330 [i.2] in order to take into account the additional delay derived from the use of Internet hosts to perform the measurements. Wire time is defined with reference to an Internet host H observing an Internet link L at a particular location. More precisely, for a given packet P, the "wire arrival time" of P at H on L is the first time (see note) T at which the first bit of P has appeared at H's observational position on L. On the other side, For a given packet P, the 'wire exit time' of P at H on L is the first time T at which all the bits of P have appeared at H's observational position on L. Wire time delay is defined as the time between the first wire arrival time, the moment in which the first bit of the packet leaves the network interface of the source and the subsequent wire exit time at the remote end, the moment at which it has arrived completely at the network interface of the destination host.

NOTE: An IP packet might arrive at the destination Dst more than once, due to retransmission.

An upper bound for the expected packet delivery is taken into account (this threshold should also be reported):

"If the packet fails to arrive within a reasonable period of time, the one-way delay is taken to be undefined (informally, infinite)."

Further RFC 2680 [i.5] states about corrupted packets:

"*If the packet arrives, but is corrupted, then it is counted as lost.*"

This is a useful approach since corrupted packets (even though they arrived) cannot be safely assigned to a flow or traffic sender, as the source identifier might be a part of the packet that was corrupted.

4.1.2 ITU-T Definition

The ITU-T Recommendation Y.1540 [i.1] defines the **IP Packet Transfer Delay (IPTD)** metric as *the one-way IP packet transfer delay for all successful and errored packet transmissions across a basic section or a Network Section Ensemble (NSE). IPTD is the time,* $(t_2 - t_1)$ *between the occurrence of two corresponding IP packet reference events, ingress event* $IPRE_1$ *at time* t_1 *and egress event* $IPRE_2$ *at time* t_2*, where* $(t_2 > t_1)$ *and* $(t_2 - t_1) \leq T_{max}$. If fragmentation occurs then the time t_2 is considered the time of the final corresponding fragment [i.1].

To understand this definition, the meaning of IP packet reference event has to be clarified. ITU-T Recommendation Y.1540 [i.1] defines it as follows:

An IP packet transfer event occurs when:

- an IP packet crosses a Measurement Point (MP); and
- standard IP procedures applied to the packet verify that the header checksum is valid; and
- the source and destination address fields within the IP packet header represent the IP addresses of the expected source host and destination host.

The IP packet transfer reference events are defined without regard to packet fragmentation.

An IP egress event is said to *correspond* to an earlier ingress event if they were created by the "same" IP packet.

Finally, the end-to-end IP packet transfer delay is the one-way delay between the measurement point at the source host and destination host.

ITU-T Recommendation Y.1540 [i.1] also mentions the mean IP packet transfer delay, which is defined as the arithmetic average of IP packet transfer delays for a packet population of interest.

4.1.3 Comparison and Recommendations

In general the ITU-T, due to its telecommunications origin, often assesses the observed delay at the endpoints of a full end-to-end connection, i.e. from Network Termination Point to Network Termination Point., while IETF is concerned with network delays among arbitrary points in the network. When setting upper limits for e.g. network planning then the desired application (e.g. VoIP) as well as effects on QoS due to codecs and de-jitter buffers should be taken into account for the assessment of the needed network QoS.

The two definitions from IETF and ITU-T for one way delay can be considered compatible, as for both definitions the relevant events are (a) the time just before the packet is being put on the wire and (b) the time after complete reception of the packet at the destination. Even though the RFC 2330 [i.2] definition of wire time delay has no direct equivalent in ITU-T terminology, it can be interpreted as "visibility of a given packet at both Measurement Points"; to be precise for the IETF the timing of the first and last bits of a packet on the wire is relevant, while ITU-T defines these events based on the moment when the packet crosses a point in the sender's or the recipient's IP stack (see note). One can state that one-way delay for both definitions yields the delay that one bit of the packet takes to travel across the network plus one time the serialisation delay of the IP packet, depending on packet length and physical link speed. From the ITU-T definition in ITU-T Recommendation Y.1540 [i.1] it is unclear what further delays inside the Measurement Points might add to the measured delay.

NOTE: However Y.1540 states "*The exact location of the IP Service measurement point in the IP protocol stack is for further study*"

In the IETF definition corrupted packets are deemed as lost, and their delay is formally defined as an undefined value of delay, which can informally be designated "infinity". In contrast to that the IP packet transport delay, defined in ITU-T Recommendation Y.1540 [i.1] applies to the combined set of successfully delivered and corrupted packets. Deciding whether a packet is corrupted or lost can be difficult to detect in certain circumstances, such as applying hashing methods to identify packets on reference and observation measurement points through an ID determined on the base of the packet content. Special care should be taken to detect corrupted packets: In case that packet header or payload fields are used in the measurement (e.g. serial packet id counter in active probing packets) then all checksums up to the respective header or payload have to be checked.

Both, RFC 3393 [i.9] and ITU-T Recommendation Y.1540 [i.1], express their metrics on populations that are conditioned on successful arrival within the waiting time, i.e. both definitions treat packets arriving after the waiting time as lost packets. It is important to select a reasonably safe threshold to decide when packets are deemed lost. It should be set so that no significant number of packets is observed as lost just because of being slightly late, i.e. a small amount over the threshold. Otherwise the derived traffic statistics would be biased. This threshold should be reported together with the measurement results. In case of fragmentation the time for the complete reception of all fragments counts. If not all fragments are received, a packet is deemed lost. A detailed comparison of packet treatment for one-way delay measurements is given in [i.19].

When computing statistics based on a series of delay values it is crucial to decide whether or not to include also the "lost" (lost or too much delayed) packets in this computation. Both standards assume that for statistics like median or other percentiles lost packets will be included (with a delay value of undefined or +infinity). For other statistics like mean delay this is not a useful approach since otherwise the average delay would be undefined or +infinity even if only one packet was lost in this measurement.

Apart from these details on the measurement of one-way delays and some differences in terminology there are no significant differences between the definitions of One Way Delay versus IP Packet Transfer Delay.

To measure one-way metrics it is needed to ensure a synchronisation of both transmission ends of the measurement device. This synchronisation may be done by GPS clocks when the two ends are distant. When ends are co-located synchronisation may be done directly by the analyser. To assess the metric, the clock accuracy of the analyzer should be better than 10 ppm.

4.1.4 Active Measurement Method

Active measurement of One-way Delay as defined by the IETF in RFC 2679 [i.4] requires the observation of test packets transmitted between two endpoints across a network, a "source" host, which sends the test packet, and a "destination" host, which receives the test packet. The One-way Delay is then calculated as the difference between the time at which the test packet is received at the destination and the time the test packet is sent.

The procedure to take a One-way Delay measurement involves the following steps:

1) The source host and the destination host are synchronized.

2) The destination host is prepared to receive the test packets sent by the source host.

3) The source host constructs a test packet. Any 'padding' portion of the packet needed only to make the test packet a given size should be filled with randomized bits to avoid a situation in which the measured delay is lower than it would otherwise be due to compression techniques along the path.

4) The source host places a timestamp in the prepared packet.

NOTE: Usually it also places a counter or identifier in the test packet, so that the destination host may also identify the number of lost test packets.

5) The source host sends the test packet to the destination host.

6) The destination host receives the test packet and records its arrival time.

7) The destination host calculates a One-way Delay by subtracting the two timestamps. If the delay between source host's timestamp and the actual sending of the packet is known, then the estimate could be adjusted by subtracting this amount. Similarly, if the delay between the actual complete receipt of the packet by the destination's network interface card and destination host's timestamp is known, then the estimate could be adjusted by subtracting this amount. If the packet fails to arrive within a reasonable period of time, the one-way delay is taken to be undefined.

8) The destination host optionally forwards the result back to the source, if the result is required at the source.

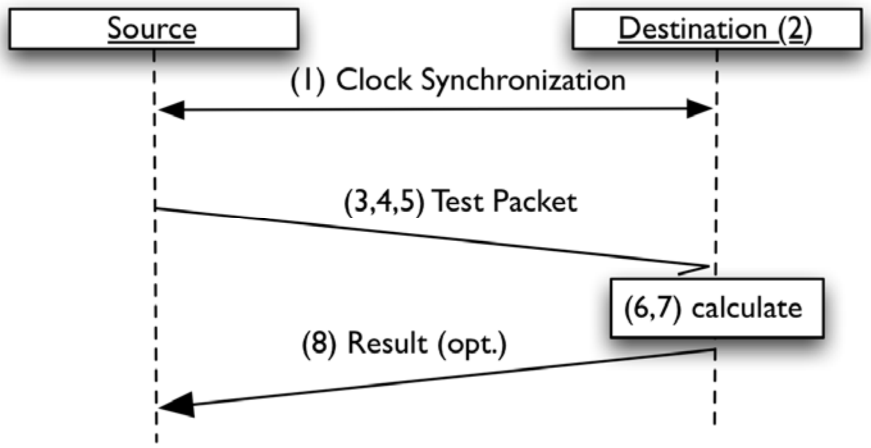

Figure 1: Active One-Way Delay Measurement Method

The IETF has defined the One-Way Active Measurement Protocol (OWAMP) [i.10] for the active measurement of the one-way delay metric. OWAMP consists of a control protocol to negotiate active performance measurement sessions, and a test protocol for transmission of actual test packets. OWAMP requires some minimal measurement infrastructure. It defines five entities, a Control Client, which specifies the test to run and passes that specification using the control protocol to a Server; this control interaction then causes a Session Sender and Session Receiver to send packets via the test protocol for the duration of the requested test. The Fetch Client then retrieves the test results from the Server. The test protocol allows accurate measurement of one-way delay without including processing overhead or other unknown timing components, and is designed to be resistant to detection and manipulation. In common usage, the Control Client, Fetch Client, and Session Sender are on one host, and the Server and Session Receiver are on another.

4.1.5 Passive Measurement Method

One-way delay may be measured passively by observing the same packet at two points in the network, and calculating the difference between the arrival times at these points. Calculating this metric requires the correlation of data from multiple Observation Points (OPs). The same packet is recognised at different Observation Points; this can be done by using invariant parts of the header or the payload or by calculating a packet identifier based on the invariant header fields and/or the packet content. Using the packet identifier reduces the amount of measurement data and can be easily obtained by calculating CRCs or hash functions. Proper clock synchronization between the observation points is important.

The passive procedure to take a One-way Delay measurement involves the following steps:

1) The source Observation Point and the destination Observation Point are synchronized.

2) The source and destination OPs are prepared to measure packets sent from the source host.

3) A source host sends a packet to the destination on a path where it is observable by source and destination OP.

4) The source OP observes this packet, records a timestamp of the observation, and identifies it.

5) The source OP forwards information about the observation event, i.e. the packet identification and timestamp to the destination OP. This operation may be performed in batches, after some specific number of packets have been observed or a given time interval has passed.

6) The destination OP recognises the packet and records its arrival time.

7) One-way delay is computed as the difference between the two arrival times. There is no strict need to perform this operation *at the destination OP*; a fifth host (not shown in figure 2) may also perform this operation when source and destination OP forward information about the observation event to it.

These steps are also shown in figure 2. The red dotted lines in the figure show the potential delay measurement error, which exists since there is a time difference between the source itself and the Observation point (OP), as well as between the destination and the destination OP. It is therefore recommended to keep the OP as close as possible to the source or destination, respectively.

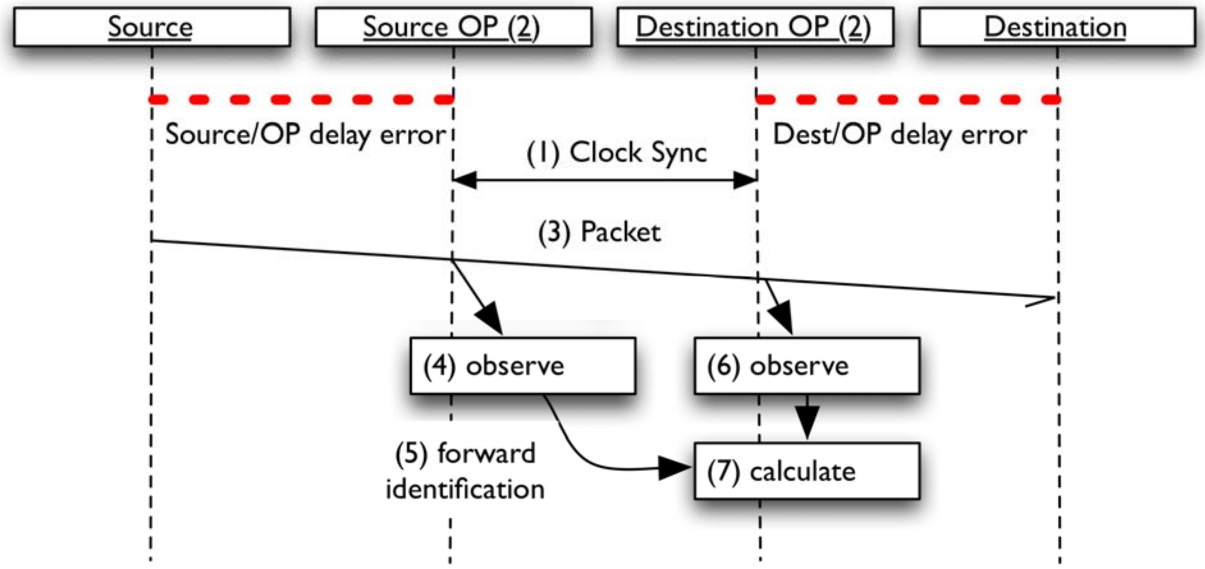

Figure 2: Passive one-way delay measurement method

Note that in contrast to active one-way delay monitoring the approach for passive one-way delay measurements does not include in the results the time for the serialisation delay due to the fact that both monitoring points are acting as packet receivers. If active and passive OWD measurement results are compared then this should be taken into account; however the serialization time becomes more and more negligible compared to the network delay the faster the net data rate of the underlying networking technology is.

The IETF IPFIX [i.12] or PSAMP protocols may be used for reporting packet contents, hashes, and timestamps from each passive observation point, as in [i.15]; each of these reports may be used to calculate a singleton in IPPM terminology or a single delay in ITU-T terminology.

4.2 Round Trip Delay

Round-trip delay is used to measure the expected time for network interaction between two hosts on a network; conceptually, it is equivalent to Delay in each direction between the two hosts. Round-trip delay is applicable to QoS for latency-sensitive protocols. The ITU-T has defined no specific round-trip delay metric, but its approach is generally compatible with the IETF's.

4.2.1 IETF Definition

RFC 2681 [i.6] defines round trip delay for a single datagram transmitted across the Internet, and, based on that, defines how this metric can be extended to produce a series of delay values to a Poisson distributed stream of sent packets. Finally it defines statistics based on the obtained series of round trip delay values such as Percentile, Median, Minimum, and Inverse-Percentile. It omits definitions for the Minimum and Average delay statistics; however these are straightforward to define. The definition assumes active probing, i.e. explicit transmission of test packets; this implies that there are no privacy considerations related to user data traffic. The IETF document also assesses the potential measurement errors and emphasises the need to note the level of uncertainty together with the results.

In IETF definitions the round trip delay (RTD) denotes only the wire time component (serialisation and transmission delay) of the total time required to send a packet back and forth. In contrast, the full round trip time (RTT) also includes the end system delay (ESD): RTT = RTD + ESD. NTP for example operates on basis of RTD, but the ping command returns the RTT (even though the ESD for an ICMP echo reply is usually very small). RFC 3611 [i.20] also notes the special case for symmetric paths:

"If the round trip delay is denoted, RTD and the end system delays associated with the two endpoints are ESD(A) and ESD(B) then: one way symmetric voice path delay = (RTD + ESD(A) + ESD(B)) / 2".

However, like noted in RFC 2681 [i.6] one has to be cautious in general when using RTT or RTD for estimating one-way delay because of potential incorrectness due to delay asymmetries.

4.2.2 ITU-T Definition

The ITU-T approach does not provide a direct definition of round trip delay for an IP packet. However it is suggested that the metric can be obtained as the sum of two IP packet transfer delays.

4.2.3 Comparison and Recommendations

With regard to round trip delay only IETF has issued a document specifically geared towards that metric. It defines the process, metric, statistics and issues to care when performing and evaluating a round-trip measurement. ITU-T work focuses on one-way delay (or more precisely IP packet transfer delay) and notes that two-way delay can be later constructed from two one-way delay measurements. We note here that this latter approach is viable but includes higher demands on the monitoring stations, especially for time synchronisation. Time synchronization is not strictly necessary with direct RTT measurements where one station only deflects the test packets or replies to them.

The statements made in the previous clause about one-way delay with respect to the treatment of long-delayed packets and their effect on delay statistics also apply to round-trip delay measurement.

4.2.4 Active Measurement Method

Active measurement of round-trip delay as defined by the IETF in RFC 2861 [i.6] requires the observation of test packets transmitted in both directions between two endpoints across a network, a "source" host, which sends the first packet, and a "destination" host, which receives the first test packet and sends a test packet back to the source in reply. The round-trip delay is then calculated as the difference between the time at which the reply is received at the source and the time at which the original test packet was sent. As with one-way delay, the measurement takes into account the delay between the transmission of the first bit of the forward packet and receipt of the last bit of the return packet.

These test packets are identifiable, so that the source can associate replies sent by the destination with test packets sent by the source. Each packet's identification may be associated with a timestamp and other timing information at the source, and may also contain timestamp information itself, in order to simplify delay calculations at the source. Of course, the identification system should allow multiple test packets and replies to be "in flight" at the same time in order to measure realistic delay measurement scenarios.

The procedure to take a round-trip measurement is then as follows:

1) The destination host is prepared to reply to test packets sent from the source host.

2) The source host constructs an identified test packet, and stores any information necessary to calculate a round-trip time from the packet's identification information. The test packet may have additional payload or header fields in order to simulate the properties of real traffic.

3) The source host sends the test packet to the destination host.

4) The destination host receives the test packet.

5) The destination host constructs a test reply packet containing the identification information back to the source host. The test reply packet may have additional payload or header fields in order to simulate the properties of real traffic. The delay here should be minimized; if it is known, it may be transmitted back to the source host, possibly in line with the reply packet, in order to allow the destination host processing delay to be taken into account in the round trip measurement.

6) The destination host sends the test reply packet back to the source host.

7) The source host receives the test reply packet from the destination host. If the source host does not receive a test reply packet in a reasonable amount of time, the packet is considered lost, and the round-trip delay is considered undefined.

8) The source host calculates a round-trip delay from the received test reply packet and any stored information linked to the test packet's identifying information. The source host calculates the delay by subtracting the time at which the source host sent the test packet from the time at which the source host received the test reply, if possible adjusted for processing delay at the source and destination hosts.

Steps 2 to 8 may be repeated to take multiple round-trip delay measurements.

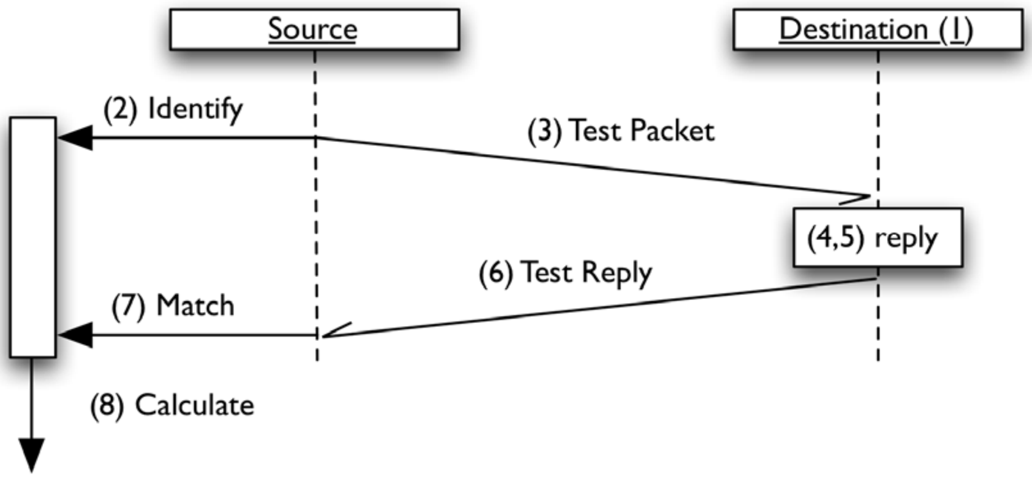

Figure 3: Active two-way delay measurement method

The IETF has defined the Two-Way Active Measurement Protocol (TWAMP) [i.23] for round trip delay measurement. TWAMP is essentially an extension of OWAMP for the IPPM round-trip delay metric. Like OWAMP, TWAMP consists of a control protocol to negotiate active performance measurement sessions, and a test protocol for transmission of actual test packets. TWAMP defines four entities, a Control Client, which specifies the test to run and passes that specification using the control protocol to a Server; this control interaction then causes a Session Sender and Session Reflector to exchange packets via the test protocol for the duration of the requested test. The test protocol allows accurate measurement of two-way delay without including processing overhead or other unknown timing components. In common usage, the Control Client and Session Sender are on one host, and the Server and Session Reflector are on another. There is no Fetch Client as in OWAMP, since the results are measurable at the Session Sender.

Note that the procedure defined in this clause is *approximated* by sending an ICMP Echo Request datagram from one host to another, and measuring the time to the receipt of the corresponding ICMP Echo Reply datagram; this approach is commonly called "pinging" after the common Unix utility which implements it. This approach is imprecise, as it does not account for the unknown processing delay at the remote host, or for different handling between ICMP packets and IP packets present within many routers, and as such is not recommended.

4.2.5 Passive Measurement Method

The procedure for passive measurement of round-trip delay is similar to those for active measurement: a packet is sent from a source to a destination, that packet causes the destination to send the packet back to the source, and the packets are identifiable at the source in order to correlate each packet of the round trip in order to calculate a delay. There are two potential architectures here; one utilizing a source Observation Point (OP) placed topologically close to the source of traffic, and one utilizing an additional destination OP placed topologically close to the destination of traffic.

The active measurement procedure in clause 4.2.4 can be modified as follows for single-OP passive two-way delay measurement:

1) The source OP is prepared to measure packets sent from the source host and replies sent to the source host.

2) The source host sends a packet to the destination.

3) The source OP observes this packet and identifies it, storing information necessary to recognize the reply packet and a timestamp it will use to calculate the delay. Identification information here consists of the source and destination addresses, transport layer protocol, transport layer ports (if applicable), and per-protocol application layer header information such as TCP sequence or acknowledgment numbers, DNS queries and answers, and so on. The source OP at this point begins waiting for the reply packet.

4) The destination host receives the packet sent from the source host and processes it. There is uncertainty in the amount of time that the destination host will wait before sending the reply which varies according to the transport and application layer information used to identify the reply in step 2. For example, for sequence and acknowledgment number matching, the destination host may wait up to 500 ms as per RFC 1122 [i.24] before sending an acknowledgment.

5) The destination host sends a reply to the packet sent by the source host.

6) The source OP observes and identifies the reply packet. If the source host does not receive a test reply packet in a reasonable amount of time, the packet is considered lost, and the round-trip delay is considered undefined.

7) The source OP calculates a round-trip delay from the received packet and reply packet. The source OP calculates the delay by subtracting the time at which the source OP received the packet from the time at which the source OP received the reply. The source OP estimates the reply delay at the destination host, as well as the delay between the source host and the source OP, in.

This method is illustrated in figure 4.

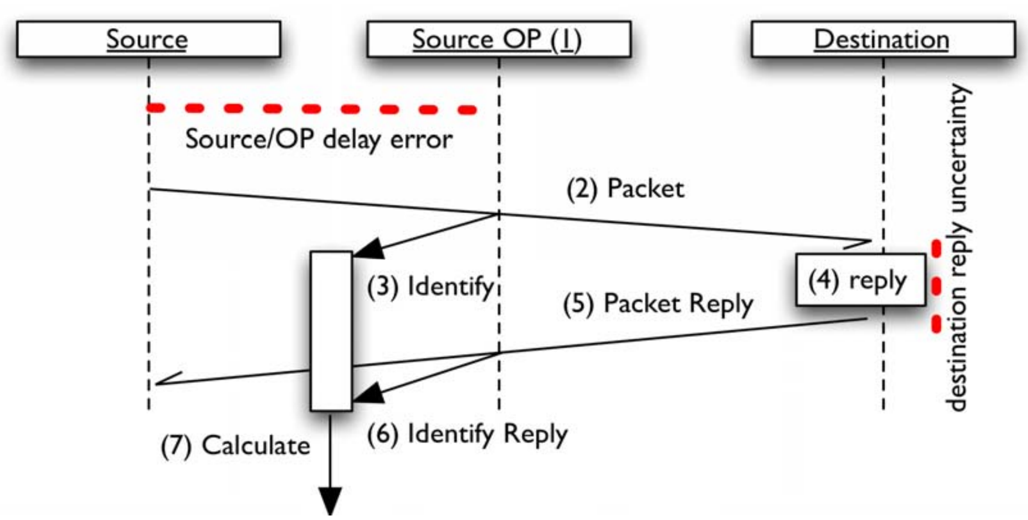

Figure 4: Single-OP variant passive two-way delay measurement method

Since this method suffers from large uncertainty in estimating the reply delay at the destination, it is only recommended in cases where the deployment of a single OP is desired. By adding a second OP close to the destination, the uncertainty in the delay at the destination host can be reduced. Here, both the source and destination OP identify the same packet and reply, and the destination OP sends the timing difference between the observation of the packet and the reply back to the source OP. While this method cannot differentiate reply delay from the network delay between the destination and the destination OP, it does allow the wire round trip time between the source OP and destination OP to be measured accurately.

This method is illustrated in figure 5.

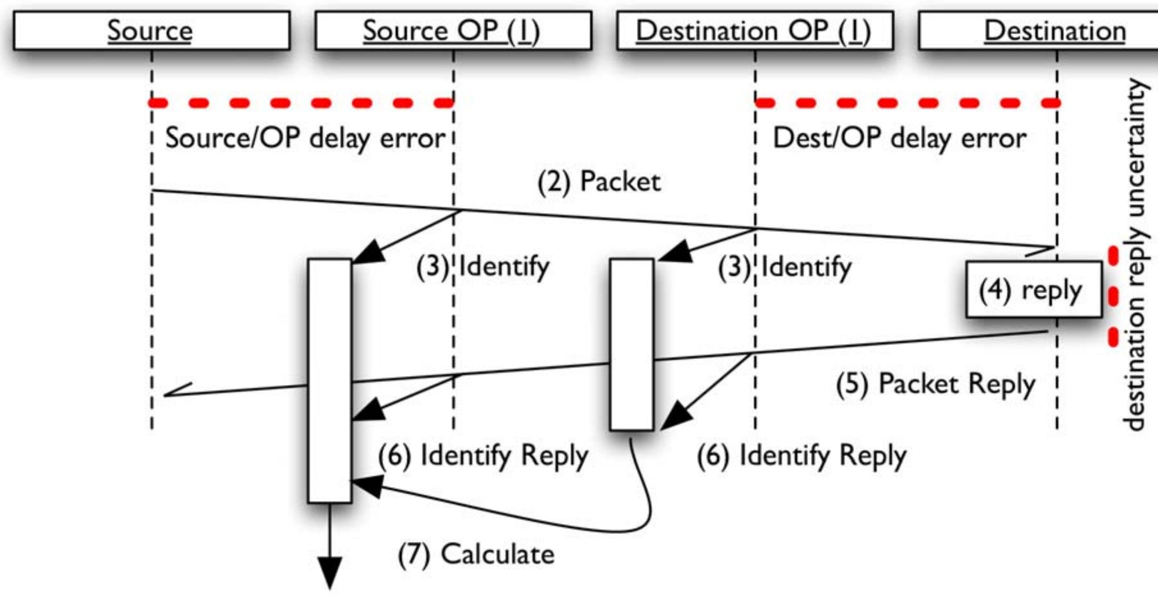

Figure 5: Dual-OP variant passive two-way delay measurement method

Note that since the ITU-T definition of round-trip delay is as a simple sum of one-way delays, round-trip delay as defined by ITU-T may also be passively measured by combining passively measured one-way delay metrics, as described in clause 4.1.5. Care should be taken with the second (or return) sample in each metric, as the upstream and downstream observation points switch roles for the reverse delay.

4.3 IP Packet Delay Variation vs. End-to-end 2-point IP Packet Delay Variation

Delay variation is used to measure the variation in a sequence of delay values over time, and is applicable to the QoS for protocols that presume or require relatively constant latency. The results of the ITU-T and IETF approaches here are not directly comparable; details are given in this clause.

4.3.1 IETF Definition

In RFC 3393 [i.9] (IP Packet Delay Variation Metric for IP Performance Metric) the **IP Packet Delay Variation** (IPDV) is presented as the difference between the One-Way-Delay two consecutive packets from a stream of selected packets across Internet paths. The definition relies on the introduction of a selection function ***F*** defining unambiguously the two packets from the stream selected for the metric.

More precisely, RFC 3393 [i.9] defines the One-way-IPDV from a source host to a destination host for two type-P packets selected by the selection function ***F*** as the difference between the value of the One-way-delay from the source host to destination host at time T2 and the value of the One-Way-Delay from the source host to the destination host at time T1. T1 is the wire arrival time at which the source host sent the first bit of the first packet, and T2 is the wire exit time at which the source host sent the first bit of the second packet. This metric is derived from the One-Way-Delay metric.

Formally, *the One-way-IPDV from a source host to a destination host at T1, T2 is equal to a real number ddT if the source host sent two packets, the first at wire-time T1 (first bit), and the second at wire-time T2 (first bit) and the packets were received by the destination host at wire-time dT1+T1 (last bit of the first packet), and at wire-time dT2+T2 (last bit of the second packet), and dT2-dT1=ddT* [i.9].

The One-way-IPDV from a source host to a destination host at T1, T2 is undefined if the source host sent the first bit of a packet at T1 and the first bit of a second packet at T2 and the destination host did not receive one or both packets.

Note that a singleton IP Packet Delay Variation measurement needs a stream of at least two packets. The purpose of the selection function is to specify exactly which two packets from the stream delimited by two interval endpoints are to be used for the singleton measurement. Examples of a selection function are:

- Consecutive packets within the specified interval.
- Packets with specified indices within the specified interval.
- Packets with the minimum and maximum one-way-delays within the specified interval.
- Packets with specified indices from the set of all defined (i.e. non-infinite) one-way-delays packets within the specified interval..

Figure 6 shows graphically how delay and IPDV relate to each other:

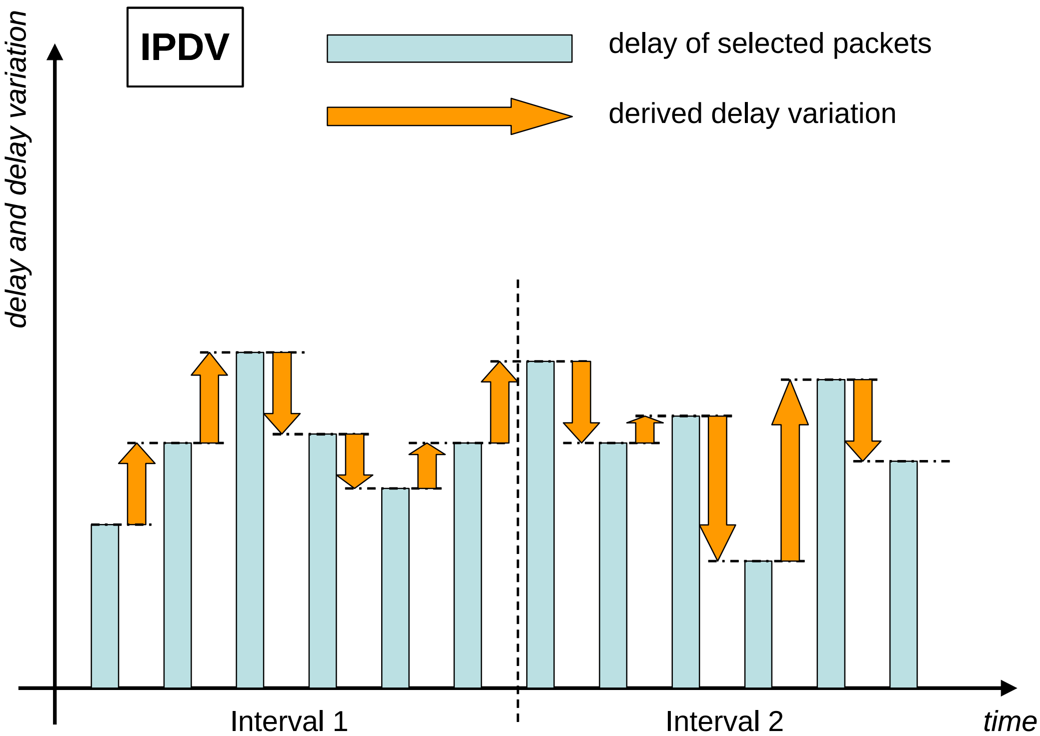

Figure 6: IPDV relation to packet delay

4.3.2 ITU-T Definition

In ITU-T Recommendation Y.1540 [i.1] (Internet Protocol Data Communication Service - IP Packet Transfer And Availability Performance Parameters) **end-to-end 2-point IP Packet Delay Variation** (IPDV) is defined based on the observations of corresponding IP packet arrivals at ingress and egress measurement points.

*"The end-to-end 2-point packet delay variation (v_k) for an IP packet **k** between a source host and a destination host is the difference between the absolute IP packet transfer delay (x_k) of the packet and a defined reference IP packet transfer delay, $d_{1,2}$, between the same measurement points: $v_k = x_k - d_{1,2}$ [i.1]".*

The reference IP packet transfer delay, $d_{1,2}$, between the source host and the destination host can be the absolute IP packet transfer delay experienced by the first IP packet between the two measurement points or any other fixed packet delay.

Positive values of end-to-end 2-point IPDV correspond to IP packet transfer delays greater than those experienced by the reference IP packet; negative values of 2-point IPDV correspond to IP packet transfer delays less than those experienced by the reference IP packet. The distribution of 2-point IPDVs is identical to the distribution of absolute IP packet transfer delays displaced by a constant value equal to $d_{1,2}$.

Therefore, the delay variation of an individual packet is naturally defined as the difference between the actual delay experienced by that packet and a reference delay. Use of the minimum delay as reference is now the primary definition of Packet Delay Variation in ITU-T Recommendation Y.1540 [i.1]. In that case the delay variation of all packets is zero or positive which simplifies the analysis of the delay variation range. Use of the average delay or the delay of the first packet of a measurement interval is possible but not recommended anymore. Figure 7 shows graphically how delay and PDV relate to each other when the minimum delay is used as a reference:

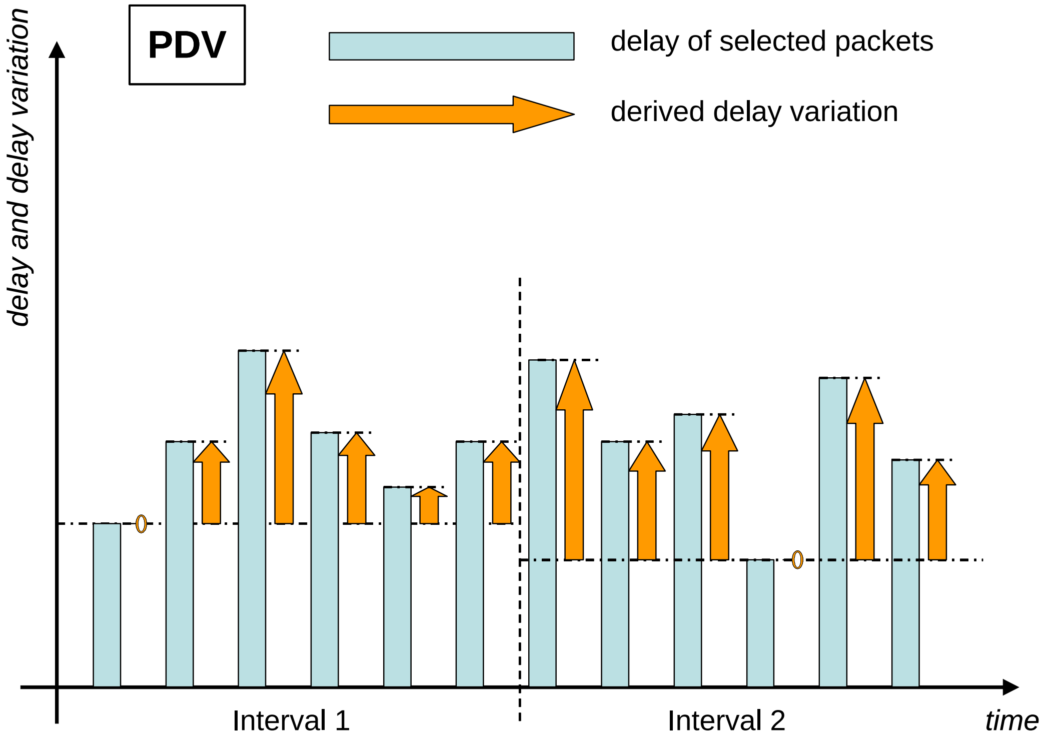

Figure 7: PDV relation to packet delay

4.3.3 Comparison and Recommendations

In RFC 3393 [i.9] it is assumed that the Type-P packet stream is generated according to a Poisson distribution.

The reason for Poisson sampling is that it ensures an unbiased and uniformly distributed sampling of times between the interval endpoints. Alternate sampling methodologies are possible. For example, continuous sampling of a constant bit rate stream (i.e. periodic packet transmission) is a possibility. However, in this case, it is necessary to avoid any "aliasing" effects that may occur with periodic samples, e.g. by varying the sampling "phase" for each measurement interval.

The measurement methodology illustrated in RFC 3393 [i.9] includes the determination of IP Packet Delay Variation in real-time as part of an active measurement. Therefore, IETF considers IPDV as a metric to be measured by using an active methodology.

Care should be taken with regard to delay variation measurements when it comes to packet size distributions, because one-way delay includes the serialization time of the packet (both IETF and ITU). Therefore, the delay of a longer packet would appear to vary from that of a short packet, simply because it takes more time to put the long packet on the wire. RFC 3393 [i.9] requires same-size packets for delay variation measurements.

Finally, it has to be noted that in the singleton metric of IP Packet Delay Variation, factors that affect the measurement are the same as those affecting the One-Way-Delay measurement:

- Errors/uncertainties due to uncertainties in the clocks of the source and destination hosts.
- Errors/uncertainties due to the difference between time measured at the host and the network observation point.

ITU-T Recommendation Y.1540 [i.1] also presents two alternative methods for summarising the delay variation experienced by a population of packets:

A first method is to pre-specify a delay variation interval and then observe the percentage of individual packet delay variations that fall inside and outside of that interval.

A second method consists of selecting upper and lower quantiles of the delay variation distribution and then measure the distance between those quantiles. It is suggested in ITU-T Recommendation Y.1540 [i.1] section 6.2.4.2 to select the 99,9th percentile and then 0th percentile (i.e. the minimum), make measurements, and observe the difference between the delay variation values at these two quantiles for each interval.

The two definitions of IP Packet Delay Variation vs. End-to-end 2-point IP Packet Delay Variation produce different statistics. While the *time series* values of IETF IPDV and ITU-T delay variation can be converted into one another, the *final statistics* based on these time series values cannot be converted into one another and are not directly comparable either. Which definition is more appropriate depends on the intended use of the results. IPDV is also influenced by the ordering of the different delay values inside an interval while PDV is not. On the other hand PDV also captures the maximum possible difference of delay in an interval while IPDV does not.

For example, the measurement of delay variation is recommended to size the play-out buffers employed by applications requiring continuous and regular delivery of packets (e.g. VoIP or multimedia streaming applications). Indeed, the size of such buffers should be determined based on the measurement of the maximum delay variation. To adjust the buffer size dynamically some dejitter buffers also use a kind of sliding average on top of the delay variation to avoid keeping large buffers for a long time: buffer_size = function(weight_factor * (current buffer size) + (1-weight_factor) * buffer_size_for_most recent delay). Furthermore, since changes in delay variation are often due to changes of the status of queues within a network, the measurement of delay variation metric can be used to understand the dynamics and the conditions of such queues.

4.3.4 Active Measurement Method

Since both the IETF and ITU-T definitions of delay variation use a sequence of one-way delay samples or measurements as the basis of the delay variation metric, delay variation may be measured by either method using the results of active one-way delay measurement as described in clause 4.1.4.

4.3.5 Passive Measurement Method

As with active measurement, since both the IETF and ITU-T definitions of delay variation use a sequence of one-way delay samples or measurements as the basis of the delay variation metric, delay variation may be measured by either method using the results of passive one-way delay measurement as described in clause 4.1.5.

4.4 One Way Packet Loss vs. IP Packet Loss Ratio

Packet loss denotes the expected probability that a transmitted packet will reach its destination, and is usually expressed as a percentage. It is widely applicable to most QoS measurements, as it is a general measurement for a network's packet delivery reliability. The IETF and ITU-T approaches are generally compatible, given assumptions about the measurement methods in use; details are given in this clause.

4.4.1 IETF Definition

RFC 2680 [i.5] provides the following definition for the One Way Packet Loss metric:

"*the One-way-Packet-Loss of a type-P packet from a source host to a destination host is 0 if the source host sent the first bit of the type-P packet to the destination host at wire-time T and the destination host received that packet. The One-way-Packet-Loss of a type-P packet from a source host to a destination host is equal to 1 if the source host sent the first bit of the type-P packet to the destination host at wire-time T and the destination host did not receive that packet.*"

RFC 2680 [i.5] also relates the definition of One-way-Packet-Loss to One-way-Delay. More precisely, it is stated that the One-way-Packet-Loss is equal to 0 exactly when the One-way-Delay assumes a finite value, whereas an infinite value of the One-way-Delay corresponds to a One-way-Packet-Loss equal to 1. Furthermore, the document outlines the need for the measurement methodology to include a way to distinguish between a packet loss and a very large, but finite, delay. As with delay measurements, the threshold time to consider a late packet lost should be reported as part of the measurement results.

4.4.2 ITU-T Definition

ITU-T Recommendation Y.1540 [i.1] defines the IP packet loss ratio (IPLR) as the ratio of total lost IP packet outcomes to total transmitted IP packets in a population of interest.

A lost packet outcome occurs when a single IP packet reference event at a permissible ingress measurement point results in a misdirected packet outcome or when some or all of the contents of that packet do not result in any IP reference event at any egress measurement point within a specified time interval.

A misdirected packet outcome occurs when a single IP packet reference event at a permissible ingress measurement point results in one (or more) corresponding reference event(s) at one (or more) egress measurement points, all within a specified time interval; and

1) the complete contents of the original packet observed at the ingress measurement point are included in the delivered packet(s); but

2) one or more of the egress measurement points where the corresponding reference events occur are not permissible egress measurement points.

4.4.3 Comparison and Recommendations

In RFC 2680 [i.5] describes a measurement methodology which explicitly relies on the use of test packets to measure One-way-Packet-Loss. ITU-T Recommendation Y.1540 [i.1], instead, mentions that measurements can be performed in active or passive fashion, but it does not specify in any way which approach should be taken. Furthermore, in RFC 2680 [i.5] the definition of the One-way-Packet-Loss-Poisson-Stream metric is given. A value of such metric is a sequence made up of (t, l) pairs, where t is a time belonging to a pseudo-random Poisson process of rate lambda, whose values fall between a beginning time T0 and a final time T1, and l is the value of the One-way-Packet-Loss at time t. Therefore, the Poisson process is used to schedule the packet sending times. Note that due to the influence of the network, test packets generally do not arrive at the destination host according to a Poisson distribution.

4.4.4 Active Measurement Method

One-way active packet loss measurements always include the sending of test packets from a sender to a packet receiver in a predefined sending schedule and/or pattern; the receiver will then count these packets. Usually the test packets contain a specific content including sequence numbers and timestamp so that analysis by the receiver can be more detailed. A sending schedule usually consists of an average rate and a distribution, e.g. a test sends 10 Packets per second on average, following a Poisson distribution for timing of the packet sending. These test specifications are part of the measurement description, in order to allow for correct interpretation of the results.

Active packet loss measurements should ensure that test packets are treated like the background traffic of interest, e.g. are sent in the same QoS class in case QoS differentiation is in effect in the network. In particular, ICMP echo request and reply packets are bad for active measurements because ICMP is often treated in IP routers in a separate processing queue or in software instead of forwarding hardware.

Packet loss as defined by RFC 2680 [i.5] may be actively measured as described in clause 4.1.4. OWAMP has built-in support for measurement of the One Way Packet Loss metric. For this purpose the information from inside the OWAMP test packets (sequence number and timestamps) is used. The test setup and configuration needs to be noted alongside the test results for correct interpretation and repetitiveness, e.g. the timeout value after which "late" arriving packets are deemed lost needs to be noted because it affects the measured results to some degree.

In any case, note that the measured loss is only accurate for the test traffic itself. Stating that this packet loss would also apply to the background traffic of the same "type" (traffic class, packet length, etc.) is applicable but introduces a level of uncertainty on the statement, which can be shrunken only with more test packets. However, the rate and volume of test traffic should be chosen so that the influence on background traffic by this additional traffic is negligible. Therefore a suitable compromise should be found for a given network and test environment.

4.4.5 Passive Measurement Method

Passive measurements of one-way packet loss are only possible if through multipoint measurement, or by observing traffic from which packet loss can be deduced from higher-layer header fields. In both cases it is important to ensure that all of the packets of interest do pass the observation point(s). In the first approach, routers or measurement probes along the path of the traffic of interest are instructed to filter and count the packets. Periodically, the probes export counters per flow, and differences along the path can be compared in order to compute and localize the packet loss.

The second approach is possible with a single point measurement, but only for protocols whose packet headers have certain fields allowing the deduction of skipped, lost, or retransmitted packets. The TCP sequence number, or the sequence number from applications such as RTP, can be used to in this way. However, there are two potential sources of inaccuracy with this method. First, these measurements can only reflect the packet loss observed by the measurement probe or between probes and may not reflect accurately the loss observed by the end systems of the traffic under observation because these loss probes are situated in the network and not directly next to the traffic sender and receiver. Note also that the required assumptions about the protocols under observation may not hold due to variations among protocol implementations.

Passive packet loss measurement may be measured by measuring one-way delay as in clause 4.1.5, and taking any packet with reasonably infinite delay to be lost.

4.5 Connectivity vs. IP Service Availability

Connectivity or service availability is a measure of the expected probability that one host can reach another, and is widely applicable to most QoS measurements, as it is a general measurement for a network's session transport reliability. The IETF and ITU-T approaches are generally compatible, given assumptions about the measurement methods in use; details are given in this clause.

4.5.1 IETF Definition

RFC 2678 [i.3] defines metrics which allow to define the level of connectivity between two hosts A1 and A2. It starts with the definition of an analytic metric, called Type-P-Instantaneous-Unidirectional-Connectivity, to define one-way connectivity at one single moment in time. This metric is extended to the notion of bi-directional connectivity (for one single moment in time), and later to connectivity within a certain period of time. Each metric delivers as a result a Boolean value which states whether the desired level of connectivity has been reached or not. A methodology for estimating the last defined metric, called Type-P1-P2-Interval-Temporal-Connectivity is sketched in RFC 2678 [i.3]. This methodology expects randomly distributed start times for sending the probe packets within a selected interval of time. RFC 2678 [i.3] also shows how to apply this testing methodology for checking TCP-based connectivity.

The following definitions are given in [i.3]:

*"**Type-P-Instantaneous-Unidirectional-Connectivity**: Src has *Type-P-Instantaneous-Unidirectional-Connectivity* to Dst at time T if a type-P packet transmitted from Src to Dst at time T will arrive at Dst."*

*"**Type-P-Instantaneous-Bidirectional-Connectivity**: Addresses A1 and A2 have *Type-P-Instantaneous-Bidirectional- Connectivity* at time T if address A1 has Type-P-Instantaneous-Unidirectional-Connectivity to address A2 and address A2 has Type-P-Instantaneous-Unidirectional-Connectivity to address A1."*

*"**Type-P-Interval-Unidirectional-Connectivity**: Address Src has *Type-P-Interval-Unidirectional-Connectivity* to address Dst during the interval [T, T+dT] if for some T' within [T, T+dT] it has Type-P-instantaneous-connectivity to Dst."*

*"**Type-P-Interval-Bidirectional-Connectivity**: Addresses A1 and A2 have *Type-P-Interval-Bidirectional-Connectivity* between them during the interval [T, T+dT] if address A1 has Type-P-Interval-Unidirectional-Connectivity to address A2 during the interval and address A2 has Type-P-Interval-Unidirectional-Connectivity to address A1 during the interval"*

*"**Type-P1-P2-Interval-Temporal-Connectivity:** Address Src has *Type-P1-P2-Interval-Temporal-Connectivity* to address Dst during the interval [T, T+dT] if there exist times T1 and T2, and time intervals dT1 and dT2, such that:*

- *T1, T1+dT1, T2, T2+dT2 are all in [T, T+dT].*
- *T1+dT1 <= T2.*
- *At time T1, Src has Type-P1 instantaneous connectivity to Dst.*
- *At time T2, Dst has Type-P2 instantaneous connectivity to Src.*
- *dT1 is the time taken for a Type-P1 packet sent by Src at time T1 to arrive at Dst.*
- *dT2 is the time taken for a Type-P2 packet sent by Dst at time T2 to arrive at Src.*

4.5.2 ITU-T Definition

"IP Service Availability" is defined in chapter 7 of ITU-T Recommendation Y.1540 [i.1]. It is based on the notion of IP packet loss ratio (IPLR), measured as a percentage of lost packets across an interval of time. If during this time interval the loss percentage is below a certain selected level c_1 then the service is called available in this time slot, else unavailable. Based on the availability results (yes/no) for consecutive time intervals the document defines the percentage of (un)availability:

"percent IP service unavailability (PIU): The percentage of total scheduled IP service time (the percentage of T_{av} intervals) that is (are) categorized as unavailable using the IP service availability function.

percent IP service availability (PIA): The percentage of total scheduled IP service time (the percentage of T_{av} intervals) that is (are) categorized as available using the IP service availability function: PIA = 100 - PIU".

ITU-T Recommendation Y.1540 [i.1] suggests 75 % as the provisional value for c_1. Other values (90 %, 99 %) are also considered. The time interval, T_{av}, is set provisionally to 5 minutes. Appendix VII of ITU-T Recommendation Y.1540 [i.1] discusses the rationale behind these values. Availability here is foremost defined as being unidirectional, but a bidirectional definition can be easily derived from that definition. In addition it is stated that "*The quantitative relationship between end-to-end IP service availability and the IP service availability of the basic section or NSE remains for further study*". This means that the experimenter still has to decide by himself which threshold to reasonably choose for his application so that "IP service availability = true" means an acceptable end-to-end service quality. In any case the source and destination of the test packets as well as the test packets themselves (type, length, and content) should be documented alongside the measurement results for proper interpretation.

4.5.3 Comparison and Recommendations

The IETF document RFC 2678 [i.3] defines analytical for connectivity as the successful transmission of *a single packet* at a given time or within a given interval. Also bidirectional metrics are defined based on a successful packet transmission in each direction. However the draft does not specify any derived metrics which could be calculated from "have/don't have connectivity" results across many measurement intervals, e.g. over the course of a day.

ITU-T Recommendation Y.1540 [i.1], chapter 7 defines "service availability" within a single interval as having an IP loss rate (IPLR) in this interval below a fixed predefined threshold c_1. This threshold should be carefully selected with the intended service (e.g. speech transmission with coded X) in mind. In contract to the IETF they also define a derived metric, called "percent IP service (un)availability", which denotes the portion of measurement intervals where the IPLR was below (or above) the threshold. This metric can be used to give a differentiated value (not only False or True) for service availability over a longer period of time, comprised of multiple measurement intervals.

Neither of the two documents specifies any suggested size for the duration of a measurement interval, nor do they suggest a probing rate or packet type to do so.

The IETF definition suffices to make a statement about the possibility of any data transmission between source host and destination host at all (i.e. IP routing works, no blocking of the test traffic by a firewall), with a potential constraint on the upper limit of transmission time. To this respect the defined connectivity metrics can only serve as a notion of "reachability", but not of suitability of the connection for any specific application service between the two hosts.

The ITU-T definition approaches this metric from an application-level view. It defines "IP service availability" solely based on an acceptable level of IP packet loss. This does not take into account other application constraints for general usefulness, e.g. enough available bandwidth between source and destination host. Therefore the defined metric cannot guarantee useful service availability.

In comparison the ITU-T definition gives a more differentiated view on the usefulness of the transmission capabilities between two hosts, and it also defines derived metrics (percent IP service availability) that allows making statements for long term monitoring of a service. Two-way service availability is not defined by ITU-T but can be easily derived.

It is recommended to further study the extension towards two-way service availability for the ITU-T and to study the applicability of the ITU-T's availability threshold approach to the IETF metrics.

4.5.4 Active Measurement Method

As ITU-T IP Service Availability is defined in terms of packet loss ratio, it can be determined actively by measuring one-way packet loss using OWAMP, and then calculating availability as defined in clause 4.5.2. For RFC 2678 [i.3] Connectivity, instantaneous unidirectional connectivity can be measured using OWAMP, and instantaneous bidirectional connectivity via TWAMP. The interval connectivity metrics can then be calculated from the instantaneous metrics.

4.5.5 Passive Measurement Method

As ITU-T IP Service Availability is defined in terms of packet loss ratio, it can be determined passively by measuring packet loss at in clause 4.4.5, and then calculating availability as defined in clause 4.5.2. For RFC 2678 [i.3] Connectivity, a passive measurement method as for delay is necessary, with observation points near each endpoint to measure connectivity between. Instantaneous unidirectional connectivity can be measured by observing a single packet passing in one direction, and instantaneous bidirectional connectivity by observing a single packet in each direction. The interval connectivity metrics can then be calculated from the instantaneous ones.

5 Other Metrics

The following metrics are also in common use in the measurement of network QoS.

5.1 Data and Packet Volume

The counting of data and packet volume as the most basic monitoring task is not standardized due to its simplicity. The only pitfall could be the question whether to count the number of bytes for each packet starting from the physical layer, MAC/LCC layer, or network (IP) layer. Usually in accounting for IP based networks only the length of the IP datagram is counted (i.e. volume of data transmitted on the network level), because the length of the layer two header may vary between two different monitoring stations on the path of a packet flow and therefore the observed data volume on the MAC/LCC layer would be different. However for network planning it is more useful to also include the layer two header bytes in the accounted data volume. Counting even the physical layer extra bits for frame synchronization (e.g. Ethernet preamble) is extremely uncommon. The IETF does only focus on accounting of packet volume starting from the network level, i.e. the IP datagram size (IP header plus payload) is used. The IETF has also standardized many aspects of how to handle (store, transmit, access, secure) accounting data in various RFC documents.

5.2 Packet Reordering

Packet reordering happens when datagrams arrive at a destination in reverse order of their initial transmission at the sender. One can distinguish between reordering of packets inside one flow and among all flows between the sender and the destination of traffic in general. Reordering inside a flow usually is a sign of an error in packet forwarding firmware, hardware, or the IP stack of sender or receiver of the traffic. Reordering of packets between flows is usually a sign of an active load balancing or features like e.g. equal cost multipath routing being active in the network. The IETF standards track document RFC 4737 [i.11] defines analytical metrics to evaluate whether a network has maintained or changed packet order. These metrics allows making a qualitative and quantitative statement about the level of packet reordering. Different possible metrics for the definition are presented and compared. Additionally defined metrics quantify the frequency of reordering and the distance between separate occurrences.

5.3 Bandwidth Capacity, Available Bandwidth, and Utilization

These metrics define network-level link and path attributes that describe how much data the link/path provides by its technical structure and how much of that (absolute or relative) measure is not yet in use at a moment in time. The IETF document RFC 5136 [i.21] provides definitions for the terms "Capacity" and "Available Capacity" related to IP traffic travelling between a source and destination in an IP network. It provides a unified nomenclature to make it easier to properly build, test, and use techniques and tools built around these constructs which produce comparable results when applied to the same network situation.

The document provides a common framework for the discussion and analysis of current and future estimation techniques for the above metrics. Specifically the document defines the following metrics: "Link Capacity", "Path Capacity" (for traffic across multiple connected links), "Link Usage" (referring to actual traffic in an interval of time), "Link Utilization" (fraction of capacity used), "Available Link Capacity" (fraction of capacity unused), and "Available Path Capacity" (minimal free capacity along multiple links). Also a comparison to the Bulk Transfer Capacity (BTC) metric is given in the document.

Although utilization is not a QoS metric by itself, it has a notable correlation with the QoS metrics described in clause 4 and, if it rises towards the bottleneck bandwidth of a set of links, may be used as an indicator for the need to monitor QoS in terms of delay more closely, since excessive delays and/or packet loss may occur more frequently.

5.4 Bulk Transport Capacity

Bulk Transport Capacity (BTC) is a measure of the ability of a network to transfer large amounts of data using a single congestion-aware transport connection (e.g. TCP) with typical applications such as FTP and HTTP-downloads or peer-to-peer data transfer operations in mind. The intuitive result of a BTC measurement is the expected long term average data rate (bits per second) of a single ideal TCP implementation along the tested network path which an application can achieve.

For the BTC metric the RFC 3148 [i.7] document defines "a framework for standardizing multiple BTC metrics that parallel the permitted transport diversity." This approach is needed because the IETF specification for congestion control (RFC 2581 [i.22]) permits considerable diversity for the practical realisation of congestion control algorithms for TCP. RFC 3148 [i.7] helps to compute comparable results for BTC from different TCP implementations. In addition the time of day, day of the week, and effects from using traffic differentiation (e.g. using integrated services (RFC 1633 [i.26] and RFC 2216 [i.27]), or differentiated services) in the network should be taken into account when evaluating BTC measurement results. BTC is also mentioned in RFC 5136 [i.21] in section 3.5 "Comparison to Bulk Transfer Capacity (BTC)", where the document provides notes on the difference between the definitions of BTC and the other capacity metrics.

5.5 Loss Patterns

The document RFC 3357 [i.8] entitled "One-way Loss Pattern Sample Metrics" targets to define metrics and their implications for differing patterns of packet loss observed. It defines "two derived metrics "loss distance" and "loss period", and the associated statistics that together capture loss patterns experienced by packet streams on the Internet." The document at first defines characteristics like bursty loss, loss distance and loss period, and based on these, two new metrics called Type-P-One-Way-Loss-Distance-Stream and Type-P-One-Way-Loss-Period-Stream. These metrics strongly reflect also the distribution and burstiness of the packet losses, in contrast to the basic loss definition from RFC 2680 [i.5].

In addition IETF draft "draft-venna-ippm-app-loss-metrics-01", which is a work-in-progress document, targets the definition of application level loss metrics, mainly for audiovisual applications using the Real-Time Transport Protocol (RTP), e.g. video streaming or Voice over IP (VoIP) applications. It defines:

"*two new metrics called Type-P-One-Way-Complete-Frame-Loss and Type-P-One-Way-Partial-Frame-Received to provide a better understanding of the affects of packet loss at the application level. The statistic Type-P-One-Way-Errored-Seconds is derived from the above metrics to compute the effect of packet loss at the application level.*"

5.6 RTCP reported metrics

RFC 3550 [i.25] illustrates the format of the Sender Report RTCP Packet, which conveys the following application level metrics:

- **sender's packet count**, which measures the total number of RTP packets transmitted by the sender since starting transmission up until the time the sender report packet was generated;
- **sender's octet count,** which refers to the total number of payload octets transmitted in RTP data packets by the sender since starting transmission up until the time this sender report packet was generated;
- **fraction lost**, which is the fraction of RTP data packets from a specific source lost since the previous sender report or receiver report was sent;
- **cumulative number of packet lost**, which represents the total number of RTP data packets from a specific source that have been lost since the beginning of reception. More precisely, this is the number of packets expected less the number of packets actually received, where the number of packets received includes any which are late or duplicates;
- **interarrival jitter**, which is defined to be the mean deviation of the difference in packet spacing at the receiver compared to the sender for a pair of packets. In other terms, the interarrival jitter is the difference in the relative transit time for two packets, where the relative transit time is the difference between a packet's RTP timestamp and the receiver's clock at the time of arrival, measured in the same units. If S_i is the RTP timestamp from packet *i*, and R_i is the time of arrival in RTP timestamps units for packet *i*, then for two packets *i* and *j* the interarrival jitter *D* can be expresses as:

$$D(i,j) = (R_j - R_i) - (S_j - S_i) = (R_j - S_j) - (R_i - S_i)$$

- delay since last Sender Report, which is the delay expressed in units of 1/65 536 seconds between receiving the last Sender Report packet from a specific source and sending the reception report.

Other and more detailed application level stream metrics are presented in RFC 3611 [i.20], where the Extended Report (XR) packet type for the RTP Control Protocol (RTCP) is defined. An XR packet consists of a header of two 32-bit words, followed by a number, possibly zero, of extended report blocks. The Statistics Summary Report Block contains statistics beyond the information carried in the standard RTCP packet format. In particular, the following metrics are reported:

- number of lost packets;
- number of duplicate packets;
- minimum jitter, which is the minimum relative transit time between two packets. All jitter values are measured as the difference between a packet's RTP timestamp and the receiver's clock at the time of arrival, measured in the same units;
- maximum jitter, which is the maximum relative transit time between two packets;
- mean jitter, which refers to the mean relative transit time between each two packet series;
- jitter deviation, which is the standard deviation of the relative transit time between each two packet series.

The VoIP Metrics Report Block provides metrics for monitoring voice over IP (VoIP) calls. These metrics include:

- loss rate, i.e. the fraction of RTP data packets from the source lost since the beginning of reception;
- discard rate, i.e. the fraction of RTP data packets from the source that have been discarded since the beginning of reception, due to late or early arrival, under-run or overflow at the receiving jitter buffer;
- round trip delay, which is the round trip time between RTP interfaces;
- end system delay, which is defined as the sum of the total sample accumulation and encoding delay associated with the sending direction and the jitter buffer, decoding, and playout buffer delay associated with the receiving direction.

6 Overview of Network Performance Relevant Standard Bodies and Working Groups

This clause summarises other network performance measurement related standards; they are briefly described and the references to the relevant documents are provided. This clause does not include IETF or ITU-T standard documents that describe metrics or measurement methods directly relevant to the scope of this current document as these have already been discussed above.

6.1 IETF

6.1.1 IPPM (IP Performance Metrics) Working Group

The aim of the IPPM WG is to develop a set of standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services. These metrics are designed such that they can be performed by network operators, end users, or independent testing groups. It is important that the metrics not represent a value judgment (i.e. define "good" and "bad"), but rather provide unbiased quantitative measures of performance.

The WG aims at producing a protocol to enable communication among test equipment that implements the one-way metrics. The intent is to create a protocol that provides a base level of functionality that allows different manufacturers' equipment that implements the metrics according to a standard to interoperate. The WG will also produce a MIB to retrieve the results of IPPM metrics, such as one-way delay and loss, to facilitate the communication of metrics to existing network management systems.

6.1.2 IPFIX (IP Flow Information eXport) Working Group

The IPFIX protocol [i.12] specifies how to export traffic flow information out of routers, network measurement probes or other devices for further processing by applications located on other devices. One of the target applications of the IPFIX protocol is QoS monitoring, intended as the non-intrusive observation of the transmission quality for single flows or traffic aggregates in the network for e.g. the validation of QoS guarantees in service level agreements (SLAs) [i.15].

In the IPFIX protocol, templates contain { type, length } pairs specifying which { value } fields are present in data records conforming to the template, giving great flexibility as to what data is transmitted. Since templates are sent very infrequently compared with data records, this results in a significant bandwidth saving. Different data records may be transmitted simply by sending new templates specifying the { type, length } pairs for the new data format Data records and templates are carried via a congestion-aware transport protocol from IPFIX exporting processes to IPFIX collecting process.

The IPFIX Information Model [i.14] defines a large number of standard Information Elements which provide the necessary { type } information for templates. The use of standard elements enables interoperability between different vendor's implementations. The list of standard elements may be extended in future. Additionally, non-interoperable enterprise specific elements may be defined for private use.

The architecture for the export of measured IP flow information from an IPFIX exporter to a collector is defined in [i.13], following the use cases given in [i.18]. Finally [i.15] describes what type of applications can use the IPFIX protocol and how they can use the information provided. It furthermore shows how the IPFIX framework relates to other architectures and frameworks.

6.1.3 PSAMP (Packet SAMPling) Working Group

The IETF Packet Sampling (PSAMP) working group was chartered "to define a standard set of capabilities for network elements to sample subsets of packets by statistical and other methods.". Their goals include devising sampling methods and algorithms and exporting capabilities for those. These operations should be applicable at line-speed and are targeted to support a wide range of measurement-based applications.

The work of this group includes the specification of operations to perform the packet selection (either statistical or deterministic), to specify per algorithm what parameters need to be configured and reported, and specify the appropriate protocols to do so. The PASMP WG works in close corporation with the IPFIX WG and has, after some thoroughly consideration, chosen to adopt the IPFIX protocol for reporting of sampling configuration data (i.e. the chosen algorithm and the selected parameters). For this purpose the PSAMP WG provides new IPFIX information elements in its information model.

The final goal is to have a set of well-understood sampling methods, and a means to perform and report sampling results to applications either co-located with or remote to the sampling network element. In addition the WG works on defining a packet sampler MIB to be placed at the network element, which includes parameters for packet selection, for packet report and stream format, and for export. This MIB allows remote entities to configure the sampler and observe its operation by writing/reading the MIB.

6.2 ITU-T

6.2.1 Study Group 12 (Performance and quality of service)

In the ITU-T, Study Group 12 is the Lead SG on Performance and Quality of Service (QoS), a role that is increasingly important with the advent of commercial VoIP and packet-based next generation networks and terminals.

With customers expecting the QoS of traditional communication services, it is crucial to be able to measure new parameters such as packet loss and jitter, and know their user impact. Thus, recent SG12 achievements include several new and revised standards on the planning and deployment of IP-based networks.

Current SG12 hot topics include the development of software tools that allow the modelling of potential network/terminal configurations and the prediction of the user impact of associated impairments. A model for voice quality prediction has been developed, while work is in progress for models for wideband speech and multimedia. The SG has also started work in other areas needing QoS guidance, such as hands-free communications in vehicles, and services based on speech technology.

SG12 is responsible for Recommendations on the end-to-end transmission performance of terminals and networks, in relation to the perceived quality and acceptance by users of text, data, speech, and multi-media applications. Although this work includes the related transmission implications of all networks and all telecommunication terminals, a special focus is given to IP QoS, interoperability and implications for NGN, and also includes work on performance and resource management.

6.2.2 Study Group 15 (Optical and other transport network infrastructures)

Study Group 15 is the home of standards for DSL. It also works on optical access and backbone technologies.

A key concern for many operators is to maximize network capacity in the last-mile (between the exchange and the customer premises). SG 15 standards on DSL are one way of helping towards this goal.

Less familiar, but equally as important are SG15 standards (ITU-T Recommendations) relating to Passive Optical Networks (PONs). PONs are a effective way of implementing fibre to the home/building etc and a crucial step towards all-optical networks. SG 15 also plays a leading role in the development of standards for the backbone including the key standard for synchronous data transmission over fibre optic networks, Synchronous Digital Hierarchy (SDH). Additionally SG 15 is working on Automatically Switched Optical Network (ASON). ASON provides quick and reliable service activation to service platforms like switches and routers.

The focus of SG15 is on the development of standards on optical and other transport network infrastructures, systems, equipment, optical fibres, and the corresponding control plane technologies to enable the evolution toward intelligent transport networks. This encompasses the development of related standards for the customer premises, access, metropolitan and long haul sections of communication networks.

Annex A: Bibliography

- ITU-T Recommendation Y.1541: "Internet protocol aspects - Quality of service and network performance Network performance objectives for IP-based services".
- IETF RFC 3432: "Network performance measurement with periodic streams". V. Raisanen, G. Grotefeld, A. Morton. November 2002.
- IETF RFC 3763: "One-way Active Measurement Protocol (OWAMP) Requirements". S. Shalunov, B. Teitelbaum. April 2004.
- IETF RFC 4148: "IP Performance Metrics (IPPM) Metrics Registry". E. Stephan. August 2005.
- ITU-T Recommendation G.810 (08/96): "Definitions and terminology for synchronization networks".
- ITU-T Recommendation I.356: "B-ISDN ATM layer cell transfer performance".
- Baek-Young Choi, Sue Moon, Rene Cruz, Zhi-Li Zhang, Christophe Diot, "Practical delay monitoring for ISPs", Proceedings of the 2005 ACM conference on Emerging network experiment and technology, October 24-27, 2005, Toulouse, France.
- I. Bilinskis and A. Mikelsons, "Randomized Signal Processing", Prentice Hall International, 1992.
- Papoulis, A. "Bernoulli Trials." §3-2 in Probability, Random Variables, and Stochastic Processes, 2nd ed. New York: McGraw-Hill, pp. 57-63, 1984.
- K. Claffy, G. Polyzos, and H-W. Braun, "Application of Sampling Methodologies to Network Traffic Characterization," Proc. SIGCOMM '93, pp. 194-203, San Francisco, September 1993.

History

Document history		
V1.1.1	December 2009	Publication
V1.1.2	July 2010	Publication