

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 20, 2008

A. Morton
G. Ramachandran
G. Maguluri
AT&T Labs
May 19, 2008

Reporting Metrics: Different Points of View
draft-morton-ippm-reporting-metrics-05

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on November 20, 2008.

Abstract

Consumers of IP network performance metrics have many different uses in mind. This memo categorizes the different audience points of view. It describes how the categories affect the selection of metric parameters and options when seeking info that serves their needs. The memo then proceeds to discuss "long-term" reporting considerations (e.g, days, weeks or months, as opposed to 10 seconds).

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",

"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Table of Contents

1.	Introduction	3
2.	Purpose and Scope	3
3.	Effect of POV on the Loss Metric	4
3.1.	Loss Threshold	4
3.1.1.	Network Characterization	4
3.1.2.	Application Performance	6
3.2.	Errored Packet Designation	6
3.3.	Causes of Lost Packets	6
3.4.	Summary for Loss	7
4.	Effect of POV on the Delay Metric	7
4.1.	Treatment of Lost Packets	7
4.1.1.	Application Performance	7
4.1.2.	Network Characterization	8
4.1.3.	Delay Variation	9
4.1.4.	Reordering	10
4.2.	Preferred Statistics	10
4.3.	Summary for Delay	11
5.	Test Streams and Sample Size	11
5.1.	Test Stream Characteristics	11
5.2.	Sample Size	11
6.	Reporting Results	12
6.1.	Overview of Metric Statistics	12
6.2.	Long-Term Reporting Considerations	13
7.	IANA Considerations	14
8.	Security Considerations	14
9.	Acknowledgements	14
10.	References	15
10.1.	Normative References	15
10.2.	Informative References	15
	Authors' Addresses	16
	Intellectual Property and Copyright Statements	17

1. Introduction

When designing measurements of IP networks and presenting the results, knowledge of the audience is a key consideration. To present a useful and relevant portrait of network conditions, one must answer the following question:

"How will the results be used?"

There are two main audience categories:

1. Network Characterization - describes conditions in an IP network for quality assurance, troubleshooting, modeling, etc. The point-of-view looks inward, toward the network, and the consumer intends their actions there.
2. Application Performance Estimation - describes the network conditions in a way that facilitates determining affects on user applications, and ultimately the users themselves. This point-of-view looks outward, toward the user(s), accepting the network as-is. This consumer intends to estimate a network-dependent aspect of performance, or design some aspect of an application's accommodation of the network. (These are **not** application metrics, they are defined at the IP layer.)

This memo considers how these different points-of-view affect both the measurement design (parameters and options of the metrics) and statistics reported when serving their needs.

The IPPM framework [[RFC2330](#)] and other RFCs describing IPPM metrics provide a background for this memo.

2. Purpose and Scope

The purpose of this memo is to clearly delineate two points-of-view

(POV) for using measurements, and describe their effects on the test design, including the selection of metric parameters and reporting the results.

The current scope of this memo is primarily limited to design and reporting of the loss and delay metrics [[RFC2680](#)] [[RFC2679](#)], but will also discuss the delay variation and reordering metrics where applicable. Sampling, or the design of the active packet stream that is the basis for the measurements, is also discussed.

Morton, et al.

Expires November 20, 2008

[Page 3]

Internet-Draft

Reporting Metrics

May 2008

[3.](#) Effect of POV on the Loss Metric

This section describes the ways in which the Loss metric can be tuned to reflect the preferences of the two audience categories, or different POV. The waiting time to declare a packet lost, or loss threshold is one area where there would appear to be a difference, but the ability to post-process the results may resolve it.

[3.1.](#) Loss Threshold

[RFC 2680](#) [[RFC2680](#)] defines the concept of a waiting time for packets to arrive, beyond which they are declared lost. The text of the RFC declines to recommend a value, instead saying that "good engineering, including an understanding of packet lifetimes, will be needed in practice." Later, in the methodology, they give reasons for waiting "a reasonable period of time", and leaving the definition of "reasonable" intentionally vague.

[3.1.1.](#) Network Characterization

Practical measurement experience has shown that unusual network circumstances can cause long delays. One such circumstance is when routing loops form during IGP re-convergence following a failure or drastic link cost change. Packets will loop between two routers until new routes are installed, or until the IPv4 Time-to-Live (TTL) field (or the IPv6 Hop Limit) decrements to zero. Very long delays on the order of several seconds have been measured [[Casner](#)] [[Cia03](#)].

Therefore, network characterization activities prefer a long waiting

time in order to distinguish these events from other causes of loss (such as packet discard at a full queue, or tail drop). This way, the metric design helps to distinguish more reliably between packets that might yet arrive, and those that are no longer traversing the network.

It is possible to calculate a worst-case waiting time, assuming that a routing loop is the cause. We model the path between Source and Destination as a series of delays in links (t) and queues (q), as these two are the dominant contributors to delay. The normal path delay across n hops without encountering a loop, D , is

$$D = t_0 + \sum_{i=1}^n t_i + q_i$$

Figure 1: Normal Path Delay

and the time spent in the loop with L hops, is

$$R = C \sum_{i=1}^{i+L-1} t_i + q_i \text{ where } C_{\max} = \frac{(TTL - n)}{L}$$

Figure 2: Delay due to Rotations in a Loop

and where C is the number of times a packet circles the loop.

If we take the delays of all links and queues as 100ms each, the TTL=255, the number of hops $n=5$ and the hops in the loop $L=4$, then

$D = 1.1$ sec and $R \approx 50$ sec, and $D + R \approx 51.1$ seconds

We note that the link delays of 100ms would span most continents, and a constant queue length of 100ms is also very generous. When a loop occurs, it is almost certain to be resolved in 10 seconds or less. The value calculated above is an upper limit for almost any realistic circumstance.

A waiting time threshold parameter, dT , set consistent with this calculation would not truncate the delay distribution (possibly causing a change in its mathematical properties), because the packets that might arrive have been given sufficient time to traverse the network.

It is worth noting that packets that are stored and deliberately forwarded at a much later time constitute a replay attack on the measurement system, and are beyond the scope of normal performance reporting.

[3.1.2.](#) Application Performance

Fortunately, application performance estimation activities are not adversely affected by the estimated worst-case transfer time. Although the designer's tendency might be to set the Loss Threshold at a value equivalent to a particular application's threshold, this specific threshold can be applied when post-processing the measurements. A shorter waiting time can be enforced by locating packets with delays longer than the application's threshold, and re-designating such packets as lost. Thus, the measurement system can use a single loss threshold and support both application and network performance POVs simultaneously.

[3.2.](#) Errored Packet Designation

[RFC 2680](#) designates packets that arrive containing errors as lost packets. Many packets that are corrupted by bit errors are discarded within the network and do not reach their intended destination.

This is consistent with applications that would check the payload integrity at higher layers, and discard the packet. However, some applications prefer to deal with errored payloads on their own, and even a corrupted payload is better than no packet at all.

To address this possibility, and to make network characterization more complete, it is recommended to distinguish between packets that do not arrive (lost) and errored packets that arrive (conditionally lost).

[3.3.](#) Causes of Lost Packets

Although many measurement systems use a waiting time to determine if a packet is lost or not, most of the waiting is in vain. The packets are no-longer traversing the network, and have not reached their destination.

There are many causes of packet loss, including:

1. Queue drop, or discard
2. Corruption of the IP header, or other essential header info
3. TTL expiration (or use of a TTL value that is too small)
4. Link or router failure

After waiting sufficient time, packet loss can probably be attributed to one of these causes.

[3.4.](#) Summary for Loss

Given that measurement post-processing is possible (even encouraged in the definitions of IPPM metrics), measurements of loss can easily serve both points of view:

- o Use a long waiting time to serve network characterization and revise results for specific application delay thresholds as

needed.

- o Distinguish between errored packets and lost packets when possible to aid network characterization, and combine the results for application performance if appropriate.

4. Effect of POV on the Delay Metric

This section describes the ways in which the Delay metric can be tuned to reflect the preferences of the two consumer categories, or different POV.

4.1. Treatment of Lost Packets

The Delay Metric [[RFC2679](#)] specifies the treatment of packets that do not successfully traverse the network: their delay is undefined.

" >>The *Type-P-One-way-Delay* from Src to Dst at T is undefined (informally, infinite)<< means that Src sent the first bit of a Type-P packet to Dst at wire-time T and that Dst did not receive that packet."

It is an accepted, but informal practice to assign infinite delay to lost packets. We next look at how these two different treatments align with the needs of measurement consumers who wish to characterize networks or estimate application performance. Also, we look at the way that lost packets have been treated in other metrics: delay variation and reordering.

4.1.1. Application Performance

Applications need to perform different functions, dependent on whether or not each packet arrives within some finite tolerance. In other words, a receivers' packet processing takes one of two directions (or "forks" in the road):

- o Packets that arrive within expected tolerance are handled by processes that remove headers, restore smooth delivery timing (as in a de-jitter buffer), restore sending order, check for errors in

payloads, and many other operations.

- o Packets that do not arrive when expected spawn other processes that attempt recovery from the apparent loss, such as retransmission requests, loss concealment, or forward error correction to replace the missing packet.

So, it is important to maintain a distinction between packets that actually arrive, and those that do not. Therefore, it is preferable to leave the delay of lost packets undefined, and to characterize the delay distribution as a conditional distribution (conditioned on arrival).

4.1.2. Network Characterization

In this discussion, we assume that both loss and delay metrics will be reported for network characterization (at least).

Assume packets that do not arrive are reported as Lost, usually as a fraction of all sent packets. If these lost packets are assigned undefined delay, then network's inability to deliver them (in a timely way) is captured only in the loss metric when we report statistics on the Delay distribution conditioned on the event of packet arrival (within the Loss waiting time threshold). We can say that the Delay and Loss metrics are Orthogonal, in that they convey non-overlapping information about the network under test.

However, if we assign infinite delay to all lost packets, then:

- o The delay metric results are influenced both by packets that arrive and those that do not.
- o The delay singleton and the loss singleton do not appear to be orthogonal (Delay is finite when Loss=0, Delay is infinite when Loss=1).
- o The network is penalized in both the loss and delay metrics, effectively double-counting the lost packets.

As further evidence of overlap, consider the Cumulative Distribution Function (CDF) of Delay when the value positive infinity is assigned to all lost packets. Figure 3 shows a CDF where a small fraction of packets are lost.




Figure 3: Cumulative Distribution Function for Delay when Loss = +Infinity

We note that a Delay CDF that is conditioned on packet arrival would not exhibit this apparent overlap with loss.

Although infinity is a familiar mathematical concept, it is somewhat disconcerting to see any time-related metric reported as infinity, in the opinion of the authors. Questions are bound to arise, and tend to detract from the goal of informing the consumer with a performance report.

[4.1.3.](#) Delay Variation

[RFC3393] excludes lost packets from samples, effectively assigning an undefined delay to packets that do not arrive in a reasonable time. [Section 4.1](#) describes this specification and its rationale (ipdv = inter-packet delay variation in the quote below).

"The treatment of lost packets as having "infinite" or "undefined" delay complicates the derivation of statistics for ipdv. Specifically, when packets in the measurement sequence are lost, simple statistics such as sample mean cannot be computed. One possible approach to handling this problem is to reduce the event space by conditioning. That is, we consider conditional statistics; namely we estimate the mean ipdv (or other derivative statistic) conditioned on the event that selected packet pairs arrive at the destination (within the given timeout). While this itself is not without problems (what happens, for example, when every other packet is lost), it offers a way to make some (valid) statements about ipdv, at the same time avoiding events with undefined outcomes."

[4.1.4.](#) Reordering

[RFC4737] defines metrics that are based on evaluation of packet arrival order, and include a waiting time to declare a packet lost (to exclude them from further processing).

If packets are assigned a delay value, then the reordering metric would declare any packets with infinite delay to be reordered, because their sequence numbers will surely be less than the "Next Expected" threshold when (or if) they arrive. But this practice would fail to maintain orthogonality between the reordering metric and the loss metric. Confusion can be avoided by designating the delay of non-arriving packets as undefined, and reserving delay values only for packets that arrive within a sufficiently long waiting time.

[4.2.](#) Preferred Statistics

Today in network characterization, the sample mean is one statistic that is almost ubiquitously reported. It is easily computed and understood by virtually everyone in this audience category. Also, the sample is usually filtered on packet arrival, so that the mean is based a conditional distribution.

The median is another statistic that summarizes a distribution, having somewhat different properties from the sample mean. The median is stable in distributions with a few outliers or without them. However, the median's stability prevents it from indicating when a large fraction of the distribution changes value. 50% or more values would need to change for the median to capture the change.

Both the median and sample mean have difficulty with bimodal distributions. The median will reside in only one of the modes, and the mean may not lie in either mode range. For this and other reasons, additional statistics such as the minimum, maximum, and 95%-ile have value when summarizing a distribution.

When both the sample mean and median are available, a comparison will sometimes be informative, because these two statistics are equal only

when the delay distribution is perfectly symmetrical.

Also, these statistics are generally useful from the Application Performance POV, so there is a common set that should satisfy audiences.

Morton, et al.

Expires November 20, 2008

[Page 10]

Internet-Draft

Reporting Metrics

May 2008

[4.3.](#) Summary for Delay

From the perspectives of:

1. application/receiver analysis, where subsequent processing depends on whether the packet arrives or times-out,
2. straightforward network characterization without double-counting defects, and
3. consistency with Delay variation and Reordering metric definitions,

the most efficient practice is to distinguish between truly lost and delayed packets with a sufficiently long waiting time, and to designate the delay of non-arriving packets as undefined.

[5.](#) Test Streams and Sample Size

This section discusses two key aspects of measurement that are sometimes omitted from the report: the description of the test stream on which the measurements are based, and the sample size.

[5.1.](#) Test Stream Characteristics

Network Characterization has traditionally used Poisson-distributed inter-packet spacing, as this provides an unbiased sample. The average inter-packet spacing may be selected to allow observation of specific network phenomena. Other test streams are designed to sample some property of the network, such as the presence of congestion, link bandwidth, or packet reordering.

If measuring a network in order to make inferences about applications or receiver performance, then there are usually efficiencies derived from a test stream that has similar characteristics to the sender. In some cases, it is essential to synthesize the sender stream, as with Bulk Transfer Capacity estimates. In other cases, it may be sufficient to sample with a "known bias", e.g., a Periodic stream to estimate real-time application performance.

[5.2.](#) Sample Size

Sample size is directly related to the accuracy of the results, and plays a critical role in the report. Even if only the sample size (in terms of number of packets) is given for each value or summary statistic, it imparts a notion of the confidence in the result.

In practice, the sample size will be selected taking both statistical and practical factors into account. Among these factors are:

1. The estimated variability of the quantity being measured
2. The desired confidence in the result (although this may be dependent on assumption of the underlying distribution of the measured quantity).
3. The effects of active measurement traffic on user traffic
4. etc.

A sample size may sometimes be referred to as "large". This is a relative, and qualitative term. It is preferable to describe what one is attempting to achieve with their sample. For example, stating an implication may be helpful: this sample is large enough such that a single outlying value at ten times the "typical" sample mean (the mean without the outlying value) would influence the mean by no more than X.

[6.](#) Reporting Results

This section gives an overview of recommendations, followed by

additional considerations for reporting results in the "long-term".

[6.1.](#) Overview of Metric Statistics

This section gives an overview of reporting recommendations for the loss, delay, and delay variation metrics based on the discussion and conclusions of the preceding sections.

The minimal report on measurements MUST include both Loss and Delay Metrics.

For Packet Loss, the loss ratio defined in [[RFC2680](#)] is a sufficient starting point, especially the guidance for setting the loss threshold waiting time. We have calculated a waiting time above that should be sufficient to differentiate between packets that are truly lost or have long finite delays under general measurement circumstances, 51 seconds. Knowledge of specific conditions can help to reduce this threshold, but 51 seconds is considered to be manageable in practice.

We note that a loss ratio calculated according to [[Y.1540](#)] would exclude errored packets from the numerator. In practice, the difference between these two loss metrics is small if any, depending

on whether the last link prior to the destination contributes errored packets.

For Packet Delay, we recommend providing both the mean delay and the median delay with lost packets designated undefined (as permitted by [[RFC2679](#)]). Both statistics are based on a conditional distribution, and the condition is packet arrival prior to a waiting time dT , where dT has been set to take maximum packet lifetimes into account, as discussed above. Using a long dT helps to ensure that delay distributions are not truncated.

For Packet Delay Variation (PDV), the minimum delay of the conditional distribution should be used as the reference delay for computing PDV according to [[Y.1540](#)] or [[RFC3393](#)]. A useful value to report is a pseudo range of delay variation based on calculating the difference between a high percentile of delay and the minimum delay. For example, the 99.9%-ile minus the minimum will give a value that can be compared with objectives in [[Y.1541](#)].

[6.2.](#) Long-Term Reporting Considerations

[[I-D.ietf-ippm-reporting](#)] describes methods to conduct measurements and report the results on a near-immediate time scale (10 seconds, which we consider to be "short-term").

Measurement intervals and reporting intervals need not be the same length. Sometimes, the user is only concerned with the performance levels achieved over a relatively long interval of time (e.g, days, weeks, or months, as opposed to 10 seconds). However, there can be risks involved with running a measurement continuously over a long period without recording intermediate results:

- o Temporary power failure may cause loss of all the results to date.
- o Measurement system timing synchronization signals may experience a temporary outage, causing sub-sets of measurements to be in error or invalid.
- o Maintenance may be necessary on the measurement system, or its connectivity to the network under test.

For these and other reasons, such as

- o the constraint to collect measurements on intervals similar to user session length, or
- o the dual-use of measurements in monitoring activities where results are needed on a period of a few minutes,

there is value in conducting measurements on intervals that are much shorter than the reporting interval.

There are several approaches for aggregating a series of measurement results over time in order to make a statement about the longer reporting interval. One approach requires the storage of all metric singletons collected throughout the reporting interval, even though the measurement interval stops and starts many times.

Another approach is described in [[I-D.ietf-ippm-framework-compagg](#)] as "temporal aggregation". This approach would estimate the results for

the reporting interval based on many individual measurement interval statistics (results) alone. The result would ideally appear in the same form as though a continuous measurement was conducted. A memo to address the details of temporal aggregation is yet to be prepared.

Yet another approach requires a numerical objective for the metric, and the results of each measurement interval are compared with the objective. Every measurement interval where the results meet the objective contribute to the fraction of time with performance as specified. When the reporting interval contains many measurement intervals it is possible to present the results as "metric A was less than or equal to objective X during Y% of time."

NOTE that numerical thresholds are not set in IETF performance work and are explicitly excluded from the IPPM charter.

7. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

8. Security Considerations

The security considerations that apply to any active measurement of live networks are relevant here as well. See [[RFC4656](#)].

9. Acknowledgements

The authors would like to thank Phil Chimento for his suggestion to employ conditional distributions for Delay, and Steve Konish Jr. for his careful review and suggestions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", [RFC 2679](#), September 1999.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", [RFC 2680](#), September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", [RFC 3393](#), November 2002.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", [RFC 4656](#), September 2006.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", [RFC 4737](#), November 2006.

[10.2](#). Informative References

- [Casner] "A Fine-Grained View of High Performance Networking, NANOG 22 Conf.; <http://www.nanog.org/mtg-0105/agenda.html>", May 20-22 2001.
- [Cia03] "Standardized Active Measurements on a Tier 1 IP Backbone, IEEE Communications Mag., pp 90-97.", June 2003.
- [I-D.ietf-ippm-framework-compagg]
Morton, A., "Framework for Metric Composition", [draft-ietf-ippm-framework-compagg-06](#) (work in progress), February 2008.
- [I-D.ietf-ippm-reporting]
Shalunov, S., "Reporting IP Performance Metrics to Users", [draft-ietf-ippm-reporting-01](#) (work in progress), October 2006.

- [Y.1540] ITU-T Recommendation Y.1540, "Internet protocol data communication service - IP packet transfer and availability performance parameters", December 2002.
- [Y.1541] ITU-T Recommendation Y.1540, "Network Performance Objectives for IP-Based Services", February 2006.

Authors' Addresses

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>

Gomathi Ramachandran
AT&T Labs
200 Laurel Avenue South
Middletown, New Jersey 07748
USA

Phone: +1 732 420 2353
Fax:
Email: gomathi@att.com
URI:

Ganga Maguluri
AT&T Labs
200 Laurel Avenue
Middletown, New Jersey 07748
USA

Phone: 732-420-2486
Fax:
Email: gmaguluri@att.com
URI:

Internet-Draft

Reporting Metrics

May 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

