

Internet Engineering Task Force (IETF)
Request for Comments: 7680
STD: 82
Obsoletes: 2680
Category: Standards Track
ISSN: 2070-1721

G. Almes
Texas A&M
S. Kalidindi
Ixia
M. Zekauskas
Internet2
A. Morton, Ed.
AT&T Labs
January 2016

A One-Way Loss Metric for IP Performance Metrics (IPPM)

Abstract

This memo defines a metric for one-way loss of packets across Internet paths. It builds on notions introduced and discussed in the IP Performance Metrics (IPPM) Framework document, RFC 2330; the reader is assumed to be familiar with that document. This memo makes RFC 2680 obsolete.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7680>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
1.1.	Motivation	5
1.2.	General Issues regarding Time	6
2.	A Singleton Definition for One-Way Packet Loss	7
2.1.	Metric Name	7
2.2.	Metric Parameters	7
2.3.	Metric Units	7
2.4.	Definition	7
2.5.	Discussion	8
2.6.	Methodologies	9
2.7.	Errors and Uncertainties	10
2.8.	Reporting the Metric	11
2.8.1.	Type-P	11
2.8.2.	Loss Threshold	11
2.8.3.	Calibration Results	11
2.8.4.	Path	12
3.	A Definition for Samples of One-Way Packet Loss	12
3.1.	Metric Name	12
3.2.	Metric Parameters	13
3.3.	Metric Units	13
3.4.	Definition	13
3.5.	Discussion	13
3.6.	Methodologies	14
3.7.	Errors and Uncertainties	15
3.8.	Reporting the Metric	15
4.	Some Statistics Definitions for One-Way Packet Loss	15
4.1.	Type-P-One-way-Packet-Loss-Ratio	15
5.	Security Considerations	16
6.	Changes from RFC 2680	17
7.	References	19
7.1.	Normative References	19
7.2.	Informative References	20
	Acknowledgements	21
	Authors' Addresses	22

1. Introduction

This memo defines a metric for one-way packet loss across Internet paths. It builds on notions introduced and discussed in the IPPM Framework document, [RFC2330]; the reader is assumed to be familiar with that document and its recent update [RFC7312].

This memo is intended to be parallel in structure to a companion document for One-way Delay ("A One-Way Delay Metric for IP Performance Metrics (IPPM)") [RFC7679]; the reader is assumed to be familiar with that document.

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 [RFC2119]. Although [RFC2119] was written with protocols in mind, the key words are used in this document for similar reasons. They are used to ensure the results of measurements from two different implementations are comparable and to note instances when an implementation could perturb the network.

Whenever a technical term from the IPPM Framework document is first used in this memo, it will be tagged with a trailing asterisk. For example, "term*" indicates that "term" is defined in the Framework document.

The structure of the memo is as follows:

- o A 'singleton*' analytic metric, called Type-P-One-way-Packet-Loss, is introduced to measure a single observation of packet transmission or loss.
- o Using this singleton metric, a 'sample*' called Type-P-One-way-Packet-Loss-Poisson-Stream is introduced to measure a sequence of singleton transmissions and/or losses measured at times taken from a Poisson process, as defined in Section 11.1.1 of [RFC2330].
- o Using this sample, several 'statistics*' of the sample will be defined and discussed.

This progression from singleton to sample to statistics, with clear separation among them, is important.

1.1. Motivation

Understanding one-way packet loss of Type-P* packets from a source host* to a destination host is useful for several reasons:

- o Some applications do not perform well (or at all) if end-to-end loss between hosts is large relative to some threshold value.
- o Excessive packet loss may make it difficult to support certain real-time applications (where the precise threshold of "excessive" depends on the application).
- o The larger the value of packet loss, the more difficult it is for transport-layer protocols to sustain high bandwidths.
- o The sensitivity of real-time applications and of transport-layer protocols to loss become especially important when very large delay-bandwidth products must be supported.

The measurement of one-way loss instead of round-trip loss is motivated by the following factors:

- o In today's Internet, the path from a source to a destination may be different than the path from the destination back to the source ("asymmetric paths"), such that different sequences of routers are used for the forward and reverse paths. Therefore, round-trip measurements actually measure the performance of two distinct paths together. Measuring each path independently highlights the performance difference between the two paths that may traverse different Internet service providers and even radically different types of networks (for example, research versus commodity networks, or networks with asymmetric link capacities, or wireless versus wireline access).
- o Even when the two paths are symmetric, they may have radically different performance characteristics due to asymmetric queuing.
- o Performance of an application may depend mostly on the performance in one direction. For example, a TCP-based communication will experience reduced throughput if congestion occurs in one direction of its communication. Troubleshooting may be simplified if the congested direction of TCP transmission can be identified.
- o In networks in which quality of service (QoS) is enabled, provisioning in one direction may be radically different than provisioning in the reverse direction and thus the QoS guarantees differ. Measuring the paths independently allows the verification of both guarantees.

It is outside the scope of this document to say precisely how loss metrics would be applied to specific problems.

1.2. General Issues regarding Time

{Comment: The terminology below differs from that defined by ITU-T documents (e.g., G.810, "Definitions and terminology for synchronization networks" and I.356, "B-ISDN ATM layer cell transfer performance") but is consistent with the IPPM Framework document. In general, these differences derive from the different backgrounds; the ITU-T documents historically have a telephony origin, while the authors of this document (and the Framework document) have a computer systems background. Although the terms defined below have no direct equivalent in the ITU-T definitions, after our definitions we will provide a rough mapping. However, note one potential confusion: our definition of "clock" is the computer operating systems definition denoting a time-of-day clock, while the ITU-T definition of clock denotes a frequency reference.}

Whenever a time (i.e., a moment in history) is mentioned here, it is understood to be measured in seconds (and fractions) relative to UTC.

As described more fully in the Framework document, there are four distinct, but related notions of clock uncertainty:

synchronization*

measures the extent to which two clocks agree on what time it is. For example, the clock on one host might be 5.4 msec ahead of the clock on a second host. {Comment: A rough ITU-T equivalent is "time error".}

accuracy*

measures the extent to which a given clock agrees with UTC. For example, the clock on a host might be 27.1 msec behind UTC. {Comment: A rough ITU-T equivalent is "time error from UTC".}

resolution*

is a specification of the smallest unit by which the clock's time is updated. It gives a lower bound on the clock's uncertainty. For example, the clock on an old Unix host might tick only once every 10 msec and thus have a resolution of only 10 msec. {Comment: A very rough ITU-T equivalent is "sampling period".}

skew*

measures the change of accuracy, or of synchronization, with time. For example, the clock on a given host might gain 1.3 msec per hour and thus be 27.1 msec behind UTC at one time and only 25.8 msec an hour later. In this case, we say that the clock of the given host has a skew of 1.3 msec per hour relative to UTC, which threatens accuracy. We might also speak of the skew of one clock relative to another clock, which threatens synchronization. {Comment: A rough ITU-T equivalent is "time drift".}

2. A Singleton Definition for One-Way Packet Loss

2.1. Metric Name

Type-P-One-way-Packet-Loss

2.2. Metric Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o T, a time
- o Tmax, a loss threshold waiting time

2.3. Metric Units

The value of a Type-P-One-way-Packet-Loss is either a zero (signifying successful transmission of the packet) or a one (signifying loss).

2.4. Definition

>>The *Type-P-One-way-Packet-Loss* from Src to Dst at T is 0<< means that Src sent the first bit of a Type-P packet to Dst at wire time* T and that Dst received that packet.

>>The *Type-P-One-way-Packet-Loss* from Src to Dst at T is 1<< 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 (within the loss threshold waiting time, Tmax).

2.5. Discussion

Thus, Type-P-One-way-Packet-Loss is 0 exactly when Type-P-One-way-Delay is a finite value, and it is 1 exactly when Type-P-One-way-Delay is undefined.

The following issues are likely to come up in practice:

- o A given methodology will have to include a way to distinguish between a packet loss and a very large (but finite) delay. As noted by Mahdavi and Paxson [RFC2678], simple upper bounds (such as the 255-second theoretical upper bound on the lifetimes of IP packets [RFC791]) could be used, but good engineering, including an understanding of packet lifetimes, will be needed in practice. {Comment: Note that, for many applications of these metrics, there may be no harm in treating a large delay as packet loss. An audio playback packet, for example, that arrives only after the playback point may as well have been lost. See Section 4.1.1 of [RFC6703] for examination of unusual packet delays and application performance estimation.}
- o If the packet arrives but is corrupted, then it is counted as lost. {Comment: One is tempted to count the packet as received since corruption and packet loss are related but distinct phenomena. If the IP header is corrupted, however, one cannot be sure about the source or destination IP addresses and is thus on shaky grounds about knowing that the corrupted received packet corresponds to a given sent test packet. Similarly, if other parts of the packet needed by the methodology to know that the corrupted received packet corresponds to a given sent test packet, then such a packet would have to be counted as lost. It would be inconsistent to count packets with corrupted methodology-specific fields as lost, and not to count packets with other corrupted aspects in the same category.} Section 15 of [RFC2330] defines the "standard-formed" packet that is applicable to all metrics. Note that at this time the definition of standard-formed packets only applies to IPv4 (see also [IPPM-UPDATES]).
- o If the packet is duplicated along the path (or paths) so that multiple non-corrupt copies arrive at the destination, then the packet is counted as received.
- o If the packet is fragmented and if, for whatever reason, reassembly does not occur, then the packet will be deemed lost.

2.6. Methodologies

As with other Type-P-* metrics, the detailed methodology will depend on the Type-P (e.g., protocol number, UDP/TCP port number, size, Differentiated Services (DS) Field [RFC2780]).

Generally, for a given Type-P, one possible methodology would proceed as follows:

- o Arrange that Src and Dst have clocks that are synchronized with each other. The degree of synchronization is a parameter of the methodology and depends on the threshold used to determine loss (see below).
- o At the Src host, select Src and Dst IP addresses and form a test packet of Type-P with these addresses.
- o At the Dst host, arrange to receive the packet.
- o At the Src host, place a timestamp in the prepared Type-P packet, and send it towards Dst (ideally minimizing time before sending).
- o If the packet arrives within a reasonable period of time, the one-way packet loss is taken to be zero (and take a timestamp as soon as possible upon the receipt of the packet).
- o If the packet fails to arrive within a reasonable period of time, T_{max} , the one-way packet loss is taken to be one. Note that the threshold of "reasonable" here is a parameter of the metric.

{Comment: The definition of reasonable is intentionally vague and is intended to indicate a value "Th" so large that any value in the closed interval [Th-delta, Th+delta] is an equivalent threshold for loss. Here, delta encompasses all error in clock synchronization and timestamp acquisition and assignment along the measured path. If there is a single value, T_{max} , after which the packet must be counted as lost, then we reintroduce the need for a degree of clock synchronization similar to that needed for one-way delay, and virtually all practical measurement systems combine methods for delay and loss. Therefore, if a measure of packet loss parameterized by a specific non-huge "reasonable" time-out value is needed, one can always measure one-way delay and see what percentage of packets from a given stream exceed a given time-out value. This point is examined in detail in [RFC6703], including analysis preferences to assign undefined delay to packets that fail to arrive with the difficulties emerging from the informal "infinite delay" assignment, and an estimation of an upper bound on waiting time for packets in transit. Further, enforcing a specific constant waiting time on stored

singletons of one-way delay is compliant with this specification and may allow the results to serve more than one reporting audience.}

Issues such as the packet format, the means by which Dst knows when to expect the test packet, and the means by which Src and Dst are synchronized are outside the scope of this document. {Comment: We plan to document the implementation techniques of our work in much more detail elsewhere; we encourage others to do so as well.}

2.7. Errors and Uncertainties

The description of any specific measurement method should include an accounting and analysis of various sources of error or uncertainty. The Framework document provides general guidance on this point.

For loss, there are three sources of error:

- o synchronization between clocks on Src and Dst.
- o the packet-loss threshold (which is related to the synchronization between clocks).
- o resource limits in the network interface or software on the receiving instrument.

The first two sources are interrelated and could result in a test packet with finite delay being reported as lost. Type-P-One-way-Packet-Loss is 1 if the test packet does not arrive, or if it does arrive and the difference between the Src timestamp and the Dst timestamp is greater than the "reasonable period of time" or loss threshold. If the clocks are not sufficiently synchronized, the loss threshold may not be "reasonable" - the packet may take much less time to arrive than its Src timestamp indicates. Similarly, if the loss threshold is set too low, then many packets may be counted as lost. The loss threshold must be high enough and the clocks synchronized well enough so that a packet that arrives is rarely counted as lost. (See the discussions in the previous two sections.)

Since the sensitivity of packet-loss measurement alone to lack of clock synchronization is less than for delay, we refer the reader to the treatment of synchronization errors in the "One-way Delay Metric for IPPM" [RFC2330] for more details.

The last source of error, resource limits, cause the packet to be dropped by the measurement instrument and counted as lost when in fact the network delivered the packet in reasonable time.

The measurement instruments should be calibrated such that the loss threshold is reasonable for application of the metrics and the clocks are synchronized enough so the loss threshold remains reasonable.

In addition, the instruments should be checked to ensure that the possibility a packet arrives at the network interface but is lost due to congestion on the interface or to other resource exhaustion (e.g., buffers) on the instrument is low.

2.8. Reporting the Metric

The calibration and context in which the metric is measured MUST be carefully considered and SHOULD always be reported along with metric results. We now present four items to consider: Type-P of the test packets, the loss threshold, instrument calibration, and the path traversed by the test packets. This list is not exhaustive; any additional information that could be useful in interpreting applications of the metrics should also be reported (see [RFC6703] for extensive discussion of reporting considerations for different audiences).

2.8.1. Type-P

As noted in Section 13 of the Framework document [RFC2330], the value of the metric may depend on the type of IP packets used to make the measurement, or "Type-P". The value of Type-P-One-way-Delay could change if the protocol (UDP or TCP), port number, size, or arrangement for special treatment (e.g., IP DS Field [RFC2780], Explicit Congestion Notification (ECN) [RFC3168], or RSVP) changes. Additional packet distinctions identified in future extensions of the Type-P definition will apply. The exact Type-P used to make the measurements MUST be accurately reported.

2.8.2. Loss Threshold

The threshold, T_{max} , between a large finite delay and loss (or other methodology to distinguish between finite delay and loss) MUST be reported.

2.8.3. Calibration Results

The degree of synchronization between the Src and Dst clocks MUST be reported. If possible, a test packet that arrives at the Dst network interface and is reported as lost due to resource exhaustion on Dst SHOULD be reported.

2.8.4. Path

Finally, the path traversed by the packet SHOULD be reported, if possible. In general, it is impractical to know the precise path a given packet takes through the network. The precise path may be known for certain Type-P on short or stable paths. If Type-P includes the record route (or loose-source route) option in the IP header, and the path is short enough, and all routers* on the path support record (or loose-source) route, then the path will be precisely recorded. This is impractical because the route must be short enough, many routers do not support (or are not configured for) record route, and use of this feature would often artificially worsen the performance observed by removing the packet from common-case processing. However, partial information is still valuable context. For example, if a host can choose between two links* (and hence, two separate routes from Src to Dst), then the initial link used is valuable context. {Comment: Backbone path selection services come and go. A historical example was Merit's NetNow setup, where a Src on one Network Access Point (NAP) can reach a Dst on another NAP by either of several different backbone networks.}

3. A Definition for Samples of One-Way Packet Loss

Given the singleton metric Type-P-One-way-Packet-Loss, we now define one particular sample of such singletons. The idea of the sample is to select a particular binding of the parameters Src, Dst, and Type-P, then define a sample of values of parameter T. The means for defining the values of T is to select a beginning time T_0 , a final time T_f , and an average rate λ , then define a pseudorandom Poisson process of rate λ , whose values fall between T_0 and T_f . The time interval between successive values of T will then average $1/\lambda$.

Note that Poisson sampling is only one way of defining a sample. Poisson has the advantage of limiting bias, but other methods of sampling will be appropriate for different situations. For example, a truncated Poisson distribution may be needed to avoid reactive network state changes during intervals of inactivity, see Section 4.6 of [RFC7312]. Sometimes the goal is sampling with a known bias, and [RFC3432] describes a method for periodic sampling with random start times.

3.1. Metric Name

Type-P-One-way-Packet-Loss-Poisson-Stream

3.2. Metric Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o T0, a time
- o Tf, a time
- o Tmax, a loss threshold waiting time
- o lambda, a rate in reciprocal seconds

3.3. Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time, and
- o L, either a zero or a one.

The values of T in the sequence are monotonic increasing. Note that T would be a valid parameter to Type-P-One-way-Packet-Loss and that L would be a valid value of Type-P-One-way-Packet-Loss.

3.4. Definition

Given T0, Tf, and lambda, we compute a pseudorandom Poisson process beginning at or before T0, with average arrival rate lambda, and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the selected times in this process, we obtain one value of Type-P-One-way-Delay. The value of the sample is the sequence made up of the resulting <time, loss> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

3.5. Discussion

The reader should be familiar with the in-depth discussion of Poisson sampling in the Framework document [RFC2330], which includes methods to compute and verify the pseudorandom Poisson process.

We specifically do not constrain the value of lambda except to note the extremes. If the rate is too large, then the measurement traffic will perturb the network and itself cause congestion. If the rate is too small, then you might not capture interesting network behavior. {Comment: We expect to document our experiences with, and suggestions

for, lambda elsewhere, culminating in a "Best Current Practice" document.}

Since a pseudorandom number sequence is employed, the sequence of times, and hence the value of the sample, is not fully specified. Pseudorandom number generators of good quality will be needed to achieve the desired qualities.

The sample is defined in terms of a Poisson process both to avoid the effects of self-synchronization and also capture a sample that is statistically as unbiased as possible. The Poisson process is used to schedule the loss measurements. The test packets will generally not arrive at Dst according to a Poisson distribution, since they are influenced by the network. Time-slotted links described in Section 3.4 [RFC7312] can greatly modify the sample characteristics. The main concern is that unbiased packet streams with randomized inter-packet time intervals will be converted to some new distribution after encountering a time-slotted link, possibly with strong periodic characteristics instead.

{Comment: there is, of course, no claim that real Internet traffic arrives according to a Poisson arrival process.}

It is important to note that, in contrast to this metric, loss ratios observed by transport connections do not reflect unbiased samples. For example, TCP transmissions both (1) occur in bursts, which can induce loss due to the burst volume that would not otherwise have been observed, and (2) adapt their transmission rate in an attempt to minimize the loss ratio observed by the connection.}

All the singleton Type-P-One-way-Packet-Loss metrics in the sequence will have the same values of Src, Dst, and Type-P.

Note also that, given one sample that runs from T_0 to T_f , and given new time values T_0' and T_f' such that $T_0 \leq T_0' \leq T_f' \leq T_f$, the subsequence of the given sample whose time values fall between T_0' and T_f' are also a valid Type-P-One-way-Packet-Loss-Poisson-Stream sample.

3.6. Methodologies

The methodologies follow directly from:

- o the selection of specific times using the specified Poisson arrival process, and
- o the methodologies discussion already given for the singleton Type-P-One-way-Packet-Loss metric.

Care must be given to correctly handle out-of-order arrival of test packets; it is possible that the Src could send one test packet at TS[i], then send a second one (later) at TS[i+1] while the Dst could receive the second test packet at TR[i+1], and then receive the first one (later) at TR[i]. Metrics for reordering may be found in [RFC4737].

3.7. Errors and Uncertainties

In addition to sources of errors and uncertainties associated with methods employed to measure the singleton values that make up the sample, care must be given to analyze the accuracy of the Poisson arrival process of the wire times of the sending of the test packets. Problems with this process could be caused by several things, including problems with the pseudorandom number techniques used to generate the Poisson arrival process. The Framework document shows how to use the Anderson-Darling test to verify the accuracy of the Poisson process over small time frames. {Comment: The goal is to ensure that the test packets are sent "close enough" to a Poisson schedule and avoid periodic behavior.}

3.8. Reporting the Metric

The calibration and context for the underlying singletons MUST be reported along with the stream (see "Reporting the Metric" (Section 2.8) for Type-P-One-way-Packet-Loss).

4. Some Statistics Definitions for One-Way Packet Loss

Given the sample metric Type-P-One-way-Packet-Loss-Poisson-Stream, we now offer several statistics of that sample. These statistics are offered mostly to be illustrative of what could be done. See [RFC6703] for additional discussion of statistics that are relevant to different audiences.

4.1. Type-P-One-way-Packet-Loss-Ratio

Given a Type-P-One-way-Packet-Loss-Poisson-Stream, the average of all the L values in the stream is the ratio of losses to total packets in the stream. In addition, the Type-P-One-way-Packet-Loss-Ratio is undefined if the sample is empty.

For example, suppose we take a sample and the results are as follows:

```
Stream1 = <
<T1, 0>
<T2, 0>
<T3, 1>
<T4, 0>
<T5, 0>
>
```

Then, the average of loss results would be 0.2, the loss ratio.

Note that, since healthy Internet paths should be operating at loss ratios below 1% (particularly if high delay-bandwidth products are to be sustained), the sample sizes needed might be larger than one would like. Thus, for example, if one wants to discriminate between various fractions of 1% over one-minute periods, then several hundred samples per minute might be needed. This would result in larger values of lambda than one would ordinarily want.

Note that although the loss threshold should be set such that any errors in loss are not significant, if the possibility that a packet that arrived is counted as lost due to resource exhaustion is significant compared to the loss ratio of interest, Type-P-One-way-Packet-Loss-Ratio will be meaningless.

5. Security Considerations

Conducting Internet measurements raises both security and privacy concerns. This memo does not specify an implementation of the metrics, so it does not directly affect the security of the Internet nor of applications that run on the Internet. However, implementations of these metrics must be mindful of security and privacy concerns.

There are two types of security concerns: potential harm caused by the measurements and potential harm to the measurements. The measurements could cause harm because they are active and inject packets into the network. The measurement parameters MUST be carefully selected so that the measurements inject trivial amounts of additional traffic into the networks they measure. If they inject

"too much" traffic, they can skew the results of the measurement and in extreme cases cause congestion and denial of service.

The measurements themselves could be harmed by routers giving measurement traffic a different priority than "normal" traffic or by an attacker injecting artificial measurement traffic. If routers can recognize measurement traffic and treat it separately, the measurements will not reflect actual user traffic. If an attacker injects artificial traffic that is accepted as legitimate, the loss ratio will be artificially lowered. Therefore, the measurement methodologies SHOULD include appropriate techniques to reduce the probability that measurement traffic can be distinguished from "normal" traffic. Authentication techniques, such as digital signatures, may be used where appropriate to guard against injected traffic attacks.

When considering privacy of those involved in measurement or those whose traffic is measured, the sensitive information available to potential observers is greatly reduced when using active techniques that are within this scope of work. Passive observations of user traffic for measurement purposes raise many privacy issues. We refer the reader to the privacy considerations described in the Large Scale Measurement of Broadband Performance (LMAP) Framework [RFC7594], which covers active and passive techniques.

Collecting measurements or using measurement results for reconnaissance to assist in subsequent system attacks is quite common. Access to measurement results or control of the measurement systems to perform reconnaissance should be guarded against. See Section 7 of [RFC7594] (the Security Considerations section of the LMAP Framework) for system requirements that help to avoid measurement system compromise.

6. Changes from RFC 2680

The text above constitutes a revision to RFC 2680, which is now an Internet Standard.

[RFC7290] provides the test plan and results supporting [RFC2680] advancement along the Standards Track, according to the process in [RFC6576]. The conclusions of [RFC7290] list four minor modifications for inclusion:

1. Section 6.2.3 of [RFC7290] asserts that the assumption of post-processing to enforce a constant waiting time threshold is compliant and that the text of the RFC should be revised slightly to include this point. The applicability of post-processing was added in the last list item of Section 2.6, above.

2. Section 6.5 of [RFC7290] indicates that the Type-P-One-way-Packet-Loss-Average statistic is more commonly called a Packet Loss Ratio, so it is renamed in this document (this small discrepancy does not affect candidacy for advancement). The renaming was implemented in Section 4.1, above.
3. The IETF has reached consensus on guidance for reporting metrics in [RFC6703], and the memo is referenced in this document to incorporate recent experience where appropriate. This reference was added in the last list item of Section 2.6, in Section 2.8, and in Section 4 above.
4. There are currently two errata with status "Verified" (EID 1528) and "Held for Document Update" (EID 3186) for [RFC2680], and these minor revisions were incorporated in Sections 1 and 2.7.

A number of updates to the [RFC2680] text have been implemented in the text to reference key IPPM RFCs that were approved after [RFC2680] (see Sections 3 and 3.6, above) and to address comments on the IPPM mailing list describing current conditions and experience.

1. Near the end of Section 1.1, there is an update of a network example using ATM, a clarification of TCP's affect on queue occupation, and discussion of the importance of one-way delay measurement.
2. Clarification of the definition of "resolution" in Section 1.2.
3. Explicit inclusion of the maximum waiting time input parameter in Sections 2.2, 2.4, and 3.2, reflecting recognition of this parameter in more recent RFCs and ITU-T Recommendation Y.1540.
4. Addition of a reference to RFC 6703 in the discussion of packet lifetime and application timeouts in Section 2.5.
5. Replaced "precedence" with updated terminology (DS Field) in Sections 2.6 and 2.8.1 (with reference).
6. Added parenthetical guidance on minimizing the interval between timestamp placement to send time or reception time in Section 2.6. Also, the text now recognizes the timestamp acquisition process and that practical systems measure both delay and loss (thus requiring the max waiting time parameter).
7. Added a reference to RFC 3432 regarding periodic sampling alongside Poisson sampling in Section 3 and also noted that a truncated Poisson distribution may be needed with modern networks as described in the IPPM Framework update [RFC7312].

8. Recognition that time-slotted links described in [RFC7312] can greatly modify the sample characteristics, in Section 3.5.
9. Added a reference to RFC 4737 regarding reordering metrics in the related discussion of Section 3.6, "Methodologies".
10. Expanded and updated the material on privacy and added cautions on use of measurements for reconnaissance in Section 5, "Security Considerations".

Section 5.4.4 of [RFC6390] suggests a common template for performance metrics partially derived from previous IPPM and Benchmarking Methodology Working Group (BMWG) RFCs, but it also contains some new items. All of the normative parts of [RFC6390] are covered, but not quite in the same section names or orientation. Several of the informative parts are covered. Maintaining the familiar outline of IPPM literature has value and minimizes unnecessary differences between this revised RFC and current/future IPPM RFCs.

7. References

7.1. Normative References

- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<http://www.rfc-editor.org/info/rfc791>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <<http://www.rfc-editor.org/info/rfc2330>>.
- [RFC2678] Mahdavi, J. and V. Paxson, "IPPM Metrics for Measuring Connectivity", RFC 2678, DOI 10.17487/RFC2678, September 1999, <<http://www.rfc-editor.org/info/rfc2678>>.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, DOI 10.17487/RFC2680, September 1999, <<http://www.rfc-editor.org/info/rfc2680>>.

- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, DOI 10.17487/RFC2780, March 2000, <<http://www.rfc-editor.org/info/rfc2780>>.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, DOI 10.17487/RFC3432, November 2002, <<http://www.rfc-editor.org/info/rfc3432>>.
- [RFC6576] Geib, R., Ed., Morton, A., Fardid, R., and A. Steinmitz, "IP Performance Metrics (IPPM) Standard Advancement Testing", BCP 176, RFC 6576, DOI 10.17487/RFC6576, March 2012, <<http://www.rfc-editor.org/info/rfc6576>>.
- [RFC7312] Fabini, J. and A. Morton, "Advanced Stream and Sampling Framework for IP Performance Metrics (IPPM)", RFC 7312, DOI 10.17487/RFC7312, August 2014, <<http://www.rfc-editor.org/info/rfc7312>>.
- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<http://www.rfc-editor.org/info/rfc7679>>.

7.2. Informative References

- [IPPM-UPDATES] Morton, A., Fabini, J., Elkins, N., Ackermann, M., and V. Hegde, "Updates for IPPM's Active Metric Framework: Packets of Type-P and Standard-Formed Packets", Work in Progress, draft-morton-ippm-2330-stdform-typep-02, December 2015.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, DOI 10.17487/RFC4737, November 2006, <<http://www.rfc-editor.org/info/rfc4737>>.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, DOI 10.17487/RFC6390, October 2011, <<http://www.rfc-editor.org/info/rfc6390>>.

- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of View", RFC 6703, DOI 10.17487/RFC6703, August 2012, <<http://www.rfc-editor.org/info/rfc6703>>.
- [RFC7290] Ciavattone, L., Geib, R., Morton, A., and M. Wieser, "Test Plan and Results for Advancing RFC 2680 on the Standards Track", RFC 7290, DOI 10.17487/RFC7290, July 2014, <<http://www.rfc-editor.org/info/rfc7290>>.
- [RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", RFC 7594, DOI 10.17487/RFC7594, September 2015, <<http://www.rfc-editor.org/info/rfc7594>>.

Acknowledgements

For [RFC2680], thanks are due to Matt Mathis for encouraging this work and for calling attention on so many occasions to the significance of packet loss. Thanks are due also to Vern Paxson for his valuable comments on early drafts and to Garry Couch and Will Leland for several useful suggestions.

For this document, thanks to Joachim Fabini, Ruediger Geib, Nalini Elkins, and Barry Constantine for sharing their measurement experience as part of their careful reviews. Brian Carpenter and Scott Bradner provided useful feedback at IETF Last Call.

Authors' Addresses

Guy Almes
Texas A&M

Email: almes@acm.org

Sunil Kalidindi
Ixia

Email: skalidindi@ixiacom.com

Matt Zekauskas
Internet2

Email: matt@internet2.edu

Al Morton (editor)
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
United States

Phone: +1 732 420 1571

Fax: +1 732 368 1192

Email: acmorton@att.com

URI: <http://home.comcast.net/~acmacm/>