

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 9, 2011

A. Morton
AT&T Labs
October 6, 2010

Round-trip Loss Metrics
draft-morton-ippm-rt-loss-01

Abstract

Many user applications and the transport protocols that make them possible require two-way communications. To address this need, and also for system simplicity, round-trip loss measurements are frequently conducted in practice. The Two-Way Active Measurement Protocol specified in RFC 5357 establishes a round-trip loss measurement capability for the Internet. However, there is currently no metric specified according to the RFC 2330 framework.

This memo proposes/adds round-trip loss to the set of IP Performance Metrics (IPPM).

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].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 9, 2011.

Copyright Notice

Copyright (c) 2010 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	3
1.1.	Motivation	3
2.	Scope	4
3.	Common Specifications for Round-trip Metrics	4
3.1.	Name: Type-P-*	4
3.2.	Metric Parameters	4
3.3.	Metric Definition	5
3.4.	Metric Units	5
4.	A Singleton Round-trip Loss Metric	5
4.1.	Name: Type-P-Round-trip-Loss	5
4.2.	Metric Parameters	6
4.3.	Definition and Metric Units	6
4.4.	Discussion and other details	7
5.	A Sample Round-trip Loss Metric	7
5.1.	Name: Type-P-Round-trip-Loss-<Sample>-Stream	7
5.2.	Metric Parameters	7
5.3.	Definition and Metric Units	7
5.4.	Discussion and other details	8
6.	Round-trip Loss Statistic	8
6.1.	Type-P-Round-trip-Loss-<Sample>-Ratio	8
7.	Round-trip Testing and One-way Reporting	8
8.	Security Considerations	9
8.1.	Denial of Service Attacks	9
8.2.	User Data Confidentiality	9
8.3.	Interference with the metrics	9
9.	IANA Considerations	10
10.	Acknowledgements	10
11.	References	10
11.1.	Normative References	10
11.2.	Informative References	11
	Author's Address	11

1. Introduction

This memo defines a metric for round-trip loss on Internet paths. It builds on the notions and conventions introduced in the IP Performance Metrics (IPPM) framework [RFC2330]. Also, the specifications of the One-way Loss metric [RFC2680] and the Round-trip Delay metric [RFC2681] are frequently referenced and modified to match the round-trip circumstances addressed here. However, this memo assumes that the reader is familiar with the references, and does not repeat material as was done in [RFC2681].

This memo uses the terms "two-way" and "round-trip" synonymously.

1.1. Motivation

Many user applications and the transport protocols that make them possible require two-way communications. For example, the TCP SYN->, <-SYN-ACK, ACK-> three-way handshake attempted billions of times each day cannot be completed without two-way connectivity in a near-simultaneous time interval. Thus, measurements of Internet round-trip loss performance provide a basis to infer application performance more easily.

Measurement system designers have also recognized advantages of system simplicity when one host simply echoes or reflects test packets to the sender. Round-trip loss measurements are frequently conducted and reported in practice. The Two-Way Active Measurement Protocol specified in [RFC5357] establishes a round-trip loss measurement capability for the Internet. However, there is currently no round-trip loss metric specified according to the [RFC2330] framework.

[RFC2681] indicates that round-trip measurements may sometimes encounter "asymmetric" paths. When loss is observed using a round-trip measurement, there is often a desire to ascertain which of the two directional paths "lost" the packet. Under some circumstances, it is possible to make this inference. The round-trip measurement method raises a few complications when interpreting the embedded one-way results, and the user should be aware of them.

[RFC2681] also points out that loss measurement conducted sequentially in both directions of a path and reported as a round-trip result may be exactly the desired metric. On the other hand, it may be difficult to derive the state of round-trip loss from one-way measurements conducted in each direction unless a method to match the appropriate one-way measurements has pre-arranged.

Finally, many measurement systems report statistics on a conditional

delay distribution, where the condition is packet arrival at the destination. This condition is encouraged in [RFC3393], [RFC5481], and [draft-ietf-ippm-reporting-metrics]. As a result, lost packets need to be reported separately, according to a standardized metric. This memo defines such a metric.

See Section 1.1 of [RFC2680] for additional motivation of the packet loss metric.

2. Scope

This memo defines a round-trip loss metric using the conventions of the IPPM framework [RFC2330].

The memo defines a singleton metric, a sample metric, and a statistic, as per [RFC2330].

The memo also investigates the topic of one-way loss inference from a two-way measurement, and lists some key considerations.

3. Common Specifications for Round-trip Metrics

To reduce the redundant information presented in the detailed metrics sections that follow, this section presents the specifications that are common to two or more metrics. The section is organized using the same subsections as the individual metrics, to simplify comparisons.

3.1. Name: Type-P-*

All metrics use the Type-P convention as described in [RFC2330]. The rest of the name is unique to each metric.

3.2. Metric Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o T, a time (start of test interval)
- o Tf, a time (end of test interval)
- o lambda, a rate in reciprocal seconds (for Poisson Streams)

- o incT, the nominal duration of inter-packet interval, first bit to first bit (for Periodic Streams)
- o T0, a time that MUST be selected at random from the interval [T, T+dT] to start generating packets and taking measurements (for Periodic Streams)
- o TstampSrc, the wire time of the packet as measured at MP(Src) as it leaves for Dst.
- o TstampDst, the wire time of the packet as measured at MP(Dst), assigned to packets that arrive within a "reasonable" time.
- o Tmax, a maximum waiting time for packets to arrive, set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost).
- o M, the total number of packets sent between T0 and Tf
- o N, the total number of packets received at Dst (sent between T0 and Tf)
- o Type-P, as defined in [RFC2330], which includes any field that may affect a packet's treatment as it traverses the network

3.3. Metric Definition

This section is specific to each metric.

3.4. Metric Units

The metric units are logical (1 or 0) when describing a single packet's loss performance, where a 0 indicates successful packet transmission and a 1 indicates packet loss.

Units of time are as specified in [RFC2330].

Other units used are defined in the associated section.

4. A Singleton Round-trip Loss Metric

4.1. Name: Type-P-Round-trip-Loss

4.2. Metric Parameters

See section 3.2.

4.3. Definition and Metric Units

Type-P-Round-trip-Loss SHALL be represented by the binary logical values (or their equivalents) when the following conditions are met:

Type-P-Round-trip-Loss = 0:

- o Src sent the first bit of a Type-P packet to Dst at wire-time $T_{stampSrc}$,
- o that Dst received that packet,
- o the Dst immediately sent a Type-P packet back to the Src, and
- o that Src received the last bit of the reflected packet at wire-time $T_{stampSrc} + T_{max}$.

Type-P-Round-trip-Loss = 1:

- o Src sent the first bit of a Type-P packet to Dst at wire-time $T_{stampSrc}$,
- o that Src did not receive the last bit of the reflected packet before the waiting time lapsed at $T_{stampSrc} + T_{max}$
- o (possibly because that Dst did not receive that packet,
- o the Dst did not immediately sent a Type-P packet back to the Src, or
- o the Src did not receive a reflected Type-P packet sent from the Dst).

Following the precedent of[RFC2681], we make the simplifying assertion:

$Type-P-Round-trip-Loss(Src \rightarrow Dst) = Type-P-Round-trip-Loss(Dst \rightarrow Src)$

(and agree with the rationale, that the ambiguity introduced is a small price to pay for measurement efficiency).

Therefore, each singleton can be represented by pairs of elements as follows:

- o TstampSrc, the wire time of the packet at the Src (beginning the round-trip journey).
- o L, either zero or one (or some logical equivalent), where L=1 indicates loss and L=0 indicates successful round-trip arrival prior to TstampSrc + Tmax.

4.4. Discussion and other details

See [RFC2680] and [RFC2681] for extensive discussion, methods of measurement, errors and uncertainties, and other fundamental considerations that need not be repeated here.

5. A Sample Round-trip Loss Metric

Given the singleton metric Type-P-Round-trip-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 TstampSrc. This can be done in several ways, including:

1. Poisson: a pseudo-random Poisson process of rate lambda, whose values fall between T and Tf. The time interval between successive values of TstampSrc will then average 1/lambda, as per [RFC2330].
2. Periodic: a periodic stream process with pseudo-random start time T0 between T and dT, and nominal inter-packet interval incT, as per [RFC3432].

In the metric name, the variable <Stream> should be replaced with the process used to define the sample, using one of the above processes (or other process, the details of which MUST be specified if used).

5.1. Name: Type-P-Round-trip-Loss-<Sample>-Stream

5.2. Metric Parameters

See section 3.2.

5.3. Definition and Metric Units

Given one of the methods for defining the test interval, the sample of times (TstampSrc) and other metric parameters, we obtain a sequence of Type-P-Round-trip-Loss singletons as defined in section 4.3.

Type-P-Round-trip-Loss-<Sample>-Stream SHALL be a sequence of pairs with elements as follows:

- o TstampSrc, as above
- o L, either zero or one (or some logical equivalent), where L=1 indicates loss and L=0 indicates successful round-trip arrival prior to TstampSrc + Tmax.

where <Sample> SHALL be replaced with "Poisson", "Periodic", or an appropriate term to designate another sample method meeting the criteria of [RFC2330].

5.4. Discussion and other details

See [RFC2680] and [RFC2681] for extensive discussion, methods of measurement, errors and uncertainties, and other fundamental considerations that need not be repeated here.

6. Round-trip Loss Statistic

This section gives the primary and overall statistic for loss performance. Additional statistics and metrics originally prepared for One-way loss MAY also be applicable.

6.1. Type-P-Round-trip-Loss-<Sample>-Ratio

Given a Type-P-Round-trip-Loss-<Sample>-Stream, the average of all the logical values, L, in the Stream is the Type-P-Round-trip-Loss-<Sample>-Ratio. This ratio is in units of lost packets per round-trip transmissions attempted.

In addition, the Type-P-Round-trip-Loss-<Sample>-Ratio is undefined if the sample is empty.

7. Round-trip Testing and One-way Reporting

This section raises considerations for results collected using a round-trip measurement architecture, such as in TWAMP [RFC5357].

The sampling process for the return path (Dst->Src) is a conditional process that depends on successful packet arrival at the Dst and correct operation at the Dst to generate the reflected packet. Therefore, the sampling process for the return path will be significantly affected when appreciable loss occurs on the Src->Dst path, making an attempt to assess the return path performance invalid

(for loss or possibly any metric).

Further, the sampling times for the return path (Dst->Src) are a random process that depends on the original sample times (TstampSrc), the one-way-delay for successful packet arrival at the Dst, and time taken at the Dst to generate the reflected packet. Therefore, the sampling process for the return path will be significantly affected when appreciable delay variation occurs on the Src->Dst path, making an attempt to assess the return path performance invalid (for loss or possibly any metric).

8. Security Considerations

8.1. Denial of Service Attacks

This metric requires a stream of packets sent from one host (source) to another host (destination) through intervening networks, and back. This method could be abused for denial of service attacks directed at the destination and/or the intervening network(s).

Administrators of source, destination, and the intervening network(s) should establish bilateral or multi-lateral agreements regarding the timing, size, and frequency of collection of sample metrics. Use of this method in excess of the terms agreed between the participants may be cause for immediate rejection or discard of packets or other escalation procedures defined between the affected parties.

8.2. User Data Confidentiality

Active use of this method generates packets for a sample, rather than taking samples based on user data, and does not threaten user data confidentiality. Passive measurement must restrict attention to the headers of interest. Since user payloads may be temporarily stored for length analysis, suitable precautions MUST be taken to keep this information safe and confidential. In most cases, a hashing function will produce a value suitable for payload comparisons.

8.3. Interference with the metrics

It may be possible to identify that a certain packet or stream of packets is part of a sample. With that knowledge at the destination and/or the intervening networks, it is possible to change the processing of the packets (e.g. increasing or decreasing delay) that may distort the measured performance. It may also be possible to generate additional packets that appear to be part of the sample metric. These additional packets are likely to perturb the results of the sample measurement.

To discourage the kind of interference mentioned above, packet interference checks, such as cryptographic hash, may be used.

9. IANA Considerations

Metrics defined in IETF are typically registered in the IANA IPPM METRICS REGISTRY as described in initial version of the registry [RFC4148]. However, areas for improvement of this registry have been identified, and the registry structure has to be revisited when there is consensus to do so.

Therefore, the metrics in this draft may be considered for registration in the future, and no IANA Action is requested at this time.

10. Acknowledgements

11. References

11.1. 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.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.

- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010.

11.2. Informative References

- [RFC5474] Duffield, N., Chiou, D., Claise, B., Greenberg, A., Grossglauser, M., and J. Rexford, "A Framework for Packet Selection and Reporting", RFC 5474, March 2009.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [Stats] McGraw-Hill NY NY, "Introduction to the Theory of Statistics, 3rd Edition,", 1974.

Author's Address

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/>

