

Internet Engineering Task Force (IETF)
Request for Comments: 6673
Category: Standards Track
ISSN: 2070-1721

A. Morton
AT&T Labs
August 2012

Round-Trip Packet Loss Metrics

Abstract

Many user applications (and the transport protocols that make them possible) require two-way communications. To assess this capability, and to achieve test 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 round-trip packet loss metric specified according to the RFC 2330 framework.

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

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/rfc6673>.

Copyright Notice

Copyright (c) 2012 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	4
1.2.	Requirements Language	5
2.	Scope	5
3.	Common Specifications for Round-Trip Metrics	5
3.1.	Name: Type-P-*	5
3.2.	Metric Parameters	5
3.3.	Metric Definition	6
3.4.	Metric Units	6
4.	A Singleton Round-Trip Loss Metric	7
4.1.	Name: Type-P-Round-trip-Loss	7
4.2.	Metric Parameters	7
4.3.	Definition and Metric Units	7
4.4.	Discussion and Other Details	8
5.	A Sample Round-Trip Loss Metric	9
5.1.	Name: Type-P-Round-trip-Loss-<Sample>-Stream	9
5.2.	Metric Parameters	9
5.3.	Definition and Metric Units	9
5.4.	Discussion and Other Details	10
6.	Round-Trip Loss Statistic	10
6.1.	Type-P-Round-trip-Loss-<Sample>-Ratio	10
7.	Round-Trip Testing and One-Way Reporting	11
8.	Measurement Considerations and Calibration	11
9.	Security Considerations	12
9.1.	Denial-of-Service Attacks	12
9.2.	User Data Confidentiality	12
9.3.	Interference with the Metrics	12
10.	IANA Considerations	13
11.	Acknowledgements	13
12.	References	13
12.1.	Normative References	13
12.2.	Informative References	14

1. Introduction

This memo defines a metric to quantify an IP network's ability to transfer packets in both directions from one host to another host. Two-way communication is almost always needed; thus, failure to transfer a packet in either direction constitutes a round-trip packet loss.

This memo defines a metric for round-trip packet 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 packet loss metric for IPPM [RFC2680] and the round-trip delay metric for IPPM [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; thus, it 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 packet 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 packet loss measurements are frequently conducted and reported in practice. The ubiquitous "ping" tools allow the measurement of round-trip packet loss and delay but usually require ICMP Echo-Request/Reply support, and ICMP packets may encounter exceptional treatment on the measurement path (see Section 2.6 of [RFC2681]). The Two-Way Active Measurement Protocol (TWAMP) specified in [RFC5357] establishes a round-trip packet loss measurement capability for the Internet. However, there is currently no round-trip packet 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 packet loss from one-way measurements conducted in each direction unless a method to match the appropriate one-way measurements has been 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 [RFC6703]. 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.

1.2. 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].

2. Scope

This memo defines a round-trip packet 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 [RFC2330] framework is for active measurement methods. Although this metric MAY be applicable in passive measurement as well, discussion of additional considerations for the passive scenario are beyond the normative scope of this memo.

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 T_f , a time (end of test interval)
- o λ , 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 T_0 , 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 T_{stampSrc} , the wire time of the packet as measured at $\text{MP}(\text{Src})$ as it leaves for Dst .
- o T_{stampDst} , the wire time of the packet as measured at $\text{MP}(\text{Dst})$, assigned to packets that arrive within a "reasonable" time (less than T_{max}).
- o T_{max} , a maximum waiting time for packets to arrive at Src , set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost).
- o M , the total number of packets sent between T_0 and T_f
- o N , the total number of packets received at Dst (sent between T_0 and T_f)
- 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 where needed (e.g., Section 6.1 in the case of $\text{Type-P-Round-trip-Loss-}<\text{Sample}>\text{-Ratio}$).

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 sent a Type-P packet back to the Src as quickly as possible (certainly less than T_{max} , and fast enough for the intended purpose), and
- o that Src received the last bit of the reflected packet prior to 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 $T_{stampSrc} + T_{max}$.

Possible causes for the Loss = 1 outcome are as follows:

- o the Dst did not receive that packet,
- o the Dst did not send 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 Section 2.4 of [RFC2681], we make the simplifying assertion that round-trip loss measured between two hosts is equal regardless of the host that originates the test:

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

(and agree with the rationale presented there -- 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.

We add the following guidance regarding the responder process to "send a Type-P packet back to the Src as quickly as possible".

A response that was not generated within Tmax is inadequate for any realistic test, and the Src will discard such responses. A responder that serves typical round-trip packet loss testing (which is relevant to higher-layer application performance) SHOULD produce a response in 1 second or less. A responder that is unable to satisfy this requirement SHOULD log the fact so that an operator can adjust the load and priorities as necessary. Analysis of responder timestamps [RFC5357] that finds responses are not generated in a timely fashion SHOULD result in operator notification, and the operator SHOULD suspend tests to the responder, since it may be overloaded. Additional measurement considerations are described in Section 8 below.

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 the following:

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 Section 11.1.1 of [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 <Sample> SHALL be replaced with the process used to define the sample, using one of the above processes (or another sample process meeting the criteria in Section 11.1 of [RFC2330], the details of which MUST be reported with the results 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

and where <Sample> SHALL be replaced with "Poisson", "Periodic", or an appropriate term to designate another sample method as described in Section 5 above.

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. However, when these references were approved, the packet reordering metrics in [RFC4737] had not yet been defined, nor had reordering been addressed in IPPM methodologies.

[RFC4737] defines packets that arrive "late" with respect to their sending order as reordered -- for example, when packets arrive with sequence numbers 4, 7, 5, 6, then packets 5 and 6 are reordered, and they are obviously not lost because they have arrived within some reasonable waiting time threshold. The presence of reordering on a round-trip path has several likely effects on the measurement.

1. Methods of measurement should continue to wait the specified time for packets and avoid prematurely declaring round-trip packet loss when a sequence gap or error is observed.
2. The time distribution of the singletons in the sample has been significantly changed.
3. Either the original packet stream or the reflected packet stream experienced path instability, and the original conditions may no longer be present.

Measurement implementations MUST address the possibility of packet reordering and avoid related errors in their processes.

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 actually 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 reverse 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 reverse path will be significantly affected when appreciable loss occurs on the Src->Dst path, making an attempt to assess the reverse path performance invalid (for loss or possibly any metric).

Further, the sampling times for the reverse 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 reverse path will be significantly affected when appreciable delay variation occurs on the Src->Dst path, making an attempt to assess the reverse path performance invalid (for loss or possibly any metric).

As discussed above in Section 5.4, packet reordering is always a possibility. In addition to the severe delay variation that usually accompanies it, reordering on the Src->Dst path will cause a misalignment of sequence numbers applied at the Dst when compared to the sender numbers. Measurement implementations MUST address this possible outcome.

8. Measurement Considerations and Calibration

Prior to conducting this measurement, the participating hosts MUST be configured to send and receive test packets of the chosen Type-P. Standard measurement protocols are capable of this task [RFC5357], but any reliable method is sufficient (e.g., if the issues with ICMP discussed in Section 2.6 of [RFC2681] can be alleviated, and the requirements of Sections 4.3 and 4.4 above are met, then ICMP could be used).

Two key features of the host that receives test packets and returns them to the originating host are described in Section 4.2 of [RFC5357]. Every received test packet MUST result in a responding packet, and the response MUST be generated as quickly as possible. This implies that interface buffers will be serviced promptly and that buffer discards will be extremely rare. These features of the

measurement equipment MUST be calibrated according to Section 3.7.3 of [RFC2679] when operating under a representative measurement load (as defined by the user). Both unexpected test packet discards, and the systematic and random errors and uncertainties, MUST be recorded.

We note that Section 4.2.1 of [RFC5357] specifies a method to collect all four significant timestamps needed to describe a packet's round-trip delay [RFC2681] and remove the processing time incurred at the responding host. This information supports the measurement of the corresponding one-way delays encountered on the round-trip path, which can identify path asymmetry or unexpected processing time at the responding host.

9. Security Considerations

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

9.2. 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.

9.3. 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) in a way 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.

Authentication or encryption techniques, such as digital signatures, MAY be used where appropriate to guard against injected traffic attacks. [RFC5357] includes both authentication and encryption features.

10. IANA Considerations

Metrics previously defined in the IETF were registered in the IANA IPPM Metrics Registry; however, this process was discontinued when the registry structure was found to be inadequate, and the registry was declared obsolete [RFC6248].

Although the metrics in this document may be considered for some form of registration in the future, no IANA action is requested at this time.

11. Acknowledgements

The author thanks Tiziano Ionta for his careful review of this memo, primarily resulting in the development of measurement considerations using TWAMP [RFC5357] as an example method. The reviews of Adrian Farrel and Benoit Claise also contributed to the clarity of the memo.

12. References

12.1. 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.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

12.2. Informative References

- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, April 2011.
- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of View", RFC 6703, August 2012.

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/>