

One-Way Metric Applicability Statement

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). 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/1id-abstracts.txt>

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

Abstract

Active traffic measurements are starting to become more widely used to ascertain network performance characteristics. All active measurement systems have the capability to measure one-way delay and one-way loss metrics, as defined in [RFC2679](#) [1] A One-way Delay Metric for IPPM and [RFC 2680](#) [2] A One-way Packet Loss Metric for IPPM, respectively. To ensure that the resulting numbers have some meaning, we attempt to characterize how the measurements are taken and what would ensure that the end numbers are indeed meaningful. This document describes an applicability statement (formerly known as best current practices) for measuring the one-way delay and one-way loss metrics in operational networks.

Overview

As more people start measuring one-way delay and one-way loss parameters it results in a large set of numbers. To ensure that these numbers have some meaning, we attempt to characterize how the measurements are taken and what would ensure that the end numbers are indeed meaningful. Much of the work relates to [RFC2679](#) [1] A One-way Delay Metric for IPPM and

[RFC2680](#)[2] A One- way Packet Loss Metric for IPPM. It is assumed that the reader is familiar with both of these documents, as well as the related framework document [RFC2330](#)[3].

Conventions used in this 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 [RFC-2119](#) [4].

1. Introduction and Terminology

Active traffic measurements are starting to become more widely used to ascertain network performance characteristics. All active measurement systems have the capability to measure one-way delay and one-way loss metrics, as defined in [RFC2679](#) [1] and [RFC 2680](#) [2], respectively. However, while these standards define how to measure quantities, there are a large number of parameters that have to be set by the operator of a measurement device. To ensure that the resulting numbers have some meaning, we attempt to characterize how the measurements are taken and what would ensure that the end numbers are indeed meaningful. This document describes best current practices for measuring the one-way delay and one-way loss metrics in operational networks.

2. Ambiguities in one-way measurement metrics

[RFC2679](#)[1] and [RFC2680](#)[2] define metrics for one-way delay and one-way loss, respectively. In practice, a large number of instances of these metrics are measured and when comparing results from different measurement entities, the numbers sometimes vary. This is partly due to ambiguities in the current documents for variables such as frequency of measurement samples, packet size, timing issues, test duration and data volumes. This draft will give recommendations for these variables for both inter-provider networks and internal networks. Inter-provider networks are those where the measurement end-points cross administrative domain boundaries, such as from one ISP to another ISP. Internal networks are those where the measurement end- points are contained within one administrative domain. This draft also discusses ambiguity issues related to reporting the metrics, such as when is a result different, alarms and sigma, average percentiles.

3. Recommendations for one way delay and loss measurements.

3.1 Measurement samples

The number of measurement samples need to be clearly defined. Specifically, we need to specify how many packets are needed to say

something about a connection. The frequency of packets should be such that one has a reasonable chance to see effects on the link but low enough that the regular traffic on the link is not affected by the measurement. In addition, it is important to ascertain what a reasonable number of packets to send, before the probability of a statistical fluke becomes small, is.

[Question: Can we benefit from packet sampling BOF work here? Ideally, math to calculate that if an effect occurs with a rate of N Hz and we send traffic with M Hz, there is a probability >X that one packet will see this effect.]

3.2 Packet size

The size of the packets is important as some devices tend to give preferential treatment to smaller packets, thus causing the delay for small packets to appear lower than for large packets, as well as overtaking or reordering. In all cases, packet sizes should be smaller than the MTU to avoid effects due to fragmentation and reassembly.

Before running any actual measurements, one should perform tests to see if delay depends on packet size other than scaling with the packet size. If this appears to be the case, one should try to estimate packet sizes for "user" data using passive measurements and adjust the packet size accordingly, or use a variable packet size according to the distribution seen in user data. These tests should be repeated when the path between source and destination changes.

Also note that some line card designs have buffer pools of different sizes. This can lead to loss being different for different packet sizes.

When packets are sent larger than the minimum size required by the measurement device, the remainder of the packet should be padded with random bits in order to avoid compression being applied to any measurement packets. The algorithm to generate these random bits as well as any seed values have to be known, in order to be able to fully understand any remaining issues with compression.

3.3 Timing issues

The measured metric should report experimental errors on the accuracy of the clocks. This has been seen to only be an issue during measurement test start-up. In the case of using NTP, it starts with an estimate and as the clock starts to stabilize it corrects the internal clock of the device.

When the IPDV metric is being measured, one use 4 time-stamps: send and

arrival time of the first packet and, send and arrival time of the second packet. The difference between these time-stamps will be small. One should take care that sufficient accuracy for the calculation is available and check that the experimental error on the overall result is still small compared to the result.

The clock should be checked for correct performance at regular intervals and measurements should be discarded when there is a problem.

One should check if the overall experimental error is small compared to the delay before further processing of the data. The errors should be recorded so they are available when calculating derived metrics such as IPDV.

[3.4](#) Test duration

The test duration can be infinitely long depending on the metric and application. In order to easily see traffic variations, measurements should run for a long time but have a limited life-time. The former requirement makes it easier to use the data for traffic engineering or load balancing.

The latter requirement allows for a easy failure detection: suppose one is measuring between A and B. At some point in time, B stops receiving packets. Until the measurement session times out, there is no way to tell if this is due to full connectivity loss between A and B, or due to a failure of the device A. When the measurement session ends, one can attempt to restart it. If one can contact the host at A, one can conservatively assume that A crashed.

How to report intermediate results while the test is in progress?

[3.5](#). Data volumes

It is important to ensure that any measurement traffic does not interfere with normal network operations. Initially, one should check if outgoing/incoming data volume for a box is small with respect to link capacity of the first few hops to avoid measurements being affected by loaded links. Also, one should check that the machine sending/receiving the data can cope with the expected offered load. Lastly, make sure that the total test traffic volume sent or received by a machine is small compared to total link capacity, a number of 3% of the total available capacity seems reasonable for routine monitoring of the performance of a link without affecting the performance of that link.

Capacity and reordering measurements that fill a link at (almost) its maximum line rate should not be used on production networks except during scheduled maintenance or test periods.

4. Reporting metrics

4.1. When is a result different?

Given 2 sets of measurements, when is set 1 statistically different from set 2?

When do you have reasonable probability that things have not changed or are OK with your network? This might vary from application to application of the data.

4.2. Alarms

From the previous paragraph, it follows when 2 results are different. This can be used to define thresholds for delay alarms.

4.3. Average/Sigma versus 2.5/median/97.5%

Since Average/Sigma for a one-way delay distribution is not well defined, and percentiles are, we should use the latter.

If it necessary to use Average/Sigma, then it should be specified how losses are treated in the calculation.

Question: what about the loss metrics: average/sigma or percentiles.

Question: Larry Dunn suggest filtering theory to get a feeling for the shape of a curve. Anybody who wants to elaborate?

5.0 Reporting the IPDV metric.

Using average/sigma for reporting the IPDV metric does not work: first of all, the average will almost always be close to zero. Then, the distribution generally is not Gaussian and the sigma is not well defined for the distributions that are being seen.

Using percentiles suffers from the same problem: the median will almost always be 0, and the 2.5 and 97.5% will be the same.

What appears to be working is 2 percentiles, for example 5 and 25%, this gives a reasonable description of the shape of the distribution.

Question: Stas: do you have some better wording?

6.0 Access to the data

Measurement results comprise of both raw data and derived results. The raw data should be kept accessible to allow for historical trend

analysis.

A minimum set of informative fields to be stored is:

- * IP address of source
- * IP address of destination
- * Time the packet was sent (or arrived)
- * Delay
- * Experimental error on sending and receiving clock
- * Packet Size
- * ...

7.0. Control/Configuration

Define maximal acceptable time to set up a measurement, latency between configuration changes and effect on measurement. No idea what the answer is, this might depend from operator to operator.

8. IANA Considerations

NONE at the moment.

9. Security Considerations

One-way delay packets can be used as a DDOS. Even if each sending box carefully checks that the outgoing rate to a destination is small, a large number of sending boxes can still be used to overflow a link. To protect against this, send configuration to receiving device before the measurements start.

Other Sanity checks? what are they?

10. References

- [1] [RFC2679](#)
- [2] [RFC2680](#)
- [3] [RFC2330](#)
- [4] [RFC2119](#)

11. Acknowledgments

Victor Reijs (HEANET) July 9's comments incorporated. Stanislav Shalunov's comments from July 26 added, Aug 8 added.

12. Authors' Addresses

Henk Uijterwaal
RIPE Network Coordination Centre
Singel 258
[1016 AB Amsterdam](#)
The Netherlands

Phone: +31.20.5354414
Fax: +31.20.5354445
Email: henk.uijterwaal@ripe.net

Merike Kaeo
Merike, Inc.
[123 Ross Street](#)
Santa Cruz, CA 95060
USA

Phone: +1 831 818 4864
Fax: +1 831 457 2654
Email: kaeo@merike.com

Full Copyright Statement Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

