

Network Working Group
Internet Draft
Expiration Date: April 1998

G. Almes, Advanced Network & Services
S. Kalidindi, Advanced Network & Services
November 1997

A Packet Loss Metric for IPPM
<[draft-ietf-ippm-loss-01.txt](#)>

1. Status of this Memo

This document is an Internet Draft. 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''.

To learn the current status of any Internet Draft, please check the ``[id-abstracts.txt](#)'' listing contained in the Internet Drafts shadow directories on [ftp.is.co.za](#) (Africa), [nic.nordu.net](#) (Europe), [munnari.oz.au](#) (Pacific Rim), [ds.internic.net](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

2. Introduction

This memo defines a metric for packet loss across Internet paths. It builds on notions introduced and discussed in the IPPM Framework document (currently 'Framework for IP Performance Metrics' <[draft-ietf-ippm-framework-01.txt](#)>); the reader is assumed to be familiar with that document.

This memo is intended to be very parallel in structure to a companion document for One-way Delay (currently 'A One-way Delay Metric for IPPM' <[draft-ietf-ippm-delay-01.txt](#)>); the reader is assumed to be familiar with that document.

The structure of the memo is as follows:

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.

ID

Packet Loss

November 1997

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.

resolution

measures the precision of a given clock. For example, the clock on an old Unix host might advance only once every 10 msec, and thus have a resolution of only 10 msec.

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, and this threatens accuracy. We might also speak of the skew of one clock relative to another clock, and this threatens synchronization.

[3. A Singleton Definition for One-way Packet Loss](#)

[3.1. Metric Name:](#)

Type-P-One-way-Packet-Loss

[3.2. Metric Parameters:](#)

- + Src, the IP address of a host
 - + Dst, the IP address of a host
 - + T, a time
 - + Path, the path* from Src to Dst; in cases where there is only one path from Src to Dst, this optional parameter can be omitted
- {Comment: the presence of path is motivated by cases such as with Merit's NetNow setup, in which a Src on one NAP can reach a Dst on another NAP by either of several different backbone networks. Generally, this optional parameter is useful only when several different

routes are possible from Src to Dst. Using the loose source route IP option is avoided since it would often artificially worsen the performance observed, and since it might not be supported along some paths.}

3.3. 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).

3.4. Definition:

>>The *Type-P-One-way-Packet-Loss* from Src to Dst at T [via path] is 0<< means that Src sent the first bit of a type-P packet [via path] 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 [via path] is 1<< means that Src sent the first bit of a type-P packet [via path] to Dst at wire-time T and that Dst did not receive that packet.

3.5. Discussion:

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

The following issues are likely to come up in practice:

- + 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, simple upper bounds (such as the 255 seconds theoretical upper bound on the lifetimes of IP packets [Postel: [RFC 791](#)]) 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.}
- + As with other 'type-P' metrics, the value of the metric may depend

on such properties of the packet as protocol, (UDP or TCP) port number, size, and arrangement for special treatment (as with IP precedence or with RSVP).

- + 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. Counting these packets as lost but packet with corruption in other parts of the packet as not lost would be confusing.}

ID

Packet Loss

November 1997

- + 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.
- + If the packet is fragmented and if, for whatever reason, reassembly does not occur, then the packet will be deemed lost.

3.6. 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, precedence).

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

- + Arrange that Src and Dst are moderately synchronized; that is, that they have clocks that are closely synchronized with each other and each fairly close to the actual time.
- + At the Src host, select Src and Dst IP addresses, and form a test packet of Type-P with these addresses.
- + Optionally, select a specific path and arrange for Src to send the packet over that path. {Comment: This could be done, for example,

- by installing a temporary host-route for Dst in Src's routing table.}
- + At the Dst host, arrange to receive the packet.
 - + At the Src host, place a timestamp in the prepared Type-P packet, and send it towards Dst [via first-hop].
 - + If the packet arrives within a reasonable period of time, the one-way packet-loss is taken to be zero.
 - + If the packet fails to arrive within a reasonable period of time, the one-way packet-loss is taken to be one. Note that the threshold of 'reasonable' here is a parameter of the methodology. {Comment: Or it could be part of the metric. If, however, we make it part of the metric, so that packets arriving after a given reasonable period must be counted as lost, then we reintroduce the need for a degree of clock synchronization similar to that needed for one-way delay. 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.}
- Issues such as the packet format, the means by which the path is ensured, 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 elsewhere our own work in describing such more detailed implementation techniques and we encourage others to do as well.}

[3.7. Errors and Uncertainties:](#)

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

Errors due to gross lack of synchronization between the Src and Dst hosts should be dealt with. Since the sensitivity of packet loss measurement to lack of synchronization is much less than for delay, we refer the reader to the treatment of synchronization errors in the One-way Delay metric.

[4. 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, path, 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 pseudo-random Poisson arrival process of rate λ , whose values fall between T_0 and T_f . The time interval between successive values of T will then average $1/\lambda$.

4.1. Metric Name:

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

4.2. Metric Parameters:

- + Src, the IP address of a host
- + Dst, the IP address of a host
- + Path, the path* from Src to Dst; in cases where there is only one path from Src to Dst, this optional parameter can be omitted
- + T_0 , a time
- + T_f , a time
- + λ , a rate in reciprocal seconds

4.3. Metric Units:

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

- + T, a time, and
- + 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.

4.4. Definition:

Given T_0 , T_f , and λ , we compute a pseudo-random Poisson process beginning at or before T_0 , with average arrival rate λ , and ending at or after T_f . Those time values greater than or equal to T_0 and less than or equal to T_f are then selected. At each of the times in this process, we obtain the value of Type-P-One-way-Packet-Loss at this time. The value of the sample is the sequence made up of the resulting $\langle \text{time}, \text{loss} \rangle$ pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

[4.5. Discussion:](#)

Note first that, since a pseudo-random number sequence is employed, the sequence of times, and hence the value of the sample, is not fully specified. Pseudo-random 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. {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 rates 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 rate 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, [path,] 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-Stream sample.

[4.6. Methodologies:](#)

The methodologies follow directly from:

- + the selection of specific times, using the specified Poisson arrival process, and
- + 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].

4.7. 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-time of the sending of the test packets. Problems with this process could be caused by either of several things, including problems with the pseudo-random number techniques used to generate the Poisson arrival process. The Framework document shows how to use an Anderson-Darling test for this.

5. Some Statistics Definitions for One-way Packet Loss

Given the sample metric Type-P-One-way-Packet-Loss-Stream, we now offer several statistics of that sample. These statistics are offered mostly to be illustrative of what could be done.

5.1. Type-P-One-way-Packet-Loss-Average

Given a Type-P-One-way-Packet-Loss-Stream, the average of all the L values in the Stream. In addition, the Type-P-One-way-Packet-Loss-Average is undefined if the sample is empty.

Example: suppose we take a sample and the results are:

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

Then the average would be 0.2.

Note that, since healthy Internet paths should be operating at loss rates 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.

6. Security Considerations

This memo raises no security issues.

7. Acknowledgements

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.

8. References

V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP Performance Metrics", Internet Draft <[draft-ietf-ippm-framework-01.txt](#)>, November 1997.

G. Almes and S. Kalidindi, "A One-way Delay Metric for IPPM", Internet Draft <[draft-ietf-ippm-delay-01.txt](#)>, November 1997.

D. Mills, "Network Time Protocol (v3)", [RFC 1305](#), April 1992.

J. Postel, "Internet Protocol", [RFC 791](#), September 1981.

ID	Packet Loss	November 1997
----	-------------	---------------

[9](#). Authors' Addresses

Guy Almes <almes@advanced.org>
Advanced Network & Services, Inc.
200 Business Park Drive
Armonk, NY 10504
USA
Phone: +1 914/273-7863

Sunil Kalidindi <kalidindi@advanced.org>
Advanced Network & Services, Inc.
200 Business Park Drive
Armonk, NY 10504
USA
Phone: +1 914/273-1219

