Reporting Metrics: Different Points of View
draft-morton-ippm-reporting-metrics-00

Status of this Memo

This Internet-Draft will expire on December 20, 2006.

Copyright Notice

Abstract

Consumers of IP network performance metrics have many different uses in mind. This memo categorizes the different audience points of view. It describes how the categories affect the selection of metric parameters and options when seeking info that serves their needs.

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].

Table of Contents

1. Introduction

When designing measurements of IP networks and presenting the results, knowledge of the audience is a key consideration. To present a useful and relevant portrait of network conditions, one must answer the following question:

"How will the results be used?"

There are two main audience categories:

1. Network Characterization - describes conditions in an IP network for quality assurance, troubleshooting, modeling, etc. The point-of-view looks inward, toward the network, and the consumer intends their actions there.

2. Application Performance Estimation - describes the network conditions in a way that facilitates determining affects on user applications, and ultimately the users themselves. This point-of-view looks outward, toward the user(s), accepting the network as-is. This consumer intends to estimate a network-dependent aspect of performance, or design some aspect of an application's accommodation of the network. (These are *not* application metrics, they are defined at the IP layer.)

This memo considers how these different points-of-view affect both the measurement design (parameters and options of the metrics) and statistics reported when serving their needs.

The IPPM framework [RFC2330] and other RFCs describing IPPM metrics provide a background for this memo.

2. Purpose and Scope

The purpose of this memo is to clearly delineate two points-of-view (POV) for using measurements, and describe their effect on the test and measurement design, including the selection of metric parameters and reporting the results.

The current scope of this memo is primarily limited to design and reporting of the loss and delay metrics [add refs] , but will also discuss the delay variation and reordering metrics where applicable. Sampling, or the design of the active packet stream that is the basis for the measurements, is also discussed.

3. Effect of POV on the Loss Metric

This section describes the ways in which the Loss metric can be tuned to reflect the preferences of the two audience categories, or different POV.

3.1. Loss Threshold

RFC 2680 [RFC2680] defines the concept of a waiting time for packets to arrive, beyond which they are declared lost. The text of the RFC declines to recommend a value, instead saying that "good engineering, including an understanding of packet lifetimes, will be needed in practice." Later, in the methodology, they give reasons for waiting "a reasonable period of time", and leaving the definition of "reasonable" intentionally vague.

Practical measurement experience has shown that unusual network circumstances can cause long delays. One such circumstance is when routing loops form during IGP re-convergence following a failure or drastic link cost change. Packets will loop between two routers until new routes are installed, or until the Time-to-Live (TTL) field decrements to zero. Very long delays on the order of several seconds have been measured [Casner] [Cia03].

Therefore, network characterization activities prefer a long waiting time in order to distinguish these events from other causes of loss (such as packet discard at a full queue, or tail drop). This way, the metric design helps to distinguish more reliably between packets that might yet arrive, and those that are no longer traversing the network.

It is possible to calculate a worst-case waiting time, assuming that a routing loop is the cause. We model the path between Source and Destination as a series of delays in links (t) and queues (q), as these two are the dominant contributors to delay. The normal path delay across n hops without encountering a loop, D, is

$$D = t_0 + \sum_{i=1}^{n} t_i + q_i$$

and the time spent in the loop with L hops, is

$$R = C \sum_{i}^{i+L} t_i + q_i \quad \text{where } C_{max} = \frac{(TTL - n)}{L}$$

and where C is the number of times a packet circles the loop.

If we take the delays of all links and queues as 100ms each, the TTL=255, the number of hops n=5 and the hops in the loop L=4, then

D = 1.1 sec and R ~= 51.2 sec, and D + R ~= 52.3 seconds

We note that the link delays of 100ms would span most continents, and a constant queue length of 100ms is also very generous. When a loop occurs, it is almost certain to be resolved in 10 seconds or less. The value calculated above is an upper limit for almost any realistic circumstance.

A waiting time threshold parameter, dT, set consistent with this calculation would not truncate the delay distribution (possibly causing a change in its mathematical properties), because the packets that might arrive have been given sufficient time to traverse the network.

It is worth noting that packets that are stored and deliberately forwarded at a much later time constitute a replay attack on the measurement system, and are beyond the scope of normal performance reporting.

Fortunately, application performance estimation activities are not adversely affected by the estimated worst-case transfer time. Although the designer's tendency might be to set the Loss Threshold at a value equivalent to a particular application's threshold, this specific threshold can be applied when post-processing the measurements. A shorter waiting time can be enforced by locating packets with delays longer than the application's threshold, and re-designating such packets as lost.

3.2. Errored Packet Designation

RFC 2680 designates packets that arrive containing errors as lost packets. Many packets that are corrupted by bit errors are discarded within the network and do not reach their intended destination.

This is consistent with applications that would check the payload

integrity at higher layers, and discard the packet. However, some applications prefer to deal with errored payloads on their own, and even a corrupted payload is better than no packet at all.

To address this possibility, and to make network characterization more complete, it is recommended to distinguish between packets that do not arrive (lost) and errored packets that arrive (conditionally lost).

3.3. Causes of Lost Packets

Although many measurement systems use a waiting time to determine if a packet is lost or not, most of the waiting is in vain. The packets are no-longer traversing the network, and have not reached their destination.

There are many causes of packet loss, including:

1. Queue drop, or discard

2. Corruption of the IP header, or other essential header info

3. TTL expiration (or use of a TTL value that is too small)

4. Link or router failure

After waiting sufficient time, packet loss can probably be attributed to one of these causes.

4. Effect of POV on the Delay Metric

This section describes the ways in which the Delay metric can be tuned to reflect the preferences of the two consumer categories, or different POV.

4.1. Treatment of Lost Packets

The Delay Metric RFC 2679[RFC2679] specifies the treatment of packets that do not successfully traverse the network: their delay is undefined.

" >>The *Type-P-One-way-Delay* from Src to Dst at T is undefined (informally, infinite)<< 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."

It is an accepted, but informal practice to assign infinite delay to

lost packets. We next look at how these two different treatments align with the needs of measurement consumers who wish to characterize networks or estimate application performance. Also, we look at the way that lost packets have been treated in other metrics: delay variation and reordering.

4.1.1. Application Performance

Applications need to perform different functions, dependent on whether or not each packet arrives within some finite tolerance. In other words, a receivers' packet processing forks on packet arrival:

- o Packets that arrive within expected tolerance are handled by processes that remove headers, restore smooth delivery timing (as in a de-jitter buffer), restore sending order, check for errors in payloads, and many other operations.

- o Packets that do not arrive when expected spawn other processes that attempt recovery from the apparent loss, such as retransmission requests, loss concealment, or forward error correction to replace the missing packet.

So, it is important to maintain a distinction between packets that actually arrive, and those that do not. Therefore, it is preferable to leave the delay of lost packets undefined, and to characterize the delay distribution as a conditional distribution (conditioned on arrival).

4.1.2. Network Characterization

In this discussion, we assume that both loss and delay metrics will be reported for network characterization (at least).

Assume packets that do not arrive are reported as Lost, usually as a percentage of all sent packets. If these lost packets are assigned undefined delay, then network's inability to deliver them in a timely way is captured only in the loss metric.

However, if we assign infinite delay to all lost packets, then:

- o The delay metric results are influenced by packets that arrive and those that do not.

- o The delay singleton and the loss singleton do not appear to be orthogonal.

- o The network is penalized in both the loss and delay metrics, effectively double-counting the lost packets.

Although infinity is a familiar mathematical concept, it is somewhat disconcerting to see any time-related metric reported as infinity, in the opinion of the authors. Questions are bound to arise, and tend to detract from the goal of informing the consumer with a performance report.

4.1.3. Delay Variation

[RFC3393] excludes lost packets from samples, effectively assigning an undefined delay to packets that do not arrive in a reasonable time. Section 4.1 describes this specification and its rationale.

"The treatment of lost packets as having "infinite" or "undefined" delay complicates the derivation of statistics for ipdv. Specifically, when packets in the measurement sequence are lost, simple statistics such as sample mean cannot be computed. One possible approach to handling this problem is to reduce the event space by conditioning. That is, we consider conditional statistics; namely we estimate the mean ipdv (or other derivative statistic) conditioned on the event that selected packet pairs arrive at the destination (within the given timeout). While this itself is not without problems (what happens, for example, when every other packet is lost), it offers a way to make some (valid) statements about ipdv, at the same time avoiding events with undefined outcomes."

4.1.4. Reordering

draft-ietf-ippm-reordering [I-D.ietf-ippm-reordering]defines metrics that are based on evaluation of packet arrival order, and include a waiting time to declare a packet lost (to exclude them from further processing).

If packets are assigned a delay value, then the reordering metric would declare any packets with infinite delay to be reordered, because their sequence numbers will surely be less than the "Next Expected" threshold when (or if) they arrive. But this practice would fail to maintain orthogonality between the reordering metric and the loss metric. Confusion can be avoided by designating the delay of non-arriving packets as undefined, and reserving delay values only for packets that arrive within a sufficiently long waiting time.

4.2. Preferred Statistics

Today in network characterization, the sample mean is one statistic that is almost ubiquitously reported. It is easily computed and understood by virtually everyone in this audience category. Also, the sample is usually filtered on packet arrival, so that the mean is

based a conditional distribution.

The median is another statistic that summarizes a distribution, having somewhat different properties from the sample mean.

When both the sample mean and median are available, a comparison will sometimes be informative, because these two statistics are equal only when the delay distribution is perfectly symmetrical.

4.3. Summary for Delay

From the perspectives of:

1. application/receiver analysis, where processing forks on packet arrival or time out,

2. straightforward network characterization without double-counting defects, and

3. consistency with Delay variation and Reordering metric definitions,

the most efficient practice is to distinguish between truly lost and delayed packets with a sufficiently long waiting time, and to designate the delay of non-arriving packets as undefined.

5. Sampling: Test Stream Characteristics

Network Characterization has traditionally used Poisson-distributed interpacket spacing, as this provides an unbiased sample. The average inter-packet spacing may be selected to allow observation of specific network phenomena. Other test streams are designed to sample some property of the network, such as the presence of congestion, link bandwidth, or packet reordering.

If measuring a network in order to make inferences about applications or receiver performance, then there are usually efficiencies derived from a test stream that has similar characteristics to the sender. In some cases, it is essential to synthesize the sender stream, as with Bulk Transfer Capacity estimates. In other cases, it may be sufficient to sample with a "known bias", e.g., a Periodic stream to estimate real-time application performance.

6. Reporting Results

>>>>>>>>Note: this section will have sub-sections that address the

different audience categories, for now it gives an overview for the loss and delay metrics discussed above.

For Packet Loss, the loss ratio defined in RFC 2680 is a sufficient starting point. We note that a loss ratio calculated according to [Y.1540] would exclude errored packets form the numerator. In practice, the difference between these two loss metrics is small if any, depending on whether the last link prior to the destination contributes errored packets.

For Packet Delay, we currently recommend providing both the mean delay and the median delay with lost packets designated as undefined. These would both be statistics on a conditional distribution, and the condition is arrival prior to a waiting time dT, where dT has been set to take maximum packet lifetimes into account.

7. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

8. Security Considerations

9. Acknowledgements

The authors would like to thank

10. References

10.1. Normative References

[I-D.ietf-ippm-reordering]
Morton, A., "Packet Reordering Metric for IPPM", draft-ietf-ippm-reordering-13 (work in progress), May 2006.

[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.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

10.2. Informative References

[Casner] "A Fine-Grained View of High Performance Networking, NANOG 22 Conf.; http://www.nanog.org/mtg-0105/agenda.html", May 20-22 2001.

[Cia03] "Standardized Active Measurements on a Tier 1 IP Backbone, IEEE Communications Mag., pp 90-97.", June 2003.

[Y.1540] ITU-T Recommendation Y.1540, "Internet protocol data communication service - IP packet transfer and availability performance parameters", December 2002.

Authors' Addresses

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/

Gomathi Ramachandran
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 2353
Fax: +1 732 368 1192
Email: gomathi@att.com
URI:

Phone:
Fax:
Email:
URI:

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.