Network Working Group A. Morton
Internet-Draft AT&T Labs
Intended status: Informational December 21, 2012
Expires: June 24, 2013

Rate Measurement Problem Statement
draft-ietf-ippm-rate-problem-01

Abstract

There is a rate measurement scenario which has wide-spread attention of Internet access subscribers and seemingly all industry players, including regulators. This memo presents an access rate-measurement problem statement for IP Performance Metrics. Key test protocol aspects require the ability to control packet size on the tested path and enable asymmetrical packet size testing in a controller-responder architecture.

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 3
2. Purpose and Scope 3
3. Active Rate Measurement 5
4. Measurement Method Categories 6
5. Test Protocol Control & Generation Requirements 8
6. Security Considerations 8
7. IANA Considerations 9
8. Acknowledgements 9
9. Appendix 9
10. References 9
  10.1. Normative References 9
  10.2. Informative References 10
Author's Address 10

1. Introduction

There are many possible rate measurement scenarios. This memo describes one rate measurement problem and presents a rate-measurement problem statement for IP Performance Metrics (IPPM).

The access-rate scenario or use case has wide-spread attention of Internet access subscribers and seemingly all Internet industry players, including regulators. This problem is being approached with many different measurement methods.

2. Purpose and Scope

The scope and purpose of this memo is to define the measurement problem statement for test protocols conducting access rate measurement on production networks. Relevant test protocols include [RFC4656] and [RFC5357]), but the problem is stated in a general way so that it can be addressed by any existing test protocol. This memo discusses possibilities for methods of measurement, but does not specify exact methods which would normally be part of the solution, not the problem.

We characterize the access rate measurement scenario as follows:

- o The Access portion of the network is the focus of this problem statement. The user typically subscribes to a service with bi-directional access partly described by rates in bits per second.

- o Rates at the edge of the network are several orders of magnitude less than aggregation and core portions.

- o Asymmetrical ingress and egress rates are prevalent.

- o Extremely large scale of access services requires low complexity devices participating at the user end of the path.

Today, the majority of widely deployed access services achieve rates less than 100 Mbit/s, and this is the order of magnitude for which a solution is sought now.

This problem statement assumes that the most-likely bottleneck device or link is adjacent to the remote (user-end) measurement device, or is within one or two router/switch hops of the remote measurement device.

Other use cases for rate measurement involve situations where the packet switching and transport facilities are leased by one operator

from another and the actual capacity available cannot be directly determined (e.g., from device interface utilization). These scenarios could include mobile backhaul, Ethernet Service access networks, and/or extensions of layer 2 or layer 3 networks. The results of rate measurements in such cases could be employed to select alternate routing, investigate whether capacity meets some previous agreement, and/or adapt the rate of traffic sources if a capacity bottleneck is found via the rate measurement. In the case of aggregated leased networks, available capacity may also be asymmetric. In these cases, the tester is assumed to have a sender and receiver location under their control. We refer to this scenario below as the aggregated leased network case.

Only active measurement methods will be addressed here, consistent with the IPPM working group's current charter. Active measurements require synthetic traffic dedicated to testing, and do not use user traffic.

The actual path used by traffic may influence the rate measurement results for some forms of access, as it may differ between user and test traffic if the test traffic has different characteristics, primarily in terms of the packets themselves (the Type-P described in [RFC2330]).

There are several aspects of Type-P where user traffic may be examined and directed to special treatment that may affect transmission rates. The possibilities include:

o Packet length

o IP addresses used

o Transport protocol used (where TCP packets may be routed differently from UDP)

o Transport Protocol port numbers used

This issue requires further discussion when specific solutions/ methods of measurement are proposed, but for this problem statement it is sufficient to Identify the problem and indicate that the solution may require an extremely close emulation of user traffic, in terms of the factors above.

Although the user may have multiple instances of network access available to them, the primary problem scope is to measure one form of access at a time. It is plausible that a solution for the single access problem will be applicable to simultaneous measurement of multiple access instances, but discussion of this is beyond the

current scope.

A key consideration is whether active measurements will be conducted with user traffic present (In-Service testing), or not present (Out-of-Service testing), such as during pre-service testing or maintenance that interrupts service temporarily. Out-of-Service testing includes activities described as "service commissioning", "service activation", and "planned maintenance". Both In-Service and Out-of-Service testing are within the scope of this problem.

It is a non-goal to solve the measurement protocol specification problem in this memo.

It is a non-goal to standardize methods of measurement in this memo. However, the problem statement will mandate that support for one or more categories of rate measurement methods and adequate control features for the methods in the test protocol.

3. Active Rate Measurement

This section lists features of active measurement methods needed to measure access rates in production networks.

Test coordination between source and destination devices through control messages and other basic capabilities described in the methods of IPPM RFCs [RFC2679][RFC2680] are taken as given (these could be listed later, if desired).

Most forms of active testing intrude on user performance to some degree. One key tenet of IPPM methods is to minimize test traffic effects on user traffic in the production network. Section 5 of [RFC2680] lists the problems with high measurement traffic rates, and the most relevant for rate measurement is the tendency for measurement traffic to skew the results, followed by the possibility of introducing congestion on the access link. Obviously, categories of rate measurement methods that use less active test traffic than others with similar accuracy SHALL be preferred for In-Service testing.

On the other hand, Out-of-Service tests where the test path shares no links with In-Service user traffic have none of the congestion or skew concerns, but these tests must address other practical concerns such as conducting measurements within a reasonable time from the tester's point of view. Out-of-Service tests where some part of the test path is shared with In-Service traffic MUST respect the In-Service constraints.

The **intended metrics to be measured** have strong influence over the categories of measurement methods required. For example, using the terminology of [RFC5136], a it may be possible to measure a Path Capacity Metric while In-Service if the level of background (user) traffic can be assessed and included in the reported result.

The measurement *architecture* MAY be either of one-way (e.g., [RFC4656]) or two-way (e.g., [RFC5357]), but the scale and complexity aspects of end-user or aggregated access measurement clearly favor two-way (with low-complexity user-end device and round-trip results collection, as found in [RFC5357]). However, the asymmetric rates of many access services mean that the measurement system MUST be able to evaluate performance in each direction of transmission. In the two-way architecture, it is expected that both end devices MUST include the ability to launch test streams and collect the results of measurements in both (one-way) directions of transmission (this requirement is consistent with previous protocol specifications, and it is not a unique problem for rate measurements).

The following paragraphs describe features for the roles of test packet SENDER, RECEIVER, and results REPORTER.

SENDER:

Generate streams of test packets with various characteristics as desired (see Section 4). The SENDER may be located at the user end of the access path, or may be located elsewhere in the production network, such as at one end of an aggregated leased network segment.

RECEIVER:

Collect streams of test packets with various characteristics (as described above), and make the measurements necessary to support rate measurement at the other end of an end-user access or aggregated leased network segment.

REPORTER:

Use information from test packets and local processes to measure delivered packet rates.

4. Measurement Method Categories

The design of rate measurement methods can be divided into two phases: test stream design and measurement (SENDER and RECEIVER), and a follow-up phase for analysis of the measurement to produce results (REPORTER). The measurement protocol that addresses this problem

MUST only serve the test stream generation and measurement functions.

For the purposes of this problem statement, we categorize the many possibilities for rate measurement stream generation as follows;

1. Packet pairs, with fixed intra-pair packet spacing and fixed or random time intervals between pairs in a test stream.

2. Multiple streams of packet pairs, with a range of intra-pair spacing and inter-pair intervals.

3. One or more packet ensembles in a test stream, using a fixed ensemble size in packets and one or more fixed intra-ensemble packet spacings (including zero spacing).

4. One or more packet chirps, where intra-packet spacing typically decreases between adjacent packets in the same chirp and each pair of packets represents a rate for testing purposes.

For all categories, the test protocol MUST support:

1. Variable payload lengths among packet streams

2. Variable length (in packets) among packet streams or ensembles

3. Variable IP header markings among packet streams

4. Choice of UDP transport and variable port numbers, OR, choice of TCP transport and variable port numbers for two-way architectures only, OR BOTH.

5. Variable number of packets-pairs, ensembles, or streams used in a test session

The items above are additional variables that the test protocol MUST be able to identify and control.

The test protocol SHALL support test packet ensemble generation (category 3), as this appears to minimize the demands on measurement accuracy. Other stream generation categories are OPTIONAL.

>>>>>>

Note: For measurement systems employing TCP Transport protocol, the ability to generate specific stream characteristics requires a sender with the ability to establish and prime the connection such that the desired stream characteristics are allowed. See Mathis' work in progress for more background [draft-mathis-ippm-model-based-metrics].

The general requirement statements needed to describe an "open-loop"

TCP sender require some additional discussion.

It may also be useful to specify a control for Bulk Transfer Capacity measurement with fully-specified TCP senders and receivers, as envisioned in [RFC3148], but this would be a brute-force assessment which does not follow the conservative tenets of IPPM measurement.

>>>>>>

Measurements for each test packet transferred between SENDER and RECEIVER MUST be compliant with the singleton measurement methods described in IPPM RFCs [RFC2679][RFC2680] (these could be listed later, if desired). The time-stamp information or loss/arrival status for each packet MUST be available for communication to the protocol entity that collects results.

5. Test Protocol Control & Generation Requirements

Essentially, the test protocol MUST support the measurement features described in the sections above. This requires:

1. Communicating all test variables to the Sender and Receiver

2. Results collection in a one-way architecture

3. Remote device control for both one-way and two-way architectures

4. Asymmetric and/or pseudo-one-way test capability in a two-way measurement architecture

The ability to control packet size on the tested path and enable asymmetrical packet size testing in a two-way architecture are REQUIRED.

The test protocol SHOULD enable measurement of the [RFC5136] Capacity metric, either Out-of-Service, In-Service, or both. Other [RFC5136] metrics are OPTIONAL.

6. Security Considerations

The security considerations that apply to any active measurement of live networks are relevant here as well. See [RFC4656] and [RFC5357].

There may be a serious issue if a proprietary Service Level Agreement involved with the access network segment provider were somehow leaked in the process of rate measurement. To address this, test protocols SHOULD NOT convey this information in a way that could be discovered by unauthorized parties.

7. IANA Considerations

This memo makes no requests of IANA.

8. Acknowledgements

Dave McDysan provided comments and text for the aggregated leased use case. Yaakov Stein suggested many considerations to address, including the In-Service vs. Out-of-Service distinction and its implication on test traffic limits and protocols.

9. Appendix

This Appendix was proposed to briefly summarize previous rate measurement experience. (There is a large body of research on rate measurement, so there is a question of what to include and what to omit. Suggestions are welcome.)

10. References

10.1. Normative References

[RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.

[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.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M.

Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

[RFC5618] Morton, A. and K. Hedayat, "Mixed Security Mode for the Two-Way Active Measurement Protocol (TWAMP)", RFC 5618, August 2009.

[RFC5938] Morton, A. and M. Chiba, "Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)", RFC 5938, August 2010.

[RFC6038] Morton, A. and L. Ciavattone, "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", RFC 6038, October 2010.

10.2. Informative References

[RFC3148] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, July 2001.

[RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, February 2008.

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/