

Internet Engineering Task Force (IETF)
Request for Comments: 7497
Category: Informational
ISSN: 2070-1721

A. Morton
AT&T Labs
April 2015

Rate Measurement Test Protocol Problem Statement and Requirements

Abstract

This memo presents a problem statement for access rate measurement for test protocols to measure IP Performance Metrics (IPPM). Key rate measurement test protocol aspects include the ability to control packet characteristics on the tested path, such as asymmetric rate and asymmetric packet size.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc7497>.

Copyright Notice

Copyright (c) 2015 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 2
 - 1.1. Requirements Language 3
- 2. Purpose and Scope 3
- 3. Active Rate Measurement 6
- 4. Measurement Method Categories 7
- 5. Test Protocol Control and Generation Requirements 9
- 6. Security Considerations 11
- 7. Operational Considerations 11
- 8. References 12
 - 8.1. Normative References 12
 - 8.2. Informative References 13
- Acknowledgements 13
- Author's Address 14

1. Introduction

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

When selecting a form of access to the Internet, subscribers are interested in the performance characteristics of the various alternatives. Standardized measurements can be a basis for comparison between these alternatives. There is an underlying need to coordinate measurements that support such comparisons and to test control protocols to fulfill this need. The figure below depicts some typical measurement points of access networks.



GW = Gateway

The access rate scenario or use case has received widespread attention of Internet-access subscribers and seemingly all Internet-industry players, including regulators. This problem is being approached with many different measurement methods. The eventual protocol solutions to this problem (and the systems that utilize the protocol) may not directly involve users, such as when tests reach from the infrastructure to a service-specific device, such as a residential gateway. However, no aspect of the problem precludes users from developing a test protocol controlled via command line

interfaces on both ends. Thus, a very wide range of test protocols, active measurement methods, and system solutions are the possible outcomes of this problem statement.

1.1. 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. 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, such as [RFC6812].

This memo discusses possible measurement methods but does not specify exact methods that would normally be part of the solution.

We are interested in access measurement scenarios with the following characteristics:

- o The access portion of the network is the focus of this problem statement. The user typically subscribes to a service with bidirectional access partly described by rates in bits per second. The rates may be expressed as raw capacity or restricted capacity, as described in [RFC6703]. These are the quantities that must be measured according to one or more standard metrics and for which measurement methods must also be agreed on as a part of the solution.
- o Referring to the reference path illustrated below and defined in [RFC7398], possible measurement points include a subscriber's host, the access service demarcation point, intra IP access (where a globally routable address is present), or the gateway between the measured access network and other networks.

```
Subsc. -- Private -- Private -- Access -- Intra IP -- GRA -- Transit
device   Net #1   Net #2   Demarc.   Access   GW     GRA GW
```

GRA = Globally Routable Address, GW = Gateway

- o Rates at some links near the edge of the provider's network can often be several orders of magnitude less than link rates in the aggregation and core portions of the network.

- o Asymmetrical access rates on ingress and egress are prevalent.
- o In many scenarios of interest, services of extremely large-scale access require low-complexity devices participating at the user end of the path, and those devices place limits on clock and control timing accuracy.

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 link 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.

This memo describes protocol support for active measurement methods consistent with the IPPM working group's traditional charter. Active measurements require synthetic traffic streams dedicated to testing and do not make measurements on user traffic. See Section 2 of [RFC2679], where the concept of a stream is first introduced in IPPM literature as the basis for collecting a sample (defined in Section 11 of [RFC2330]).

As noted in [RFC2330], the focus of access traffic management 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 (see Section 13 of [RFC2330] for the considerations on packet type, or Type-P).

There are several aspects of Type-P where user traffic may be examined and selected for special treatment that may affect transmission rates. Various aspects of Type-P are known to influence

Equal-Cost Multipath (ECMP) routing with possible rate measurement variability across parallel paths. Without being exhaustive, the possibilities include:

- o Packet length
- o IP addresses
- o Transport protocol (e.g., where TCP packets may be routed differently from UDP)
- o Transport-protocol port numbers

This issue requires further discussion when specific solutions/methods of measurement are proposed; 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 one or more 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 treatment of this scenario is beyond the current scope this document.

A key consideration is whether or not active measurements will be conducted with user traffic present. In-Service testing takes place with user traffic present. Out-of-Service testing occurs during pre-service assessment or during maintenance that interrupts service temporarily. Out-of-Service testing includes activities described as "service commissioning", "service activation", and "planned maintenance". Opportunistic In-Service testing (when there is no user traffic present throughout the test interval, such as outside normal business hours) is essentially equivalent to Out-of-Service testing. 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 mandates support for one category of rate measurement methods in the test protocol and adequate control features for the methods in the control protocol (assuming the control and test protocols are separate).

3. Active Rate Measurement

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

Coordination between source and destination devices through control messages and other basic capabilities described in the methods of IPPM RFCs [RFC2679] [RFC2680], and assumed for test protocols such as [RFC5357] and [RFC4656], are taken as given.

Most forms of active testing intrude on user performance to some degree, especially In-Service testing. 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 ("too much" traffic); 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. Section 4 of [RFC3148] provides additional considerations. The user of protocols for In-Service testing MUST respect these traffic constraints. Obviously, categories of rate measurement methods that use less active test traffic than others with similar accuracy are preferred for In-Service testing, and the specifications of this memo encourage traffic reduction through asymmetric control capabilities.

Out-of-Service tests where the test path shares no links with In-Service user traffic, have none of the congestion or skew concerns. Both types should address practical matters common to all test efforts, such as conducting measurements within a reasonable time from the tester's point of view and ensuring that timestamp accuracy is consistent with the precision needed for measurement [RFC2330]. Out-of-Service tests where some part of the test path is shared with In-Service traffic MUST respect the In-Service constraints described above.

The intended metrics to be measured have strong influence over the categories of measurement methods required. For example, using the terminology of [RFC5136], 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 a 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, 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 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 receiving end of an access or aggregated leased-network segment.

REPORTER:

Use information from test packets and local processes to measure delivered packet rates and prepare results in the required format (the REPORTER role may be combined with another role, most likely the SENDER).

4. Measurement Method Categories

A protocol that addresses the rate measurement problem MUST serve the test stream generation and measurement functions (SENDER and RECEIVER). The follow-up phase of analyzing the measurement results to produce a report is outside the scope of this problem and memo (REPORTER).

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, meaning that back-to-back burst ensembles and constant rate ensembles fall in this category).
4. One or more packet chirps (a set of packets with specified characteristics), where inter-packet spacing typically decreases between adjacent packets in the same chirp and each pair of packets represents a rate for testing purposes.

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.

For all supported categories, the following is a list of additional variables that the protocol(s) MUST be able to specify, control, and generate:

- a. variable payload lengths among packet streams;
- b. variable length (in packets) among packet streams or ensembles;
- c. variable IP header markings among packet streams;
- d. choice of UDP transport and variable port numbers, or choice of TCP transport and variable port numbers for two-way architectures only, or both (see below for additional requirements on TCP transport generation); and
- e. variable number of packet pairs, ensembles, or streams used in a test session.

The ability to revise these variables during an established test session is OPTIONAL, as multiple test sessions could serve the same purpose. Another OPTIONAL feature is the ability to generate streams with VLAN tags and other markings.

For measurement systems employing TCP as the 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 [IPPM-METRICS] for more background.

Beyond a simple connection handshake and the options establishment, an "open-loop" TCP sender requires the SENDER ability to:

- o generate TCP packets with well-formed headers (all fields valid), including Acknowledgement aspects;
- o produce packet streams at controlled rates and variable inter-packet spacings, including packet ensembles (back-to-back at server rate); and
- o continue the configured sending stream characteristics despite all control indications except receive-window exhaust.

The corresponding TCP RECEIVER performs normally, having some ability to configure the receive window sufficiently large so as to allow the SENDER to transmit at will (up to a configured target).

It may also be useful (for diagnostic purposes) to provide a control for the bulk transfer capacity measurement with fully-specified (and congestion-controlled) 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 [RFC2330].

Measurements for each UDP test packet transferred between SENDER and RECEIVER MUST be compliant with the singleton measurement methods described in IPPM RFCs [RFC2679][RFC2680]. The timestamp information or loss/arrival status for each packet MUST be available for communication to the REPORTER function.

5. Test Protocol Control and Generation Requirements

In summary, 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; and
4. Asymmetric packet rates in a two-way measurement architecture, or coordinated one-way test capabilities with the same effect (asymmetric rates may be achieved through directional control of packet rate or packet size).

The ability to control and generate asymmetric rates in a two-way architecture is REQUIRED. Two-way architectures are RECOMMENDED to include control and generation capability for both asymmetric and symmetric packet sizes because packet size often matters in the scope of this problem and test systems SHOULD be equipped to detect directional size dependency through comparative measurements.

Asymmetric packet size control is indicated when the result of a measurement may depend on the size of the packets used in each direction, i.e., when any of the following conditions hold:

- o there is a link in the path with asymmetrical capacity in opposite directions (in combination with one or more of the conditions below, but their presence or specific details may be unknown to the tester);
- o there is a link in the path that aggregates (or divides) packets into link-level frames and may have a capacity that depends on packet size, rate, or timing;
- o there is a link in the path where transmission in one direction influences performance in the opposite direction;
- o there is a device in the path where transmission capacity depends on packet header processing capacity (in other words, the capacity is sensitive to packet size);
- o the target application stream is nominally MTU size packets in one direction versus ACK stream in the other (noting that there are a vanishing number of symmetrical rate application streams for which rate measurement is wanted or interesting but such streams might have some relevance at this time);
- o the distribution of packet losses is critical to rate assessment;

and possibly other circumstances revealed by measurements comparing streams with symmetrical size and asymmetrical size.

Implementations may support control and generation for only symmetric packet sizes when none of the above conditions hold.

The test protocol SHOULD enable measurement of the capacity metric [RFC5136] either Out-of-Service, In-Service, or both; other metrics [RFC5136] 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].

Privacy considerations for measurement systems, particularly when Internet users participate in the tests in some way, are described in [LMAP-FRAMEWORK].

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. Operational Considerations

All forms of testing originate network traffic, either through their communications for control and results collection, from dedicated measurement packet streams, or from a combination of both types of traffic. Testing traffic primarily falls in one of two categories: subscriber traffic or network management traffic. There is an ongoing need to engineer networks so that various forms of traffic are adequately served, and publication of this memo does not change this need. Service subscribers and authorized users SHOULD obtain their network operator's or service provider's permission before conducting tests. Likewise, a service provider or third party SHOULD obtain the subscriber's permission to conduct tests, since they might temporarily reduce service quality. The protocol SHOULD communicate the permission status once the overall system has obtained it, either explicitly or through other means.

Subscribers, their service providers and network operators, and sometimes third parties, all seek to measure network performance. Capacity testing with active traffic often affects the packet transfer performance of streams traversing shared components of the test path, to some degree. The degradation can be minimized by scheduling such tests infrequently and restricting the amount of measurement traffic required to assess capacity metrics. As a result, occasional short-duration estimates with minimal traffic are preferred to measurements based on frequent file transfers of many megabytes with similar accuracy. New measurement methodologies intended for standardization should be evaluated individually for potential operational issues. However, the scheduled frequency of testing is as important as the methods used (and schedules are not typically submitted for standardization).

The new test protocol feature of asymmetrical packet size generation in two-way testing is recommended in this memo. It can appreciably reduce the load and packet processing demands of each test and therefore reduce the likelihood of degradation in one direction of the tested path. Current IETF standardized test protocols (e.g., [RFC5357] and [RFC6812]) do not possess the asymmetric size generation capability with two-way testing.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998, <<http://www.rfc-editor.org/info/rfc2330>>.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999, <<http://www.rfc-editor.org/info/rfc2679>>.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999, <<http://www.rfc-editor.org/info/rfc2680>>.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006, <<http://www.rfc-editor.org/info/rfc4656>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008, <<http://www.rfc-editor.org/info/rfc5357>>.
- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of View", RFC 6703, August 2012, <<http://www.rfc-editor.org/info/rfc6703>>.

8.2. Informative References

[IPPM-METRICS]

Mathis, M. and A. Morton, "Model Based Bulk Performance Metrics", Work in Progress, draft-ietf-ippm-model-based-metrics-04, March 2015.

[LMAP-FRAMEWORK]

Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for Large-Scale Measurement of Broadband Performance (LMAP)", Work in Progress, draft-ietf-lmap-framework-12, March 2015.

[RFC3148] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, July 2001, <<http://www.rfc-editor.org/info/rfc3148>>.

[RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, February 2008, <<http://www.rfc-editor.org/info/rfc5136>>.

[RFC6812] Chiba, M., Clemm, A., Medley, S., Salowey, J., Thombare, S., and E. Yedavalli, "Cisco Service-Level Assurance Protocol", RFC 6812, January 2013, <<http://www.rfc-editor.org/info/rfc6812>>.

[RFC7398] Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and A. Morton, "A Reference Path and Measurement Points for Large-Scale Measurement of Broadband Performance", RFC 7398, February 2015, <<http://www.rfc-editor.org/info/rfc7398>>.

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. Bill Cervený, Marcelo Bagnulo, Kostas Pentikousis (a persistent reviewer), and Joachim Fabini have contributed insightful, clarifying comments that made this a better document. Barry Constantine also provided suggestions for clarification.

Author's Address

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
United States

Phone: +1 732 420 1571
Fax: +1 732 368 1192
EMail: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>