

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 18, 2014

J. Fabini
Vienna University of Technology
A. Morton
AT&T Labs
October 15, 2013

Advanced Stream and Sampling Framework for IPPM
draft-ietf-ippm-2330-update-01

Abstract

To obtain repeatable results in modern networks, test descriptions need an expanded stream parameter framework that also augments aspects specified as Type-P for test packets. This memo proposes to update the IP Performance Metrics (IPPM) Framework with advanced considerations for measurement methodology and testing. The existing framework mostly assumes deterministic connectivity, and that a single test stream will represent the characteristics of the path when it is aggregated with other flows. Networks have evolved and test stream descriptions must evolve with them, otherwise unexpected network features may dominate the measured performance. This memo describes new stream parameters for both network characterization and support of application design using IPPM metrics.

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].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 18, 2014.

Copyright Notice

Copyright (c) 2013 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	3
1.1.	Definition: Reactive Path Behavior	3
2.	Scope	4
3.	New or Revised Stream Parameters	4
3.1.	Test Packet Type-P	6
3.1.1.	Multiple Test Packet Lengths	6
3.1.2.	Test Packet Payload Content Optimization	6
3.2.	Packet History	7
3.3.	Access Technology Change	7
3.4.	Time-Slotted Randomness Cancellation	7
4.	Quality of Metrics and Methodologies	9
4.1.	Repeatability	9
4.2.	Continuity	10
4.3.	Actionable	10
4.4.	Conservative	11
4.5.	Spatial and Temporal Composition	12
4.6.	Poisson Sampling	12
5.	Conclusions	12
6.	Security Considerations	12
7.	IANA Considerations	13
8.	Acknowledgements	13
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	14
	Authors' Addresses	14

1. Introduction

The IETF IP Performance Metrics (IPPM) working group first created a framework for metric development in [RFC2330]. This framework has stood the test of time and enabled development of many fundamental metrics, while only being updated once in a specific area [RFC5835].

The IPPM framework [RFC2330] generally relies on several assumptions, one of which is not explicitly stated but assumed: lightly loaded paths conform to the linear "delay = packet size / capacity" equation, being state/history-less (with some exceptions, firewalls are mentioned). However, this does not hold true for many modern network technologies, such as reactive paths (those with demand-driven resource allocation) and links with time-slotted operation. Per-flow state can be observed on test packet streams, and such treatment will influence network characterization if it is not taken into account. Flow history will also affect the performance of applications and be perceived by their users.

Moreover, Sections 4 and 6.2 of [RFC2330] explicitly recommend repeatable measurement metrics and methodologies. Measurements in today's access networks illustrate that methodological guidelines of [RFC2330] must be extended to capture the reactive nature of these networks. Although the proposed extensions can support methodologies to fulfill the continuity requirement stated in section 6.2 of [RFC2330], there is no guarantee. Practical measurements confirm that some link types exhibit distinct responses to repeated measurements with identical stimulus, i.e., identical traffic patterns. If feasible, appropriate fine-tuning of measurement traffic patterns can improve measurement continuity and repeatability for these link types as shown in [IBD].

We stress that this update of [RFC2330] does not invalidate or require changes to the analytic metric definitions prepared in the IPPM working group to date. Rather, it adds considerations for active measurement methodologies and expands the importance of existing conventions and notions in [RFC2330], such as "packets of Type-P".

Among the evolutionary networking changes is a phenomenon we call "reactive behavior", defined below.

1.1. Definition: Reactive Path Behavior

Reactive path behavior will be observable by the test packet stream as a repeatable phenomenon where packet transfer performance characteristics **change** according to prior observations of the packet flow of interest (at the reactive host or link). Therefore,

reactive path behavior is nominally deterministic with respect to the flow of interest. Other flows or traffic load conditions may result in additional performance-affecting reactions, but these are external to the characteristics of the flow of interest.

In practice, a sender may not have absolute control of the ingress packet stream characteristics at a reactive host or link, but this does not change the deterministic reactions present there. If we measure a path and the arrival characteristics at the reactive host/link depend on both the sending characteristics and the transfer characteristics of intervening hosts and links, then the reaction will appear less deterministic owing to the noise in the pattern at the reactive host/link.

Other than the size of the payload at the layer of interest and the header itself, packet content does not influence the measurement. Reactive behavior at the IP layer is not influenced by the TCP ports in use, for example. Therefore, the indication of reactive behavior must include the layer at which measurements are instituted.

Examples include links with Active/In-active state detectors, and hosts or links that revise their traffic serving and forwarding rates (up or down) based on packet arrival history.

Although difficult to handle from a measurement point of view, reactive paths entities are usually designed to improve overall network performance and user experience, for example by making capacity available to an active user. Reactive behavior may be an artifact of solutions to allocate scarce resources according to the demands of users, thus it is an important problem to solve for measurement and other disciplines, such as application design.

2. Scope

The scope of this memo is to describe useful stream parameters in addition to the information in Section 11.1 of [RFC2330] and described in [RFC3432] for periodic streams. The purpose is to foster repeatable measurement results in modern networks by highlighting the key aspects of test streams and packets and make them part of the IPPM performance metric framework.

3. New or Revised Stream Parameters

There are several areas where measurement methodology definition and test result interpretation will benefit from an increased understanding of the stream characteristics and the (possibly

unknown) network condition that influence the measured metrics.

1. Network treatment depends on the fullest extent on the "packet of Type-P" definition in [RFC2330], and has for some time.
 - * State is often maintained on the per-flow basis at various points in the path, where "flows" are determined by IP and other layers. Significant treatment differences occur with the simplest of Type-P parameters: packet length. Use of multiple lengths is RECOMMENDED.
 - * Payload content optimization (compression or format conversion) in intermediate segments. This breaks the convention of payload correspondence when correlating measurements made at different points in a path.
2. Packet history (instantaneous or recent test rate or inactivity, also for non-test traffic) profoundly influences measured performance, in addition to all the Type-P parameters described in [RFC2330].
3. Access technology may change during testing. A range of transfer capacities and access methods may be encountered during a test session. When different interfaces are used, the host seeking access will be aware of the technology change which differentiates this form of path change from other changes in network state. Section 14 of [RFC2330] treats the possibility that a host may have more than one attachment to the network, and also that assessment of the measurement path (route) is valid for some length of time (in Section 5 and Section 7 of [RFC2330]). Here we combine these two considerations under the assumption that changes may be more frequent and possibly have greater consequences on performance metrics.
4. Paths including links or nodes with time-slotted service opportunities represent several challenges to measurement (when service time period is appreciable):
 - * Random/unbiased sampling is not possible beyond one such link in the path.
 - * The above encourages a segmented approach to end to end measurement, as described in [RFC6049] for Network Characterization (as defined in [RFC6703]) to understand the full range of delay and delay variation on the path. Alternatively, if application performance estimation is the goal (also defined in [RFC6703]), then a stream with un-biased or known-bias properties [RFC3432] may be sufficient.

- * Multi-modal delay variation makes central statistics unimportant, others must be used instead.

Each of these topics is treated in detail below.

3.1. Test Packet Type-P

We recommend two Type-P parameters to be added to the factors which have impact on path performance measurements, namely packet length and payload type. Carefully choosing these parameters can improve measurement methodologies in their continuity and repeatability when deployed in reactive paths.

3.1.1. Multiple Test Packet Lengths

Many instances of network characterization using IPPM metrics have relied on a single test packet length. When testing to assess application performance or an aggregate of traffic, benchmarking methods have used a range of fixed lengths and frequently augmented fixed size tests with a mixture of sizes, or IMIX as described in [RFC6985].

Test packet length influences delay measurements, in that the IPPM one-way delay metric [RFC2679] includes serialization time in its first-bit to last bit time stamping requirements. However, different sizes can have a larger effect on link delay and link delay variation than serialization would explain alone. This effect can be non-linear and change the instantaneous network performance when a different size is used, or the performance of packets following the size change.

Repeatability is a main measurement methodology goal as stated in section 6.2 of [RFC2330]. To eliminate packet length as a potential measurement uncertainty factor, successive measurements must use identical traffic patterns. In practice a combination of random payload and random start time can yield representative results as illustrated in [IRR].

3.1.2. Test Packet Payload Content Optimization

The aim for efficient network resource use has resulted in deployment of server-only or client-server lossless or lossy payload compression techniques on some links or paths. These optimizers attempt to compress high-volume traffic in order to reduce network load. Files are analyzed by application-layer parsers and parts (like comments) might be dropped. Although typically acting on HTTP or JPEG files, compression might affect measurement packets, too. In particular measurement packets are qualified for efficient compression when they

use standard plain-text payload.

IPPM-conforming measurements should add packet payload content as a Type-P parameter which can help to improve measurement determinism. Some packet payloads are more susceptible to compression than others, but optimizers in the measurement path can be out ruled by using incompressible packet payload. This payload content could be either generated by a random device or by using part of a compressed file (e.g., a part of a ZIP compressed archive).

3.2. Packet History

Recent packet history and instantaneous data rate influence measurement results for reactive links supporting on-demand capacity allocation. Measurement uncertainty may be reduced by knowledge of measurement packet history and total host load. Additionally, small changes in history, e.g., because of lost packets along the path, can be the cause of large performance variations.

For instance delay in reactive 3G networks like High Speed Packet Access (HSPA) depends to a large extent on the test traffic data rate. The reactive resource allocation strategy in these networks affects the uplink direction in particular. Small changes in data rate can be the reason of more than 200% increase in delay, depending on the specific packet size.

3.3. Access Technology Change

[RFC2330] discussed the scenario of multi-homed hosts. If hosts become aware of access technology changes (e.g., because of IP address changes or lower layer information) and make this information available, measurement methodologies can use this information to improve measurement representativeness and relevance.

However, today's various access network technologies can present the same physical interface to the host. A host may or may not become aware when its access technology changes on such an interface. Measurements for paths which support on-demand capacity allocation are therefore challenging in that it is difficult to differentiate between access technology changes (e.g., because of mobility) and reactive path behavior (e.g., because of data rate change).

3.4. Time-Slotted Randomness Cancellation

Time-Slotted operation of path entities - interfaces, routers or links - in a network path is a particular challenge for measurements, especially if the time slot period is substantial. The central observation as an extension to Poisson stream sampling in [RFC2330]

is that the first such time-slotted component cancels unbiased measurement stream sampling. In the worst case, time-slotted operation converts an unbiased, random measurement packet stream into a periodic packet stream. Being heavily biased, these packets may interact with periodic behavior of subsequent time-slotted network entities[TSRC].

Time-slotted randomness cancellation (TSRC) sources can be found in virtually any system, network component or path, their impact on measurements being a matter of the order of magnitude when compared to the metric under observation. Examples of TSRC sources include but are not limited to system clock resolution, operating system ticks, time-slotted component or network operation, etc. The amount of measurement bias is determined by the particular measurement stream, relative offset between allocated time-slots in subsequent path entities, delay variation in these paths, and other sources of variation. Measurement results might change over time, depending on how accurately the sending host, receiving host, and time-slotted components in the measurement path are synchronized to each other and to global time. If path segments maintain flow state, flow parameter change or flow re-allocations can cause substantial variation in measurement results.

Practical measurements confirm that such interference limits delay measurement variation to a sub-set of theoretical value range. Measurement samples for such cases can aggregate on artificial limits, generating multi-modal distributions as demonstrated in [IRR]. In this context, the desirable measurement sample statistics differentiate between multi-modal delay distributions caused by reactive path behavior and the ones due to time-slotted interference.

Measurement methodology selection for time-slotted paths depends to a large extent on the respective viewpoint. End-to-end metrics can provide accurate measurement results for short-term sessions and low likelihood of flow state modifications. Applications or services which aim at approximating path performance for a short time interval (in the order of minutes) and expect stable path conditions should therefore prefer end-to-end metrics. Here stable path conditions refer to any kind of global knowledge concerning measurement path flow state and flow parameters.

However, if long-term forecast of time-slotted path performance is the main measurement goal, a segmented approach relying on measurement of sub-path metrics is preferred. Re-generating unbiased measurement traffic at any hop can help to reveal the true range of path performance for all path segments.

4. Quality of Metrics and Methodologies

[RFC6808] proposes repeatability and continuity as one of the metric and methodology properties to infer on measurement quality. Depending mainly on the set of controlled measurement parameters, measurements repeated for a specific network path using a specific methodology may or may not yield repeatable results. Challenging measurement scenarios for adequate parameter control include wireless, reactive, or time-slotted networks as discussed earlier in this document. This section presents an expanded definition of "repeatability" beyond the definition in [RFC2330] and an expanded examination of the [RFC2330] concept of "continuity" and its limited applicability.

4.1. Repeatability

[RFC2330] defines repeatability in a general way:

"A methodology for a metric should have the property that it is repeatable: if the methodology is used multiple times under identical conditions, the same measurements should result in the same measurements."

The challenge is to develop this definition further, such that it becomes an objective measurable criterion (and does not depend on the concept of continuity discussed below). Fortunately, this topic has been treated in other IPPM work. In BCP 176 [RFC6576], the criteria of equivalent results was agreed as the surrogate for interoperability when assessing metric RFCs for standards track advancement. The criteria of equivalence were expressed as objective statistical requirements for comparison across same implementations and independent implementations in the test plans specific to each RFC evaluated ([RFC2679] in the test plan of [RFC6808]).

The tests of [RFC6808] rely on nearly identical conditions to be present for analysis, but accept that these conditions cannot be exactly identical in the production network paths used. The test plans allow some correction factors to be applied (some statistical tests are hyper-sensitive to differences in the mean of distributions), and recognize the original findings of [RFC2330] regarding excess sample sizes.

One way to view the reliance on identical conditions is to view it as a challenge: how few parameters and path conditions need to be controlled and still produce repeatable methods/measurements?

Although the [RFC6808] test plan documented numerical criteria for equivalence, we cannot specify the exact numerical criteria for

repeatability *in general*. The process in the BCP [RFC6576] and statistics in [RFC6808] have been used successfully, and the numerical criteria to declare a metric repeatable should be agreed by all interested parties prior to measurement.

We revise the definition slightly, as follows:

"A methodology for a metric should have the property that it is repeatable: if the methodology is used multiple times under identical conditions, the methods should produce equivalent measurement results."

4.2. Continuity

In the original framework [RFC2330], the concept of continuity was introduced to provide a relaxed criteria for judging repeatability, and was described in section 6.2 of [RFC2330] as follows:

"...a methodology for a given metric exhibits continuity if, for small variations in conditions, it results in small variations in the resulting measurements."

Although there are conditions where metrics may exhibit continuity, there are others where this criteria would fail for both user traffic and active measurement traffic. Consider link fragmentation, and the non-linear increase in delay when we increase packet size just beyond the limit of a single fragment. An active measurement packet would see the same delay increase when exceeding the fragment size.

The Bulk Transfer Capacity (BTC) [RFC3148] gives another example at bottom of page 2:

"There is also evidence that most TCP implementations exhibit non-linear performance over some portion of their operating region. It is possible to construct simple simulation examples where incremental improvements to a path (such as raising the link data rate) results in lower overall TCP throughput (or BTC) [Mat98]."

Clearly, the time-slotted network elements described in section 3.4 above also qualifies as a new exception to the ideal of continuity. Therefore, we deprecate continuity as an alternate criterion on metrics, and prefer the more exact evaluation of repeatability instead.

4.3. Actionable

The IP Performance Metrics Framework [RFC2330] includes usefulness as a metric criterion:

"...The metrics must be useful to users and providers in understanding the performance they experience or provide..."

When considering measurements as part of a maintenance process, evaluation of measurement results for a path under observation can draw attention to potential performance problems "somewhere" on the path. Anomaly detection is therefore an important phase and first step which already satisfies the usefulness criterion for many metrics.

This concept of usefulness can be extended, becoming a sub-set of what we refer to as "actionable" criterion in the following. Central to maintenance is the isolation of the root cause of reported anomalies down to a specific sub-path, link or host, and metrics should support this second step as well. While detection of path anomaly may be the result of an on-going monitoring process, the second step of cause isolation consists of specific, directed on-demand measurements on components and sub-paths. Metrics must support users in this directed search, becoming actionable:

Metrics must enable users and operators to understand path performance and SHOULD help to direct corrective actions when warranted, based on the measurement results.

Besides characterizing metrics, usefulness and actionable properties are also applicable to methodologies and measurements.

4.4. Conservative

[RFC2330] adopts the term "conservative" for measurement methodologies for which:

"... the act of measurement does not modify, or only slightly modifies, the value of the performance metric the methodology attempts to measure."

It should be noted that this definition of "conservative" in the sense of [RFC2330] depends to a large extent on the measurement path's technology and characteristics. In particular, when deployed on reactive paths, sub-paths, links or hosts conforming to the definition in Section 1.1 of this document, measurement packets can originate capacity (re)allocations. In addition, small measurement flow variations can result in other users on the same path perceiving significant variations in measurement results.

4.5. Spatial and Temporal Composition

Concepts related to temporal and spatial composition of metrics in Section 9 of [RFC2330] have been extended in [RFC5835]. [RFC5835] defines multiple new types of metrics, including Spatial Composition, Temporal Aggregation, and Spatial Aggregation. So far, only the metrics for Spatial Composition have been standardized [RFC6049], providing the ability to estimate the performance of a complete path from subpath metrics. Spatial Composition aligns with the finding of [TSRC], that unbiased sampling is not possible beyond the first time-slotted link within a measurement path. In cases where measurement of subpaths is not feasible, restoring randomness of measurement samples when necessary is recommended as presented in [TSRC].

4.6. Poisson Sampling

Section 11.1.1 of [RFC2330] describes Poisson sampling, where the inter-packet send times have a Poisson distribution. A path element with reactive behavior sensitive to flow inactivity could change state if the random inter-packet time is too long. It is recommended to truncate the tail of Poisson distribution to avoid reactive element state changes. Truncation has been used without issue to ensure that minimum sample sizes can be attained in a fixed test interval.

5. Conclusions

Safeguarding repeatability as a key property of measurement methodologies is highly challenging and sometimes impossible in reactive paths. Measurements in paths with demand-driven allocation strategies must use a prototypical application packet stream to infer a specific application's performance. Measurement repetition with unbiased network and flow states (e.g., by rebooting measurement hosts) can help to avoid interference with periodic network behavior, randomness being a mandatory feature for avoiding correlation with network timing. Inferring the path performance between one measurement session or packet stream and other streams with alternate characteristics is generally discouraged with reactive paths because of the huge set of global parameters which have influence instantaneous path performance.

6. Security Considerations

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

7. IANA Considerations

This memo makes no requests of IANA.

8. Acknowledgements

The authors thank Rudiger Geib, Matt Mathis and Konstantinos Pentikousis for their helpful comments on this draft.

9. References

9.1. Normative References

- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [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.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.
- [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.
- [RFC5657] Dusseault, L. and R. Sparks, "Guidance on Interoperation and Implementation Reports for Advancement to Draft Standard", BCP 9, RFC 5657, September 2009.
- [RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric

Composition", RFC 5835, April 2010.

- [RFC6049] Morton, A. and E. Stephan, "Spatial Composition of Metrics", RFC 6049, January 2011.
- [RFC6576] Geib, R., Morton, A., Fardid, R., and A. Steinmitz, "IP Performance Metrics (IPPM) Standard Advancement Testing", BCP 176, RFC 6576, March 2012.
- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of View", RFC 6703, August 2012.

9.2. Informative References

- [IBD] Fabini, J., Karner, W., Wallentin, L., and T. Baumgartner, "The Illusion of Being Deterministic - Application-Level Considerations on Delay in 3G HSPA Networks", Lecture Notes in Computer Science, Springer, Volume 5550, 2009, pp 301-312, May 2009.
- [IRR] Fabini, J., Wallentin, L., and P. Reichl, "The Importance of Being Really Random: Methodological Aspects of IP-Layer 2G and 3G Network Delay Assessment", ICC'09 Proceedings of the 2009 IEEE International Conference on Communications, doi: 10.1109/ICC.2009.5199514, June 2009.
- [RFC3148] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, July 2001.
- [RFC6808] Ciavattone, L., Geib, R., Morton, A., and M. Wieser, "Test Plan and Results Supporting Advancement of RFC 2679 on the Standards Track", RFC 6808, December 2012.
- [RFC6985] Morton, A., "IMIX Genome: Specification of Variable Packet Sizes for Additional Testing", RFC 6985, July 2013.
- [TSRC] Fabini, J. and M. Abmayer, "Delay Measurement Methodology Revisited: Time-slotted Randomness Cancellation", IEEE Transactions on Instrumentation and Measurement doi: 10.1109/TIM.2013.2263914, October 2013.

Authors' Addresses

Joachim Fabini
Vienna University of Technology
Gusshausstrasse 25/E389
Vienna, 1040
Austria

Phone: +43 1 58801 38813
Fax: +43 1 58801 38898
Email: Joachim.Fabini@tuwien.ac.at
URI: <http://www.tc.tuwien.ac.at/about-us/staff/joachim-fabini/>

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/>