

Internet Engineering Task Force (IETF)
Request for Comments: 7799
Category: Informational
ISSN: 2070-1721

A. Morton
AT&T Labs
May 2016

Active and Passive Metrics and Methods
(with Hybrid Types In-Between)

Abstract

This memo provides clear definitions for Active and Passive performance assessment. The construction of Metrics and Methods can be described as either "Active" or "Passive". Some methods may use a subset of both Active and Passive attributes, and we refer to these as "Hybrid Methods". This memo also describes multiple dimensions to help evaluate new methods as they emerge.

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/rfc7799>.

Copyright Notice

Copyright (c) 2016 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. Terms and Definitions	3
3.1. Performance Metric	3
3.2. Method of Measurement	4
3.3. Observation Point	4
3.4. Active Methods	4
3.5. Active Metric	5
3.6. Passive Methods	5
3.7. Passive Metric	6
3.8. Hybrid Methods and Metrics	6
4. Discussion	8
4.1. Graphical Representation	8
4.2. Discussion of PDM	10
4.3. Discussion of "Coloring" Method	11
4.4. Brief Discussion of OAM Methods	11
5. Security Considerations	12
6. References	12
6.1. Normative References	12
6.2. Informative References	13
Acknowledgements	14
Author's Address	14

1. Introduction

The adjectives "Active" and "Passive" have been used for many years to distinguish between two different classes of Internet performance assessment. The first Passive and Active Measurement (PAM) Conference was held in 2000, but the earliest proceedings available online are from the second PAM conference in 2001
<<https://www.ripe.net/ripe/meetings/pam-2001>>.

The notions of "Active" and "Passive" are well-established. In general:

- o An Active Metric or Method depends on a dedicated measurement packet stream and observations of the stream.
- o A Passive Metric or Method depends **solely** on observation of one or more existing packet streams. The streams only serve measurement when they are observed for that purpose, and are present whether or not measurements take place.

As new techniques for assessment emerge, it is helpful to have clear definitions of these notions. This memo provides more-detailed definitions, defines a new category for combinations of traditional Active and Passive techniques, and discusses dimensions to evaluate new techniques as they emerge.

This memo provides definitions for Active and Passive Metrics and Methods based on long usage in the Internet measurement community, and especially the Internet Engineering Task Force (IETF). This memo also describes the combination of fundamental Active and Passive categories that are called Hybrid Methods and Metrics.

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 [RFC2119].

2. Purpose and Scope

The scope of this memo is to define and describe Active and Passive versions of metrics and methods that are consistent with the long-time usage of these adjectives in the Internet measurement community and especially the IETF. Since the science of measurement is expanding, we provide a category for combinations of the traditional extremes, treating Active and Passive as a continuum and designating combinations of their attributes as Hybrid Methods.

Further, this memo's purpose includes describing multiple dimensions to evaluate new methods as they emerge.

3. Terms and Definitions

This section defines the key terms of the memo. Some definitions use the notion of "stream of interest", which is synonymous with "population of interest" defined in clause 6.1.1 of ITU-T Recommendation Y.1540 [Y.1540]. These definitions will be useful for any work in progress, such as [PASSIVE] (with which there is already good consistency).

3.1. Performance Metric

The standard definition of a quantity, produced in an assessment of performance and/or reliability of the network, which has an intended utility and is carefully specified to convey the exact meaning of a measured value. (This definition is consistent with that of Performance Metric in [RFC2330] and [RFC6390]).

3.2. Method of Measurement

The procedure or set of operations having the object of determining a Measured Value or Measurement Result.

3.3. Observation Point

See Section 2 of [RFC7011] for the definition of Observation Point (a location in the network where packets can be observed), and related definitions. The comparable term defined in IETF literature on Active measurement is "Measurement Point" (see Section 4.1 of [RFC5835]). Both of these terms have come into use describing similar actions at the identified point in the network path.

3.4. Active Methods

Active Methods of Measurement have the following attributes:

- o Active Methods generate packet streams. Commonly, the packet stream of interest is generated as the basis of measurement. Sometimes, the adjective "synthetic" is used to categorize Active measurement streams [Y.1731]. An accompanying packet stream or streams may be generated to increase overall traffic load, though the loading stream(s) may not be measured.
- o The packets in the stream of interest have fields or field values (or are augmented or modified to include fields or field values) that are dedicated to measurement. Since measurement usually requires determining the corresponding packets at multiple measurement points, a sequence number is the most common information dedicated to measurement, and it is often combined with a timestamp.
- o The Source and Destination of the packet stream of interest are usually known a priori.
- o The characteristics of the packet stream of interest are known at the Source (at least), and may be communicated to the Destination as part of the method. Note that some packet characteristics will normally change during packet forwarding. Other changes along the path are possible, see [STDFORM].

When adding traffic to the network for measurement, Active Methods influence the quantities measured to some degree, and those performing tests should take steps to quantify the effect(s) and/or minimize such effects.

3.5. Active Metric

An Active Metric incorporates one or more of the aspects of Active Methods in the metric definition.

For example, IETF metrics for IP performance (developed according to the framework described in [RFC2330]) include the Source-packet stream characteristics as metric-input parameters, and also specify the packet characteristics (Type-P) and Source and Destination IP addresses (with their implications on both stream treatment and interfaces associated with measurement points).

3.6. Passive Methods

Passive Methods of Measurement are:

- o based solely on observations of an undisturbed and unmodified packet stream of interest (in other words, the method of measurement MUST NOT add, change, or remove packets or fields or change field values anywhere along the path).
- o dependent on the existence of one or more packet streams to supply the stream of interest.
- o dependent on the presence of the packet stream of interest at one or more designated Observation Points.

Some Passive Methods simply observe and collect information on all packets that pass Observation Point(s), while others filter the packets as a first step and only collect information on packets that match the filter criteria, and thereby narrow the stream of interest.

It is common that Passive Methods are conducted at one or more Observation Points. Passive Methods to assess Performance Metrics often require multiple Observation Points, e.g., to assess the latency of packet transfer across a network path between two Observation Points. In this case, the observed packets must include enough information to determine the corresponding packets at different Observation Points.

Communication of the observations (in some form) to a collector is an essential aspect of Passive Methods. In some configurations, the traffic load generated when communicating (or exporting) the Passive Method results to a collector may itself influence the measured network's performance. However, the collection of results is not unique to Passive Methods, and the load from management and operations of measurement systems must always be considered for potential effects on the measured values.

3.7. Passive Metric

Passive Metrics apply to observations of packet traffic (traffic flows in [RFC7011]).

Passive performance metrics are assessed independently of the packets or traffic flows, and solely through observation. Some refer to such assessments as "out of band".

One example of Passive Performance Metrics for IP packet transfer can be found in ITU-T Recommendation Y.1540 [Y.1540], where the metrics are defined on the basis of reference events generated as packets pass reference points. The metrics are agnostic to the distinction between Active and Passive when the necessary packet correspondence can be derived from the observed stream of interest as required.

3.8. Hybrid Methods and Metrics

Hybrid Methods are Methods of Measurement that use a combination of Active Methods and Passive Methods, to assess Active Metrics, Passive Metrics, or new metrics derived from the a priori knowledge and observations of the stream of interest. ITU-T Recommendation Y.1540 [Y.1540] defines metrics that are also applicable to the hybrid categories, since packet correspondence at different observation/reference points could be derived from "fields or field values which are dedicated to measurement", but otherwise the methods are Passive.

There are several types of Hybrid Methods, as categorized below.

With respect to a *single* stream of interest, Hybrid Type I methods fit in the continuum as follows, in terms of what happens at the Source (or Observation Point nearby):

- o Generation of the stream of interest => Active
- o Augmentation or modification of the stream of interest, or employment of methods that modify the treatment of the stream => Hybrid Type I
- o Observation of a stream of interest => Passive

As an example, consider the case where the method generates traffic load stream(s), and observes an existing stream of interest according to the criteria for Passive Methods. Since loading streams are an aspect of Active Methods, the stream of interest is not "solely observed", and the measurements involve a single stream of interest whose treatment has been modified by the presence of the load. Therefore, this is a Hybrid Type I method.

We define Hybrid Type II as follows: Methods that employ two or more different streams of interest with some degree of mutual coordination (e.g., one or more Active streams and one or more undisturbed and unmodified packet streams) to collect both Active and Passive Metrics and enable enhanced characterization from additional joint analysis. [HYBRID] presents a problem statement for Hybrid Type II Methods and Metrics. Note that one or more Hybrid Type I streams could be substituted for the Active streams or undisturbed streams in the mutually coordinated set. It is the Type II Methods where unique Hybrid Metrics are anticipated to emerge.

Methods based on a combination of a single (generated) Active stream and Passive observations applied to the stream of interest at intermediate Observation Points are also Hybrid Methods. However, [RFC5644] already defines these as Spatial Metrics and Methods. It is possible to replace the Active stream of [RFC5644] with a Hybrid Type I stream and measure Spatial Metrics (but this was unanticipated when [RFC5644] was developed).

The table below illustrates the categorization of methods (where "Synthesis" refers to a combination of Active and Passive Method attributes).

	Single Stream of Interest	Multiple Simultaneous Streams of Interest from Different Methods
Single Fundamental Method	Active or Passive	
Synthesis of Fundamental Methods	Hybrid Type I	
Multiple Methods	Spatial Metrics [RFC5644]	Hybrid Type II

There may be circumstances where results measured with Hybrid Methods can be considered equivalent to those measured with Passive Methods. This notion references the possibility of a "class C" where packets of different Type-P are treated equally in network implementation, as described in Section 13 of [RFC2330] and using the terminology for paths from Section 5 of [RFC2330]:

Hybrid Methods of measurement that augment or modify packets of a "class C" in a host should produce results equivalent to Passive Methods of Measurement when hosts accessing and links transporting these packets along the path (other than those performing augmentation/modification) treat packets from both categories of

methods (with and without the augmentation/modification) as the same "class C". The Passive Methods of Measurement represent the Ground Truth when comparing results between Passive and Hybrid Methods, and this comparison should be conducted to confirm the "class C" treatment.

4. Discussion

This section illustrates the definitions and presents some examples.

4.1. Graphical Representation

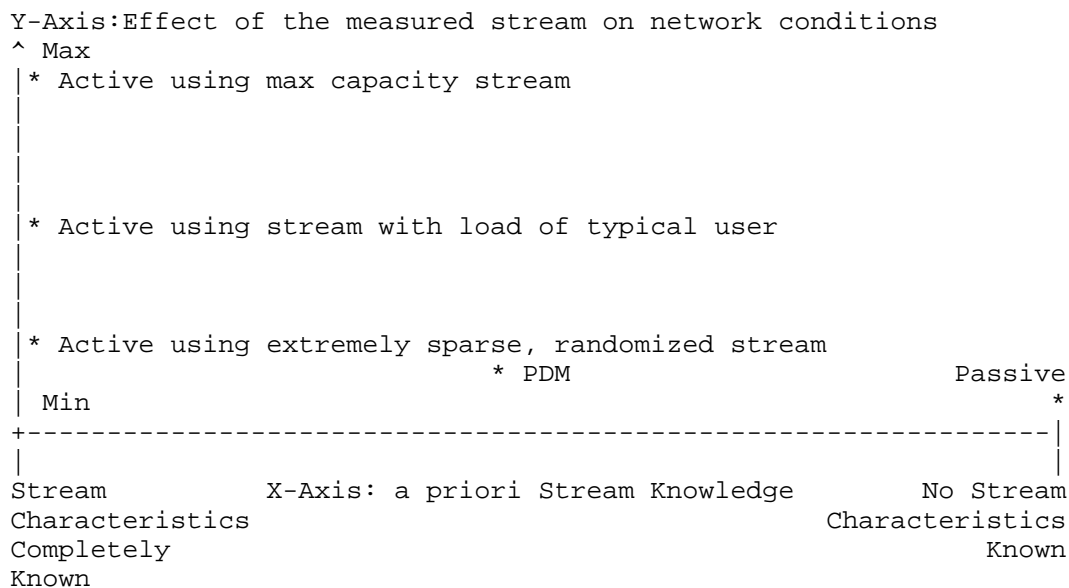
If we compare the Active and Passive Methods, there are at least two dimensions on which methods can be evaluated. This evaluation space may be useful when a method is a combination of the two alternative methods.

The two dimensions (initially chosen) are:

Y-Axis: "Effect of the measured stream on network conditions". The degree to which the stream of interest biases overall network conditions experienced by that stream and other streams. This is a key dimension for Active measurement error analysis. (Comment: There is also the notion of time averages -- a measurement stream may have significant effect while it is present, but the stream is only generated 0.1% of the time. On the other hand, observations alone have no effect on network performance. To keep these dimensions simple, we consider the stream effect only when it is present, but note that reactive networks defined in [RFC7312] may exhibit bias for some time beyond the life of a stream.)

X-Axis: "a priori Stream Knowledge". The degree to which stream characteristics are known a priori. There are methodological advantages of knowing the source stream characteristics, and having complete control of the stream characteristics. For example, knowing the number of packets in a stream allows more-efficient operation of the measurement receiver, and so is an asset for Active Methods of Measurement. Passive Methods (with no sample filter) have few clues available to anticipate what protocol the first packet observed will use or how many packets will comprise the flow; once the standard protocol of a flow is known, the possibilities narrow (for some compliant flows). Therefore, this is a key dimension for Passive measurement error analysis.

There are a few examples we can plot on a two-dimensional space. We can anchor the dimensions with reference point descriptions.



(In the graph above, "PDM" refers to [PDMOPTION], an IPv6 Option Header for Performance and Diagnostic Measurements, described in Section 4.2.)

We recognize that method categorization could be based on additional dimensions, but this would require a different graphical approach.

For example, "effect of stream of interest on network conditions" could easily be further qualified into:

1. effect on the performance of the stream of interest itself: for example, choosing a packet marking or Differentiated Services Code Point (DSCP) resulting in domain treatment as a real-time stream (as opposed to default/best-effort marking).
2. effect on unmeasured streams that share the path and/or bottlenecks: for example, an extremely sparse measured stream of minimal size packets typically has little effect on other flows (and itself), while a stream designed to characterize path capacity may affect all other flows passing through the capacity bottleneck (including itself).
3. effect on network conditions resulting in network adaptation: for example, a network monitoring load and congestion conditions might change routing, placing some flows on alternate paths to mitigate the congestion.

We have combined 1 and 2 on the Y-axis, as examination of examples indicates strong correlation of the effects in this pair, and network adaptation is not addressed.

It is apparent that different methods of IP network measurement can produce different results, even when measuring the same path at the same time. The two dimensions of the graph help us to understand how the results might change with the method chosen. For example, an Active Method to assess throughput adds some amount of traffic to the network, which might result in lower throughput for all streams. However, a Passive Method to assess throughput can also err on the low side due to unknown limitations of the hosts providing traffic, competition for host resources, limitations of the network interface, or private sub-networks that are not an intentional part of the path, etc. Hybrid Methods could easily suffer from both forms of error. Another example of potential errors stems from the pitfalls of using an Active stream with known a bias, such as a periodic stream defined in [RFC3432]. The strength of modeling periodic streams (like Voice over IP (VoIP)) is a potential weakness when extending the measured results to other application whose streams are non-periodic. The solutions are to model the application streams more exactly with an Active Method or to accept the risks and potential errors with the Passive Method discussed above.

4.2. Discussion of PDM

In [PDMOPTION], an IPv6 Option Header for Performance and Diagnostic Measurements (PDM) is described which, when added to the stream of interest at strategic interfaces, supports performance measurements. This method processes a user traffic stream and adds "fields which are dedicated to measurement" (the measurement intent is made clear in the title of this option). Thus:

- o The method intends to have a minor effect on the measured stream and other streams in the network. There are conditions where this intent may not be realized.
- o The measured stream has unknown characteristics until it is processed to add the PDM Option header. Note that if the packet MTU is exceeded after adding the header, the intent to have a minor effect will not be realized.

We conclude that this is a Hybrid Type I method, having at least one characteristic of both Active and Passive Methods for a single stream of interest.

4.3. Discussion of "Coloring" Method

[OPSAWG], proposed to color packets by re-writing a field of the stream at strategic interfaces to support performance measurements (noting that this is a difficult operation at an intermediate point on an encrypted Virtual Private Network). This method processes a user traffic stream and inserts "fields or values which are dedicated to measurement". Thus:

- o The method intends to have a minor effect on the measured stream and other streams in the network (less than PDM above). There are conditions where this intent may not be realized.
- o The measured stream has unknown characteristics until it is processed to add the coloring in the header, and the stream could be measured and time-stamped during that process.

We note that [COLORING] proposes a method similar to [OPSAWG], as discussion on the IPPM mailing list revealed.

We conclude that this is a Hybrid Type I method, having at least one characteristic of both Active and Passive Methods for a single stream of interest.

4.4. Brief Discussion of OAM Methods

Many Operations, Administration, and Management (OAM) methods exist beyond the IP layer. For example, [Y.1731] defines several different measurement methods that we would classify as follows:

- o Loss Measurement (LM) occasionally injects frames with a count of previous frames since the last LM message. We conclude LM is Hybrid Type I, because this method processes a user traffic stream and augments the stream of interest with frames having "fields which are dedicated to measurement".
- o Synthetic Loss Measurement (SLM) and Delay Measurement (DM) methods both inject dedicated measurement frames, so the "stream of interest is generated as the basis of measurement". We conclude that SLM and DM methods are Active Methods.

We also recognize the existence of alternate terminology used in OAM at layers other than IP. Readers are encouraged to consult [RFC6374] for MPLS Loss and Delay measurement terminology, for example.

5. Security Considerations

When considering the security and privacy of those involved in measurement or those whose traffic is measured, there is sensitive information communicated and observed at observation and measurement points described above, and protocol issues to consider. We refer the reader to the security and privacy considerations described in the Large-Scale Measurement of Broadband Performance (LMAP) Framework [RFC7594], which covers Active and Passive measurement techniques and supporting material on measurement context.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, 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, DOI 10.17487/RFC2330, May 1998, <<http://www.rfc-editor.org/info/rfc2330>>.
- [RFC3432] Raisanen, V., Grotfeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, DOI 10.17487/RFC3432, November 2002, <<http://www.rfc-editor.org/info/rfc3432>>.
- [RFC5644] Stephan, E., Liang, L., and A. Morton, "IP Performance Metrics (IPPM): Spatial and Multicast", RFC 5644, DOI 10.17487/RFC5644, October 2009, <<http://www.rfc-editor.org/info/rfc5644>>.
- [RFC5835] Morton, A., Ed. and S. Van den Berghe, Ed., "Framework for Metric Composition", RFC 5835, DOI 10.17487/RFC5835, April 2010, <<http://www.rfc-editor.org/info/rfc5835>>.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, DOI 10.17487/RFC6390, October 2011, <<http://www.rfc-editor.org/info/rfc6390>>.

- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<http://www.rfc-editor.org/info/rfc7011>>.
- [RFC7312] Fabini, J. and A. Morton, "Advanced Stream and Sampling Framework for IP Performance Metrics (IPPM)", RFC 7312, DOI 10.17487/RFC7312, August 2014, <<http://www.rfc-editor.org/info/rfc7312>>.
- [RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", RFC 7594, DOI 10.17487/RFC7594, September 2015, <<http://www.rfc-editor.org/info/rfc7594>>.

6.2. Informative References

- [COLORING] Chen, M., Ed., Zheng, L., Ed., Mirsky, G., Ed., Fioccola, G., Ed., and T. Mizrahi, Ed., "IP Flow Performance Measurement Framework", Work in Progress, draft-chen-ippm-coloring-based-ipfpm-framework-06, March 2016.
- [HYBRID] Trammell, B., Zheng, L., Berenguer, S., and M. Bagnulo, "Hybrid Measurement using IPPM Metrics", Work in Progress, draft-trammell-ippm-hybrid-ps-01, February 2014.
- [OPSAWG] Capello, A., Cociglio, M., Castaldelli, L., and A. Bonda, "A packet based method for passive performance monitoring", Work in Progress, draft-tempia-opsawg-p3m-04, February 2014.
- [PASSIVE] Zheng, L., Elkins, N., Lingli, D., Ackermann, M., and G. Mirsky, "Framework for IP Passive Performance Measurements", Work in Progress, draft-zheng-ippm-framework-passive-03, February 2015.
- [PDMOPTION] Elkins, N. and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", Work in Progress, draft-ietf-ippm-6man-pdm-option-02, April 2016.

- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <<http://www.rfc-editor.org/info/rfc6374>>.
- [STDFORM] Morton, A., Fabini, J., Elkins, N., Ackermann, M., and V. Hegde, "Updates for IPPM's Active Metric Framework: Packets of Type-P and Standard-Formed Packets", Work in Progress, draft-morton-ippm-2330-stdform-typep-02, December 2015.
- [Y.1540] ITU-T, "Internet protocol data communication service - IP packet transfer and availability performance parameters", March 2011, <<https://www.itu.int/rec/T-REC-Y.1540-201103-I/en>>.
- [Y.1731] ITU-T, "Operation, administration and management (OAM) functions and mechanisms for Ethernet-based networks", August 2015, <<https://www.itu.int/rec/T-REC-G.8013-201508-I/en>>.

Acknowledgements

Thanks to Mike Ackermann for asking the right question, and for several suggestions on terminology. Brian Trammell provided key terms and references for the Passive category, and suggested ways to expand the Hybrid description and types. Phil Eardley suggested some hybrid scenarios for categorization as part of his review. Tiziano Ionta reviewed the document and suggested the classification for the "coloring" Method of Measurement. Nalini Elkins identified several areas for clarification following her review. Bill Jouris, Stenio Fernandes, and Spencer Dawkins suggested several editorial improvements. Tal Mizrahi, Joachim Fabini, Greg Mirsky, and Mike Ackermann raised many key considerations in their Working Group Last Call (WGLC) reviews, based on their broad measurement experience.

Author's Address

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

Email: acmorton@att.com