

IP Performance Measurement (ippm)
Internet-Draft
Intended status: Informational
Expires: August 18, 2014

B. Trammell
ETH Zurich
L. Zheng
Huawei
S. Silva
LACNIC
M. Bagnulo
UC3M
February 14, 2014

Hybrid Measurement using IPPM Metrics
draft-trammell-ippm-hybrid-ps-01

Abstract

Hybrid measurement is the combination of metrics derived from passive and active measurement to produce a measurement result. This document discusses use cases for hybrid measurement using metrics defined within the IPPM framework, and discusses requirements for the definition of passive methodologies for selected IPPM metrics.

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 August 18, 2014.

Copyright Notice

Copyright (c) 2014 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.

1. Introduction

Hybrid measurement is the combination of metrics derived from passive and active measurement to produce a measurement result. This combination can be either spatial or temporal. For example, one way delay to a given endpoint could be derived from passive measurements from a sample of remote endpoints with which traffic is frequently exchanged, and supplemented with active measurements from endpoints with less frequent traffic, to build a "delay map" to a certain point in the network. On the temporal side, loss or delay metrics could be passively measured and stored over time to provide a baseline against which actively-measured loss or delay metrics could be compared during troubleshooting, in order to determine whether a specific path or path segment is contributing to an observed performance problem.

The IPPM working group has produced a framework [RFC2330] for a rich set of well-defined metrics (e.g. [RFC2679], [RFC2680]) for IP performance measurement using active methods, and protocols for measuring them. These metrics could form the basis of a platform for hybrid measurement, provided that passively-derived metrics were defined to be compatible with the corresponding actively-derived metrics; or alternately, provided that methodologies for passive measurement can be defined for each of the existing active metrics to be used, such that those methodologies produce values for the metrics equivalent to the active methodology for the same metric parameters, given some assumptions about the packet stream to be observed to perform the passive measurement, and given tolerances for uncertainty in the results.

2. Problem Statement

Complicating the definition of hybrid measurements is that passive measurement must make do with the traffic that is observable, while active measurement has some control over the traffic observed. Measurements for some set of parameters are not possible if no suitable traffic is observed, and the timing of the measurement cannot be controlled. Placement of the observation points for passive measurement along a path additionally introduces uncertainty in the results. For example, passive one-way delay measurement could be performed using two measurement points, one close to each endpoint, with synchronized clocks, comparing the observation times of packets via their hashes. This will not produce a value which is

directly comparable to a Type-P-One-way-Delay measured as specified in section 3.6 of [RFC2679], because it will not account for the one-way-delay from the source to the source-side observation point, or from the destination-side observation point to the destination. Any specification of hybrid measurement using IPPM metrics must handle these complications.

The proposed specification entails:

- o Definition of scenarios and requirements for hybrid measurement.
- o Selection of existing IPPM metrics to be used for the active side of hybrid measurements to meet these requirements.
- o Definition of equivalent passive measurement methodologies for each selected metric, specifically addressing the assumptions about the observed packet stream which must hold for the metric to be valid, and with a specific allowance for the measurement and/or estimation of uncertainty due to uncontrollable conditions or observation point placement.
- o Definition of metrics based on these passive methodologies, or modification of the definition of existing metrics to add passive methodologies.
- o Definition of methods for comparison between active and passive metrics allowing for estimated uncertainty.
- o Definition of methods for spatial and temporal composition of active and passive metrics together allowing for estimated uncertainty.

3. Selected IPPM Metrics

[EDITOR'S NOTE: this section will contain information on the metrics selected for passive measurement, and initial discussion of passive measurement methodologies for them. Metric definition will presumably be left for a future document.]

3.1. Packet Loss

In order to perform packet loss measurements on a live traffic flow, different approaches exist. An approach is to count the number of packets sent on one end, and the number of packets received on the other end. Packet loss over a path is the difference between the number of packets transmitted at the starting interface of the path and received at the ending interface of this path.

3.1.1.1. Passive Measurement Method

In order to count the number of packets sent and received and to compare two counters, it is required that the two counters refer exactly to the same set of packets. One difficulty is it is hard to determine exactly when to read the counter since a flow is continuous. A possible solution is to virtually split the flow in consecutive blocks by periodically inserting a delimiting packet, so that each counter refers exactly to the same block of packets. However, delimiting the flow using specific packets requires generating additional packets within the flow and requires the equipment to be able to process those packets. In addition, the method is vulnerable to out of order reception and the loss of delimiting packets.

An existing method by "coloring" IP packets for performance measurement is introduced in [I-D.templia-opsawg-p3m]. This "colored" based approach doesn't use delimiting packets. Instead, it "colours" the packets so that the packets belonging to the same block will have the same "colour", while consecutive blocks will have different colours. Each change of colour represents a sort of auto-synchronization signal that guarantees the consistency of measurements taken by different devices along the path.

3.2. One-way Delay

IPPM has defined a protocol for active one way delay measurement OWAMP in [RFC4656] It consists of a control protocol for negotiating measurement sessions and a data plane protocol for test packets. OWAMP is an active protocol meaning that the one delay is measured for artificial packets that are generated for this purpose.

It would be natural to pursue passive and/or hybrid approaches for measuring one way delay. In this case, the goal would be to measure one way delay for packets that are flowing through the network. This can be achieved by defining two observation points that will record the packets they see and the corresponding timestamps. This information will be used to determine the one way delay of the observed packets, similarly than in the active measurement approach. In order to do that, it is necessary to identify which packets are the ones that the measurement will be performed with. One way to do this is to define a certain flow of packets and then record some fields of the packets that are unlikely to change during its journey through their journey between the observation points. Once the packets have been properly identified, and the timestamp information about them has been recorded in the observation points, it should be possible to calculate the one way delay for the observed packets.

If defining a passive metric for one way delay is deemed interesting, it would be then needed to perform a gap analysis for the additional protocols that are needed for this. As the passive approach would also need to negotiate measurement sessions, it may be worth exploring the re use of OWAMP for this. Similarly, both observation points should agree what packet flow will be used for the measurement, so additional negotiation is needed. Finally, IPFIX could be used to report the results so that the actual delay can be calculated.

An additional exercise that would be then relevant is to understand how comparable are measurements obtained through the active and passive measurements. In particular, depending on the packet frequency, it may or may not be possible to achieve the different packet streams available in active measurements.

A hybrid approach for measuring one way delay seems attractive as it would be possible to measure reliably one way delay reusing the packets available in the network when they exist and generating artificial traffic when they don't exist. This requires careful consideration in order to obtain the desired packet streams and it is likely to require additional control protocol to specific the hybrid measurement.

3.3. Round-trip Delay

Round-trip delay is used to measure the expected time for network interaction between two hosts on a network; conceptually, it is equivalent to Delay in each direction between the two hosts.

Active measurement of round-trip delay as defined in [RFC2681] requires the observation of test packets transmitted in both directions between two endpoints across a network, a "source" host, which sends the first packet, and a "destination" host, which receives the first test packet and sends a test packet back to the source in reply. The round-trip delay is then calculated as the difference between the time at which the reply is received at the source and the time at which the original test packet was sent from this same source.

IPPM has defined the Two-Way Active Measurement Protocol (TWAMP) [RFC5357] for round trip delay measurement. TWAMP is essentially an extension of OWAMP for the IPPM round-trip delay metric. Like OWAMP, TWAMP consists of a control protocol to negotiate active performance measurement sessions, and a test protocol for transmission of actual test packets.

TWAMP is defined for active performance metrics, which means that the Round-trip delay is measured for packets that are generated specifically for this purpose.

3.3.1. Passive Measurement Method

The passive approach for measuring Round-trip delay would consist on measuring this delay for existent packets in contrast with the active approach in which test packets are generated. Similarly to the method used for measuring One-way Delay, for Round-trip Delay it would be needed to define two observation points that will record the packets they see and the corresponding timestamps.

The procedure for passive measurement of round-trip delay is similar to the procedure for active measurement: a packet sent from a source to a destination is recorded; that packet causes the destination to send a reply back to the source. This reply is also recorded. The packets are identifiable at the source in order to correlate each packet of the round trip in order to calculate a delay.

There are two potential architectures here; one utilizing a source Observation Point (OP) placed topologically close to the source of traffic, and one utilizing an additional destination OP placed topologically close to the destination of traffic.

In order to be able to measure the Round-trip Delay of the observed packets, it would be necessary to identify which packets will be used to perform the measurement.

4. Methodology

For certain performance metrics, many passive measurement methodologies may exist. This section give the functional requirements and design considerations of the passive measurement methodology.

4.1. Measurement Session Management

A measurement session refers to the period of time in which measurement for certain performance metrics is enabled over a forwarding path. When an interface on the measurement node is activated, the interfaces start collecting statistics. When both the upstream and downstream measurement interfaces are activated, the measurement session starts. During a measurement session, data from two active interfaces are periodically collected and the performance metrics, such as loss rate or delay, are derived. A measurement session SHOULD be started either proactively or on demand.

4.1.1. Measurement Configuration

A measurement session can be configured statically. In this case, network operators activate the two interfaces or configure their parameter settings on the relevant nodes either manually or automatically through agents of network management system (NMS). Alternatively, a measurement session can be configured dynamically. In this case, an interface may coordinate another interface on its forwarding path to start or stop a session. Accordingly, the format and process routines of the measurement session control packets need to be specified. The delivery of such packets SHOULD be reliable and it MUST be possible to secure the delivery of such packets.

4.2. Measurement Result Report

Performance reports contain streams of measurement data over a period of time. A data collection agent MAY actively poll the monitoring nodes and collect the measurement reports from all active interfaces. Alternatively, the monitoring nodes might be configured to upload the reports to the specific data collection agents once the data become available. To save bandwidth, the content of the reports might be aggregated and compressed. The period of reporting SHOULD be able to be configured or controlled by rate limitation mechanisms.

4.3. Synchronization

During a measurement session, data from the active upstream and downstream interfaces are periodically collected and the performance metrics are derived. Certain synchronization mechanism is required to ensure the data are correlated. This may further requires that the upstream and downstream interfaces having a certain time synchronization capability (e.g., supporting the Network Time Protocol (NTP) [RFC5905], or the IEEE 1588 Precision Time Protocol (PTP) [IEEE1588].) For packet delay measurement, this requirement for time synchronization is already present.

4.4. Scalability

The measurement methodology MUST be scalable. A service provider production network usually comprises of thousands of nodes. Given the scale, the collecting, processing and reporting overhead of performance measurement data SHOULD NOT overwhelm either monitoring nodes or management nodes. The volume of reporting traffic should be reasonable and not cause any network congestion.

4.5. Robustness

The measurements MUST be independent of the failure of the underlying network. For example, the correct measurement result SHOULD be generated even if some measurement coordinating packets are lost; invalid performance reports should be able to be identified in case that the underlying network is undergoing drastic changes. If dynamic measurement configuration is supported, the delivery of measurement session control packets SHOULD be reliable so that the measurement sessions can be started, ended and performed in a predictable manner.

4.6. Security

The measurement methodology MUST not impose security risks on the network. For example, the monitoring nodes should be prevented from being exploited by third parties to control measurement sessions arbitrarily, which might make the nodes vulnerable for DDoS attacks. If dynamic configuration is supported, the measurement session control packets need to be encrypted and authenticated.

5. Security Considerations

[EDITOR'S NOTE: this section will discuss general security considerations of using passive measurement for performance, both on the potential for attacks against the measurement system as well as the potential for privacy or security threats posed by the measurement system itself.]

6. IANA Considerations

This document contains no considerations for IANA.

7. References

7.1. Normative References

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

7.2. Informative References

[I-D.tempia-opsawg-p3m] Capello, A., Cociglio, M., Castaldelli, L., and A. Bonda, "A packet based method for passive performance monitoring", draft-tempia-opsawg-p3m-04 (work in progress), February 2014.

- [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.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, 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. Babiarez, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

Authors' Addresses

Brian Trammell
Swiss Federal Institute of Technology Zurich
Gloriastrasse 35
8092 Zurich
Switzerland

Phone: +41 44 632 70 13
Email: trammell@tik.ee.ethz.ch

Lianshu Zheng
Huawei Technologies
China

Email: vero.zheng@huawei.com

Sofia Silva
LACNIC
Uruguay

Email: sofia@lacnic.net

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes
Spain

Email: bagnulo@it.uc3m.es