Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 15, 2015

L. Zheng
Huawei Technologies
N. Elkins
Inside Products, Inc.
L. Deng
China Mobile
M. Ackermann
Blue Cross Blue Shield of Michigan
G. Mirsky
Ericsson

February 11, 2015

Framework for IP Passive Performance Measurements
draft-zheng-ippm-framework-passive-03

Abstract

This document describes the framework for passive measurement. In particular, the differences between passive and active measurements are analyzed, general considerations for both metric definition and measurement methodology are discussed, and requirements for various entities performing a given passive measurement task are described according to a reference model.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This document describes the framework for passive measurement. In particular, the differences between passive and active measurements are analyzed, general considerations for both metric definition and measurement methodology are discussed, and requirements for various entities performing a given passive measurement task are described according to a reference model.

The IETF IP Performance Metrics (IPPM) working group first created a framework for metric development in [RFC2330], which enabled development of many fundamental metrics. [RFC2330] has been updated once by [RFC5835], which describes a detailed framework for composing and aggregating metrics originally defined in [RFC2330].

2. Terminology

TBD

3. Measurement Methods

3.1. Active Measurement Method

Active Measurement Method: The process of measuring performance or reliability parameters by the examination of traffic (IP Packets) injected into the network, expressly for the purpose of measurement by intended Measurement Point(s).

The packets in an Active Measurement Stream typically have fields which are dedicated to and customized for measurement purposes. As an example, a sequence number is a common information field used for dedicated measurements, potentially at multiple measurement points.

Packet stream characteristics (e.g. Protocol Type) and specific field information (e.g. IP Address), are known at the source and usually communicated to the measurement point(s) as well.

Because traffic stream characteristics (e.g. number of packets), and traffic type (e.g. protocol) are known to the Active Measurement Point receivers, more efficient and focused operations are possible.

3.2. Passive Measurement Method

Passive Measurement Method: The process of measuring some performance or reliability parameter associated with the existing traffic (packets) on the network.

[Note: There are definitions for both active and passive measurement methods in [I-D.ietf-ippm-metric-registry]. Further discussion and coordination may be needed.]

Some passive methods observe and collect information on all packets that pass the observation or Measurement Point(s), while other Passive Methods filter the packets as a first step and only process information on packets that match the filter criteria.

Passive Methods may be conducted at one or more Measurement Points. Certain Metrics (e.g. latency across a particular network path), require multiple Measurement Points and observed packets must include sufficient information (e.g. sequence number), to correlate packets from different observation points.

Passive traffic may be observed/measured at any point in an IP session path, including source host, destination host and middleboxes. Passive traffic may also be observed/measured by ""Out of band"" devices, which do not participate in processing the actual session traffic. This parallel approach typically has the least effect upon network conditions and the session traffic being measured.

3.3. Hybrid Measurement Method

Hybrid Measurement Method: Methods of Measurement which use a combination of Active Methods and Passive Methods.

Hybrid Methods are not fully defined or delineated at this time. Details and examples will be forthcoming. As this occurs, this section will be expanded upon accordingly.

4. Measured Metrics

The de facto focus of RFC2330 is on active measurement. Although many of the concepts discussed in RFC2330, metrics, measurement methodology, errors with time apply to both passive and active methods of measurement techniques, there are considerable differences in terms of metric definition and measurement methodology for passive measurement.

4.1. Active Metrics

Active Metrics: A set of standard measurements for evaluating network performance or reliability, based upon the results of active traffic (IP Packets), injected into the network by a source node, expressly for the purpose of measurement and examined by one or more Measurement Points.

Examples of Active Metrics include: Latency, Throughput, errors, etc.

4.2. Passive Metrics

Passive Metrics: A set of standard measurements for evaluating network performance or reliability, based upon the results of Passive traffic (IP Packets), existing on the network and examined by one or more Measurement Points.

[Editor Note]: While Active and Passive Methods differ considerably, the Metrics requirements and definitions for Active and Passive are similar if not identical. Both can be described as defined reference events, as packets pass defined reference points. These concepts are consistent with and further elucidated by ITU-T Recommendation Y.1540 [Y.1540.2011].

Therefore it makes sense to be agnostic to the distinction between active and passive, with respect to Metrics. Distinctions or different definitions for Active and Passive Metrics, should only be created as needed, consistent with the IPPM Metric Registry [I-D.ietf-ippm-metric-registry].

Passive measurements may be used in scenarios where active measurement alone is not enough or applicable. Since no extra in-band traffic which may alter service and performance behavior is introduced, passive measurement may be done during peak traffic.

Passive measurement is not without cost. In the best scenario, the passive measurement point is external to the devices participating in the network traffic. For example, a passive network TAP may be placed at a switch to capture traffic. This would create very little, if any, interference with in-band traffic. Alternatively, care must be taken if a passive measurement technique creates load on a participant in the network. For example, a packet trace taken at one of the end host points may add load to the device thus potentially changing the environment which it is measuring. The benefits of this method for measurement and diagnostics must be weighed with the costs.

For networks where charges are based on the amount of data sent, passive measurement may be the first choice for end-to-end measurement, as it does not introduce any extra expense to the subscriber. In terms of Quality of Experience (QoE) measurement, passive measurement is expected to be more accurate and helpful in troubleshooting as it reflects the status of real application traffic.

For passive measurement, the concepts of singleton, sample and statistical, as defined in [RFC2330], also apply. However, there are some differences. The singleton, sample, and statistical measurements are those taken within the boundaries of captured traffic.

4.2.1. Passive Measurement Metric Elements

In passive measurement, the most important aspects have to do with the portion of reality which is actually measured at any point in time. So, it may be useful to define some terms for passive measurement. These are as follows:

1. Capture content: this is the type(s) of packet or metric found.

2. Capture distribution: this is the actual pattern of data in the collected packets. The pattern or distribution may be Poisson but it may also be bimodal, uniform, or skewed. For example, one might see an FTP transfer as a relatively uniform distribution, a TCP connection with a windowing issue may display a skewed distribution, etc.

3. Capture limits: this is the way the set of packets or metrics are selected. For example, one may decide to take a trace that consists of 1,000 packets. Alternatively, one might take a packet capture for 5 minutes with no regard to how many packets are found.

4. Capture methodology: this is the area in which passive differs most greatly from active methods. For example, [RFC2679], section 3.6. Methodologies discusses the various techniques of injecting test packets into the network. This is not applicable to passive measurement. Passive measurement simply collects that which exists.

5. Unruly Nature of Capture: With reality, there are no guarantees. That is, if one imagines a passive sample to be a packet trace taken at a host. If the metric one is looking for is IP/TCP connectivity measured by a TCP three way handshake, then in active measurement, one can be guaranteed to find that metric because one has injected packets of that type into the stream. In passive measurement, the capture may contain anywhere from zero occurrences of the desired metric to many instances of the desired metric.

6. Capture Selection: With active measurement, one may create 500 packets of a certain type and pick according to the sampling distribution desired.

For example, [RFC2330] in the discussion of generating Poisson distributions (11.1.3), discusses a method:

Method 1 is to proceed as follows:

1. Generate E1 and wait that long.

2. Perform a measurement.

3. Generate E2 and wait that long.

4. Perform a measurement.

5. Generate E3 and wait that long.

6. Perform a measurement ...

With passive measurement, one has no way of knowing if a particular desired packet or packet sequence exists at all in the set of packets captured.

Having said that, if there do exist many such packets, one may use a random (or another) sampling method to pick the instances desired. That is, if one has 100,000 instances of TCP three-way handshakes, one may decide to randomly choose 50 to examine more closely.

Inherent Inequality of Active and Passive Measurements: due to the nature of data traffic, depending on what metric is measured, it is unlikely that it will have a random or Poisson distribution. Hence, metrics created using Active methods and those generated using Passive methods are likely to differ. It is not known at this point whether that difference is significant or not.

[TBD: More discussion here on distributions and inequality]

. Point of View: In passive measurement, it matters greatly where the measurement is being done. Point of view is critical. Passive measurement only knows what it sees from its own perspective.

In troubleshooting problems using passive measurement, it is often necessary to get multiple points of view. Let us take a simple case of diagnosing packet loss from an end user perspective. If one takes a packet trace at the client host, one sees that certain packets are not being received. If one takes two packet traces at the same time at the server and client, one sees that the server sends these packets yet the client does not receive them. Hence, the problem must be at a middle box. So, then, one must start taking traces at client, server, and a trace point after the first middle box, etc.

The measurement techniques for passive measurement must accommodate and facilitate such tasks.

Active measurement techniques know clearly the measurement point and path because that is a part of the definition of the Active measurement task.

5. Reference Model

This section describes the main functional components of the passive measurement system, and the interactions between the components. Some new terms are defined in this document and some are borrowed from the LMAP Framework [I-D.ietf-lmap-framework].

```
           +---------------+                +---------------+
           |  Measurement  |  Coordination  |  Measurement  |
           |   Agent A     |--------------|   Agent B     |
           +---------------+                +---------------+
                     ^  |                       ^    |
            Control  |  | Report        Control |    | Report
                     |  |     +-----------------+    |
                     |  +-----|-------------------+  |
                     |        |                   |  |
                     v        v                   v  v
                   +------------+           +------------+
                   | Controller |---------|  Collector |
                   +------------+           +------------+
```

Figure 1: Passive Measurement Reference Model

Although there are considerable similarities between the proposed reference model and the LMAP framework [I-D.ietf-lmap-framework], it should be noted that the above architecture is provided as a more general outline of an integral collection of functional components collaborating in performing a specific instance of passive measurement method. Various functions from LMAP framework in performing a passive measurement task represent a specific way of realizing the general model.

Controller: A entity that exchanges the Control of the Measurement Task with the Measurement Entity, receives the Report from the Collector and conducts the value calculation/derivation for the metrics measured of the Measurement Task. When multiple Measurement Entities are involved for a certain Measurement Task, Controller may only have Control exchanged with one or some of the Measurement Entities.

Collector: A entity that receives a Report from a Measurement Entity and provides the Report to the Controller for metric calculation / derivation.

Measurement Agent: An entity that exchanges the Control of the Measurement Task with the Controller, performs Measurement Tasks and sends the Report to Collector. When multiple Measurement Agents are involved for a certain Measurement Task, Coordination may be required between Measurement Entities.

Control: The collective description of information exchanged between Controller and Measurement Agent, i.e. configurations, instructions, states, etc. for a Measurement Agent to perform and Report Measurement Tasks.

Coordination: [TBD. Discuss coordination with MAs and Controller]

Report: The set of Measurement Results and other associated information as defined by the Control.

[Measurement Task]: The act that consists of the single operation of the Measurement Method at a particular time and with all its Input Parameters set to specific values.

[Measurement Result]: The output of a single Measurement Task (the value obtained for the parameter of interest or Metric).

[Note: further discussion and clarifications regarding these borrowed terms from LMAP framework are to be expected, with coordination with [I-D.ietf-lmap-framework].]

6. Methodology

For a given set of well-defined metrics, a number of distinct measurement methodologies may exist. Let us take One-way Packet Loss as example. 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. In order to perform packet loss measurements on a live traffic flow, different methodologies exist. A partial list includes:

1. observation, e.g. Sequence Number, pros and cons

2. inserting a delimiting packet: Y.1731, RFC6374, pros and cons

3. altering the packet:

Note: This list is by no means exhaustive. The purpose is to point out the variety of measurement techniques.

Note: A methodology for a metric should have the property that it is repeatable: if the methodology is used multiple times under identical

conditions, it should result in consistent measurements. A methodology for a metric should be scalable, robust and secured.

Following sections list the functional requirements and design considerations of any passive measurement methodology.

6.1. Discussion of Errors / Unintended Consequences

As discussed in Section 6.3 Measurements, Uncertainties and Errors of RFC2330, the measurement technique itself can introduce errors.

"consider the timing error due to measurement overheads within the computer making the measurement, as opposed to delays due to the Internet component being measured. The former is a measurement error, while the latter reflects the metric of interest. Note that one technique that can help avoid this overhead is the use of a packet filter/sniffer, running on a separate computer that records network packets and timestamps them accurately."

With some types of passive measurement, changing the packet may create extra load on the network, change the characteristics of network traffic, or change the nature of the problem itself. Obviously, the benefits of the measurement must be such as to offset the potential unintended consequences.

6.2. Control Protocol

As depicted by the reference model, there are different functional components residing along an end-to-end path or within an ISP's domain that cooperate to perform a specific passive measurement task. This section describes the high level function requirements for the control protocol between these collaborating components.

Note: LMAP is developing the control protocol between MA and controller, here will be the discussion for control protocol between measurement parties, i.e. MA to MA or MA to MP.

6.3. 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. A measurement session may be started either proactively or on demand. The methodology must indicate how the measurement session is to be started.

6.4. Data Collected Correlation

When there is no coordination between MAs during a measurement session, data collected on the upstream MA and downstream MA, e.g. packet counts or timestamps, may be periodically report to the Controller. And the value of the performance metrics are calculated/ derived on the Controller. Certain synchronization mechanism is required to ensure the data collected on upstream and downstream are correlated. This may further require that the upstream and downstream MEs have a certain time synchronization capability (e.g., supporting the Network Time Protocol (NTP) [RFC5905], or the IEEE 1588 Precision Time Protocol (PTP) [IEEE.1588.2008].)

6.5. Measurement Configuration

A measurement session can be configured statically or dynamically. The methods must be discussed.

6.6. Scalability and Robustness

[TBD]

6.7. Privacy Issues

[TBD]

7. Security Considerations

This document does not bring new security issues to IPPM.

8. Acknowledgements

The authors would like to thank Al Morton, Brian Trammell and Robert Hamilton for their valuable comments.

9. References

9.1. Normative References

[I-D.li-mpls-seamless-mpls-mbb]
Li, Z., Li, L., Morillo, M., and T. Yang, "Seamless MPLS for Mobile Backhaul", draft-li-mpls-seamless-mpls-mbb-01 (work in progress), February 2014.

[IEEE.1588.2008]
"Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Standard 1588, March 2008.

[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.

[RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010.

[RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

9.2. Informational References

[I-D.ietf-ippm-metric-registry]
Bagnulo, M., Claise, B., Eardley, P., Morton, A., and A. Akhter, "Registry for Performance Metrics", draft-ietf-ippm-metric-registry-01 (work in progress), September 2014.

[I-D.ietf-lmap-framework]
Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for large-scale measurement platforms (LMAP)", draft-ietf-lmap-framework-10 (work in progress), January 2015.

[Y.1540.2011]
"Internet protocol data communication service - IP packet transfer and availability performance parameters", ITU-T Y.1540, March 2011.

Authors' Addresses

Lianshu Zheng
Huawei Technologies
China

Email: vero.zheng@huawei.com

Nalini Elkins
Inside Products, Inc.
USA

Email: nalini.elkins@insidethestack.com

Lingli Deng
China Mobile
China

Email: denglingli@chinamobile.com

Michael Ackermann
Blue Cross Blue Shield of Michigan
USA

Email: mike.ackermann@bcbsmi.com

Greg Mirsky
Ericsson
USA

Email: gregory.mirsky@ericsson.com