Network Working Group E. Stephan
Internet-Draft France Telecom
Intended status: Informational L. Liang
Expires: January 7, 2008 University of Surrey
A. Morton
AT&T Labs
July 6, 2007

IP Performance Metrics (IPPM) for spatial and multicast
draft-ietf-ippm-multimetrics-04

Status of this Memo

Copyright Notice

Abstract

The IETF IP Performance Metrics (IPPM) working group has standardized metrics for measuring end-to-end performance between 2 points. This memo defines 2 sets of metrics to extend these end-to-end ones. It defines spatial metrics for measuring the performance of segments along a path and metrics for measuring the performance of a group of

users in multiparty communications.

Table of Contents

1. Introduction

The IPPM WG defined a framework for metric definitions and end-to-end measurements:

- o A general framework for defining performance metrics, described in the Framework for IP Performance Metrics [RFC2330];
- o A One-way Active Measurement Protocol Requirements [RFC3763];
- o A One-way Active Measurement Protocol (OWAMP) [RFC4656];
- o An IP Performance Metrics Registry [RFC4148];

It specified a set of end-to-end metrics, which conform to this framework:

- o The IPPM Metrics for Measuring Connectivity [RFC2678];
- o The One-way Delay Metric for IPPM [RFC2679];
- o The One-way Packet Loss Metric for IPPM [RFC2680];
- o The Round-trip Delay Metric for IPPM [RFC2681];
- o A Framework for Defining Empirical Bulk Transfer Capacity Metrics [RFC3148];
- o One-way Loss Pattern Sample Metrics [RFC3357];

o IP Packet Delay Variation Metric for IPPM [RFC3393];

o Network performance measurement for periodic streams [RFC3432];

o Packet Reordering Metric for IPPM [RFC4737];

This memo defines spatial and one-to-group metrics based on the framework and on the end-to-end metrics defined in these documents.

Firstly it defines spatial metrics:

o A 'vector', called Type-P-Spatial-One-way-Delay-Vector, will be introduced to divide an end-to-end Type-P-One-way-Delay [RFC2679] in a spatial sequence of one-way delays.

o A 'vector', called Type-P-Spatial-One-way-Packet-Loss-Vector, will be introduced to divide an end-to-end Type-P-One-way-Packet-Loss [RFC2680] in a spatial sequence of packet loss.

o Using the Type-P-Spatial-One-way-Delay-Vector metric, a 'vector', called Type-P-Spatial-One-way-ipdv-Vector, will be introduced to divide an end-to-end Type-P-One-way-ipdv in a spatial sequence of ipdv.

o Using the Type-P-Spatial-One-way-Delay-Vector metric, a 'sample', called Type-P-Segment-One-way-Delay-Stream, will be introduced to define a set of one-way delays between a pair of host of the path;

o Using the Type-P-Spatial-Packet-Loss-Vector metric, a 'sample', called Type-P-Segment-Packet-Loss-Stream, will be introduced to define a set of packet losses between a pair of host of the path;

o Using the Type-P-Spatial-ipdv-Vector metric, a 'sample', called Type-P-Segment-ipdv-Stream, will be introduced to define a set of ipdvs between a pair of host of the path;

Then it defines one-to-group metrics.

o Using one test packet sent from one sender to a group of receivers, a 'sample', called Type-P-one-to-group-One-way-Delay-Vector, will be introduced to define the list of Type-P-one-way-delay [RFC2679] between this sender and the group of receivers.

- o Using one test packet sent from one sender to a group of receivers, a 'sample', called Type-P-one-to-group-One-way-Packet-Loss-Vector, will be introduced to define the list of Type-P-One-way-Packet-Loss [RFC2680] between this sender and the group of receivers

- o Using one test packet sent from one sender to a group of receivers, a 'sample', called Type-P-one-to-group-One-way-ipdv-Vector, will be introduced to define the list of Type-P-One-way-ipdv between this sender and the group of receivers

- o Then a discussion section presents the set of statistics that may be computed using these metrics to present the network performance in the view of a group of users. The statistics may be the basis for requirements (e.g. fairness) on multiparty communications. .

Reporting of metrics is defined in the section "Manageability Consideration".

2. Terminology

2.1. Path Digest Hosts

The list of the hosts of a path from a source to a destination.

2.2. Multiparty metric

A metric is said to be multiparty if the topology involves more than one source or destination in the measurements. All multiparty metrics define a set of hosts called "points of interest", where one host is the source and other hosts are the measurement collection points. For example, if the set of points of interest is < ha, hb, hc, ..., hn >, where ha is the source and < hb, hc, ..., hn > are the destinations, then measurements may be conducted between < ha, hb>, < ha, hc>, ..., <ha, hn >.

For the purposes of this memo (reflecting the scope of a single source), the only multiparty metrics are one-to-group metrics.

2.3. Spatial metric

A metric is said to be spatial if one of the hosts involved is neither the source nor the destination of the measured packet.

2.4. One-to-group metric

A metric is said to be one-to-group if the measured packet is sent by one source and (potentially) received by several destinations. Thus, the topology of the communication group can be viewed as a centre-distributed or server-client topology with the source as the centre/server in the topology.

2.5. Points of interest

Points of interest are the set of hosts* (as per RFC2330 definition, that is including nodes) of the set of hosts involved in the delivery of the packets (in addition to the source itself). Note that the set of the points of interest are (a possibly arbitrary) subset of all the hosts involved in the path. Points of interest of One-to-group metrics are the hosts receiving packets from the source (in addition to the source itself).


```
                         Src                  Recv
                         `.          ,-.
                           `.      ,'   `...... 1
                             `.    ;       :
                               `.  ;       :
                                 ;          :... 2
                                 |          |
                                 :          ;
```

```
                                  :       ;.... 3
                                  :       ;
                                   `.   ,'
                                     `-'....... N
```

Figure 1: One-to-group points of interest

A points of interest of spatial metrics is a host of the set of hosts involved in the delivery of the packets from the source.

```
                        Src ------.           Hosts
                                   \
                                    `---X   ... 1
                                        \
                                         x
                                        /
                             .---------X   .... 2
                           /
                          x
                           \
                           `---X           .... 3
                                  \
                                   \
                                    \
                                     X     .... N
                                      \
                                       \
                                        \
                                         `---- Dst
```

Note: 'x' are nodes which are not points of interest

Figure 2: Spatial points of interest

2.6. Reference point

The centre/server in the multimetrics measurement that is controlled by network operators can be a very good reference point where measurement data can be collected for further processing although the actual measurements have to be carried out at all points of interest. Thus, we can define the reference point as the server where the statistic calculation will be carried out.

2.7. Vector

A Vector is a set of singletons, which are a set of results of the observation of the behaviour of the same packet at different places of a network at different time. For instance, if One-way delay singletons observed at N receivers for Packet P sent by the source Src are dT1, dT2,..., dTN, it can be say that a vector V with N elements can be organized as {dT1, dT2,..., dTN}. The elements in one vector are singletons distinct with each other in terms of both measurement point and time. Given the vector V as an example, the element dT1 is distinct from the rest by measured at receiver 1 at time T1. Additional to a singleton, Vector gives information over a space dimension.

2.8. Matrix

Several vectors can form up a Matrix, which contains results observed in a sampling interval at different place of a network at different time. For instance, given One-way delay vectors V1={dT11, dT12,..., dT1N}, V2={dT21, dT22,..., dT2N},..., Vm={dTm1, dTm2,..., dTmN} for Packet P1, P2,...,Pm, we can have a One-way delay Matrix {V1, V2,...,Vm}. Additional to the information given by a Vector, a Matrix is more powerful to present network performance in both space and time dimensions. It normally corresponds to a sample.

The relation among Singleton, Vector and Matrix can be shown in the following Figure 3.

```
                 Point of          Singleton
                 interest            /          Samples
                  ,----.    ^      /
                 /   R1.....|  / R1dT1   R1dT2   R1dT3 ... R3dTk \
                /         \ | |                                   |
               ;  R2........| |  R2dT1   R2dT2   R2dT3 ... R3dTk  |
          Src  |           || |                                   |
               |      R3....| |  R3dT1   R3dT2   R3dT3 ... R3dTk  |
               |           || |                                   |
               :           ;| |                                   |
                \         / | |                                   |
                 \  Rn......|  \ RndT1   RndT2   RndT3 ... RndTk /
                  `-----'   +-------------------------------------> time

                                Vector           Matrix
                               (space)           (time)
```

Figure 3: Relation beween Singletons, vectors and matrix

3. Motivations

All IPPM metrics are defined for end-to-end measurement. These metrics provide very good guides for measurement in the pair communications. However, further efforts should be put to define metrics for multiparty measurements such as one to one trajectory metrics and one to multipoint metrics.

3.1. Motivations for spatial metrics

Decomposition of instantaneous end-to-end measures is needed:

- o Decomposing the performance of interdomain path is desirable in interdomain to qualify per AS contribution to the performance. So it is necessary to define standard spatial metrics before going further in the computation of inter domain path with QoS constraint.

- o Traffic engineering and troubleshooting applications require spatial views of the one-way delay consumption, identification of the location of the lost of packets and the decomposition of the ipdv over the path.

- o Monitoring the QoS of a multicast tree, of MPLS point-to-multipoint and inter-domain communication require spatial decomposition of the one-way delay, of the packet loss and of the ipdv.

- o Composition of metrics is a need to scale in the measurement plane. Spatial measure give typically the individual performance of an intra domain segment. It is the elementary piece of information to exchange for measuring interdomain performance based on composition of metrics.

- o The PSAMP WG defines capabilities to sample packets in a way to to support instantaneous measurement respecful of the IPPM framework [RFC2330]. Consequently it is necessary to define a set of spatial metrics for passive and active techniques.

3.2. Motivations for One-to-group metrics

While the node-to-node based spatial measures can provide very useful data in the view of each connection, we also need measures to present the performance of a multiparty communication in the view of a group with consideration that it involves a group of people rather than two. As a consequence a simple one-way metric cannot describe the multi-connection situation. We need some new metrics to collect performance of all the connections for further statistics analysis. A group of metrics are proposed in this stage named one-to-group performance metrics based on the unicast metrics defined in IPPM WG. One-to-group metrics are trying to composite one-way metrics from one source to a group of destinations to make up new metrics. The compositions are necessary for judging the network performance of multiparty communications and can also be used to describe the difference of the QoS served among a group of users.

One-to-group performance metrics are needed for several reasons:

- o For designing and engineering multicast trees and MPLS point-to-multipoint LSP;

- o For evaluating and controlling of the quality of the multicast services;

- o For controlling the performance of the inter domain multicast services;

o For presenting and evaluating the relative QoS requirements for the multiparty communications.

To understand the connection situation between one source and any one receiver in the multiparty communication group, we need the collection of instantaneous end-to-end measures. It will give us very detailed insight into each branch of the multicast tree in terms of end-to-end absolute QoS. It can provide clear and helpful information for engineers to identify the connection with problems in

a complex multiparty routing tree.

The one-to-group metrics described in this memo introduce one-to-many concerns to the IPPM working group to measure the performance of a group of users who receiving data from the same source. The concept extends the "path" in the one-way measurement to "path tree" to cover both one-to-one and one-to-many communications. Nevertheless, applied to one-to-one communications they provide exactly the same results as the corresponding one-to-one metrics.

3.3. Discussion on Group-to-one and Group-to-group metrics

We note that points of interest can also be selected to define measurements on Group-to-one and Group-to-group topologies. These topologies are currently beyond the scope of this memo, because they would involve multiple packets launched from different sources. However, we can give some clues here on these two cases.

The measurements for group-to-one topology can be easily derived from the one-to-group measurement. The measurement point is the reference point that is acting as a receiver while all of clients/receivers defined for one-to-group measurement act as sources in this case.

For the group-to-group connection topology, we can hardly define the reference point and, therefore, have difficulty to define the measurement points. However, we can always avoid this confusion by treating the connections as one-to-group or group-to-one in our measurements without consideration on how the real communication will be carried out. For example, if one group of hosts < ha, hb, hc, ..., hn > are acting as sources to send data to another group of hosts < Ha, Hb, Hc, ..., Hm >, we can always decompose them into n one-to-group communications as < ha, Ha, Hb, Hc, ..., Hm >, < hb, Ha, Hb, Hc, ..., Hm >, <hc, Ha, Hb, Hc, ..., Hm >, ..., < hn, Ha, Hb, Hc, ..., Hm >.

4. Spatial vectors metrics definitions

This section defines vectors for the decomposition of end-to-end singleton metrics over a path.

Spatial vectors metrics are based on the decomposition of standard end-to-end metrics defined by the IPPM WG in [RFC2679], [RFC2680], [RFC3393] and [RFC3432].

Definitions are coupled with the corresponding end-to-end metrics. Methodology specificities are common to all the vectors defined and are consequently discussed in a common section.

4.1. A Definition for Spatial One-way Delay Vector

This section is coupled with the definition of Type-P-One-way-Delay. When a parameter from section 3 of [RFC2679] is first used in this section, it will be tagged with a trailing asterisk.

Sections 3.5 to 3.8 of [RFC2679] give requirements and applicability statements for end-to-end one-way-delay measurements. They are applicable to each point of interest Hi involved in the measure. Spatial one-way-delay measurement SHOULD be respectful of them, especially those related to methodology, clock, uncertainties and reporting.

Following we adapt some of them and introduce points specific to spatial measurement.

4.1.1. Metric Name

Type-P-Spatial-One-way-Delay-Vector

4.1.2. Metric Parameters

o Src*, the IP address of the sender.

o Dst*, the IP address of the receiver.

o i, An integer if the list <1,2,...,n> which ordered the hosts in the path.

o Hi, A host* of the path digest.

o T*, a time, the sending (or initial observation) time for a measured packet.

o dT* a delay, the one-way delay for a measured packet.

o <dT1,..., dTn> a list of delay.

o P*, the specification of the packet type.

o <H1, H2,..., Hn>, hosts path digest.

4.1.3. Metric Units

A sequence of times.

4.1.4. Definition

Given a Type-P packet sent by the sender Src at wire-time (first bit) T to the receiver Dst in the path <H1, H2,..., Hn>. Given the sequence of values <T+dT1,T+dT2,...,T+dTn,T+dT> such that dT is the Type-P-One-way-Delay from Src to Dst and such that for each Hi of the path, T+dTi is either a real number corresponding to the wire-time the packet passes (last bit received) Hi, or undefined if the packet never passes Hi.

Type-P-Spatial-One-way-Delay-Vector metric is defined for the path <Src, H1, H2,..., Hn, Dst> as the sequence of values <T,dT1,dT2,...,dTn,dT>.

4.1.5. Discussion

Following are specific issues which may occur:

o the delay looks to decrease: dTi > DTi+1. This seem typically du to some clock synchronisation issue. This point is discussed in the section 3.7.1. "Errors or uncertainties related to Clocks" of of [RFC2679]. One consequence of these uncertainties is that times of a measure at different hosts shall not be used to order hosts on the path of a measure;

o The location of the point of interest in the device influences the result. If the packet is not observed on the input interface the delay includes buffering time and consequently an uncertainty due to the difference between 'wire time' and 'host time';

4.2. A Definition for Spatial One-way Packet Loss Vector

This section is coupled with the definition of Type-P-One-way-Packet-Loss. Then when a parameter from the section 2 of [RFC2680] is first used in this section, it will be tagged with a trailing asterisk.

Sections 2.5 to 2.8 of [RFC2680] give requirements and applicability statements for end-to-end one-way-Packet-Loss measurements. They are applicable to each point of interest Hi involved in the measure. Spatial packet loss measurement SHOULD be respectful of them, especially those related to methodology, clock, uncertainties and reporting.

Following we define the spatial metric, then we adapt some of the points above and introduce points specific to spatial measurement.

4.2.1. Metric Name

Type-P-Spatial-One-way-Packet-Loss-Vector

4.2.2. Metric Parameters

+ Src*, the IP address of the sender.

+ Dst*, the IP address of the receiver.

+ i, An integer which ordered the hosts in the path.

+ Hi, exchange points of the path digest.

+ T*, a time, the sending (or initial observation) time for a measured packet.

+ dT1,..., dTn, dT, a list of delay.

+ P*, the specification of the packet type.

+ <SH1, H2,..., Hn>, hosts path digest.

+ B1, B2, ..., Bi, ..., Bn, a list of Boolean values.

4.2.3. Metric Units

A sequence of Boolean values.

4.2.4. Definition

Given a Type-P packet sent by the sender Src at time T to the receiver Dst in the path <H1, H2, ..., Hn>. Given the sequence of times <T+dT1,T+dT2,...,T+dTn,T+dT> the packet passes <H1, H2 ..., Hn, Dst>,

Type-P-One-way-Packet-Lost-Vector metric is defined as the sequence of values <B1, B2, ..., Bn> such that for each Hi of the path, a value of Bi of 0 means that dTi is a finite value, and a value of 1 means that dTi is undefined.

4.2.5. Discussion

Following are specific issues which may occur:

o the result includes the sequence 1,0. This case means that the packet was seen by a host but not by it successor on the path;

The location of the point of interest in the device influences the result:

o Even if the packet is received by a host, it may be not observed by the point of interest located after a buffer;

4.3. A Definition for Spatial One-way Ipdv Vector

This section uses parameters from the definition of Type-P-One-way-ipdv. When a parameter from section 2 of [RFC3393] is first used in this section, it will be tagged with a trailing asterisk.

Following we adapt some of them and introduce points specific which are to spatial measurement.

4.3.1. Metric Name

Type-P-Spatial-One-way-ipdv-Vector

4.3.2. Metric Parameters

+ Src*, the IP address of the sender.

+ Dst*, the IP address of the receiver.

+ i, An integer which ordered the hosts in the path.

+ Hi, exchange points of the path digest.

+ T1*, the time the first packet was sent.

+ T2*, the time the second packet was sent.

+ P, the specification of the packet type.

+ P1, the first packet sent at time T1.

+ P2, the second packet sent at time T2.

+ <H1, H2,..., Hn>, host path digest.

+ <T1,dT1.1, dT1.2,..., dT1.n,dT1>, the Type-P-Spatial-One-way-Delay-Vector for packet sent at time T1;

+ <T2,dT2.1, dT2.2,..., dT2.n,dT2>, the Type-P-Spatial-One-way-Delay-Vector for packet sent at time T2;

+ L*, a packet length in bits. The packets of a Type P packet stream from which the Type-P-Spatial-One-way-Delay-Vector metric is taken MUST all be of the same length.

4.3.3. Metric Units

A sequence of times.

4.3.4. Definition

Given the Type-P packet having the size L and sent by the sender Src at wire-time (first bit) T1 to the receiver Dst in the path <H1, H2,..., Hn>.

Given the Type-P packet having the size L and sent by the sender Src at wire-time (first bit) T2 to the receiver Dst in the same path.

Given the Type-P-Spatial-One-way-Delay-Vector <T1,dT1.1, dT1.2,..., dT1,n,dT1> of the packet P1.

Given the Type-P-Spatial-One-way-Delay-Vector <T2,dT2.1, dT2.2,..., dT2,n,dT2> of the packet P2.

Type-P-Spatial-One-way-ipdv-Vector metric is defined as the sequence of values <T2-T1,dT2.1-dT1.1,dT2.2-dT1.2,...,dT2.n-dT1.n,dT2-dT1> Such that for each Hi of the path <H1, H2,..., Hn>, dT2.i-dT1.i is either a real number if the packets P1 and P2 passes Hi at wire-time (last bit) dT1.i, respectively dT2.i, or undefined if at least one of them never passes Hi. T2-T1 is the inter-packet emission interval and dT2-dT1 is ddT* the Type-P-One-way-ipdv at T1,T2*.

4.4. Spatial Methodology

Methodology, reporting and uncertainties points specified in section 3 of [RFC2679][RFC2679] applies to each point of interest Hi measuring a element of a spatial delay vector.

Methodology, reporting and uncertainties points specified in section 2 of [RFC2680][RFC2680] applies to each point of interest Hi measuring a element of a spatial packet loss vector.

Sections 3.5 to 3.7 of [RFC3393] give requirements and applicability statements for end-to-end One-way ipdv measurements. They are applicable to each point of interest Hi involved in the measure. Spatial One-way ipdv measurement SHOULD be respectful of methodology, clock, uncertainties and reporting aspects given in this section.

Passive and active measurement approaches of the metrology are considered as orthogonal. Active measure differs mainly from passive measure by the knowledge of the time the packet was send by the source and received by the destination, making the RFC2330 framework difficults to apply to passive measurement. On the other hand, spatial metrics rely on passive observation of the packets involved in the measure.

In fact each approach compliments the other setting the base of spatial measurement methodology: Active points of interest provide information observed at the source and at the destination while Passive points of interests provide information observed at intermediary hosts of the path.

Generally, for a given Type-P of length L, in a given Hi, the methodology for spatial vector metrics would proceed as follows:

o At each Hi, points of interest prepare to capture the packet sent a time T, take a timestamp Ti', determine the internal delay correction dTi' (See section 3.7.1. "Errors or uncertainties related to Clocks" of [RFC2679]),

o Each Hi extracts the path ordering information from the packet (e.g. time-to-live);

o Each Hi compute the wiretime from Src to Hi: Ti = Ti' - dTi'. This arrival time is undefined (infinite) if the packet is not detected after the 'loss threshold' duration;

o Each Hi extracts the timestamp T from the packet,

o Each Hi computes the one-way-delay from Src to Hi: dTi = Ti - T;

o The reference point gathers the result and time-to-live of each Hi and order them according to the path to build the Type-P-Spatial-One-way-Delay-Vector metric <T,dT1,dT2,...,dTn,dT> over the path <Src,H1, H2,..., Hn, Dst>.

4.4.1. Loss threshold

Loss threshold is the centrality of any methodology because it determines the presence the packet in the measurement process of the point of interest and consequently determines any ground truth metric result. It determines the presence of an effective delay, and bias the measure of ipdv, of packet loss and of the statistics.

This is consistent for end-to-end but impacts spatial measure: depending on the consistency of the Loss threshold among the points of interest, a packet may be considered loss a one host but present in another one, or may be observed by the last host (last hop) of the path but considered lost by Dst. The analysis of such results is not deterministic: has the path change? Does the packet arrive at destination or was it lost during the last mile? The same applies, of course, for one-way-delay measures: a delay measured may be infinite at one host but a real value in another one, or may be measured as a real value by the last host of the path but observed as infinite by Dst. The Loss threshold should be set up with the same value in each host of the path and in the destination. The Loss threshold must be systematically reported to permit careful introspection and to avoid the introduction of any contradiction in the statistic computation process.

4.4.2. Host Path Digest

The methodology given above adds the order of the points of interest over the path to [RFC2679] one's.

A perfect Host Path Digest (hum! of cource from the measurement point of view only, that is, corresponding to the real path the test packet experimented) may include several times several hosts:

- <Ha,..., Ha> coresponds to a loop in the path;
- <Ha,..,Hb,..., Ha,...,Hb> coresponds to a loop in the path which may occurs during rerouting phases;

These cases MUST be identified before statistics computation to avoid corrupted results' production. This applies especially to the measure of segments which are build from results of a measure of a vector metric.

5. Spatial Segments metrics definitions

This section defines samples to measure a sequence of delays, a sequence of lost and a sequence of ipdv between 2 hosts of the path, a segment. Singletons are taken from segments of vectors defined above.

5.1. A Definition of a sample of One-way Delay of a segment of the path

This metric defines a sample of One-way delays between a pair of hosts of a path.

5.1.1. Metric Name

Type-P-Segment-One-way-Delay-Stream

5.1.2. Metric Parameters

+ Src*, the IP address of the sender.

+ Dst*, the IP address of the receiver.

+ P*, the specification of the packet type;

+ i, An integer which orders exchange points in the path.

+ k, An integer which orders the packets sent.

+ Hi, a host of the path digest;

+ <H1, H2,..., Hn>, host path digest.

+ Ha, a host of the path digest different from Dst and Hb;

+ Hb, a host of the path digest different from Src and Ha.

Hb order in the path must greater that Ha;

+ <T1, ..., Tk>, a list of time ordered by k;

+ dT1,..., dTn a list of delay;

5.1.3. Metric Units

A sequence of delay

5.1.4. Definition

Given 2 hosts Ha and Hb of the path <Src, H1, H2,..., Hn, Dst>, given a flow of packets of Type-P sent from Src to Dst at the times T1, T2... Tn. At each of these times, we obtain a Type-P-Spatial-One-way-Delay-Vector <T1,dT1.a, ..., dT1.b,...,, dT1.n,dT1>. We define the value of the sample Type-P-segment-One-way-Delay-Stream as the sequence made up of the delays dTk.b - dTk.a. dTk.a is the delay between Src and Ha. dTk.b is the delay between Src and Hb. 'dTk.b - dTk.a' is the one-way delay experienced by the packet sent by Src at the time Tk when going from Ha to Hb.

5.1.5. Discussion

Following are specific issues which may occur:

o When a is Src <Tk,dTk.b - dTk.a> is the measure of the first hop.

o When b is Dst <Tk,dTk.b - dTk.a> is the measure of the last hop.

o the delay looks to decrease: dTi > DTi+1:

  * This is typically du to clock synchronisation issue. this point is discussed in the section 3.7.1. "Errors or uncertainties

related to Clocks" of of [RFC2679];

* This may occurs too when the clock resolution of one probe is bigger than the minimum delay of a path. As an example this happen when measuring the delay of a path which is 500 km long with one probe synchronized using NTP having a clock resolution of 8ms.

o The location of the point of interest in the device influences the result. If the packet is not observed on the input interface the delay includes buffering time and consequently an uncertainty due to the difference between 'wire time' and 'host time';

o dTk.b may be observed and not dTk.a.

o Tk is unknown if the flow is made of end user packets, that is pure passive measure. In this case Tk may be forced to Tk+dTk.a. This motivate separate metrics names for pure passive measurement or specific reporting information.

o Pure passive measure should consider packets of the same size and of the same Type-P.

5.2. A Definition of a sample of Packet Loss of a segment of the path

This metric defines a sample of Packet lost between a pair of hosts of a path.

5.2.1. Metric Name

Type-P-segment-Packet-loss-Stream

5.2.2. Metric Parameters

+ Src*, the IP address of the sender.

+ Dst*, the IP address of the receiver.

+ P*, the specification of the packet type.

+ k, An integer which orders the packets sent.

+ n, An integer which orders the hosts on the path.

+ <H1, H2,..., Hn>, hosts path digest.

+ Ha, a host of the path digest different from Dst and Hb;

+ Hb, a host of the path digest different from Src and Ha. Hb order in the path must greater that Ha;

+ Hi, exchange points of the path digest.

+ <B1, B2, ..., Bn> a list of bits.

+ <L1, L2, ..., Ln> a list of integers.

5.2.3. Metric Units

A sequence of integers <L1, L2,..., Lk>

5.2.4. Definition

Given 2 hosts Ha and Hb of the path <Src, H1, H2,..., Hn, Dst>, given a flow of packets of Type-P sent from Src to Dst at the times T1, T2... Tn. At each of these times, we obtain a Type-P-Spatial-Packet-Lost-Vector <B1.1, B1.2,..., B1.n>. We define the value of the sample Type-P-segment-Packet-Lost-Stream between Ha and Hb as the sequence made up of the integer <L1, L2,..., Lk> such that for each Tk:

o a value of Lk of 0 means that Bk.a has a value of 0 (observed) and that Bk.b have a value of 0 (observed);

o a value of Lk of 1 means that Bk.a has a value of 0 (observed) and that Bk.b have a value of 1 (not observed);

o a value of Lk of 2 means that Bk.a has a value of 1 (not observed) and that Bk.b have a value of 0 (observed);

o a value of Lk of 3 means that Bk.a has a value of 1 (not observed) and that Bk.b have a value of 1 (not observed).

5.2.5. Discussion

The semantic of a Type-P-segment-Packet-loss-Stream is similar to the one of Type-P-Packet-loss-Stream:

o a value of 0 means that the packet was observed by Ha (similar to 'send by Src') and not observed by Hb (similar to 'not received by Dst');

o a value of 1 means that it was observed by Ha (similar to 'send by Src') and observed by Hb (similar to 'received by Dst').

This definition of Type-P-segment-Packet-loss-Stream is similar to the Type-P-Packet-loss-Stream defined in [RFC2680] excepted that in a Type-P-segment-Packet-loss-Stream the rules of the point of interests Ha and Hb are symetrical: The asumption that a set of packets are going from Ha to Hb does not apply to Type-P-segment-Packet-loss-Stream because as the host path digest is dynamic each packet has its own host path digest.

Making the asumption that the host path digest of a Type-P-spatial-Packet-loss-vector does not change and that the set of (Hk, Hk+1) tuples is mostly stable over time lead to unusable results and to the introduction of mistakes in the metrics aggregation processes. The right approach consists in carefully scrutening the path ordering information to build sample with elements of vectors sharing the same properties in term of Ha and Hb and 'Ha to Hb'. So a measure of Type-P-spatial-Packet-loss-vector differs from a Type-P-Packet-loss one in that it produces different samples of packet loss over time.

The semantic of a Type-P-segment-Packet-loss-Stream defines 2 new results:

o A value of Lk of 2 (1,0) corresponds to a mistake in the ordering of Ha and Hb over the path coming either from the configuration (asumption on the path) or from the processing of the vectors: bad scrutening of the path ordering information, or some other mistake in the measure or the reporting. It is not in the scope of this document to go in further details which are mostly implementation dependent. This value MUST not be used to compute packet lost statistics.

o A value of Lk of 3 (1,1) corresponds to a lost of the packet in upper segment of the path.

5.3. A Definition of a sample of One-way Ipdv of a segment of the path

This metric defines a sample of ipdv between a pair of hosts of a path.

Editor note: work in progress

5.3.1. Metric Name

Type-P-Segment-Ipdv-Stream

5.3.2. Metric Parameters

5.3.3. Metric Units

5.3.4. Definition

5.3.5. Discussion

5.4. Discussion on Passive Segment Metrics

A pure passive spatial measure is the measure of a spatial metric based on the observation of user traffic instead of packets dedicated to the measure.

This section discusses the applicability of pure passive measurement methodology (e.g. without injecting test traffic as described by PSAMP documents [draft-ietf-psamp-sample-tech-10.txt]) to measure spatial metrics.

To permit comparison and discussion, we firstly define pure passive measurement methodology following the spirit of IPPM framework [RFC2330] and the methodology of [RFC2679]. Then we propose several passive metrics complying to this framework.

5.4.1. A methodoloy for passive segment metrics

The following starts from the point that the time a packet is sent is not needed to measure the delay from one host Ha of the path to another host Hb.

Generally, for the packets of Type-P and length L sent a time <T1,

T2,..., Tn> by the source Src pure passive methodology might proceed as follows:

o Each point of interest Ha and Hb prepares to capture these packets;

o Each point of interest Ha and Hb takes a timestamp Ti', determines the internal delay correction dTi' (See section 3.7.1. "Errors or uncertainties related to Clocks" of [RFC2679]),

o Each points of interest Ha and Hb extracts the path ordering information from the packet (e.g. time-to-live)

o Each points of interest Ha and Hb computes the wiretime fSrom Src to Hi: Ti = Ti' - dTi'. ;

o The reference point gathers individual information for the packets sent a time <T1, T2,..., Tn> from each point of interest Ha and Hb and proceeds as follow:

  * Orders them according to the path ordering information;

  * Extracts the timestamps Ti.a and Ti+1.b. This arrival time is undefined (infinite) if the packet is not detected;

  * Computes the one-way-delay from Ha to Hb as (Ti+1.b - Ti.a). The delay is undefined (infinite) if the packet is not detected in Ha or Hb;

o The reference point builds the segment sample <T1.b - T1.a, T2.b - T2.a,..., Tn.b - Tn.a> from Ha to Hb;

5.4.2. Discussion on passive methodoloy

Intrinsically passive methodoloy does not known (neither in the points of interest nor in the point of reference) the time each packet is sent <T1, T2,..., Tn> and the time each packet it received.

Section 5.4.1 shows that a passive segment one-way delay measure does not rely on the time T the packet is sent to compute the delay from a host Ha to a host Hb.

Intuitively, packets loss measurement does not require any time information and only expects the packet was sent. Passive packet loss methodology relies on the detection of the packet by one point of interest and not by another. This relies on asumptions similar to spatial methodology:

o The knowledge of the path and the order of the points of interest over the path;

o The packet is observed by one point of interest and not by another;

Nevertheless, passive packet loss measure is limited by the fact that information that neither a packet has be sent nor that the packet was received is never available:

when the path changes and the packet is not observed it is not deterministic to state that the packet is lost because the measure does not known that the packet is received by Dst.

when the measure does not observe any packets it is not possible to state that all packets are lost because the measure does not known that any packets were sent.

The drawback is that monitoring finely these events is crucial for troubleshooting workflow.

IPPM framework relies on the mesure of the behavior of packets of the same size. Consequently a passive metric sample MUST not mix information of packets of different sizes.

Segment metrics may be measured using pure passive technics. Passive segment metrics definitions are very closed to spatial segment metrics definitions. Therefore below we just name passive segment metrics to distinguish the methodology used. Having distinct metrics identifiers for spatial metrics and passive spatial metrics in the [RFC4148] will avoid interoperability issues especially during composition of metrics the IPPM WG is currently defining.

5.4.3. Passive Segment metrics

5.4.3.1. Passive Segment One way Delay metric

This metric definition is based on the top of the Type-P-spatial-segment-One-way-Delay-Stream metric definition.

name: Type-P-Passive-Segment-One-way-Delay-Stream

5.4.3.2. Passive Segment Packet Loss metric

This metric definition is based on the top of the Type-P-spatial-segment-Packet-Loss-Stream metric definition.

name: Type-P-Passive-Segment-Packet-Loss-Stream

5.4.3.3. Passive Segment One-way Ipdv metric

This metric definition is based on the top of the Type-P-Segment-Ipdv-Stream metric definition.

name: Type-P-Passive-Segment-One-way-Ipdv-Stream

6. One-to-group metrics definitions

6.1. A Definition for one-to-group One-way Delay

6.1.1. Metric Name

Type-P-one-to-group-One-way-Delay-Vector

6.1.2. Metric Parameters

- o Src, the IP address of a host acting as the source.
- o Recv1,..., RecvN, the IP addresses of the N hosts acting as receivers.
- o T, a time.
- o dT1,...,dTn a list of time.

o P, the specification of the packet type.

o Gr, the multicast group address (optional). The parameter Gr is the multicast group address if the measured packets are transmitted by multicast. This parameter is to identify the measured traffic from other unicast and multicast traffic. It is set to be optional in the metric to avoid losing any generality, i.e. to make the metric also applicable to unicast measurement where there is only one receivers.

6.1.3. Metric Units

The value of a Type-P-one-to-group-One-way-Delay-Vector is a set of singletons metrics Type-P-One-way-Delay [RFC2679].

6.1.4. Definition

Given a Type P packet sent by the source Src at Time T, given the N hosts { Recv1,...,RecvN } which receive the packet at the time { T+dT1,...,T+dTn }, a Type-P-one-to-group-One-way-Delay-Vector is defined as the set of the Type-P-One-way-Delay singleton between Src and each receiver with value of { dT1, dT2,...,dTn }.

6.2. A Definition for one-to-group One-way Packet Loss

6.2.1. Metric Name

Type-P-one-to-group-One-way-Packet-Loss-Vector

6.2.2. Metric Parameters

o Src, the IP address of a host acting as the source.

o Recv1,..., RecvN, the IP addresses of the N hosts acting as receivers.

o T, a time.

o T1,...,Tn a list of time.

o P, the specification of the packet type.

o Gr, the multicast group address (optional).

6.2.3. Metric Units

The value of a Type-P-one-to-group-One-way-Packet-Loss-Vector is a set of singletons metrics Type-P-One-way-Packet-Loss [RFC2680].

6.2.4. Definition

Given a Type P packet sent by the source Src at T and the N hosts, Recv1,...,RecvN, which should receive the packet at T1,...,Tn, a Type-P-one-to-group-One-way-Packet-Loss-Vector is defined as a set of the Type-P-One-way-Packet-Loss singleton between Src and each of the receivers {<T1,0|1>,<T2,0|1>,..., <Tn,0|1>}.

6.3. A Definition for one-to-group One-way Ipdv

6.3.1. Metric Name

Type-P-One-to-group-One-way-ipdv-Vector

6.3.2. Metric Parameters

+ Src, the IP address of a host acting as the source.

+ Recv1,..., RecvN, the IP addresses of the N hosts acting as receivers.

+ T1, a time.

+ T2, a time.

+ ddT1,...,ddTn, a list of time.

+ P, the specification of the packet type.

+ F, a selection function defining unambiguously the two packets from the stream selected for the metric.

+ Gr, the multicast group address (optional)

6.3.3. Metric Units

The value of a Type-P-One-to-group-One-way-ipdv-Vector is a set of singletons metrics Type-P-One-way-ipdv [RFC3393].

6.3.4. Definition

Given a Type P packet stream, Type-P-one-to-group-One-way-ipdv-Vector is defined for two packets from the source Src to the N hosts {Recv1,...,RecvN },which are selected by the selection function F, as the difference between the value of the Type-P-one-to-group-One-way-Delay-Vector from Src to { Recv1,..., RecvN } at time T1 and the value of the Type-P-one-to-group- One-way-Delay-Vector from Src to { Recv1,...,RecvN } at time T2. T1 is the wire-time at which Src sent the first bit of the first packet, and T2 is the wire-time at which Src sent the first bit of the second packet. This metric is derived from the Type-P-one-to- group-One-way-Delay-Vector metric.

Therefore, for a set of real number {ddT1,...,ddTn},Type-P-one- to-group-One-way-ipdv-Vector from Src to { Recv1,...,RecvN } at T1, T2 is {ddT1,...,ddTn} means that Src sent two packets, the first at wire-time T1 (first bit), and the second at wire-time T2 (first bit) and the packets were received by { Recv1,...,RecvN } at wire-time {dT1+T1,...,dTn+T1}(last bit of the first packet), and at wire-time {dT'1+T2,...,dT'n+T2} (last bit of the second packet), and that {dT'1-dT1,...,dT'n-dTn} ={ddT1,...,ddTn}.

7. One-to-Group Sample Statistics

The defined one-to-group metrics above can all be directly achieved from the relevant unicast one-way metrics. They managed to collect all unicast measurement results of one-way metrics together in one profile and sort them by receivers and packets in a multicast group. They can provide sufficient information regarding the network performance in terms of each receiver and guide engineers to identify potential problem happened on each branch of a multicast routing tree. However, these metrics can not be directly used to conveniently present the performance in terms of a group and neither to identify the relative performance situation.

From the performance point of view, the multiparty communication services not only require the absolute performance support but also the relative performance. The relative performance means the difference between absolute performance of all users. Directly using the one-way metrics cannot present the relative performance situation. However, if we use the variations of all users one-way parameters, we can have new metrics to measure the difference of the absolute performance and hence provide the threshold value of relative performance that a multiparty service might demand. A very good example of the high relative performance requirement is the online gaming. A very light difference in delay might result in failure in the game. We have to use multicast specific statistic metrics to define exactly how small the relative delay the online gaming requires. There are many other services, e.g. online biding, online stock market, etc., that require multicast metrics in order to evaluate the network against their requirements. Therefore, we can see the importance of new, multicast specific, statistic metrics to feed this need.

We might also use some one-to-group statistic conceptions to present and report the group performance and relative performance to save the report transmission bandwidth. Statistics have been defined for One-way metrics in corresponding RFCs. They provide the foundation of definition for performance statistics. For instance, there are definitions for minimum and maximum One-way delay in [RFC2679]. However, there is a dramatic difference between the statistics for one-to-one communications and for one-to-many communications. The former one only has statistics over the time dimension while the later one can have statistics over both time and space dimensions. This space dimension is introduced by the Matrix concept as illustrated in Figure 9. For a Matrix M each row is a set of One-way singletons spreading over the time dimension and each column is another set of One-way singletons spreading over the space dimension.

```
          Receivers
           Space
             ^
           1 |    / R1dT1   R1dT2     R1dT3 ... R3dTk \
             |   |                                     |
           2 |   |  R2dT1   R2dT2     R2dT3 ... R3dTk  |
             |   |                                     |
           3 |   |  R3dT1   R3dT2     R3dT3 ... R3dTk  |
           . |   |                                     |
           . |   |                                     |
           . |   |                                     |
           n |    \ RndT1   RndT2     RndT3 ... RndTk /
             +---------------------------------------------> time
            T0
```

Figure 9: Matrix M (n*m)

In Matrix M, each element is a One-way delay singleton. Each column is a delay vector contains the One-way delays of the same packet observed at M points of interest. It implies the geographical factor of the performance within a group. Each row is a set of One-way delays observed during a sampling interval at one of the points of interest. It presents the delay performance at a receiver over the time dimension.

Therefore, one can either calculate statistics by rows over the space dimension or by columns over the time dimension. It's up to the operators or service provides which dimension they are interested in. For example, a TV broadcast service provider might want to know the statistical performance of each user in a long term run to make sure their services are acceptable and stable. While for an online gaming service provider, he might be more interested to know if all users are served fairly by calculating the statistics over the space dimension. This memo does not intend to recommend which of the statistics are better than the other.

To save the report transmission bandwidth, each point of interest can send statistics in a pre-defined time interval to the reference point rather than sending every One-way singleton it observed. As long as an appropriate time interval is decided, appropriate statistics can represent the performance in a certain accurate scale. How to decide the time interval and how to bootstrap all points of interest and the reference point depend on applications. For instance, applications with lower transmission rate can have the time interval longer and ones with higher transmission rate can have the time interval shorter. However, this is out of the scope of this memo.

Moreover, after knowing the statistics over the time dimension, one

might want to know how this statistics distributed over the space dimension. For instance, a TV broadcast service provider had the performance Matrix M and calculated the One-way delay mean over the time dimension to obtain a delay Vector as {V1,V2,..., VN}. He then calculated the mean of all the elements in the Vector to see what level of delay he has served to all N users. This new delay mean gives information on how good the service has been delivered to a group of users during a sampling interval in terms of delay. It needs twice calculation to have this statistic over both time and space dimensions. We name this kind of statistics 2-level statistics to distinct with those 1-level statistics calculated over either space or time dimension. It can be easily prove that no matter over which dimension a 2-level statistic is calculated first, the results are the same. I.e. one can calculate the 2-level delay mean using the Matrix M by having the 1-level delay mean over the time dimension first and then calculate the mean of the obtained vector to find out the 2-level delay mean. Or, he can do the 1-level statistic calculation over the space dimension first and then have the 2-level delay mean. Both two results will be exactly the same. Therefore, when define a 2-level statistic, there is no need to specify in which procedure the calculation should follow.

Comment: The above statement depends on whether the order of operations has any affect on the outcome.

Many statistics can be defined for the proposed one-to-group metrics over either the space dimension or the time dimension or both. This memo treats the case where a stream of packets from the Source results in a sample at each of the Receivers in the Group, and these samples are each summarized with the usual statistics employed in one-to-one communication. New statistic definitions are presented, which summarize the one-to-one statistics over all the Receivers in the Group.

7.1. Discussion on the Impact of packet loss on statistics

The packet loss does have effects on one-way metrics and their statistics. For example, the lost packet can result an infinite one-way delay. It is easy to handle the problem by simply ignoring the infinite value in the metrics and in the calculation of the corresponding statistics. However, the packet loss has so strong impact on the statistics calculation for the one-to-group metrics that it can not be solved by the same method used for one-way metrics. This is due to the complex of building a Matrix, which is needed for calculation of the statistics proposed in this memo.

The situation is that measurement results obtained by different end users might have different packet loss pattern. For example, for

User1, packet A was observed lost. And for User2, packet A was successfully received but packet B was lost. If the method to overcome the packet loss for one-way metrics is applied, the two singleton sets reported by User1 and User2 will be different in terms of the transmitted packets. Moreover, if User1 and User2 have different number of lost packets, the size of the results will be different. Therefore, for the centralized calculation, the reference point will not be able to use these two results to build up the group Matrix and can not calculate the statistics. In an extreme situation, no single packet arrives all users in the measurement and the Matrix will be empty. One of the possible solutions is to replace the infinite/undefined delay value by the average of the two adjacent values. For example, if the result reported by user1 is { R1dT1 R1dT2 R1dT3 ... R1dTK-1 UNDEF R1dTK+1... R1DM } where "UNDEF" is an undefined value, the reference point can replace it by R1dTK = {(R1dTK-1)+(R1dTK+1)}/2. Therefore, this result can be used to build up the group Matrix with an estimated value R1dTK. There are other possible solutions such as using the overall mean of the whole result to replace the infinite/undefined value, and so on. It is out of the scope of this memo.

For the distributed calculation, the reported statistics might have different "weight" to present the group performance, which is especially true for delay and ipdv relevant metrics. For example, User1 calculates the Type-P-Finite-One-way-Delay-Mean R1DM as shown in Figure. 8 without any packet loss and User2 calculates the R2DM with N-2 packet loss. The R1DM and R2DM should not be treated with equal weight because R2DM was calculated only based on 2 delay values in the whole sample interval. One possible solution is to use a weight factor to mark every statistic value sent by users and use this factor for further statistic calculation.

7.2. General Metric Parameters

- Src, the IP address of a host
- G, the Group IP address
- N, the number of Receivers (Recv1, Recv2, ... RecvN)
- T, a time (start of test interval)
- Tf, a time (end of test interval)
- K, the number of packets sent from the source during the test interval

- J[n], the number of packets received at a particular Receiver, n, where 1<=n<=N
- lambda, a rate in reciprocal seconds (for Poisson Streams)
- incT, the nominal duration of inter-packet interval, first bit to first bit (for Periodic Streams)
- T0, a time that MUST be selected at random from the interval [T, T+I] to start generating packets and taking measurements (for Periodic Streams)
- TstampSrc, the wire time of the packet as measured at MP(Src) (the Source Measurement Point)
- TstampRecv, the wire time of the packet as measured at MP(Recv), assigned to packets that arrive within a "reasonable" time
- Tmax, a maximum waiting time for packets at the destination, set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost), thus the distribution of delay is not truncated

o dT, shorthand notation for a one-way delay singleton value

o L, shorthand notation for a one-way loss singleton value, either zero or one, where L=1 indicates loss and L=0 indicates arrival at the destination within TstampSrc + Tmax, may be indexed over n Receivers

o DV, shorthand notation for a one-way delay variation singleton value

7.3. One-to-Group one-way Delay Statistics

This section defines the overall one-way delay statistics for an entire Group or receivers. For example, we can define the group mean delay, as illustrated below. This is a metric designed to summarize the whole matrix.


```
     Recv    /----------- Sample -------------\   Stats     Group Stat

      1      R1dT1   R1dT2     R1dT3 ... R1dTk     R1DM  \
                                                         |
      2      R2dT1   R2dT2     R2dT3 ... R2dTk     R2DM  |
                                                         |
      3      R3dT1   R3dT2     R3dT3 ... R3dTk     R2DM   >  GMD
      .                                                  |
      .                                                  |
      .                                                  |
      n      RndT1   RndT2     RndT3 ... RndTk     RnDM  /
```

Figure 10: One-to-GroupGroup Mean Delay

where:

R1dT1 is the Type-P-Finite-One-way-Delay singleton evaluated at Receiver 1 for packet 1.

R1DM is the Type-P-Finite-One-way-Delay-Mean evaluated at Receiver 1 for the sample of packets (1,...K).

GMD is the mean of the sample means over all Receivers (1, ...N).

7.3.1. Definition and Metric Units

Using the parameters above, we obtain the value of Type-P-One-way-Delay singleton for all packets sent during the test interval at each Receiver (Destination), as per [RFC2679]. For each packet that arrives within Tmax of its sending time, TstampSrc, the one-way delay singleton (dT) will be a finite value in units of seconds. Otherwise, the value of the singleton is Undefined.

For each packet [i] that has a finite One-way Delay at Receiver n (in other words, excluding packets which have undefined one-way delay):

Type-P-Finite-One-way-Delay-Receiver-n-[i] =

= TstampRecv[i] - TstampSrc[i]

The units of Finite one-way delay are seconds, with sufficient resolution to convey 3 significant digits.

7.3.2. Sample Mean Statistic

This section defines the Sample Mean at each of N Receivers.


```
Type-P-Finite-One-way-Delay-Mean-Receiver-n = RnDM =
              J[n]
               ---
        1      \
       --- *    >   Type-P-Finite-One-way-Delay-Receiver-n-[i]
       J[n]    /
               ---
              i = 1
```

Figure 11: Type-P-Finite-One-way-Delay-Mean-Receiver-n

where all packets i= 1 through J[n] have finite singleton delays.

7.3.3. One-to-Group Mean Delay Statistic

This section defines the Mean One-way Delay calculated over the entire Group (or Matrix).

```
Type-P-One-to-Group-Mean-Delay = GMD =
                                  N
                                 ---
                            1    \
                            - *   >    RnDM
                            N    /
                                 ---
                                n = 1
```

Figure 12: Type-P-One-to-Group-Mean-Delay

Note that the Group Mean Delay can also be calculated by summing the Finite one-way Delay singletons in the Matrix, and dividing by the number of Finite One-way Delay singletons.

7.3.4. One-to-Group Range of Mean Delays

This section defines a metric for the range of mean delays over all N receivers in the Group, (R1DM, R2DM,...RnDM).

Type-P-One-to-Group-Range-Mean-Delay = GRMD = max(RnDM) - min(RnDM)

7.3.5. One-to-Group Maximum of Mean Delays

This section defines a metrics for the maximum of mean delays over all N receivers in the Group, (R1DM, R2DM,...RnDM).

Type-P-One-to-Group-Max-Mean-Delay = GMMD = max(RnDM)

7.4. One-to-Group one-way Loss Statistics

This section defines the overall 1-way loss statistics for an entire Group. For example, we can define the group loss ratio, as illustrated below. This is a metric designed to summarize the entire Matrix.

```
     Recv    /----------- Sample ----------\   Stats     Group Stat

       1       R1L1    R1L2     R1L3 ... R1Lk     R1LR \
                                                        |
       2       R2L1    R2L2     R2L3 ... R2Lk     R2LR  |
                                                        |
       3       R3L1    R3L2     R3L3 ... R3Lk     R3LR    >  GLR
       .                                                |
       .                                                |
       .                                                |
       n       RnL1    RnL2     RnL3 ... RnLk     RnLR /
```

Figure 13: One-to-Group Loss Ratio

where:

R1L1 is the Type-P-One-way-Loss singleton (L) evaluated at Receiver 1 for packet 1.

R1LR is the Type-P-One-way-Loss-Ratio evaluated at Receiver 1 for the sample of packets (1,...K).

GLR is the loss ratio over all Receivers (1, ..., N).

7.4.1. One-to-Group Loss Ratio

The overall Group loss ratio id defined as

Type-P-One-to-Group-Loss-Ratio =

```
                                K,N
                                ---
                          1     \
                      = --- *    >   L(k,n)
                         K*N    /
                                ---
                               k,n = 1
```

Figure 14

ALL Loss ratios are expressed in units of packets lost to total packets sent.

7.4.2. One-to-Group Loss Ratio Range

Given a Matrix of loss singletons as illustrated above, determine the Type-P-One-way-Packet-Loss-Average for the sample at each receiver, according to the definitions and method of [RFC2680]. The Type-P-One-way-Packet-Loss-Average, RnLR for receiver n, and the Type-P-One-way-Loss-Ratio illustrated above are equivalent metrics. In terms of the parameters used here, these metrics definitions can be expressed as

```
Type-P-One-way-Loss-Ratio-Receiver-n = RnLR =
                                K
                               ---
                          1    \
                          - *   >   RnLk
                          K    /
                               ---
                              k = 1
```

Figure 15: Type-P-One-way-Loss-Ratio-Receiver-n

The One-to-Group Loss Ratio Range is defined as

Type-P-One-to-Group-Loss-Ratio-Range = max(RnLR) - min(RnLR)

It is most effective to indicate the range by giving both the max and minimum loss ratios for the Group, rather than only reporting the difference between them.

7.4.3. Comparative Loss Ratio

Usually, the number of packets sent is used in the denominator of packet loss ratio metrics. For the comparative metrics defined here, the denominator is the maximum number of packets received at any receiver for the sample and test interval of interest.

The Comparative Loss Ratio is defined as

```
   Type-P-Comp-Loss-Ratio-Receiver-n = RnCLR =
                                 K
                                ---
                                \
                                 >   Ln(k)
                                /
                                ---
                                k=1
                       = -----------------------------
                                 /    K         \
                                 |   ---        |
                                 |   \          |
                         K - Min |    >   Ln(k) |
                                 |   /          |
                                 |   ---        |
                                 \   k=1        / N
```

Figure 16: Type-P-Comp-Loss-Ratio-Receiver-n

7.5. One-to-Group one-way Delay Variation Statistics

There are two delay variation (DV) statistics that summarize the performance over the Group: the maximum DV over all receivers and the minimum DV over all receivers (where DV is a point-to-point metric). For each receiver, the DV is usually expressed as the 1-10^(-3) quantile of one-way delay minus the minimum one-way delay.

8. Measurement Methods: Scaleability and Reporting

Virtually all the guidance on measurement processes supplied by the earlier IPPM RFCs (such as [RFC2679] and [RFC2680]) for one-to-one scenarios is applicable here in the spatial and multiparty measurement scenario. The main difference is that the spatial and multiparty configurations require multiple measurement points where a stream of singletons will be collected. The amount of information requiring storage grows with both the number of metrics and the number of measurement points, so the scale of the measurement architecture multiplies the number of singleton results that must be collected and processed.

It is possible that the architecture for results collection involves a single aggregation point with connectivity to all the measurement points. In this case, the number of measurement points determines both storage capacity and packet transfer capacity of the host acting

as the aggregation point. However, both the storage and transfer capacity can be reduced if the measurement points are capable of computing the summary statistics that describe each measurement interval. This is consistent with many operational monitoring architectures today, where even the individual singletons may not be stored at each measurement point.

In recognition of the likely need to minimize form of the results for storage and communication, the Group metrics above have been constructed to allow some computations on a per-Receiver basis. This means that each Receiver's statistics would normally have an equal weight with all other Receivers in the Group (regardless of the number of packets received).

8.1. Computation methods

The scalability issue can be raised when there are thousands of points of interest in a group who are trying to send back the measurement results to the reference point for further processing and analysis. The points of interest can send either the whole measured sample or only the calculated statistics. The former one is a centralized statistic calculation method and the latter one is a distributed statistic calculation method. The sample should include all metrics parameters, the values and the corresponding sequence numbers. The transmission of the whole sample can cost much more bandwidth than the transmission of the statistics that should include all statistic parameters specified by policies and the additional information about the whole sample, such as the size of the sample, the group address, the address of the point of interest, the ID of the sample session, and so on. Apparently, the centralized calculation method can require much more bandwidth than the distributed calculation method when the sample size is big. This is especially true when the measurement has huge number of the points of interest. It can lead to a scalability issue at the reference point by over load the network resources. The distributed calculation method can save much more bandwidth and release the pressure of the scalability issue at the reference point side. However, it can result in the lack of information because not all measured singletons are obtained for building up the group matrix. The performance over time can be hidden from the analysis. For example, the loss pattern can be missed by simply accepting the loss ratio as well as the delay pattern. This tradeoff between the bandwidth consuming and the information acquiring has to be taken into account when design the measurement campaign to optimize the measurement results delivery. The possible solution could be to transit the statistic parameters to the reference point first to obtain the general information of the group performance. If the detail results are required, the reference point should send the requests to the points of interest, which could

be particular ones or the whole group. This procedure can happen in the off peak time and can be well scheduled to avoid delivery of too many points of interest at the same time. Compression techniques can also be used to minimize the bandwidth required by the transmission. This could be a measurement protocol to report the measurement results. It is out of the scope of this memo.

8.2. Measurement

To prevent any bias in the result, the configuration of a one-to-many measure must take in consideration that implicitly more packets will to be routed than send and selects a test packets rate that will not impact the network performance.

8.3. Effect of Time and Space Aggregation Order on Stats

This section presents the impact of the aggregation order on the scalability of the reporting and of the computation. It makes the hypothesis that receivers are managed remotely and not co-located.

multimetrics samples represented a matrix as illustrated below

```
   Point of
   interest
     1      R1S1   R1S1     R1S1 ... R1Sk    \
```

```
                                                 |
      2      R2S1   R2S2     R2S3 ... R2Sk       |
                                                 |
      3      R3S1   R3S2     R3S3 ... R3Sk        >  sample over space
      .                                          |
      .                                          |
      .                                          |
      n      RnS1   RnS2     RnS3 ... RnSk     /

             S1M    S2M      S3M  ... SnM      Stats over space

             \-------------  ------------/
                           \/
               Stat over space and time
```

Figure 17: Impact of space aggregation on multimetrics Stat

2 methods are available to compute statistics on the resulting matrix:

o metric is computed over time and then over space;

o metric is computed over space and then over time.

They differ only by the order of the time and of the space aggregation. View as a matrix this order is neutral as does not impact the result, but the impact on a measurement deployment is critical.

In both cases the volume of data to report is proportional to the number of probes. But there is a major difference between these 2 methods:

method2: In space and time aggregation mode the volume of data to collect is proportional to the number of test packets received; Each received packet RiSi triggers out a block of data that must be reported to a common place for computing the stat over space;

method1: In time and space aggregation mode the volume of data to collect is proportional to the period of aggregation, so it does not depend on the number of packet received;

Method 2 property has severe drawbacks in terms of security and dimensioning:

The increasing of the rate of the test packets may result in a sort of DoS toward the computation points;

The dimensioning of a measurement system is quite impossible to validate.

The time aggregation interval provides the reporting side with a control of various collecting aspects such as bandwidth and computation and storage capacities. So this draft defines metrics based on method 1.

Note: In some specific cases one may need sample of singletons over space. To address this need it is suggested firstly to limit the number of test and the number of test packets per seconds. Then reducing the size of the sample over time to one packet give sample of singleton over space..

8.3.1. Impact on group stats

2 methods are available to compute group statistics:

- o method1: Figure 10 andFigure 13 illustrate the method chosen: the one-to-one statistic is computed per interval of time before the computation of the mean over the group of receivers;

- o method2: Figure 17 presents the second one, metric is computed over space and then over time.

8.3.2. Impact on spatial stats

2 methods are available to compute group statistics:

- o method 1: spatial segment metrics and statistics are preferably computed over time by each points of interest;

- o method 2: Vectors metrics are intrinsecally instantaneous space metrics which must be reported using method2 whenever instantaneous metrics information is needed. Pure passive measurement approach has no choice but to use this method because delay and losses may not be computed in each point of interest.

9. Manageability Considerations

Usually IPPM WG documents defines each metric reporting within its definition. This document defines the reporting of all the metrics introduced in a single section to provide consistent information while avoiding repetitions. the aim is to contribute to the work of the WG on the reporting and to satisfy IESG recommendation of gathering manageability considerations in a dedicated section.

Data models of spatial and one-to-group metrics are similar excepted that points of interests of spatial vectors must be ordered.

The complexity of the reporting relies on the number of points of interests.

9.1. Reporting spatial metric

The reporting of spatial metrics shares a lot of aspects with RFC2679-80. New ones are common to all the definitions and are mostly related to the reporting of the path and of methodology parameters that may bias raw results analysis. This section presents these specific parameters and then lists exhaustively the parameters that shall be reported.

9.1.1. Path

End-to-end metrics can't determine the path of the measure despite IPPM RFCs recommend it to be reported (Section 3.8.4 of [RFC2679]). Spatial metrics vectors provide this path. The report of a spatial vector must include the points of interests involved: the sub set of the hosts of the path participating to the instantaneous measure.

9.1.2. Host order

A spatial vector must order the points of interest according to their order in the path. It is highly suggested to use the TTL in IPv4, the Hop Limit in IPv6 or the corresponding information in MPLS.

The report of a spatial vector must include the ordered list of the hosts involved in the instantaneous measure.

9.1.3. Timestamping bias

The location of the point of interest inside a node influences the timestamping skew and accuracy. As an example, consider that some internal machinery delays the timestamping up to 3 milliseconds then the minimal uncertainty reported be 3 ms if the internal delay is unknown at the time of the timestamping.

The report of a spatial vector must include the uncertainty of the timestamping compared to wire time.

9.1.4. Reporting spatial One-way Delay

The reporting includes information to report for one-way-delay as perthe Section 3.6 of [RFC2679].the same apply for packet loss and ipdv

9.2. Reporting One-to-group metric

9.3. Metric identification

IANA assigns each metric defined by the IPPM WG with a unique identifier as per [RFC4148] in the IANA-IPPM-METRICS-REGISTRY-MIB.

To avoid misunderstanding and to address specific reporting constraints, section Section 5.4.3 of this memo gives distinct names to passive metrics and Section 13 requests a distinct metric identifier for each metrics the memo defines.

As it is crucial for composition of metrics to know the methodology used (e.g. generation method, detection method...), the report of a metric result used in composition of metrics MUST alway include its metric identifier.

9.4. Reporting data model

This section presents the elements of the datamodel and the usage of the information reported for real network performance analysis. It is out of the scope of this section to define how the information is reported.

The data model is build with pieces of information introduced and explained in one-way delay definitions [RFC2679], in packet loss definitions [RFC2680] and in IPDV definitions[RFC3393][RFC3432]. It includes not only information given by "Reporting the metric" sections but by sections "Methodology" and "Errors and Uncertainties" sections.

Following are the elements of the datamodel taken from end-to-end definitions referred in this memo and from spatial and multicast metrics it defines:

o Packet_type, The Type-P of test packets (Type-P);

o Packet_length, a packet length in bits (L);

o Src_host, the IP address of the sender;

o Dst_host, the IP address of the receiver;

o Hosts_serie: <H1, H2,..., Hn>, a list of points of interest;

o Loss_threshold: The threshold of infinite delay;

o Systematic_error: constant delay between wire time and timestamping;

o Calibration_error: maximal uncertainty;

o Src_time, the sending time for a measured packet;

o Dst_time, the receiving time for a measured packet;

o Result_status : an indicator of usability of a result 'Resource exhaustion' 'infinite', 'lost';

o Delays_serie: <dT1,..., dTn> a list of delays;

o Losses_serie: <B1, B2, ..., Bi, ..., Bn>, a list of Boolean values (spatial) or a set of Boolean values (one-to-group);

o Result_status_serie: a list of results status;

o dT: a delay;

o Singleton_number: a number of singletons;

- o Observation_duration: An observation duration;

- o metric_identifier.

Following is the information of each vector that should be available to compute samples:

- o Packet_type;

- o Packet_length;

- o Src_host, the sender of the packet;

- o Dst_host, the receiver of the packet, apply only for spatial vectors;

- o Hosts_serie: not ordered for one-to-group;

- o Src_time, the sending time for the measured packet;

- o dT, the end-to-end one-way delay for the measured packet, apply only for spatial vectors;

- o Delays_serie: apply only for delays and ipdv vector, not ordered for one-to-group;

- o Losses_serie: apply only for packets loss vector, not ordered for one-to-group;

- o Result_status_serie;

- o Observation_duration: the difference between the time of the last singleton and the time of the first singleton.

- o Following is the context information (measure, points of interests) that should be available to compute samples :

  - * Loss threshold;

  - * Systematic error: constant delay between wire time and timestamping;

* Calibration error: maximal uncertainty;

A spatial or a one-to-group sample is a collection of singletons giving the performance from the sender to a single point of interest. Following is the information that should be available for each sample to compute statistics:

- Packet_type;
- Packet_length;
- Src_host, the sender of the packet;
- Dst_host, the receiver of the packet;
- Start_time, the sending time of the first packet;
- Delays_serie: apply only for delays and ipdv samples;
- Losses_serie: apply only for packets loss samples;
- Result_status_serie;
- Observation_duration: the difference between the time of the last singleton of the last sample and the time of the first singleton of the first sample.
- Following is the context information (measure, points of interests) that should be available to compute statistics :
  * Loss threshold;
  * Systematic error: constant delay between wire time and timestamping;
  * Calibration error: maximal uncertainty;

Following is the information of each statistic that should be reported:

o Result;

o Start_time;

o Duration;

o Result_status;

o Singleton_number, the number of singletons the statistic is computed on;

10. Open issues

Do we define min, max, avg of for each segment metrics ?

having the maximum loss metric value could be interesting. Say, the segment between router A and B always contributes loss metric value of "1" means it could be the potential problem segment.

Uploading dTi of each Hi consume a lot of bandwidth. Computing statistics (min, max and avg) of dTi locally in each Hi reduce the bandwidth consumption.

11. Security Considerations

Spatial and one-to-group metrics are defined on the top of end-to-end metrics. Security considerations discussed in One-way delay metrics definitions of [RFC2679] , in packet loss metrics definitions of[RFC2680] and in IPDV metrics definitions of[RFC3393] and [RFC3432] apply to multimetrics.

11.1. Spatial metrics

Malicious generation of packets with spoofing addresses may corrupt the results without any possibility to detect the spoofing.

Malicious generation of packets which match systematically the hash function used to detect the packets may lead to a DoS attack toward the point of reference.

11.2. one-to-group metric

The reporting of measurement results from a huge number of probes may overload the network the reference point is attach to, the reference point network interfaces and the reference point computation capacities.

The configuration of a measure must take in consideration that implicitly more packets will to be routed than send and selects a test packets rate accordingly. Collecting statistics from a huge number of probes may overload any combination of the network where the measurement controller is attach to, measurement controller network interfaces and measurement controller computation capacities.

one-to-group metrics measurement should consider using source authentication protocols, standardized in the MSEC group, to avoid fraud packet in the sampling interval. The test packet rate could be negotiated before any measurement session to avoid deny of service

attacks.

12. Acknowledgments

Lei would like to acknowledge Prof. Zhili Sun from CCSR, University of Surrey, for his instruction and helpful comments on this work.

13. IANA Considerations

Metrics defined in this memo Metrics defined in this memo are designed to be registered in the IANA IPPM METRICS REGISTRY as described in initial version of the registry [RFC4148] :

IANA is asked to register the following metrics in the IANA-IPPM-METRICS-REGISTRY-MIB :

Spatial-One-way-Delay-Vector OBJECT-IDENTITY

```
      STATUS current

      DESCRIPTION

         "Type-P-Spatial-One-way-Delay-Vector"

      REFERENCE

         "Reference "RFCyyyy, section 4.1."

         -- RFC Ed.: replace yyyy with actual RFC number & remove this
         note

      := { ianaIppmMetrics nn } -- IANA assigns nn


   Spatial-Packet-Loss-Vector OBJECT-IDENTITY

      STATUS current

      DESCRIPTION

         "Type-P-Spatial-Packet-Loss-Vector"

      REFERENCE
```



```
         "Reference "RFCyyyy, section 4.2."

         -- RFC Ed.: replace yyyy with actual RFC number & remove this
         note

      := { ianaIppmMetrics nn } -- IANA assigns nn

   Spatial-One-way-ipdv-Vector OBJECT-IDENTITY

      STATUS current

      DESCRIPTION
```

```
      "Type-P-Spatial-One-way-ipdv-Vector"

    REFERENCE

      "Reference "RFCyyyy, section 4.3."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

    := { ianaIppmMetrics nn } -- IANA assigns nn

  Spatial-Segment-One-way-Delay-Stream OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

      "Type-P-Spatial-Segment-One-way-Delay-Stream"

    REFERENCE

      "Reference "RFCyyyy, section 5.1."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

    := { ianaIppmMetrics nn } -- IANA assigns nn

  Spatial-Segment-Packet-Loss-Stream OBJECT-IDENTITY

    STATUS current

    DESCRIPTION
```



```
      "Type-P-Spatial-Segment-Packet-Loss-Stream"

    REFERENCE

      "Reference "RFCyyyy, section 5.2."
```

-- RFC Ed.: replace yyyy with actual RFC number & remove this note

:= { ianaIppmMetrics nn } -- IANA assigns nn

Spatial-Segment-One-way-ipdv-Stream OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-Spatial-Segment-ipdv-Stream"

REFERENCE

"Reference "RFCyyyy, section 5.3."

-- RFC Ed.: replace yyyy with actual RFC number & remove this note

:= { ianaIppmMetrics nn } -- IANA assigns nn

Passive-Segment-One-way-Delay-Stream OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-Passive-Segment-One-way-Delay-Stream"

REFERENCE

"Reference "RFCyyyy, section 5.4.1."

-- RFC Ed.: replace yyyy with actual RFC number & remove this note

:= { ianaIppmMetrics nn } -- IANA assigns nn

Passive-Segment-Packet-Loss-Stream OBJECT-IDENTITY

```
   STATUS current

   DESCRIPTION

      "Type-P-Passive-Segment-Packet-Loss-Stream"

   REFERENCE

      "Reference "RFCyyyy, section 5.4.2."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

   := { ianaIppmMetrics nn } -- IANA assigns nn

Passive-Segment-One-way-ipdv-Stream OBJECT-IDENTITY

   STATUS current

   DESCRIPTION

      "Type-P-Passive-Segment-One-way-ipdv-Stream"

   REFERENCE

      "Reference "RFCyyyy, section 5.4.3."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

   := { ianaIppmMetrics nn } -- IANA assigns nn

-- One-to-group metrics

one-to-group-One-way-Delay-Vector OBJECT-IDENTITY

   STATUS current

   DESCRIPTION

      "Type-P-one-to-group-One-way-Delay-Vector"

   REFERENCE

      "Reference "RFCyyyy, section 5.1."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note
```

```
      := { ianaIppmMetrics nn } -- IANA assigns nn

   one-to-group-One-way-Packet-Loss-Vector OBJECT-IDENTITY

      STATUS current

      DESCRIPTION

         "Type-P-one-to-group-One-way-Packet-Loss-Vector"

      REFERENCE

         "Reference "RFCyyyy, section 5.2."

         -- RFC Ed.: replace yyyy with actual RFC number & remove this
         note

      := { ianaIppmMetrics nn } -- IANA assigns nn

   one-to-group-One-way-ipdv-Vector OBJECT-IDENTITY

      STATUS current

      DESCRIPTION

         "Type-P-one-to-group-One-way-ipdv-Vector"

      REFERENCE

         "Reference "RFCyyyy, section 5.3."

         -- RFC Ed.: replace yyyy with actual RFC number & remove this
         note

      := { ianaIppmMetrics nn } -- IANA assigns nn

   One-to-Group-Mean-Delay OBJECT-IDENTITY

      STATUS current

      DESCRIPTION
```

"Type-P-One-to-Group-Mean-Delay"

REFERENCE

"Reference "RFCyyyy, section 6.3.3."


```
       -- RFC Ed.: replace yyyy with actual RFC number & remove this
       note

    := { ianaIppmMetrics nn } -- IANA assigns nn

 One-to-Group-Range-Mean-Delay OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

       "Type-P-One-to-Group-Range-Mean-Delay"

    REFERENCE

       "Reference "RFCyyyy, section 6.3.4."

       -- RFC Ed.: replace yyyy with actual RFC number & remove this
       note

    := { ianaIppmMetrics nn } -- IANA assigns nn

 One-to-Group-Max-Mean-Delay OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

       "Type-P-One-to-Group-Max-Mean-Delay"

    REFERENCE

       "Reference "RFCyyyy, section 6.3.5."

       -- RFC Ed.: replace yyyy with actual RFC number & remove this
```

```
      note

   := { ianaIppmMetrics nn } -- IANA assigns nn

One-to-Group-Loss-Ratio OBJECT-IDENTITY

   STATUS current

   DESCRIPTION

      "Type-P-One-to-Group-Loss-Ratio"
```



```
   REFERENCE

      "Reference "RFCyyyy, section 6.4.1."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

   := { ianaIppmMetrics nn } -- IANA assigns nn

--

One-to-Group-Loss-Ratio-Range OBJECT-IDENTITY

   STATUS current

   DESCRIPTION

      "Type-P-One-to-Group-Loss-Ratio-Range"

   REFERENCE

      "Reference "RFCyyyy, section 6.4.2."

      -- RFC Ed.: replace yyyy with actual RFC number & remove this
      note

   := { ianaIppmMetrics nn } -- IANA assigns nn
```

--

14. References

14.1. Normative References

[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.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

[RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.

14.2. Informative References

[RFC2678] Mahdavi, J. and V. Paxson, "IPPM Metrics for Measuring Connectivity", RFC 2678, September 1999.

[RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.

[RFC3148] Mathis, M. and M. Allman, "A Framework for Defining Empirical Bulk Transfer Capacity Metrics", RFC 3148, July 2001.

[RFC3357] Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", RFC 3357, August 2002.

[RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.

[RFC3763] Shalunov, S. and B. Teitelbaum, "One-way Active Measurement Protocol (OWAMP) Requirements", RFC 3763, April 2004.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.

Authors' Addresses

Stephan Emile
France Telecom Division R&D
2 avenue Pierre Marzin
Lannion, F-22307

Fax: +33 2 96 05 18 52
Email: emile.stephan@orange-ftgroup.com

Lei Liang
CCSR, University of Surrey
Guildford
Surrey, GU2 7XH

Fax: +44 1483 683641
Email: L.Liang@surrey.ac.uk

Al Morton
200 Laurel Ave. South
Middletown, NJ 07748
USA

Phone: +1 732 420 1571
Email: acmorton@att.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).