

Network Working Group
Request for Comments: 5644
Category: Standards Track

E. Stephan
France Telecom
L. Liang
University of Surrey
A. Morton
AT&T Labs
October 2009

IP Performance Metrics (IPPM): Spatial and Multicast

Abstract

The IETF has standardized IP Performance Metrics (IPPM) for measuring end-to-end performance between two points. This memo defines two new categories of metrics that extend the coverage to multiple measurement points. It defines spatial metrics for measuring the performance of segments of a source to destination path, and metrics for measuring the performance between a source and many destinations in multiparty communications (e.g., a multicast tree).

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow

modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction and Scope	3
2. Terminology	4
3. Brief Metric Descriptions	7
4. Motivations	10
5. Spatial Vector Metrics Definitions	12
6. Spatial Segment Metrics Definitions	19
7. One-to-Group Metrics Definitions	27
8. One-to-Group Sample Statistics	30
9. Measurement Methods: Scalability and Reporting	40
10. Manageability Considerations	44
11. Security Considerations	49
12. Acknowledgments	50
13. IANA Considerations	50
14. References	56
14.1. Normative References	56
14.2. Informative References	57

1. Introduction and Scope

IETF has standardized IP Performance Metrics (IPPM) for measuring end-to-end performance between two points. This memo defines two new categories of metrics that extend the coverage to multiple measurement points. It defines spatial metrics for measuring the performance of segments of a source to destination path, and metrics for measuring the performance between a source and many destinations in multiparty communications (e.g., a multicast tree).

The purpose of this memo is to define metrics to fulfill the new requirements of measurement involving multiple measurement points. Spatial metrics measure the performance of each segment along a path. One-to-group metrics measure the performance for a group of users. These metrics are derived from one-way end-to-end metrics, all of which follow the IPPM framework [RFC2330].

This memo is organized as follows: Section 2 introduces new terms that extend the original IPPM framework [RFC2330]. Section 3 briefly introduces the new metrics, and Section 4 motivates each metric category. Sections 5 through 8 develop each category of metrics with definitions and statistics. Then the memo discusses the impact of the measurement methods on the scalability and proposes an information model for reporting the measurements. Finally, the memo discusses security aspects related to measurement and registers the metrics in the IANA IP Performance Metrics Registry [RFC4148].

The scope of this memo is limited to metrics using a single source packet or stream, and observations of corresponding packets along the path (spatial), at one or more destinations (one-to-group), or both. Note that all the metrics defined herein are based on observations of packets dedicated to testing, a process that is called active measurement. Passive measurement (for example, a spatial metric based on the observation of user traffic) is beyond the scope of this memo.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

2.1. Naming of the Metrics

The names of the metrics, including capitalized letters, are as close as possible of the names of the one-way end-to-end metrics they are derived from.

2.2. Terms Defined Elsewhere

host: section 5 of RFC 2330

router: section 5 of RFC 2330

loss threshold: section 2.8.2 of RFC 2680

path: section 5 of RFC 2330

sample: section 11 of RFC 2330

singleton: section 11 of RFC 2330

2.3. Routers Digest

The list of the routers on the path from the source to the destination that act as points of interest, also referred to as the routers digest.

2.4. Multiparty Metric

A metric is said to be multiparty if the topology involves more than one measurement collection point. All multiparty metrics designate a set of hosts as "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 $\langle ha, hb, hc, \dots, hn \rangle$, where ha is the source and $\langle hb, hc, \dots, hn \rangle$ are the destinations, then measurements may be conducted between $\langle ha, hb \rangle$, $\langle ha, hc \rangle$, ..., $\langle ha, hn \rangle$.

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

2.5. Spatial Metric

A metric is said to be spatial if one of the hosts (measurement collection points) involved is neither the source nor a destination of the measured packet(s). Such measurement hosts will usually be routers that are members of the routers digest.

2.6. 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 more than one destination. Thus, the topology of the communication group can be viewed as a center-distributed or server-client topology with the source as the center/server in the topology.

2.7. Points of Interest

Points of interest are the hosts (as per the RFC 2330 definition, "hosts" include routing nodes) that are measurement collection points, which are a sub-set of the set of hosts involved in the delivery of the packets (in addition to the source itself).

For spatial metrics, points of interest are a (possibly arbitrary) sub-set of all the routers involved in the path.

Points of interest of one-to-group metrics are the intended destination hosts for packets from the source (in addition to the source itself).





Figure 1: One-to-Group Points of Interest

A candidate point of interest for spatial metrics is a router from the set of routers involved in the delivery of the packets from source to destination.



Note: 'X' are nodes that are points of interest,
 'x' are nodes that are not points of interest

Figure 2: Spatial Points of Interest

2.8. Reference Point

A reference point is defined as the server where the statistical calculations will be carried out. It is usually a centralized server in the measurement architecture that is controlled by a network operator, where measurement data can be collected for further processing. The reference point is distinctly different from hosts at measurement collection points, where the actual measurements are carried out (e.g., points of interest).

2.9. Vector

A vector is a set of singletons (single atomic results) comprised of observations corresponding to a single source packet at different hosts in a network. For instance, if the one-way delay singletons observed at N receivers for Packet P sent by the source Src are dt_1, dt_2, \dots, dt_N , then a vector V with N elements can be organized as $\{dt_1, dt_2, \dots, dt_N\}$. The element dt_1 is distinct from all others as the singleton at receiver 1 in response to a packet sent from the source at a specific time. The complete vector gives information over the dimension of space, a set of N receivers in this example.

The singleton elements of any vector are distinctly different from each other in terms of their measurement collection point. Different vectors for common measurement points of interest are distinguished by the source packet sending time.

2.10. Matrix

Several vectors form a matrix, which contains results observed over a sampling interval at different places in a network at different times. For example, the one-way delay vectors $V1=\{dT11, dT12, \dots, dT1N\}$, $V2=\{dT21, dT22, \dots, dT2N\}$, ..., $Vm=\{dTm1, dTm2, \dots, dTmN\}$ for Packet $P1, P2, \dots, Pm$, form a one-way delay Matrix $\{V1, V2, \dots, Vm\}$. The matrix organizes the vector information to present network performance in both space and time.

A one-dimensional matrix (row) corresponds to a sample in simple point-to-point measurement.

The relationship among singleton, sample, vector, and matrix is illustrated in Figure 3.

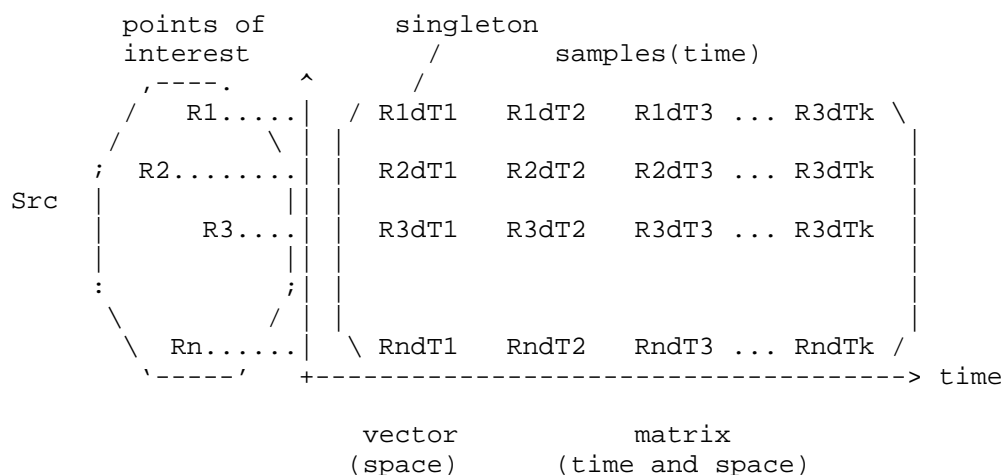


Figure 3: Relationship between Singletons, Samples, Vectors, and Matrix

3. Brief Metric Descriptions

The metrics for spatial and one-to-group measurement are based on the source-to-destination, or end-to-end metrics defined by IETF in [RFC2679], [RFC2680], [RFC3393], and [RFC3432].

This memo defines seven new spatial metrics using the [RFC2330] framework of parameters, units of measure, and measurement methodologies. Each definition includes a section that describes measurement constraints and issues, and provides guidance to increase the accuracy of the results.

The spatial metrics are:

- o Type-P-Spatial-One-way-Delay-Vector divides the end-to-end Type-P-One-way-Delay [RFC2679] into a spatial vector of one-way delay singletons.
- o Type-P-Spatial-One-way-Packet-Loss-Vector divides an end-to-end Type-P-One-way-Packet-Loss [RFC2680] into a spatial vector of packet loss singletons.
- o Type-P-Spatial-One-way-ipdv-Vector divides an end-to-end Type-P-One-way-ipdv into a spatial vector of ipdv (IP Packet Delay Variation) singletons.
- o Using elements of the Type-P-Spatial-One-way-Delay-Vector metric, a sample called Type-P-Segment-One-way-Delay-Stream collects one-way delay metrics between two points of interest on the path over time.
- o Likewise, using elements of the Type-P-Spatial-Packet-Loss-Vector metric, a sample called Type-P-Segment-Packet-Loss-Stream collects one-way delay metrics between two points of interest on the path over time.
- o Using the Type-P-Spatial-One-way-Delay-Vector metric, a sample called Type-P-Segment-ipdv-prev-Stream will be introduced to compute ipdv metrics (using the previous packet selection function) between two points of interest on the path over time.
- o Again using the Type-P-Spatial-One-way-Delay-Vector metric, a sample called Type-P-Segment-ipdv-min-Stream will define another set of ipdv metrics (using the minimum delay packet selection function) between two points of interest on the path over time.

The memo also defines three one-to-group metrics to measure the one-way performance between a source and a group of receivers. They are:

- o Type-P-One-to-group-Delay-Vector which collects the set of Type-P-One-way-Delay singletons between one sender and N receivers;

- o Type-P-One-to-group-Packet-Loss-Vector which collects the set of Type-P-One-way-Packet-Loss singletons between one sender and N receivers; and
- o Type-P-One-to-group-ipdv-Vector which collects the set of Type-P-One-way-ipdv singletons between one sender and N receivers.

Finally, based on the one-to-group vector metrics listed above, statistics are defined to capture single receiver performance, group performance, and the relative performance for a multiparty communication:

- o Using the Type-P-One-to-group-Delay-Vector, a metric called Type-P-One-to-group-Receiver-n-Mean-Delay, or RnMD, presents the mean of delays between one sender and a single receiver 'n'. From this metric, three additional metrics are defined to characterize the mean delay over the entire group of receivers during the same time interval:
 - * Type-P-One-to-group-Mean-Delay, or GMD, presents the mean of delays;
 - * Type-P-One-to-group-Range-Mean-Delay, or GRMD, presents the range of mean delays; and
 - * Type-P-One-to-group-Max-Mean-Delay, or GMMD, presents the maximum of mean delays.
- o Using the Type-P-One-to-group-Packet-Loss-Vector, a metric called Type-P-One-to-group-Receiver-n-Loss-Ratio, or RnLR, captures the packet loss ratio between one sender and a single receiver 'n'. Based on this definition, two more metrics are defined to characterize packet loss over the entire group during the same time interval:
 - * Type-P-One-to-group-Loss-Ratio, or GLR, captures the overall packet loss ratio for the entire group of receivers; and
 - * Type-P-One-to-group-Range-Loss-Ratio, or GRLR, presents the comparative packet loss ratio during the test interval between one sender and N receivers.
- o Using the Type-P-One-to-group-Packet-Loss-Vector, a metric called Type-P-One-to-group-Receiver-n-Comp-Loss-Ratio, or RnCLR, computes a packet loss ratio using the maximum number of packets received at any receiver.

- o Using Type-P-One-to-group-ipdv-Vector, a metric called Type-P-One-to-group-Range-Delay-Variation, or GRDV, presents the range of delay variation between one sender and a group of receivers.

4. Motivations

All existing IPPM metrics are defined for end-to-end (source-to-destination) measurement of point-to-point paths. It is logical to extend them to multiparty situations such as one-to-one trajectory metrics and one-to-multipoint metrics.

4.1. Motivations for Spatial Metrics

Spatial metrics are needed for:

- o Decomposing the performance of an inter-domain path to quantify the per-AS (Autonomous System) contribution to the end-to-end performance.
- o Traffic engineering and troubleshooting, which benefit from spatial views of one-way delay and ipdv consumption, or identification of the path segment where packets were lost.
- o Monitoring the decomposed performance of a multicast tree based on MPLS point-to-multipoint communications.
- o Dividing end-to-end metrics, so that some segment measurements can be re-used and help measurement systems reach large-scale coverage. Spatial measures could characterize the performance of an intra-domain segment and provide an elementary piece of information needed to estimate inter-domain performance to another destination using Spatial Composition metrics [SPATIAL].

4.2. 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 topology. A simple point-to-point metric cannot completely describe the multiparty situation. New one-to-group metrics assess performance of the multiple paths for further statistical analysis. The new metrics are named one-to-group performance metrics, and they are based on the unicast metrics defined in IPPM RFCs. One-to-group metrics are one-way metrics from one source to a group of destinations or receivers. The metrics are helpful for judging the overall performance of a multiparty communications network and for describing the performance variation across a group of destinations.

One-to-group performance metrics are needed for:

- o Designing and engineering multicast trees and MPLS point-to-multipoint Label Switched Paths (LSPs).
- o Evaluating and controlling the quality of multicast services, including inter-domain multicast.
- o Presenting and evaluating the performance requirements for multiparty communications and overlay multicast.

To understand the packet transfer performance between one source and any one receiver in the multiparty communication group, we need to collect instantaneous end-to-end metrics, or singletons. This gives a very detailed view into the performance of each branch of the multicast tree, and can provide clear and helpful information for engineers to identify the branch with problems in a complex multiparty routing tree.

The one-to-group metrics described in this memo introduce the multiparty topology into the IPPM framework, and they describe the performance delivered to a group receiving packets from the same source. The concept extends the "path" of the point-to-point measurement to "path tree" to cover one-to-many topologies. If applied to one-to-one topology, the one-to-group metrics provide exactly the same results as the corresponding one-to-one metrics.

4.3. 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 beyond the scope of this memo, because they would involve multiple packets launched from different sources. However, this section gives some insights 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 host that is acting as a receiver while all other hosts act as sources in this case.

The group-to-group communication topology has no obvious focal point: the sources and the measurement collection points can be anywhere. However, it is possible to organize the problem by applying measurements in one-to-group or group-to-one topologies for each host in a uniform way (without taking account of how the real

communication might be carried out). For example, one group of hosts $\langle ha, hb, hc, \dots, hn \rangle$ might act as sources to send data to another group of hosts $\langle Ha, Hb, Hc, \dots, Hm \rangle$, and they can be organized into n sets of points of interest for one-to-group communications:

$\langle ha, Ha, Hb, Hc, \dots, Hm \rangle, \langle hb, Ha, Hb, Hc, \dots, Hm \rangle, \langle hc, Ha, Hb, Hc, \dots, Hm \rangle, \dots, \langle hn, Ha, Hb, Hc, \dots, Hm \rangle$.

5. Spatial Vector Metrics Definitions

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

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

The spatial vector definitions are coupled with the corresponding end-to-end metrics. Measurement methodology aspects are common to all the vectors defined and are consequently discussed in a common section.

5.1. A Definition for Spatial One-Way Delay Vector

This section is coupled with the definition of Type-P-One-way-Delay in section 3 of [RFC2679]. When a parameter from the definition in [RFC2679] is re-used in this section, the first instance 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, H_i , involved in the measure. Spatial one-way delay measurements MUST respect them, especially those related to methodology, clock, uncertainties, and reporting.

5.1.1. Metric Name

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

5.1.2. Metric Parameters

- o Src*, the IP address of the sender.
- o Dst*, the IP address of the receiver.
- o i , an integer in the ordered list $\langle 1, 2, \dots, n \rangle$ of routers in the path.

- o H_i , a router of the routers 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 dT_i , a delay, the one-way delay for a measured packet from the source to router H_i .
- o $\langle dT_1, \dots, dT_i, \dots, dT_n \rangle$ a list of n delay singletons.
- o Type-P*, the specification of the packet type.
- o $\langle H_1, H_2, \dots, H_n \rangle$ the routers digest.

5.1.3. Metric Units

The value of Type-P-Spatial-One-way-Delay-Vector is a sequence of times (a real number in the dimension of seconds with sufficient resolution to convey the results).

5.1.4. Definition

Given a Type-P packet sent by the Src at wire-time (first bit) T to the receiver Dst on the path $\langle H_1, H_2, \dots, H_n \rangle$. There is a sequence of values $\langle T+dT_1, T+dT_2, \dots, T+dT_n, T+dT \rangle$ such that dT is the Type-P-One-way-Delay from Src to Dst, and for each H_i of the path, $T+dT_i$ is either a real number corresponding to the wire-time the packet passes (last bit received) H_i , or undefined if the packet does not pass H_i within a specified loss threshold* time.

Type-P-Spatial-One-way-Delay-Vector metric is defined for the path $\langle \text{Src}, H_1, H_2, \dots, H_n, \text{Dst} \rangle$ as the sequence of values $\langle T, dT_1, dT_2, \dots, dT_n, dT \rangle$.

5.1.5. Discussion

Some specific issues that may occur are as follows:

- o the delay singletons "appear" to decrease: $dT_i > dT_{i+1}$. This may occur despite being physically impossible with the definition used.

- * This is frequently due to a measurement clock synchronization issue. This point is discussed in section 3.7.1 "Errors or uncertainties related to Clocks" of [RFC2679]. Consequently, the values of delays measured at multiple routers may not match the order of those routers on the path.
- * The actual order of routers on the path may change due to reconvergence (e.g., recovery from a link failure).
- * The location of the measurement collection point in the device influences the result. If the packet is not observed directly on the input interface, the delay includes buffering time and consequently an uncertainty due to the difference between 'wire-time' and 'host time'.

5.2. A Definition for Spatial Packet Loss Vector

This section is coupled with the definition of Type-P-One-way-Packet-Loss. When a parameter from section 2 of [RFC2680] is used in this section, the first instance 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, H_i , involved in the measure. Spatial packet loss measurement MUST respect them, especially those related to methodology, clock, uncertainties, and reporting.

The following sections define the spatial loss vector, adapt some of the points above, and introduce points specific to spatial loss measurement.

5.2.1. Metric Name

Type-P-Spatial-Packet-Loss-Vector

5.2.2. Metric Parameters

- o Src*, the IP address of the sender.
- o Dst*, the IP address of the receiver.
- o i, an integer in the ordered list $\langle 1, 2, \dots, n \rangle$ of routers in the path.
- o H_i , a router of the routers digest.
- o T*, a time, the sending time for a measured packet.

- o dTi , a delay, the one-way delay for a measured packet from the source to host H_i .
- o $\langle dT1, \dots, dTn \rangle$, list of n delay singletons.
- o Type-P*, the specification of packet type.
- o $\langle H1, H2, \dots, Hn \rangle$, the routers digest.
- o $\langle L1, L2, \dots, Ln \rangle$, a list of Boolean values.

5.2.3. Metric Units

The value of Type-P-Spatial-Packet-Loss-Vector is a sequence of Boolean values.

5.2.4. Definition

Given a Type-P packet sent by the Src at time T to the receiver Dst on the path $\langle H1, H2, \dots, Hn \rangle$. For the sequence of times $\langle T+dT1, T+dT2, \dots, T+dTi, \dots, T+dTn \rangle$ the packet passes in $\langle H1, H2, \dots, Hi, \dots, Hn \rangle$, define the Type-P-Packet-Loss-Vector metric as the sequence of values $\langle T, L1, L2, \dots, Ln \rangle$ such that for each H_i of the path, a value of 0 for L_i means that dTi is a finite value, and a value of 1 means that dTi is undefined.

5.2.5. Discussion

Some specific issues that may occur are as follows:

- o The result might include the sequence of values 1,0. Although this appears physically impossible (a packet is lost, then reappears later on the path):
 - * The actual routers on the path may change due to reconvergence (e.g., recovery from a link failure).
 - * The order of routers on the path may change due to reconvergence.
 - * A packet may not be observed in a router due to some buffer or CPU overflow at the measurement collection point.

5.3. A Definition for Spatial One-Way ipdv Vector

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

The following sections define the spatial ipdv vector, adapt some of the points above, and introduce points specific to spatial ipdv measurement.

5.3.1. Metric Name

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

5.3.2. Metric Parameters

- o Src*, the IP address of the sender.
- o Dst*, the IP address of the receiver.
- o i, an integer in the ordered list $\langle 1, 2, \dots, n \rangle$ of routers in the path.
- o Hi, a router of the routers digest.
- o T1*, a time, the sending time for a first measured packet.
- o T2*, a time, the sending time for a second measured packet.
- o dT*, a delay, the one-way delay for a measured packet.
- o dTi, a delay, the one-way delay for a measured packet from the source to router Hi.
- o Type-P*, the specification of the packet type.
- o P1, the first packet sent at time T1.
- o P2, the second packet sent at time T2.
- o $\langle H1, H2, \dots, Hn \rangle$, the routers digest.
- o $\langle T1, dT1.1, dT1.2, \dots, dT1.n, dT1 \rangle$, the Type-P-Spatial-One-way-Delay-Vector for a packet sent at time T1.
- o $\langle T2, dT2.1, dT2.2, \dots, dT2.n, dT2 \rangle$, the Type-P-Spatial-One-way-Delay-Vector for a packet sent at time T2.
- o 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.

5.3.3. Metric Units

The value of Type-P-Spatial-One-way-ipdv-Vector is a sequence of times (a real number in the dimension of seconds with sufficient resolution to convey the results).

5.3.4. Definition

Given P1 the Type-P packet sent by the sender Src at wire-time (first bit) T1 to the receiver Dst. Given $\langle T1, dT1.1, dT1.2, \dots, dT1.n, dT1 \rangle$ the Type-P-Spatial-One-way-Delay-Vector of P1 over the sequence of routers $\langle H1, H2, \dots, Hn \rangle$.

Given P2 the Type-P packet sent by the sender Src at wire-time (first bit) T2 to the receiver Dst. Given $\langle T2, dT2.1, dT2.2, \dots, dT2.n, dT2 \rangle$ the Type-P-Spatial-One-way-Delay-Vector of P2 over the same path.

The Type-P-Spatial-One-way-ipdv-Vector metric is defined as the sequence of values $\langle T1, T2, dT2.1-dT1.1, dT2.2-dT1.2, \dots, dT2.n-dT1.n, dT2-dT1 \rangle$ such that for each H_i of the sequence of routers $\langle H1, H2, \dots, Hn \rangle$, $dT2.i-dT1.i$ is either a real number if the packets P1 and P2 pass H_i at wire-time (last bit) $dT1.i$ and $dT2.i$ respectively, or undefined if at least one of them never passes H_i (and the respective one-way delay is undefined). The $T1, T2^*$ pair indicates the inter-packet emission interval and $dT2-dT1$ is ddT^* the Type-P-One-way-ipdv.

5.4. Spatial Methodology

The methodology, reporting specifications, and uncertainties specified in section 3 of [RFC2679] apply to each point of interest (or measurement collection point), H_i , measuring an element of a spatial delay vector.

Likewise, the methodology, reporting specifications, and uncertainties specified in section 2 of [RFC2680] apply to each point of interest, H_i , measuring an 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, H_i , involved in the measure. Spatial One-way ipdv measurement MUST respect the methodology, clock, uncertainties, and reporting aspects given there.

Generally, for a given Type-P packet of length L at a specific Hi, the methodology for spatial vector metrics may proceed as follows:

- o At each Hi, points of interest/measurement collection points prepare to capture the packet sent at time T, record a timestamp Ti', and 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 (TTL));
- o Each Hi computes the corrected wire-time from Src to Hi: $T_i = T_i' - dTi'$. This arrival time is undefined 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: $dTi = T_i - T$;
- o The reference point gathers the result of each Hi and arranges them according to the path ordering information received to build the Type-P spatial one-way vector (e.g., Type-P-Spatial-One-way-Delay-Vector metric $\langle T, dT_1, dT_2, \dots, dT_n, dT \rangle$ over the path $\langle Src, H_1, H_2, \dots, H_n, Dst \rangle$ at time T.

5.4.1. Packet Loss Detection

In a pure end-to-end measurement, packet losses are detected by the receiver only. A packet is lost when Type-P-One-way-Delay is undefined or very large (see sections 2.4 and 2.5 of [RFC2680] and section 3.5 of [RFC2680]). A packet is deemed lost by the receiver after a duration that starts at the time the packet is sent. This timeout value is chosen by a measurement process. It determines the threshold between recording a long packet transfer time as a finite value or an undefined value.

In a spatial measurement, packet losses may be detected at several measurement collection points. Depending on the consistency of the packet loss detections among the points of interest, a packet may be considered as lost at one point despite having a finite delay at another, or it may be observed by the last measurement collection point of the path but considered lost by Dst.

There is a risk of misinterpreting such results: has the path changed? Did the packet arrive at the destination or was it lost on the very last link?

The same concern applies to one-way delay measures: a delay measured may be computed as infinite by one observation point but as a real value by another one, or may be measured as a real value by the last observation point of the path but designated as undefined by Dst.

The observation/measurement collection points and the destination SHOULD use consistent methods to detect packets losses. The methods and parameters must be systematically reported to permit careful comparison and to avoid introducing any confounding factors in the analysis.

5.4.2. Routers Digest

The methodology given above relies on knowing the order of the router/measurement collection points on the path [RFC2330].

Path instability might cause a test packet to be observed more than once by the same router, resulting in the repetition of one or more routers in the routers digest.

For example, repeated observations may occur during rerouting phases that introduce temporary micro loops. During such an event, the routers digest for a packet crossing Ha and Hb may include the pattern <Hb, Ha, Hb, Ha, Hb>, meaning that Ha ended the computation of the new path before Hb and that the initial path was from Ha to Hb, and that the new path is from Hb to Ha.

Consequently, duplication of routers in the routers digest of a vector MUST be identified before computation of statistics to avoid producing corrupted information.

6. Spatial Segment Metrics Definitions

This section defines samples to measure the performance of a segment of a path over time. The definitions rely on the matrix of the spatial vector metrics defined above.

First, this section defines a sample of one-way delay, Type-P-Segment-One-way-Delay-Stream, and a sample of packet loss, Type-P-Segment-Packet-Loss-Stream.

Then, it defines two different samples of ipdv: Type-P-Segment-ipdv-prev-Stream uses the current and previous packets as the selection function, and Type-P-Segment-ipdv-min-Stream uses the minimum delay as one of the selected packets in every pair.

6.1. A Definition of a Sample of One-Way Delay of a Segment of the Path

This metric defines a sample of one-way delays over time between a pair of routers on a path. Since it is very close semantically to the metric Type-P-One-way-Delay-Poisson-Stream defined in section 4 of [RFC2679], sections 4.5 to 4.8 of [RFC2679] are integral parts of the definition text below.

6.1.1. Metric Name

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

6.1.2. Metric Parameters

- o Src, the IP address of the sender.
- o Dst, the IP address of the receiver.
- o Type-P, the specification of the packet type.
- o i, an integer in the ordered list $\langle 1, 2, \dots, n \rangle$ of routers in the path.
- o k, an integer that orders the packets sent.
- o a and b, two integers where $b > a$.
- o H_i , a router of the routers digest.
- o $\langle H_1, \dots, H_a, \dots, H_b, \dots, H_n \rangle$, the routers digest.
- o $\langle T_1, T_2, \dots, T_m \rangle$, a list of times.

6.1.3. Metric Units

The value of a Type-P-Segment-One-way-Delay-Stream is a pair of:

A list of times $\langle T_1, T_2, \dots, T_m \rangle$; and

A sequence of delays.

6.1.4. Definition

Given two routers, H_a and H_b , of the path $\langle H_1, H_2, \dots, H_a, \dots, H_b, \dots, H_n \rangle$, and the matrix of Type-P-Spatial-One-way-Delay-Vector for the packets sent from Src to Dst at times $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$:

```

<T1, dT1.1, dT1.2, ..., dT1.a, ..., dT1.b,..., dT1.n, dT1>;
<T2, dT2.1, dT2.2, ..., dT2.a, ..., dT2.b,..., dT2.n, dT2>;
...
<Tm, dTm.1, dTm.2, ..., dTm.a, ..., dTm.b,..., dTm.n, dTm>.

```

We define the sample Type-P-Segment-One-way-Delay-Stream as the sequence $\langle dT1.ab, dT2.ab, \dots, dTk.ab, \dots, dTm.ab \rangle$ such that for each time T_k , 'dTk.ab' is either the real number 'dTk.b - dTk.a', if the packet sent at the time T_k passes H_a and H_b , or is undefined if this packet never passes H_a or (inclusive) never passes H_b .

6.1.5. Discussion

Some specific issues that may occur are as follows:

- o the delay singletons "appear" to decrease: $dT_i > dT_{i+1}$, and is discussed in section 5.1.5.
- * This could also occur when the clock resolution of one measurement collection point is larger than the minimum delay of a path. For example, the minimum delay of a 500 km path through optical fiber facilities is 2.5 ms, but the measurement collection point has a clock resolution of 8 ms.

The metric SHALL be invalid for times $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$ if the following conditions occur:

- o H_a or H_b disappears from the path due to some routing change.
- o The order of H_a and H_b changes in the path.

6.2. A Definition of a Sample of Packet Loss of a Segment of the Path

This metric defines a sample of packet loss over time between a pair of routers of a path. Since it is very close semantically to the metric Type-P-Packet-loss-Stream defined in section 3 of [RFC2680], sections 3.5 to 3.8 of [RFC2680] are integral parts of the definition text below.

6.2.1. Metric Name

Type-P-Segment-Packet-Loss-Stream

6.2.2. Metric Parameters

- o Src, the IP address of the sender.
- o Dst, the IP address of the receiver.
- o Type-P, the specification of the packet type.
- o k, an integer that orders the packets sent.
- o n, an integer that orders the routers on the path.
- o a and b, two integers where $b > a$.
- o $\langle H_1, H_2, \dots, H_a, \dots, H_b, \dots, H_n \rangle$, the routers digest.
- o H_i , a router of the routers digest.
- o $\langle T_1, T_2, \dots, T_m \rangle$, a list of times.
- o $\langle L_1, L_2, \dots, L_n \rangle$, a list of Boolean values.

6.2.3. Metric Units

The value of a Type-P-Segment-Packet-Loss-Stream is a pair of:

The list of times $\langle T_1, T_2, \dots, T_m \rangle$; and

A sequence of Boolean values.

6.2.4. Definition

Given two routers, H_a and H_b , of the path $\langle H_1, H_2, \dots, H_a, \dots, H_b, \dots, H_n \rangle$ and the matrix of Type-P-Spatial-Packet-Loss-Vector for the packets sent from Src to Dst at times $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$:

$\langle T_1, L_{1.1}, L_{1.2}, \dots, L_{1.a}, \dots, L_{1.b}, \dots, L_{1.n}, L \rangle$,
 $\langle T_2, L_{2.1}, L_{2.2}, \dots, L_{2.a}, \dots, L_{2.b}, \dots, L_{2.n}, L \rangle$,
 \dots ,
 $\langle T_m, L_{m.1}, L_{m.2}, \dots, L_{m.a}, \dots, L_{m.b}, \dots, L_{m.n}, L \rangle$.

We define the value of the sample Type-P-Segment-Packet-Loss-Stream from H_a to H_b as the sequence of Booleans $\langle L_{1.ab}, L_{2.ab}, \dots, L_{k.ab}, \dots, L_{m.ab} \rangle$ such that for each T_k :

- o A value of Lk of 0 means that Ha and Hb observed the packet sent at time Tk (both Lk.a and Lk.b have a value of 0).
- o A value of Lk of 1 means that Ha observed the packet sent at time Tk (Lk.a has a value of 0) and that Hb did not observe the packet sent at time Tk (Lk.b has a value of 1).
- o The value of Lk is undefined when neither Ha nor Hb observed the packet (both Lk.a and Lk.b have a value of 1).

6.2.5. Discussion

Unlike Type-P-Packet-loss-Stream, Type-P-Segment-Packet-Loss-Stream relies on the stability of the routers digest. The metric SHALL be invalid for times $\langle T1, T2, \dots, Tm-1, Tm \rangle$ if the following conditions occur:

- o Ha or Hb disappears from the path due to some routing change.
- o The order of Ha and Hb changes in the path.
- o Lk.a or Lk.b is undefined.
- o Lk.a has the value 1 (not observed) and Lk.b has the value 0 (observed).
- o L has the value 0 (the packet was received by Dst) and Lk.ab has the value 1 (the packet was lost between Ha and Hb).

6.3. A Definition of a Sample of ipdv of a Segment Using the Previous Packet Selection Function

This metric defines a sample of ipdv [RFC3393] over time between a pair of routers using the previous packet as the selection function.

6.3.1. Metric Name

Type-P-Segment-ipdv-prev-Stream

6.3.2. Metric Parameters

- o Src, the IP address of the sender.
- o Dst, the IP address of the receiver.
- o Type-P, the specification of the packet type.
- o k, an integer that orders the packets sent.

- o n , an integer that orders the routers on the path.
- o a and b , two integers where $b > a$.
- o $\langle H1, H2, \dots, Ha, \dots, Hb, \dots, Hn \rangle$, the routers digest.
- o $\langle T1, T2, \dots, Tm-1, Tm \rangle$, a list of times.
- o $\langle Tk, dTk.1, dTk.2, \dots, dTk.a, \dots, dTk.b, \dots, dTk.n, dTk \rangle$, a Type-P-Spatial-One-way-Delay-Vector.

6.3.3. Metric Units

The value of a Type-P-Segment-ipdv-prev-Stream is a pair of:

The list of $\langle T1, T2, \dots, Tm-1, Tm \rangle$; and

A list of pairs of interval of times and delays;

6.3.4. Definition

Given two routers, Ha and Hb , of the path $\langle H1, H2, \dots, Ha, \dots, Hb, \dots, Hn \rangle$ and the matrix of Type-P-Spatial-One-way-Delay-Vector for the packets sent from Src to Dst at times $\langle T1, T2, \dots, Tm-1, Tm \rangle$:

$\langle T1, dT1.1, dT1.2, \dots, dT1.a, \dots, dT1.b, \dots, dT1.n, dT1 \rangle$,
 $\langle T2, dT2.1, dT2.2, \dots, dT2.a, \dots, dT2.b, \dots, dT2.n, dT2 \rangle$,
 ...
 $\langle Tm, dTm.1, dTm.2, \dots, dTm.a, \dots, dTm.b, \dots, dTm.n, dTm \rangle$.

We define the Type-P-Segment-ipdv-prev-Stream as the sequence of packet time pairs and delay variations

$\langle (T1, T2, dT2.ab - dT1.ab), \dots,$
 $(Tk-1, Tk, dTk.ab - dTk-1.ab), \dots,$
 $(Tm-1, Tm, dTm.ab - dTm-1.ab) \rangle$

For any pair, $Tk, Tk-1$ in $k=1$ through m , the difference $dTk.ab - dTk-1.ab$ is undefined if:

- o the delay $dTk.a$ or the delay $dTk-1.a$ is undefined, OR
- o the delay $dTk.b$ or the delay $dTk-1.b$ is undefined.

6.3.5. Discussion

This metric belongs to the family of inter-packet delay variation metrics (IPDV in uppercase) whose results are extremely sensitive to the inter-packet interval in practice.

The inter-packet interval of an end-to-end IPDV metric is under the control of the source (ingress point of interest). In contrast, the inter-packet interval of a segment IPDV metric is not under the control the ingress point of interest of the measure, Ha. The interval will certainly vary if there is delay variation between the Source and Ha. Therefore, the ingress inter-packet interval must be known at Ha in order to fully comprehend the delay variation between Ha and Hb.

6.4. A Definition of a Sample of ipdv of a Segment Using the Minimum Delay Selection Function

This metric defines a sample of ipdv [RFC3393] over time between a pair of routers on a path using the minimum delay as one of the selected packets in every pair.

6.4.1. Metric Name

Type-P-Segment-One-way-ipdv-min-Stream

6.4.2. Metric Parameters

- o Src, the IP address of the sender.
- o Dst, the IP address of the receiver.
- o Type-P, the specification of the packet type.
- o k, an integer that orders the packets sent.
- o i, an integer that identifies a packet sent.
- o n, an integer that orders the routers on the path.
- o a and b, two integers where $b > a$.
- o $\langle H1, H2, \dots, Ha, \dots, Hb, \dots, Hn \rangle$, the routers digest.
- o $\langle T1, T2, \dots, Tm-1, Tm \rangle$, a list of times.
- o $\langle Tk, dTk.1, dTk.2, \dots, dTk.a, \dots, dTk.b, \dots, dTk.n, dTk \rangle$, a Type-P-Spatial-One-way-Delay-Vector.

6.4.3. Metric Units

The value of a Type-P-Segment-One-way-ipdv-min-Stream is a pair of:

The list of $\langle T1, T2, \dots, Tm-1, Tm \rangle$; and

A list of times.

6.4.4. Definition

Given two routers, H_a and H_b , of the path $\langle H1, H2, \dots, H_a, \dots, H_b, \dots, H_n \rangle$ and the matrix of Type-P-Spatial-One-way-Delay-Vector for the packets sent from Src to Dst at times $\langle T1, T2, \dots, Tm-1, Tm \rangle$:

```

<T1, dT1.1, dT1.2, ..., dT1.a, ..., dT1.b, ..., dT1.n, dT1>,
<T2, dT2.1, dT2.2, ..., dT2.a, ..., dT2.b, ..., dT2.n, dT2>,
...
<Tm, dTm.1, dTm.2, ..., dTm.a, ..., dTm.b, ..., dTm.n, dTm>.

```

We define the Type-P-Segment-One-way-ipdv-min-Stream as the sequence of times $\langle dT1.ab - \min(dTi.ab), \dots, dTk.ab - \min(dTi.ab), \dots, dTm.ab - \min(dTi.ab) \rangle$ where:

- o $\min(dTi.ab)$ is the minimum value of the tuples $(dT_k.b - dT_k.a)$;
- o for each time T_k , $dT_k.ab$ is undefined if $dT_k.a$ or (inclusive) $dT_k.b$ is undefined, or the real number $(dT_k.b - dT_k.a)$ is undefined.

6.4.5. Discussion

This metric belongs to the family of packet delay variation metrics (PDV). PDV distributions have less sensitivity to inter-packet interval variations than IPDV values, as discussed above.

In principle, the PDV distribution reflects the variation over many different inter-packet intervals, from the smallest inter-packet interval, up to the length of the evaluation interval, $Tm - T1$. Therefore, when delay variation occurs and disturbs the packet spacing observed at H_a , the PDV results will likely compare favorably to a PDV measurement where the source is H_a and the destination is H_b , because a wide range of spacings are reflected in any PDV distribution.

7. One-to-Group Metrics Definitions

This section defines performance metrics between a source and a group of receivers.

7.1. A Definition for One-to-Group Delay

This section defines a metric for one-way delay between a source and a group of receivers.

7.1.1. Metric Name

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

7.1.2. 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 times.
- o Type-P, the specification of the packet type.
- o Gr, the receiving group identifier. The parameter Gr is the multicast group address if the measured packets are transmitted over IP multicast. This parameter is to differentiate the measured traffic from other unicast and multicast traffic. It is OPTIONAL for this metric to avoid losing any generality, i.e., to make the metric also applicable to unicast measurement where there is only one receiver.

7.1.3. Metric Units

The value of a Type-P-One-to-group-Delay-Vector is a set of Type-P-One-way-Delay singletons [RFC2679], that is a sequence of times (a real number in the dimension of seconds with sufficient resolution to convey the results).

7.1.4. Definition

Given a Type-P packet sent by the source Src at time T, and the N hosts { Recv1,...,RecvN } which receive the packet at the time { T+dT1,...,T+dTn }, or the packet does not pass a receiver within a specified loss threshold time, then the Type-P-One-to-group-Delay-

Vector is defined as the set of the Type-P-One-way-Delay singletons between Src and each receiver with value of { dT1, dT2,...,dTn }, where any of the singletons may be undefined if the packet did not pass the corresponding receiver within a specified loss threshold time.

7.2. A Definition for One-to-Group Packet Loss

7.2.1. Metric Name

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

7.2.2. 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 Type-P, the specification of the packet type.
- o Gr, the receiving group identifier, OPTIONAL.

7.2.3. Metric Units

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

- o T, time the source packet was sent.
- o L1,...,LN a list of Boolean values.

7.2.4. Definition

Given a Type-P packet sent by the source Src at T and the N hosts, Recv1,...,RecvN, the Type-P-One-to-group-Packet-Loss-Vector is defined as a set of the Type-P-One-way-Packet-Loss singletons between Src and each of the receivers:

{T, <L1=0|1>,<L2=0|1>,..., <LN=0|1>},

where the Boolean value 0|1 depends on receiving the packet at a particular receiver within a loss threshold time.

7.3. A Definition for One-to-Group ipdv

7.3.1. Metric Name

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

7.3.2. 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 T1, a time.
- o T2, a time.
- o ddT1, ...,ddTn, a list of times.
- o Type-P, the specification of the packet type.
- o F, a selection function non-ambiguously defining the two packets from the stream selected for the metric.
- o Gr, the receiving group identifier. The parameter Gr is the multicast group address if the measured packets are transmitted over IP multicast. This parameter is to differentiate the measured traffic from other unicast and multicast traffic. It is 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 receiver.

7.3.3. Metric Units

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

7.3.4. Definition

Given a Type-P packet stream, Type-P-One-to-group-ipdv-Vector is defined for two packets transferred 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-Delay-Vector from Src to { Recv1,..., RecvN } at time T1 and the value of the Type-P-One-to-group-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 T_2 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-Delay-Vector metric.

For a set of real numbers $\{ddT_1, \dots, ddT_n\}$, the Type-P-One-to-group-ipdv-Vector from Src to $\{Recv_1, \dots, Recv_N\}$ at T_1, T_2 is $\{ddT_1, \dots, ddT_n\}$ means that Src sent two packets, the first at wire-time T_1 (first bit), and the second at wire-time T_2 (first bit) and the packets were received by $\{Recv_1, \dots, Recv_N\}$ at wire-time $\{dT_1 + T_1, \dots, dT_n + T_1\}$ (last bit of the first packet), and at wire-time $\{dT'_1 + T_2, \dots, dT'_n + T_2\}$ (last bit of the second packet), and that $\{dT'_1 - dT_1, \dots, dT'_n - dT_n\} = \{ddT_1, \dots, ddT_n\}$.

For any pair of selected packets, the difference $dT'_n - dT_n$ is undefined if:

- o the delay dT_n to Receiver n is undefined, OR
- o the delay dT'_n to Receiver n is undefined.

8. One-to-Group Sample Statistics

The one-to-group metrics defined above are directly achieved by collecting relevant unicast one-way metrics measurements results and by gathering them per group of receivers. They produce network performance information that guides engineers toward potential problems that may have happened on any branch of a multicast routing tree.

The results of these metrics are not directly usable to present the performance of a group because each result is made of a huge number of singletons that are difficult to read and analyze. As an example, delays are not comparable because the distance between receiver and sender differs. Furthermore, they don't capture relative performance situations in a multiparty communication.

From the performance point of view, the multiparty communication services not only require the support of absolute performance information but also information on "relative performance". "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 online gaming. A very small difference in delay might result in failure in the game. We have to use multicast-

specific statistic metrics to define the relative delay required by online gaming. There are many other services, e.g., online bidding, 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 4. 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.

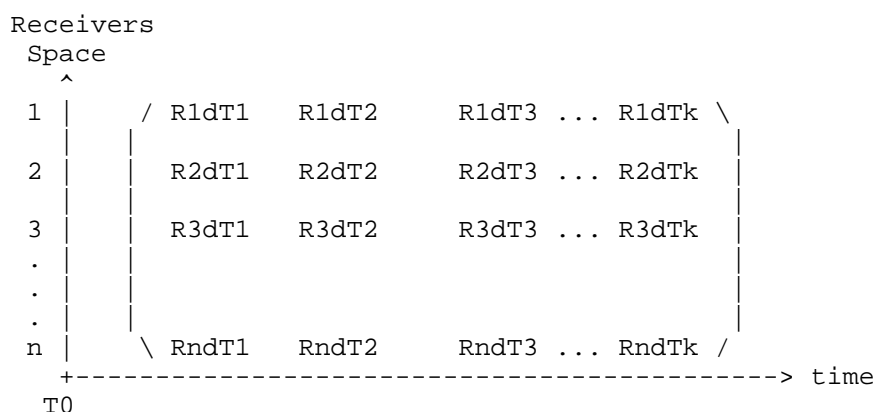


Figure 4: Matrix M (n*m)

In Matrix M, each element is a one-way delay singleton. Each column is a delay vector. It contains the one-way delays of the same packet observed at n 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 providers in which dimension they are interested. 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 in knowing 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 others.

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 a lower transmission rate can have the time interval be longer, and ones with higher transmission rate can have the time interval be 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 these statistics are 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 $\{V_1, V_2, \dots, V_N\}$. 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 well the service has been delivered to a group of users during a sampling interval in terms of delay. It requires twice as much calculation to have this statistic over both time and space dimensions. These kinds of statistics are referred to as 2-level statistics to distinguish them from 1-level statistics calculated over either space or time dimension. It can be easily proven that no matter over which dimension a 2-level statistic is calculated first, the results are the same. That is, 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 results will be exactly the same. Therefore, when defining a 2-level statistic, there is no need to specify the order in which the calculation is executed.

Many statistics can be defined for the proposed one-to-group metrics over the space dimension, 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.

8.1. Discussion on the Impact of Packet Loss on Statistics

Packet loss does have effects on one-way metrics and their statistics. For example, a lost packet can result in 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 such a 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 complexity 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 to be 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 a 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 cannot calculate the statistics. The extreme situation being the case when no packets arrive at any user. 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 { RldT1 RldT2 RldT3 ... RldTK-1 UNDEF RldTK+1... RlMD } where "UNDEF" is an undefined value, the reference point can replace it by $RldTK = \{(RldTK-1) + (RldTK+1)\} / 2$. Therefore, this result can be used to build up the group Matrix with an estimated value RldTK. There are other possible solutions, such as using the overall mean of the whole result to replace the infinite/undefined value, and so on. However, this 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 R1MD as shown in Figure 7 without any packet loss, and User2 calculates the R2MD with N-2 packet loss. The R1MD and R2MD should not be treated with equal weight because R2MD was calculated only based on two 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.

8.2. General Metric Parameters

- o Src, the IP address of a host.
- o G, the receiving group identifier.
- o N, the number of Receivers (Recv1, Recv2, ... RecvN).
- o T, a time (start of test interval).
- o Tf, a time (end of test interval).
- o K, the number of packets sent from the source during the test interval.
- o J[n], the number of packets received at a particular Receiver, n, where $1 \leq n \leq N$.
- o lambda, a rate in reciprocal seconds (for Poisson Streams).
- o incT, the nominal duration of inter-packet interval, first bit to first bit (for Periodic Streams).
- o 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).
- o TstampSrc, the wire-time of the packet as measured at MP(Src) (the Source Measurement Point).
- o TstampRecv, the wire-time of the packet as measured at MP(Recv), assigned to packets that arrive within a "reasonable" time.
- o 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.

8.3. One-to-Group Delay Statistics

This section defines the overall one-way delay statistics for a receiver and for an entire group as illustrated by the matrix below.

Recv	/----- Sample -----\					Stats	Group Stat
1	R1dT1	R1dT2	R1dT3	...	R1dT _k	R1MD \	} Group Delay
2	R2dT1	R2dT2	R2dT3	...	R2dT _k	R2MD	
3	R3dT1	R3dT2	R3dT3	...	R3dT _k	R3MD	
.							
.							
n	RndT1	RndT2	RndT3	...	RndT _k	RnMD /	

Receiver-n
Delay

Figure 5: One-to-Group Mean Delay

Statistics are computed on the finite one-way delays of the matrix above.

All one-to-group delay statistics are expressed in seconds with sufficient resolution to convey three significant digits.

8.3.1. Type-P-One-to-group-Receiver-n-Mean-Delay

This section defines Type-P-One-to-group-Receiver-n-Mean-Delay, the Delay Mean, at each Receiver N, also named RnMD.

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 the finite value TstampRecv[i] - TstampSrc[i] in units of seconds. Otherwise, the value of the singleton is Undefined.

$$\text{RnMD} = \frac{1}{J[n]} * \sum_{i=1}^{J[n]} \text{TstampRecv}[i] - \text{TstampSrc}[i]$$

Note: RnMD value is Undefined when $J[n] = 0$ for all n .

Figure 6: Type-P-One-to-group-Receiver-n-Mean-Delay

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

8.3.2. Type-P-One-to-group-Mean-Delay

This section defines Type-P-One-to-group-Mean-Delay, the Mean one-way Delay calculated over the entire Group, also named GMD.

$$\text{GMD} = \frac{1}{N} * \sum_{n=1}^N \text{RnMD}$$

Figure 7: 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.

8.3.3. Type-P-One-to-group-Range-Mean-Delay

This section defines a metric for the Range of Mean Delays over all N receivers in the Group ($R1MD, R2MD \dots RnMD$).

Type-P-One-to-group-Range-Mean-Delay = GRMD = $\max(\text{RnMD}) - \min(\text{RnMD})$

8.3.4. Type-P-One-to-group-Max-Mean-Delay

This section defines a metric for the Maximum of Mean Delays over all N receivers in the Group ($R1MD, R2MD, \dots RnMD$).

Type-P-One-to-group-Max-Mean-Delay = GMMD = $\max(\text{RnMD})$

8.4. One-to-Group Packet Loss Statistics

This section defines the overall one-way loss statistics for a receiver and for an entire group as illustrated by the matrix below.

Recv	/----- Sample -----\				Stats	Group Stat
1	R1L1	R1L2	R1L3	... R1Lk	R1LR \	> Group Loss Ratio
2	R2L1	R2L2	R2L3	... R2Lk	R2LR	
3	R3L1	R3L2	R3L3	... R3Lk	R3LR	
.	
n	RnL1	RnL2	RnL3	... RnLk	RnLR /	
Receiver-n Loss Ratio						

Figure 8: One-to-Group Loss Ratio

Statistics are computed on the sample of Type-P-One-way-Packet-Loss [RFC2680] of the matrix above.

All loss ratios are expressed in units of packets lost to total packets sent.

8.4.1. Type-P-One-to-group-Receiver-n-Loss-Ratio

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 and the Type-P-One-to-group-Receiver-n-Loss-Ratio, also named RnLR, are equivalent metrics. In terms of the parameters used here, these metrics definitions can be expressed as

$$RnLR = \frac{1}{K} * \sum_{k=1}^K RnLk$$

Figure 9: Type-P-One-to-group-Receiver-n-Loss-Ratio

8.4.2. Type-P-One-to-group-Receiver-n-Comp-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 numerator is the sum of the losses at receiver n.

The Comparative Loss Ratio, also named, RnCLR, is defined as

$$RnCLR = \frac{\sum_{k=1}^K \ln(k)}{K - \min \left(\sum_{k=1}^K \ln(k), N \right)}$$

Note: Ln is a set of one-way loss values at receiver n.
There is one value for each of the K packets sent.

Figure 10: Type-P-One-to-group-Receiver-n-Comp-Loss-Ratio

8.4.3. Type-P-One-to-group-Loss-Ratio

Type-P-One-to-group-Loss-Ratio, the overall Group Loss Ratio, also named GLR, is defined as:

$$GLR = \frac{1}{K \cdot N} \sum_{k,n=1}^{K,N} \ln(k)$$

Figure 11: Type-P-One-to-group-Loss-Ratio

Where the sum includes all of the Loss singletons, $Ln(k)$, over the N receivers and K packets sent, in a ratio with the total packets over all receivers.

8.4.4. Type-P-One-to-group-Range-Loss-Ratio

The One-to-group Loss Ratio Range is defined as:

Type-P-One-to-group-Range-Loss-Ratio = $\max(RnLR) - \min(RnLR)$

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

8.5. One-to-group Delay Variation Statistics

This section defines one-way delay variation (DV) statistics for an entire group as illustrated by the matrix below.

Recv	/----- Sample -----\					Stats
1	R1ddT1	R1ddT2	R1ddT3 ... R1ddTk	R1DV	\	> Group Stat
2	R2ddT1	R2ddT2	R2ddT3 ... R2ddTk	R2DV		
3	R3ddT1	R3ddT2	R3ddT3 ... R3ddTk	R3DV		
.						
.						
n	RnddT1	RnddT2	RnddT3 ... RnddTk	RnDV	/	

Figure 12: One-to-group Delay Variation Matrix (DVMa)

Statistics are computed on the sample of Type-P-One-way-ipdv singletons of the group delay variation matrix above where $RnddTk$ is the Type-P-One-way-ipdv singleton evaluated at Receiver n for the packet k and where $RnDV$ is the point-to-point one-way packet delay variation for Receiver n .

All One-to-group delay variation statistics are expressed in seconds with sufficient resolution to convey three significant digits.

8.5.1. Type-P-One-to-group-Range-Delay-Variation

This section defines a metric for the Range of Delay Variation over all N receivers in the Group.

Maximum DV and minimum DV over all receivers summarize the performance over the Group (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.

Type-P-One-to-group-Range-Delay-Variation = GRDV =

= $\max(R_nDV) - \min(R_nDV)$ for all n receivers

This range is determined from the minimum and maximum values of the point-to-point one-way IP Packet Delay Variation for the set of Destinations in the group and a population of interest, using the Packet Delay Variation expressed as the $1-10^{-3}$ quantile of one-way delay minus the minimum one-way delay. If a more demanding service is considered, one alternative is to use the $1-10^{-5}$ quantile, and in either case, the quantile used should be recorded with the results. Both the minimum and the maximum delay variation are recorded, and both values are given to indicate the location of the range.

9. Measurement Methods: Scalability 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 points of interest where a stream of singletons will be collected. The amount of information requiring storage grows with both the number of metrics and the points of interest, 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 reference point with connectivity to all the points of interest. In this case, the number of points of interest determines both storage capacity and packet transfer capacity of the host acting as the reference point. However, both the storage and transfer capacity can be reduced if the points of interest 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 point of interest.

In recognition of the likely need to minimize the 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).

9.1. 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 is a centralized statistic calculation method and the latter 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 a very large number of the points of interest. It can lead to a scalability issue at the reference point by overloading the network resources.

The distributed calculation method can save much more bandwidth and mitigate issues arising from scalability at the reference point side.

However, it may result in a loss of information. As not all measured singletons are available for building up the group matrix, the real performance over time can be hidden from the result. For example, the loss pattern can be missed by simply accepting the loss ratio. This tradeoff between bandwidth consumption and information acquisition has to be taken into account when designing the measurement approach.

One possible solution could be to transmit the statistic parameters to the reference point first to obtain the general information of the group performance. If detailed 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. However, this is out of the scope of this memo.

9.2. Measurement

To prevent any bias in the result, the configuration of a one-to-many measure must take into consideration that more packets will be routed than sent (copies of a packet sent are expected to arrive at many destination points) and select a test packet rate that will not impact the network performance.

9.3. 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 not co-located and that results are gathered in a point of reference for further usages.

Multimetric samples are represented in a matrix as illustrated below



Figure 13: Impact of Space Aggregation on Multimetrics Stats

Two methods are available to compute statistics on a matrix:

- o Method 1: The statistic metric is computed over time and then over space; or
- o Method 2: The statistic metric is computed over space and then over time.

These two methods differ only by the order of the aggregation. The order does not impact the computation resources required. It does not change the value of the result. However, it impacts severely the minimal volume of data to report:

- o Method 1: Each point of interest periodically computes statistics over time to lower the volume of data to report. They are reported to the reference point for subsequent computations over the spatial dimension. This volume no longer depends on the number of samples. It is only proportional to the computation period.
- o Method 2: The volume of data to report is proportional to the number of samples. Each sample, RiSi, must be reported to the reference point for computing statistic over space and statistic over time. The volume increases with the number of samples. It is proportional to the number of test packets;

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

- o Increasing the rate of the test packets may result in a Denial of Service (DoS) toward the points of reference;
- o The dimensioning of a measurement system is quite impossible to validate because any increase of the rate of the test packets will increase the bandwidth requested to collect the raw results.

The computation period over time period (commonly named the aggregation period) provides the reporting side with a control of various collecting aspects such as bandwidth, computation, and storage capacities. So this document defines metrics based on method 1.

9.3.1. Impact on Spatial Statistics

Two methods are available to compute spatial statistics:

- o Method 1: Spatial segment metrics and statistics are preferably computed over time for each points of interest;
- o Method 2: Vectors metrics are intrinsically instantaneous space metrics, which must be reported using Method 2 whenever instantaneous metrics information is needed.

9.3.2. Impact on One-to-Group Statistics

Two methods are available to compute group statistics:

- o Method 1: Figure 5 and Figure 8 illustrate the method. The one-to-one statistic is computed per interval of time before the computation of the mean over the group of receivers.
- o Method 2: Figure 13 presents the second method. The metric is computed over space and then over time.

10. Manageability Considerations

This section defines the reporting of all the metrics introduced in the document.

Information models of spatial metrics and of one-to-group metrics are similar except that points of interests of spatial vectors MUST be ordered.

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

10.1. Reporting Spatial Metric

The reporting of spatial metrics shares a lot of aspects with RFC 2679 and RFC 2680. 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 SHOULD be reported.

10.1.1. Path

End-to-end metrics can't determine the path of the measure despite the fact that IPPM RFCs recommend it be reported (see 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 routers of the path participating to the instantaneous measure.

10.1.2. Host Order

A spatial vector MUST order the points of interest according to their order in the path. The ordering MAY be based on information from 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.

10.1.3. 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 three 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.

10.1.4. Reporting Spatial One-Way Delay

The reporting includes information to report for one-way delay as section 3.6 of [RFC2679]. The same applies for packet loss and ipdv.

10.2. Reporting One-to-Group Metric

All reporting rules described in [RFC2679] and [RFC2680] apply to the corresponding One-to-group metrics. The following are specific parameters that SHOULD be reported.

10.2.1. Path

As suggested by [RFC2679] and [RFC2680], the path traversed by the packet SHOULD be reported, if possible. For One-to-group metrics, the path tree between the source and the destinations or the set of paths between the source and each destination SHOULD be reported.

The path tree might not be as valuable as individual paths because an incomplete path might be difficult to identify in the path tree. For example, how many points of interest are reached by a packet traveling along an incomplete path?

10.2.2. Group Size

The group size SHOULD be reported as one of the critical management parameters. One-to-group metrics, unlike spatial metrics, don't require the ordering of the points of interests because group members receive the packets in parallel.

10.2.3. Timestamping Bias

It is the same as described in section 10.1.3.

10.2.4. Reporting One-to-group One-way Delay

It is the same as described in section 10.1.4.

10.2.5. Measurement Method

As explained in section 9, the measurement method will have impact on the analysis of the measurement result. Therefore, it SHOULD be reported.

10.3. Metric Identification

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

10.4. Information Model

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

The information model is built with pieces of information introduced and explained in the sections of [RFC2679] , [RFC2680] , [RFC3393], and [RFC3432] that define the IPPM metrics and from any of the sections named "Reporting the metric" , "Methodology", and "Errors and Uncertainties" whenever they exist in these documents.

The following are the elements of information taken from end-to-end metrics definitions referred to 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_series: <H1, H2,..., Hn>, a list of points of interest participating in the instantaneous measure. They are routers in the case of spatial metrics or receivers in the case of one-to-group metrics.
- 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_series`, `<dT1,..., dTn>`, a list of delays.
- o `Losses_series`, `<B1, B2, ..., Bi, ..., Bn>`, a list of Boolean values (spatial) or a set of Boolean values (one-to-group).
- o `Result_status_series`, 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`.

The 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_series`, 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_series`, apply only for delays and ipdv vector, not ordered for one-to-group;

- o `Losses_series`, apply only for packets loss vector, not ordered for one-to-group;
- o `Result_status_series`;
- o `Observation_duration`, the difference between the time of the last singleton and the time of the first singleton.

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

- o Loss threshold;
- o Systematic error, constant delay between wire-time and timestamping;
- o 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.

The following is the information that SHOULD be available for each sample to compute statistics:

- o `Packet_type`;
- o `Packet_length`;
- o `Src_host`, the sender of the packet;
- o `Dst_host`, the receiver of the packet;
- o `Start_time`, the sending time of the first packet;
- o `Delays_series`, apply only for delays and ipdv samples;
- o `Losses_series`, apply only for packets loss samples;
- o `Result_status_series`;
- o `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.

The following is the context information (measure, points of interests) that SHOULD be available to compute statistics:

- o Loss threshold;

- o Systematic error, constant delay between wire-time and timestamping;
- o Calibration error, maximal uncertainty;

The 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 on which the statistic is computed;

11. Security Considerations

Spatial and one-to-group metrics are defined on the top of end-to-end metrics. Security considerations discussed in the 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 metrics defined in this memo.

Someone may spoof the identity of a point of interest identity and intentionally send corrupt results in order to remotely orient the traffic engineering decisions.

A point of interest could intentionally corrupt its results in order to remotely orient the traffic engineering decisions.

11.1. Spatial Metrics

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

Spatial measurement results carry the performance of individual segments of the path and the identity of nodes. An attacker may infer from this information the points of weakness of a network (e.g., congested node) that would require the least amount of additional attacking traffic to exploit. Therefore, monitoring information should be carried in a way that prevents unintended

recipients from inspecting the measurement reports. A straightforward solution is to restrict access to the reports using encrypted sessions or secured networks.

11.2. One-to-Group Metrics

Reporting of measurement results from a huge number of probes may overload reference point resources (network, network interfaces, computation capacities, etc.).

The configuration of a measurement must take into consideration that implicitly more packets will be routed than sent and select a test packet rate accordingly. Collecting statistics from a huge number of probes may overload any combination of the network to which the measurement controller is attached, measurement controller network interfaces, and measurement controller computation capacities.

One-to-group metric measurements 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 denial-of-service attacks.

A point of interest could intentionally degrade its results in order to remotely increase the quality of the network on the branches of the multicast tree to which it is connected.

12. Acknowledgments

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

13. IANA Considerations

Metrics defined in this memo have been registered in the IANA IPPM METRICS REGISTRY as described in the initial version of the registry [RFC4148]:

IANA has registered the following metrics in the IANA-IPPM-METRICS-REGISTRY-MIB:

ietfSpatialOneWayDelayVector OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 5.1."

:= { ianaIppmMetrics 52 }

ietfSpatialPacketLossVector OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 5.2."

:= { ianaIppmMetrics 53 }

ietfSpatialOneWayIpdvVector OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 5.3."

:= { ianaIppmMetrics 54 }

ietfSegmentOneWayDelayStream OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 6.1."

```
:= { ianaIppmMetrics 55 }

ietfSegmentPacketLossStream OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

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

    REFERENCE

        "RFC 5644, section 6.2."

:= { ianaIppmMetrics 56 }

ietfSegmentIpdvPrevStream OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

        "Type-P-Segment-ipdv-prev-Stream"

    REFERENCE

        "RFC 5644, section 6.3."

:= { ianaIppmMetrics 57 }

ietfSegmentIpdvMinStream OBJECT-IDENTITY

    STATUS current

    DESCRIPTION

        "Type-P-Segment-ipdv-min-Stream"

    REFERENCE

        "RFC 5644, section 6.4."

:= { ianaIppmMetrics 58 }

-- One-to-group metrics

ietfOneToGroupDelayVector OBJECT-IDENTITY
```

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 7.1."

:= { ianaIppmMetrics 59 }

ietfOneToGroupPacketLossVector OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 7.2."

:= { ianaIppmMetrics 60 }

ietfOneToGroupIpdvVector OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 7.3."

:= { ianaIppmMetrics 61 }

-- One to group statistics

--

ietfOnetoGroupReceiverNMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-One-to-group-Receiver-n-Mean-Delay"

REFERENCE

"RFC 5644, section 8.3.1."

:= { ianaIppmMetrics 62 }

ietfOneToGroupMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 8.3.2."

:= { ianaIppmMetrics 63 }

ietfOneToGroupRangeMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 8.3.3."

:= { ianaIppmMetrics 64 }

ietfOneToGroupMaxMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

```
"RFC 5644, section 8.3.4."

:= { ianaIppmMetrics 65 }

ietfOneToGroupReceiverNLossRatio OBJECT-IDENTITY
    STATUS current
    DESCRIPTION
        "Type-P-One-to-group-Receiver-n-Loss-Ratio"
    REFERENCE
        "RFC 5644, section 8.4.1."
    := { ianaIppmMetrics 66 }

--

ietfOneToGroupReceiverNCompLossRatio OBJECT-IDENTITY
    STATUS current
    DESCRIPTION
        "Type-P-One-to-group-Receiver-n-Comp-Loss-Ratio"
    REFERENCE
        "RFC 5644, section 8.4.2."
    := { ianaIppmMetrics 67 }

ietfOneToGroupLossRatio OBJECT-IDENTITY
    STATUS current
    DESCRIPTION
        "Type-P-One-to-group-Loss-Ratio"
    REFERENCE
        "RFC 5644, section 8.4.3."
    := { ianaIppmMetrics 68 }
```

--

ietfOneToGroupRangeLossRatio OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"RFC 5644, section 8.4.4."

:= { ianaIppmMetrics 69 }

ietfOneToGroupRangeDelayVariation OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-One-to-group-Range-Delay-Variation"

REFERENCE

"RFC 5644, section 8.5.1."

:= { ianaIppmMetrics 70 }

--

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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

- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.
- [SPATIAL] Morton, A. and E. Stephan, "Spatial Composition of Metrics", Work in Progress, June 2009.

Authors' Addresses

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

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

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

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

