

Network Working Group	E. Stephan	TOC
Internet-Draft	France Telecom	
Intended status: Standards Track	L. Liang	
Expires: October 30, 2009	University of Surrey	
	A. Morton	
	AT&T Labs	
	April 28, 2009	

IP Performance Metrics (IPPM) for spatial and multicast draft-ietf-ippm-multimetrics-11

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79. 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.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 30, 2009.

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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

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).

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 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

Table of Contents

- [1. Introduction and Scope](#)
- [2. Terminology](#)
- [3. Brief Metric Descriptions](#)
- [4. Motivations](#)
- [5. Spatial vector metrics definitions](#)
- [6. Spatial Segment Metrics Definitions](#)
- [7. One-to-group metrics definitions](#)
- [8. One-to-group Sample Statistics](#)
- [9. Measurement Methods: Scalability and Reporting](#)
- [10. Manageability Considerations](#)
- [11. Security Considerations](#)
- [12. Acknowledgments](#)
- [13. IANA Considerations](#)
- [14. References](#)
 - [14.1. Normative References](#)
 - [14.2. Informative References](#)
- [§ Authors' Addresses](#)

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 the 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 [\[RFC2330\] \(Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.\)](#). This memo is organized as follows: Section 2 introduces new terms that extend the original IPPM framework [\[RFC2330\] \(Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.\)](#). Section 3 motivates each metric category and briefly introduces the new metrics. Sections 4 through 7 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 [\[RFC4148\] \(Stephan, E., "IP Performance Metrics \(IPPM\) Metrics Registry," August 2005.\)](#).

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 which 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.

2. Terminology

[TOC](#)

2.1. Naming of the metrics

[TOC](#)

The names of the metrics, including capitalization 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

[TOC](#)

host: section 5 of RFC 2330
loss threshold: section 2.8.2 of RFC 2680
path: section 5 of RFC 2330
path digest: section 5 of RFC 2330
sample: section 11 of RFC 2330
singleton: section 11 of RFC 2330

2.3. Path Digest Hosts

[TOC](#)

The list of the hosts on a path from the source to the destination, also referred to as the host path digest.

2.4. Multiparty metric

[TOC](#)

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

[TOC](#)

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 members of the path digest.

2.6. One-to-group metric

[TOC](#)

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

[TOC](#)

Points of interest are the hosts (as per the RFC 2330 definition, "hosts" include routing nodes) that are measurement collection points, 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 hosts 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).

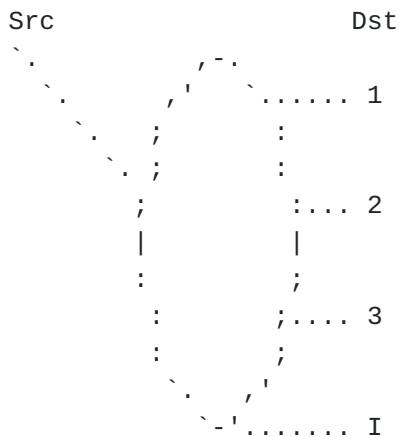
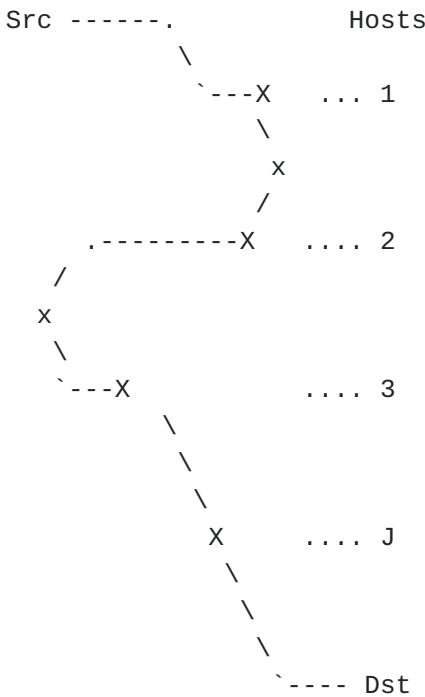


Figure 1: One-to-group points of interest

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



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

Figure 2: Spatial points of interest

2.8. Reference point

[TOC](#)

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

[TOC](#)

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 $dT1, dT2, \dots, dTN$, then a vector V with N elements can be organized as $\{dT1, dT2, \dots, dTN\}$. The element $dT1$ 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

[TOC](#)

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 $V1=\{dT11, dT12, \dots, dT1N\}$, $V2=\{dT21, dT22, \dots, dT2N\}, \dots, Vm=\{dTm1, dTm2, \dots, dTmN\}$ 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 the following [Figure 3 \(Relationship between singletons, samples, vectors and matrix\)](#).

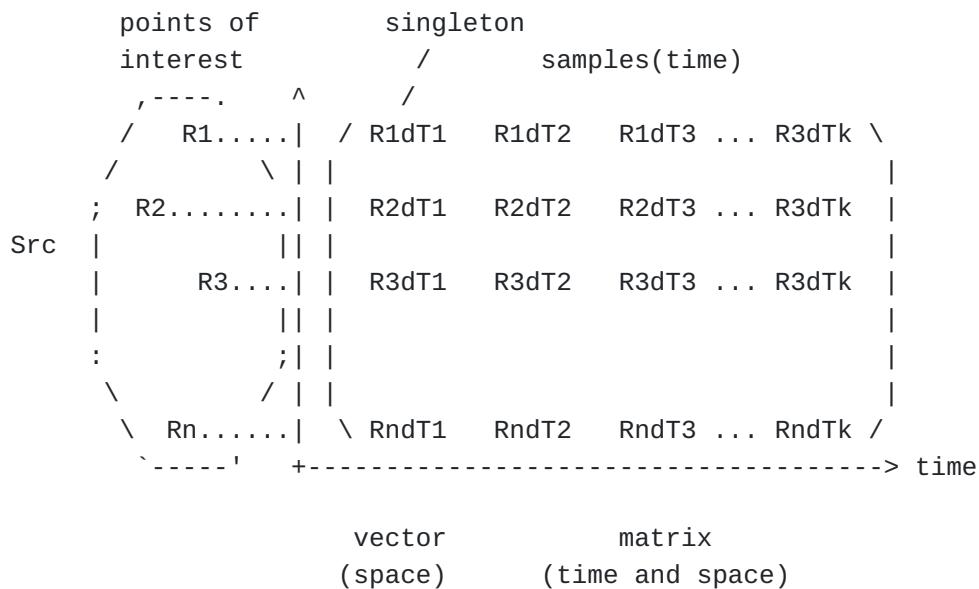


Figure 3: Relationship between singletons, samples, vectors and matrix

3. Brief Metric Descriptions

[TOC](#)

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](#)] ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.](#)), [[RFC2680](#)] ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.](#)), [[RFC3393](#)] ([Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.](#)), [[RFC3432](#)] ([Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams," November 2002.](#)).

This memo defines seven new spatial metrics using the [[RFC2330](#)] ([Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.](#)) framework of parameters, units of measure, and measurement methodologies. Each definition includes a section that describes measurements constraints and issues, and provides guidance to increase the accuracy of the results.

The spatial metrics are:

*Type-P-Spatial-One-way-Delay-Vector divides the end-to-end Type-P-One-way-Delay [[RFC2679](#)] ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.](#)) into a spatial vector of one-way delay singletons.

*Type-P-Spatial-One-way-Packet-Loss-Vector divides an end-to-end Type-P-One-way-Packet-Loss [[RFC2680](#)] ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.](#)) into a spatial vector of packet loss singletons.

*Type-P-Spatial-One-way-ipdv-Vector divides an end-to-end Type-P-One-way-ipdv into a spatial vector of ipdv singletons.

*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.

*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.

*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.

*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:

*Type-P-One-to-group-Delay-Vector collects the set of Type-P-one-way-delay singletons between one sender and N receivers.

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

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

*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, 3 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;

-Type-P-One-to-group-Max-Mean-Delay or GMMD, presents the maximum of mean delays.

*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, 2 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;
- 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.

*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.

*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

[TOC](#)

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

[TOC](#)

Spatial metrics are needed for:

*Decomposing the performance of an inter-domain path to quantify the per-AS contribution to the end-to-end performance.

*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.

*Monitoring the decomposed performance of a multicast tree based on of MPLS point-to-multipoint communications.

*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 [[I-D.ietf-ipmm-spatial-composition](#)] (Morton, A. and E. Stephan, "Spatial Composition of Metrics," April 2010.).

4.2. Motivations for One-to-group metrics

[TOC](#)

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:

- *Designing and engineering multicast trees and MPLS point-to-multipoint LSPs.
- *Evaluating and controlling the quality of multicast services, including inter-domain multicast.
- *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 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

[TOC](#)

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 $\langle h_a, h_b, h_c, \dots, h_n \rangle$ might act as sources to send data to another group of hosts $\langle H_a, H_b, H_c, \dots, H_m \rangle$, and they can be organized into n sets of points of interest for one-to-group communications:

$\langle h_a, H_a, H_b, H_c, \dots, H_m \rangle, \langle h_b, H_a, H_b, H_c, \dots, H_m \rangle, \langle h_c, H_a, H_b, H_c, \dots, H_m \rangle, \dots, \langle h_n, H_a, H_b, H_c, \dots, H_m \rangle$.

5. Spatial vector metrics definitions

[TOC](#)

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 [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM, September 1999.\)](#), [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM, September 1999.\)](#), [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\), November 2002.\)](#) and [\[RFC3432\] \(Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams, November 2002.\)](#).

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

[TOC](#)

This section is coupled with the definition of Type-P-One-way-Delay of the section 3 of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas,](#)

["A One-way Delay Metric for IPPM," September 1999.](#)). When a parameter from the definition in [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) is reused in this section, the first instance will be tagged with a trailing asterisk.

Sections 3.5 to 3.8 of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) 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 measurement MUST respect them, especially those related to methodology, clock, uncertainties and reporting.

5.1.1. Metric Name

[TOC](#)

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

5.1.2. Metric Parameters

[TOC](#)

Src, the IP address of the sender.

Dst, the IP address of the receiver.

* i , an integer in the ordered list $\langle 1, 2, \dots, n \rangle$ of hosts in the path.

* H_i , a host in the path digest.

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

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

* dTi , a delay, the one-way delay for a measured packet from the source to host H_i .

* $\langle dT_1, \dots, dT_i, \dots, dT_n \rangle$ a list of n delay singletions.

Type-P, the specification of the packet type.

* $\langle H_1, H_2, \dots, H_n \rangle$, a path host digest.

5.1.3. Metric Units

[TOC](#)

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

[TOC](#)

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 $\langle Src, H_1, H_2, \dots, H_n, Dst \rangle$ as the sequence of values $\langle T, dT_1, dT_2, \dots, dT_n, dT \rangle$.

5.1.5. Discussion

[TOC](#)

Some specific issues that may occur are as follows:

*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 the section 3.7.1. "Errors or uncertainties related to Clocks" of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#). Consequently, the values of delays measured at multiple hosts may not match the order of those hosts on the path.

-The actual order of hosts 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

[TOC](#)

This section is coupled with the definition of Type-P-One-way-Packet-Loss. When a parameter from the section 2 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) is used in this section, the first instance will be tagged with a trailing asterisk.

Sections 2.5 to 2.8 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) 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

[TOC](#)

Type-P-Spatial-Packet-Loss-Vector

5.2.2. Metric Parameters

[TOC](#)

* Src^* , the IP address of the sender.

* Dst^* , the IP address of the receiver.

* i , an integer in the ordered list $<1, 2, \dots, n>$ of hosts in the path.

* Hi , points of interest from the path digest.

* T^* , a time, the sending time for a measured packet.

* dTi , a delay, the one-way delay for a measured packet from the source to host Hi .

* $<dT1, \dots, dTn>$, list of n delay singletions.

* $Type-P^*$, the specification of packet type.

* $\langle H_1, H_2, \dots, H_n \rangle$, a host path digest.

* $\langle L_1, L_2, \dots, L_n \rangle$, a list of Boolean values.

5.2.3. Metric Units

[TOC](#)

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

5.2.4. Definition

[TOC](#)

Given a Type-P packet sent by the Src at time T to the receiver Dst on the path $\langle H_1, H_2, \dots, H_n \rangle$. For the sequence of times $\langle T+dT_1, T+dT_2, \dots, T+dT_i, \dots, T+dT_n \rangle$ the packet passes in $\langle H_1, H_2, \dots, H_i, \dots, H_n \rangle$, define the Type-P-Packet-Loss-Vector metric as the sequence of values $\langle L_1, L_2, \dots, L_n \rangle$ such that for each H_i of the path, a value of 0 for L_i means that dT_i is a finite value, and a value of 1 means that dT_i is undefined.

5.2.5. Discussion

[TOC](#)

Some specific issues that may occur are as follows:

*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 hosts on the path may change due to reconvergence (e.g., recovery from a link failure).
 - The order of hosts on the path may change due to reconvergence.
 - A packet may not be observed in a host due to some buffer or CPU overflow at the measurement collection point.
-

[TOC](#)

5.3. A Definition for Spatial One-way Ipdv Vector

When a parameter from section 2 of [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.\)](#) (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

[TOC](#)

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

5.3.2. Metric Parameters

[TOC](#)

Src, the IP address of the sender.

Dst, the IP address of the receiver.

*i, an integer in the ordered list <1,2,...,n> of hosts in the path.

*Hi, a host of the path digest.

T1, a time, the sending time for a first measured packet.

T2, a time, the sending time for a second measured packet.

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

*dT_i, a delay, the one-way delay for a measured packet from the source to host Hi.

Type-P, the specification of the packets type.

*P1, the first packet sent at time T1.

*P2, the second packet sent at time T2.

*<H1, H2, ..., Hn>, a 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.

$* <T2, dT2.1, dT2.2, \dots, dT2.n, dT2>$, 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.

5.3.3. Metric Units

[TOC](#)

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

[TOC](#)

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

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

Type-P-Spatial-One-way-ipdv-Vector metric is defined as the sequence of values $<T1, T2, dT2.1-dT1.1, dT2.2-dT1.2, \dots, dT2.n-dT1.n, dT2-dT1>$ such that for each H_i of the path $<H1, H2, \dots, Hn>$, $dT2.i-dT1.i$ 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

[TOC](#)

The methodology, reporting specifications, and uncertainties specified in section 3 of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) 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 [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#)

apply to each point of interest, H_i , measuring an element of a spatial packet loss vector.

Sections 3.5 to 3.7 of [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.\)](#) 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 H_i , the methodology for spatial vector metrics may proceed as follows:

- *At each H_i , points of interest/measurement collection points prepare to capture the packet sent at time T , record a timestamp T_i' , and determine the internal delay correction dT_i' (See section 3.7.1. "Errors or uncertainties related to Clocks" of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#));
- *Each H_i extracts the path ordering information from the packet (e.g. time-to-live);
- *Each H_i computes the corrected wiretime from Src to H_i : $T_i = T_i' - dT_i'$. This arrival time is undefined if the packet is not detected after the 'loss threshold' duration;
- *Each H_i extracts the timestamp T from the packet;
- *Each H_i computes the one-way-delay from Src to H_i : $dT_i = T_i - T$;
- *The reference point gathers the result of each H_i 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 $\langle \text{Src}, H_1, H_2, \dots, H_n, \text{Dst} \rangle$ at time T .

5.4.1. Packet Loss Detection

[TOC](#)

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 section 2.4 and 2.5 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) and section 3.5 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#)). A packet is deemed lost by the receiver after a duration which 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 one, or 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. Host Path Digest

[TOC](#)

The methodology given above relies on knowing the order of the hosts/ measurement collection points on the path [[RFC2330](#)] ([Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.](#)).

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

For example, repeated observations may occur during rerouting phases which introduce temporary micro loops. During such an event the host path 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 hosts in the path digest of a vector MUST be identified before computation of statistics to avoid producing corrupted information.

[TOC](#)

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.

Firstly 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 2 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

[TOC](#)

This metric defines a sample of One-way delays over time between a pair of hosts on a path. Since it is very close semantically to the metric Type-P-One-way-Delay-Poisson-Stream defined in section 4 of [\[RFC2679\]](#) ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.](#)), sections 4.5 to 4.8 of [\[RFC2679\]](#) ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.](#)) are integral parts of the definition text below.

6.1.1. Metric Name

[TOC](#)

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

6.1.2. Metric Parameters

[TOC](#)

*Src, the IP address of the sender.

*Dst, the IP address of the receiver.

*Type-P, the specification of the packet type.

*i, an integer in the ordered list <1,2,...,n> of hosts in the path.

*k, an integer which orders the packets sent.

*a and b, two integers where b > a.

* H_i , a host of the path digest.

* $\langle H_1, \dots, H_a, \dots, H_b, \dots, H_n \rangle$, a host path digest.

* $\langle T_1, T_2, \dots, T_m \rangle$, a list of times.

6.1.3. Metric Units

[TOC](#)

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

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

A sequence of delays.

6.1.4. Definition

[TOC](#)

Given 2 hosts, 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 $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$:

$\langle T_1, dT_{1.1}, dT_{1.2}, \dots, dT_{1.a}, \dots, dT_{1.b}, \dots, dT_{1.n}, dT_1 \rangle$;

$\langle T_2, dT_{2.1}, dT_{2.2}, \dots, dT_{2.a}, \dots, dT_{2.b}, \dots, dT_{2.n}, dT_2 \rangle$;

...

$\langle T_m, dT_{m.1}, dT_{m.2}, \dots, dT_{m.a}, \dots, dT_{m.b}, \dots, dT_{m.n}, dT_m \rangle$.

We define the sample Type-P-segment-One-way-Delay-Stream as the sequence $\langle dT_{1.ab}, dT_{2.ab}, \dots, dT_{k.ab}, \dots, dT_{m.ab} \rangle$ such that for each time T_k , ' $dT_{k.ab}$ ' is either the real number ' $dT_{k.b} - dT_{k.a}$ ' if the packet sent at time T_k passes H_a and H_b or undefined if this packet never passes H_a or (inclusive) never passes H_b .

[TOC](#)

6.1.5. Discussion

Some specific issues that may occur are as follows:

*the delay singletons "appear" to decrease: $dTi > DTi+1$, 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 2.5ms, but the measurement collection point has a clock resolution of 8ms.

The metric SHALL be invalid for times $< T1 , T2, \dots, Tm-1, Tm>$ if the following conditions occur:

*Ha or Hb disappears from the path due to some routing change.

*The order of Ha and Hb changes in the path.

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

[TOC](#)

This metric defines a sample of packet loss over time between a pair of hosts of a path. Since it is very close semantically to the metric Type-P-Packet-loss-Stream defined in section 3 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#), sections 3.5 to 3.8 of [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) are integral parts of the definition text below.

6.2.1. Metric Name

[TOC](#)

Type-P-segment-Packet-Loss-Stream

6.2.2. Metric Parameters

[TOC](#)

*Src, the IP address of the sender.

*Dst, the IP address of the receiver.

*Type-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.

*a and b, two integers where b > a.

* $\langle H_1, H_2, \dots, H_a, \dots, H_b, \dots, H_n \rangle$, a host path digest.

* H_i , exchange points of the path digest.

* $\langle T_1, T_2, \dots, T_m \rangle$, a list of times.

* $\langle L_1, L_2, \dots, L_n \rangle$, a list of Boolean values.

6.2.3. Metric Units

[TOC](#)

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

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

A sequence of Boolean values.

6.2.4. Definition

[TOC](#)

Given two hosts, 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 $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$:

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

We define the value of the sample Type-P-segment-Packet-Lost-Stream from Ha to Hb 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 :

*A value of L_k of 0 means that Ha and Hb observed the packet sent at time T_k (both $L_k.a$ and $L_k.b$ have a value of 0).

*A value of L_k of 1 means that Ha observed the packet sent at time T_k ($L_k.a$ has a value of 0) and that Hb did not observe the packet sent at time T_k ($L_k.b$ has a value of 1).

*The value of L_k is undefined when neither Ha nor Hb observed the packet (both $L_k.a$ and $L_k.b$ have a value of 1).

6.2.5. Discussion

[TOC](#)

Unlike Type-P-Packet-loss-Stream, Type-P-Segment-Packet-Loss-Stream relies on the stability of the host path digest. The metric SHALL be invalid for times $\langle T_1, T_2, \dots, T_{m-1}, T_m \rangle$ if the following conditions occur:

*Ha or Hb disappears from the path due to some routing change.

*The order of Ha and Hb changes in the path.

* $L_k.a$ or $L_k.b$ is undefined.

* $L_k.a$ has the value 1 (not observed) and $L_k.b$ has the value 0 (observed);

*L has the value 0 (the packet was received by Dst) and L_{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

[TOC](#)

This metric defines a sample of ipdv [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.\)](#) over time between a pair of hosts using the previous packet as the selection function.

6.3.1. Metric Name

[TOC](#)

Type-P-Segment-ipdv-prev-Stream

6.3.2. Metric Parameters

[TOC](#)

*Src, the IP address of the sender.

*Dst, the IP address of the receiver.

*Type-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.

*a and b, two integers where b > a.

*<H1, H2, ..., Ha, ..., Hb, ..., Hn>, the hosts path digest.

*<T1, T2, ..., Tm-1, Tm>, a list of times.

*<Tk, dTk.1, dTk.2, ..., dTk.a, ..., dTk.b, ..., dTk.n, dTk>, a Type-P-Spatial-One-way-Delay-Vector.

6.3.3. Metric Units

[TOC](#)

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

The list of <T1, T2, ..., Tm-1, Tm>;

A list of pairs of interval of times and delays;

6.3.4. Definition

[TOC](#)

Given two hosts, Ha and Hb, of the path <H1, H2, ..., Ha, ..., Hb, ..., Hn>, and the matrix of Type-P-Spatial-One-way-Delay-Vector for the packets sent from Src to Dst at times <T1, T2, ..., Tm-1, Tm> :

<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-ipdv-prev-Stream as the sequence of packet time pairs and delay variations

```
<(T1, T2 , dT2.ab - dT1.ab) ,...,  
(Tk-1, Tk, dTk.ab - dTk-1.ab), ...,  
(Tm-1, Tm, dTm.ab - dTm-1.ab)>
```

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

*the delay dTk.a or the delay dTk-1.a is undefined, OR

*the delay dTk.b or the delay dTk-1.b is undefined.

6.3.5. Discussion

[TOC](#)

This metric belongs to the family of inter packet delay variation metrics (IPDV in upper case) 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 of 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

[TOC](#)

This metric defines a sample of ipdv [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.\)](#) over time between a pair of hosts on a path using the minimum delay as one of the selected packets in every pair.

[TOC](#)

6.4.1. Metric Name

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

6.4.2. Metric Parameters

[TOC](#)

*Src, the IP address of the sender.

*Dst, the IP address of the receiver.

*Type-P, the specification of the packet type.

*k, an integer which orders the packets sent.

*i, an integer which identifies a packet sent.

*n, an integer which orders the hosts on the path.

*a and b, two integers where b > a.

*<H1, H2, ..., Ha, ..., Hb, ..., Hn>, the host path digest.

*<T1, T2, ..., Tm-1, Tm>, a list of times.

*<Tk, dTk.1, dTk.2, ..., dTk.a, ..., dTk.b, ..., dTk.n, dTk>, a Type-P-Spatial-One-way-Delay-Vector.

6.4.3. Metric Units

[TOC](#)

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

The list of <T1, T2, ..., Tm-1, Tm>;

A list of times.

6.4.4. Definition

[TOC](#)

Given two hosts, Ha and Hb, of the path <H1, H2, ..., Ha, ..., Hb, ..., Hn>, and the matrix of Type-P-Spatial-One-way-Delay-Vector for the packets sent from Src to Dst at times <T1, T2, ..., Tm-1, Tm> :

```

<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:

* $\min(dTi.ab)$ is the minimum value of the tuples $(dTk.b - dTk.a)$;

*for each time Tk , $dTk.ab$ is undefined if $dTk.a$ or (inclusive) $dTk.b$ is undefined, or the real number $(dTk.b - dTk.a)$ is undefined.

6.4.5. Discussion

[TOC](#)

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 Ha, the PDV results will likely compare favorably to a PDV measurement where the source is Ha and the destination is Hb, because a wide range of spacings are reflected in any PDV distribution.

7. One-to-group metrics definitions

[TOC](#)

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

7.1. A Definition for One-to-group Delay

[TOC](#)

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

7.1.1. Metric Name

[TOC](#)

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

7.1.2. Metric Parameters

[TOC](#)

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

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

*T, a time.

*dT1, ..., dTn a list of times.

*Type-P, the specification of the packet type.

*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

[TOC](#)

The value of a Type-P-One-to-group-Delay-Vector is a set of Type-P-One-way-Delay singletons [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#), which is a sequence of times (a real number in the dimension of seconds with sufficient resolution to convey the results).

7.1.4. Definition

[TOC](#)

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

[TOC](#)

7.2.1. Metric Name

[TOC](#)

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

7.2.2. Metric Parameters

[TOC](#)

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

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

*T, a time.

*Type-P, the specification of the packet type.

*Gr, the receiving group identifier, OPTIONAL.

7.2.3. Metric Units

[TOC](#)

The value of a Type-P-One-to-group-Packet-Loss-Vector is a set of Type-P-One-way-Packet-Loss singletons [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#).

*T, time the source packet was sent

*L1, ..., LN a list of boolean values

7.2.4. Definition

[TOC](#)

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

[TOC](#)

7.3.1. Metric Name

[TOC](#)

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

7.3.2. Metric Parameters

[TOC](#)

*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 times.

*Type-P, the specification of the packet type.

*F, a selection function non-ambiguously defining the two packets from the stream selected for the metric.

*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

[TOC](#)

The value of a Type-P-One-to-group-ipdv-Vector is a set of Type-P-One-way-ipdv singletons [\[RFC3393\]](#) ([Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.](#)).

7.3.4. Definition

[TOC](#)

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 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-Delay-Vector metric.

For a set of real numbers {ddT1,...,ddTn}, the Type-P-One-to-group-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}.

For any pair of selected packets, the difference dT'n-dTn is undefined if:

*the delay dTn to Receiver n is undefined, OR

*the delay dT'n to Receiver n is undefined.

8. One-to-group Sample Statistics

[TOC](#)

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 which guides engineers toward potential problems which 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 which 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 situation 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". 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 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 \(Matrix M \(n*m\)\)](#). 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 4: 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 provider 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 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 {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 requires twice as much calculation to have this statistic over both time and space dimensions. This kind of statistics is 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. 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 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 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.

8.1. Discussion on the Impact of packet loss on statistics

[TOC](#)

The 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 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. 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 { 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. 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 $R1DM$ 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.

8.2. General Metric Parameters

[TOC](#)

*Src, the IP address of a host;

*G, the receiving group identifier;

*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 $1 \leq n \leq N$;

*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);

* T_0 , a time that MUST be selected at random from the interval $[T, T+I]$ to start generating packets and taking measurements (for Periodic Streams);

* $T_{stampSrc}$, the wire time of the packet as measured at $MP(Src)$ (the Source Measurement Point);

* $T_{stampRecv}$, the wire time of the packet as measured at $MP(Recv)$, assigned to packets that arrive within a "reasonable" time;

* T_{max} , 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;

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

* 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 $T_{stampSrc} + T_{max}$, may be indexed over n Receivers;

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

8.3. One-to-group Delay Statistics

[TOC](#)

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

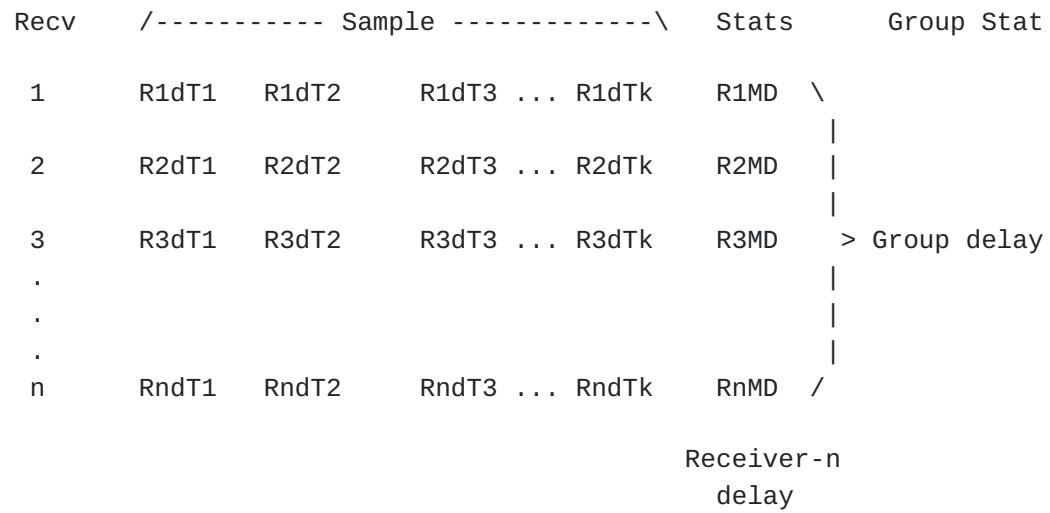


Figure 5: One-to-group Mean Delay

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 3 significant digits.

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

[TOC](#)

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

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 [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#). 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.

```

J[n]
---
1   \
RnMD = --- * > TstampRecv[i] - TstampSrc[i]
J[n]  /
---
i = 1

```

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

[TOC](#)

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

```

N
---
1   \
GMD = - * > RnDM
N   /
---
n = 1

```

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

[TOC](#)

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 = $\max(RnDM) - \min(RnDM)$

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

[TOC](#)

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

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

8.4. One-to-group Packet Loss Statistics

[TOC](#)

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	\
2	R2L1	R2L2	R2L3	...	R2Lk	R2LR	
3	R3L1	R3L2	R3L3	...	R3Lk	R3LR	> Group Loss Ratio
.							
.							
.							
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 [\[RFC2680\]](#) ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.](#)) 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

[TOC](#)

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 [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#). 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

$$RnLR = \frac{\sum_{k=1}^K \frac{1 - RnL_k}{RnL_k}}{K}$$

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

8.4.2. Type-P-One-to-group-Receiver-n-Comp-Loss-Ratio

[TOC](#)

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, also named, RnCLR, is defined as

$$RnCLR = \frac{K}{\sum_{k=1}^N \ln(k)}$$

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

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

[TOC](#)

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

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

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

[TOC](#)

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

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

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

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.

8.5. One-to-group Delay Variation Statistics

[TOC](#)

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 \
2	R2ddT1	R2ddT2	R2ddT3	...	R2ddTk	R2DV
3	R3ddT1	R3ddT2	R3ddT3	...	R3ddTk	R3DV > Group Stat
.						
.						
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-Delay-Variation singletons of the group delay variation matrix above where RnddTk is the Type-P-One-way-Delay-Variation 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 3 significant digits.

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

[TOC](#)

This section defines a metric for the range of delays 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 $1 \cdot 10^{-3}$ quantile of one-way delay minus the minimum one-way delay.

Type-P-One-to-group-Range-Delay-Variation = GRDV =
= $\max(RnDV) - \min(RnDV)$ 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 $1 \cdot 10^{-3}$ quantile of one-way delay minus the minimum one-way delay. If a more demanding service is considered, one alternative is to use the $1 \cdot 10^{-5}$ 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

[TOC](#)

Virtually all the guidance on measurement processes supplied by the earlier IPPM RFCs (such as [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) and [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#)) 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

[TOC](#)

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 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 all measured singletons are not 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 transit 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

[TOC](#)

To prevent any bias in the result, the configuration of a one-to-many measure must take in consideration that intrinsically more packets will be routed than sent (copies of a packet sent are expected to arrive at

many destination points) and selects a test packets rate that will not impact the network performance.

9.3. Effect of Time and Space Aggregation Order on Stats

[TOC](#)

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.

Multimetrics samples are represented in a matrix as illustrated below

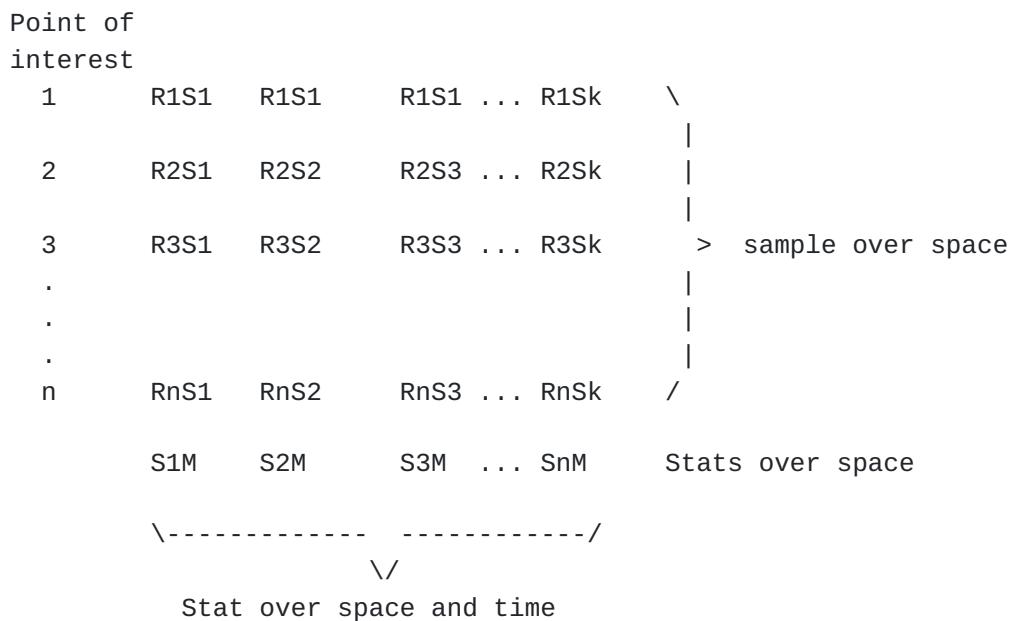


Figure 13: Impact of space aggregation on multimetrics Stat

Two methods are available to compute statistics on a matrix:

*Method 1: The statistic metric is computed over time and then over space;

*Method 2: The statistic metric is computed over space and then over time.

These 2 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:

*Method 1: Each point of interest computes periodically 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;

*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:

*Increasing the rate of the test packets may result in a Denial of Service toward the points of reference;

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

9.3.1. Impact on spatial statistics

[TOC](#)

Two methods are available to compute spatial statistics:

*Method 1: spatial segment metrics and statistics are preferably computed over time for each points of interest;

*Method 2: Vectors metrics are intrinsically instantaneous space metrics which must be reported using Method2 whenever instantaneous metrics information is needed.

[TOC](#)

9.3.2. Impact on one-to-group statistics

Two methods are available to compute group statistics:

*Method1: [Figure 5 \(One-to-group Mean Delay\)](#) and [Figure 8 \(One-to-group Loss Ratio\)](#) 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;

*Method2: [Figure 13 \(Impact of space aggregation on multimetrics Stat\)](#) presents the second one, metric is computed over space and then over time.

10. Manageability Considerations

[TOC](#)

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, to avoid repetitions and to conform to IESG recommendation of gathering manageability considerations in a dedicated section.

Information models of spatial metrics and of 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.

10.1. Reporting spatial metric

[TOC](#)

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.

10.1.1. Path

[TOC](#)

End-to-end metrics can't determine the path of the measure despite IPPM RFCs recommend it to be reported (See Section 3.8.4 of [\[RFC2679\]](#) ([Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.](#))). 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.

10.1.2. Host order

[TOC](#)

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.

10.1.3. Timestamping bias

[TOC](#)

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.

10.1.4. Reporting spatial One-way Delay

[TOC](#)

The reporting includes information to report for one-way-delay as the Section 3.6 of [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#). The same apply for packet loss and ipdv.

10.2. Reporting One-to-group metric

[TOC](#)

All reporting rules described in [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) and [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) apply to the corresponding One-to-group metrics. Following are specific parameters that should be reported.

10.2.1. Path

[TOC](#)

As suggested by the [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#) and [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#), 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. 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 travelling along an incomplete path?

10.2.2. Group size

[TOC](#)

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

[TOC](#)

It is the same as described in section 10.1.3.

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

[TOC](#)

It is the same as described in section 10.1.4.

10.2.5. Measurement method

[TOC](#)

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

[TOC](#)

10.3. Metric identification

IANA assigns each metric defined by the IPPM WG with a unique identifier as per [\[RFC4148\] \(Stephan, E., "IP Performance Metrics \(IPPM\) Metrics Registry," August 2005.\)](#) in the IANA-IPPM-METRICS-REGISTRY-MIB.

10.4. Information model

[TOC](#)

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 one-way delay definitions [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#), in packet loss definitions [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#) and in IPDV definitions of [\[RFC3393\] \(Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\)," November 2002.\)](#) and [\[RFC3432\] \(Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams," November 2002.\)](#). It includes not only information given by "Reporting the metric" sections but by sections "Methodology" and "Errors and Uncertainties".

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

- *Packet_type, The Type-P of test packets (Type-P);
- *Packet_length, a packet length in bits (L);
- *Src_host, the IP address of the sender;
- *Dst_host, the IP address of the receiver;
- *Hosts_serie: <H1, H2, ..., Hn>, a list of points of interest;
- *Loss_threshold: The threshold of infinite delay;
- *Systematic_error: constant delay between wire time and timestamping;
- *Calibration_error: maximal uncertainty;
- *Src_time, the sending time for a measured packet;

```

*Dst_time, the receiving time for a measured packet;

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

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

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

*Result_status_serie: a list of results status;

*dT: a delay;

*Singleton_number: a number of singletons;

*Observation_duration: An observation duration;

*metric_identifier.
```

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

```

*Packet_type;

*Packet_length;

*Src_host, the sender of the packet;

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

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

*Src_time, the sending time for the measured packet;

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

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

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

*Result_status_serie;

*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 :

- 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:

- *Result;
- *Start_time;

*Duration;
*Result_status;
*Singleton_number, the number of singletons the statistic is computed on;

11. Security Considerations

[TOC](#)

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 [\[RFC2679\]](#) (Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.) , in packet loss metrics definitions of [\[RFC2680\]](#) (Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.) and in IPDV metrics definitions of [\[RFC3393\]](#) (Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," November 2002.) and [\[RFC3432\]](#) (Raisanen, V., Grotfeld, G., and A. Morton, "Network performance measurement with periodic streams," November 2002.) apply to metrics defined in this memo.

11.1. Spatial metrics

[TOC](#)

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 metrics

[TOC](#)

Reporting of measurement results from a huge number of probes may overload reference point resources (network, network interfaces, computation capacities).
The configuration of a measurement must take in consideration that implicitly more packets will be routed than sent 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 attached 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

[TOC](#)

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

[TOC](#)

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 [\[RFC4148\] \(Stephan, E., "IP Performance Metrics \(IPPM\) Metrics Registry," August 2005.\)](#) :

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

`ietfSpatialOneWayDelayVector OBJECT-IDENTITY`

`STATUS current`

`DESCRIPTION`

`"Type-P-Spatial-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`

`ietfSpatialPacketLossVector OBJECT-IDENTITY`

`STATUS current`

`DESCRIPTION`

`"Type-P-Spatial-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  
  
ietfSpatialOneWayIpdvVector OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-Spatial-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  
  
ietfSegmentOneWayDelayStream OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-Segment-One-way-Delay-Stream"  
  
REFERENCE  
  
"Reference "RFCyyyy, section 6.1."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
ietfSegmentPacketLossStream OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-Segment-Packet-Loss-Stream"  
  
REFERENCE
```

```
"Reference "RFCyyyy, section 6.2."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
ietfSegmentIpdvPrevStream OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-Segment-ipdv-prev-Stream"  
  
REFERENCE  
  
"Reference "RFCyyyy, section 6.3."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
ietfSegmentIpdvMinStream OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-Segment-ipdv-min-Stream"  
  
REFERENCE  
  
"Reference "RFCyyyy, section 6.4."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
-- One-to-group metrics  
ietfOneToGroupDelayVector OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-One-to-group-Delay-Vector"  
  
REFERENCE
```

```
"Reference "RFCyyyy, section 7.1."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
ietfOneToGroupPacketLossVector OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-One-to-group-Packet-Loss-Vector"  
  
REFERENCE  
  
"Reference "RFCyyyy, section 7.2."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
ietfOneToGroupIpdvVector OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-One-to-group-ipdv-Vector"  
  
REFERENCE  
  
"Reference "RFCyyyy, section 7.3."  
-- RFC Ed.: replace yyyy with actual RFC number & remove this  
note  
  
:= { ianaIppmMetrics nn } -- IANA assigns nn  
  
-- One to group statistics  
--  
ietfOneToGroupReceiverNMeanDelay OBJECT-IDENTITY  
  
STATUS current  
  
DESCRIPTION  
  
"Type-P-One-to-group-Receiver-n-Mean-Delay"
```

REFERENCE

"Reference "RFCyyyy, section 8.3.1."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.3.2."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupRangeMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.3.3."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupMaxMeanDelay OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.3.4."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupReceiverNLossRatio OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-One-to-group-Receiver-n-Loss-Ratio"

REFERENCE

"Reference "RFCyyyy, section 8.4.1."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

--

ietfOneToGroupReceiverNCompLossRatio OBJECT-IDENTITY

STATUS current

DESCRIPTION

"Type-P-One-to-group-Receiver-n-Comp-Loss-Ratio"

REFERENCE

"Reference "RFCyyyy, section 8.4.2."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupLossRatio OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.4.3."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

--

ietfOneToGroupRangeLossRatio OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.4.4."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

ietfOneToGroupRangeDelayVariation OBJECT-IDENTITY

STATUS current

DESCRIPTION

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

REFERENCE

"Reference "RFCyyyy, section 8.5.1."

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

:= { ianaIppmMetrics nn } -- IANA assigns nn

--

14. References

14.1. Normative References

[TOC](#)

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[RFC2679]	Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," RFC 2679, September 1999 (TXT).
[RFC2680]	Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," RFC 2680, September 1999 (TXT).
[RFC3393]	Demichelis, C. and P. Chimento, " IP Packet Delay Variation Metric for IP Performance Metrics (IPPM), " RFC 3393, November 2002 (TXT).
[RFC4148]	Stephan, E., " IP Performance Metrics (IPPM) Metrics Registry, " BCP 108, RFC 4148, August 2005 (TXT).

14.2. Informative References

[TOC](#)

[I-D.ietf-ippm-spatial-composition]	Morton, A. and E. Stephan, " Spatial Composition of Metrics, " draft-ietf-ippm-spatial-composition-11 (work in progress), April 2010 (TXT).
[RFC2330]	Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," RFC 2330, May 1998 (TXT , HTML , XML).
[RFC3432]	Raisanen, V., Grotefeld, G., and A. Morton, " Network performance measurement with periodic streams, " RFC 3432, November 2002 (TXT).

Authors' Addresses

[TOC](#)

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