

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

A. Morton
AT&T Labs
M. Bagnulo
UC3M
P. Eardley
BT
K. D'Souza
AT&T Labs
October 22, 2018

Initial Performance Metric Registry Entries
draft-ietf-ippm-initial-registry-08

Abstract

This memo defines the Initial Entries for the Performance Metrics Registry. This version includes:

- * resolved open items on UDP one-way measurements.
- * resolved questions on DNS measurement method(s)

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on April 25, 2019.

Internet-Draft

Initial Registry

October 2018

Copyright Notice

Copyright (c) 2018 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	7
2.	Scope	8
3.	Registry Categories and Columns	8
4.	UDP Round-trip Latency and Loss Registry Entries	9
4.1.	Summary	10
4.1.1.	ID (Identifier)	10
4.1.2.	Name	10
4.1.3.	URIs	10
4.1.4.	Description	10
4.1.5.	Change Controller	10
4.1.6.	Version (of Registry Format)	10
4.2.	Metric Definition	11
4.2.1.	Reference Definition	11
4.2.2.	Fixed Parameters	12
4.3.	Method of Measurement	12
4.3.1.	Reference Method	13
4.3.2.	Packet Stream Generation	14
4.3.3.	Traffic Filtering (observation) Details	14
4.3.4.	Sampling Distribution	15
4.3.5.	Run-time Parameters and Data Format	15
4.3.6.	Roles	15
4.4.	Output	16
4.4.1.	Type	16
4.4.2.	Reference Definition	16
4.4.3.	Metric Units	17
4.4.4.	Calibration	17

4.5.	Administrative items	18
4.5.1.	Status	18
4.5.2.	Requestor (keep?)	18
4.5.3.	Revision	18
4.5.4.	Revision Date	18

4.6.	Comments and Remarks	18
5.	Packet Delay Variation Registry Entry	18
5.1.	Summary	18
5.1.1.	ID (Identifier)	18
5.1.2.	Name	18
5.1.3.	URIs	19
5.1.4.	Description	19
5.1.5.	Change Controller	19
5.1.6.	Version (of Registry Format)	19
5.2.	Metric Definition	19
5.2.1.	Reference Definition	19
5.2.2.	Fixed Parameters	20
5.3.	Method of Measurement	21
5.3.1.	Reference Method	21
5.3.2.	Packet Stream Generation	22
5.3.3.	Traffic Filtering (observation) Details	22
5.3.4.	Sampling Distribution	23
5.3.5.	Run-time Parameters and Data Format	23
5.3.6.	Roles	23
5.4.	Output	23
5.4.1.	Type	24
5.4.2.	Reference Definition	24
5.4.3.	Metric Units	24
5.4.4.	Calibration	25
5.5.	Administrative items	25
5.5.1.	Status	25
5.5.2.	Requestor (keep?)	26
5.5.3.	Revision	26
5.5.4.	Revision Date	26
5.6.	Comments and Remarks	26
6.	DNS Response Latency and Loss Registry Entries	26
6.1.	Summary	26
6.1.1.	ID (Identifier)	27
6.1.2.	Name	27
6.1.3.	URI	27
6.1.4.	Description	27

6.1.5.	Change Controller	27
6.1.6.	Version (of Registry Format)	27
6.2.	Metric Definition	27
6.2.1.	Reference Definition	28
6.2.2.	Fixed Parameters	28
6.3.	Method of Measurement	30
6.3.1.	Reference Method	30
6.3.2.	Packet Stream Generation	32
6.3.3.	Traffic Filtering (observation) Details	32
6.3.4.	Sampling Distribution	32
6.3.5.	Run-time Parameters and Data Format	33
6.3.6.	Roles	34

6.4.	Output	34
6.4.1.	Type	34
6.4.2.	Reference Definition	34
6.4.3.	Metric Units	35
6.4.4.	Calibration	35
6.5.	Administrative items	35
6.5.1.	Status	35
6.5.2.	Requestor	36
6.5.3.	Revision	36
6.5.4.	Revision Date	36
6.6.	Comments and Remarks	36
7.	UDP Poisson One-way Delay and Loss Registry Entries	36
7.1.	Summary	36
7.1.1.	ID (Identifier)	36
7.1.2.	Name	37
7.1.3.	URI and URL	37
7.1.4.	Description	37
7.2.	Metric Definition	38
7.2.1.	Reference Definition	38
7.2.2.	Fixed Parameters	39
7.3.	Method of Measurement	40
7.3.1.	Reference Method	40
7.3.2.	Packet Stream Generation	41
7.3.3.	Traffic Filtering (observation) Details	41
7.3.4.	Sampling Distribution	42
7.3.5.	Run-time Parameters and Data Format	42
7.3.6.	Roles	42
7.4.	Output	42
7.4.1.	Type	43

7.4.2.	Reference Definition	43
7.4.3.	Metric Units	46
7.4.4.	Calibration	46
7.5.	Administrative items	47
7.5.1.	Status	47
7.5.2.	Requestor (keep?)	47
7.5.3.	Revision	47
7.5.4.	Revision Date	47
7.6.	Comments and Remarks	47
8.	UDP Periodic One-way Delay and Loss Registry Entries	47
8.1.	Summary	47
8.1.1.	ID (Identifier)	47
8.1.2.	Name	48
8.1.3.	URIs	48
8.1.4.	Description	48
8.2.	Metric Definition	49
8.2.1.	Reference Definition	49
8.2.2.	Fixed Parameters	50
8.3.	Method of Measurement	51

8.3.1.	Reference Method	51
8.3.2.	Packet Stream Generation	52
8.3.3.	Traffic Filtering (observation) Details	52
8.3.4.	Sampling Distribution	53
8.3.5.	Run-time Parameters and Data Format	53
8.3.6.	Roles	53
8.4.	Output	53
8.4.1.	Type	54
8.4.2.	Reference Definition	54
8.4.3.	Metric Units	57
8.4.4.	Calibration	57
8.5.	Administrative items	58
8.5.1.	Status	58
8.5.2.	Requestor (keep?)	58
8.5.3.	Revision	58
8.5.4.	Revision Date	58
8.6.	Comments and Remarks	58
9.	ICMP Round-trip Latency and Loss Registry Entries	58
9.1.	Summary	59
9.1.1.	ID (Identifier)	59
9.1.2.	Name	59
9.1.3.	URIs	59

9.1.4.	Description	59
9.1.5.	Change Controller	60
9.1.6.	Version (of Registry Format)	60
9.2.	Metric Definition	60
9.2.1.	Reference Definition	60
9.2.2.	Fixed Parameters	61
9.3.	Method of Measurement	62
9.3.1.	Reference Method	62
9.3.2.	Packet Stream Generation	63
9.3.3.	Traffic Filtering (observation) Details	64
9.3.4.	Sampling Distribution	64
9.3.5.	Run-time Parameters and Data Format	64
9.3.6.	Roles	65
9.4.	Output	65
9.4.1.	Type	65
9.4.2.	Reference Definition	65
9.4.3.	Metric Units	67
9.4.4.	Calibration	67
9.5.	Administrative items	68
9.5.1.	Status	68
9.5.2.	Requestor (keep?)	68
9.5.3.	Revision	68
9.5.4.	Revision Date	68
9.6.	Comments and Remarks	68
10.	TCP Round-Trip Delay and Loss Registry Entries	68
10.1.	Summary	68

10.1.1.	ID (Identifier)	69
10.1.2.	Name	69
10.1.3.	URIs	69
10.1.4.	Description	69
10.1.5.	Change Controller	70
10.1.6.	Version (of Registry Format)	70
10.2.	Metric Definition	70
10.2.1.	Reference Definitions	70
10.2.2.	Fixed Parameters	72
10.3.	Method of Measurement	73
10.3.1.	Reference Methods	73
10.3.2.	Packet Stream Generation	75
10.3.3.	Traffic Filtering (observation) Details	75
10.3.4.	Sampling Distribution	75
10.3.5.	Run-time Parameters and Data Format	75

10.3.6.	Roles	76
10.4.	Output	76
10.4.1.	Type	76
10.4.2.	Reference Definition	77
10.4.3.	Metric Units	79
10.4.4.	Calibration	79
10.5.	Administrative items	79
10.5.1.	Status	79
10.5.2.	Requestor (keep?)	79
10.5.3.	Revision	79
10.5.4.	Revision Date	79
10.6.	Comments and Remarks	79
11.	ver08 BLANK Registry Entry	80
11.1.	Summary	80
11.1.1.	ID (Identifier)	80
11.1.2.	Name	80
11.1.3.	URIs	80
11.1.4.	Description	80
11.1.5.	Reference	80
11.1.6.	Change Controller	80
11.1.7.	Version (of Registry Format)	80
11.2.	Metric Definition	80
11.2.1.	Reference Definition	81
11.2.2.	Fixed Parameters	81
11.3.	Method of Measurement	81
11.3.1.	Reference Method	81
11.3.2.	Packet Stream Generation	81
11.3.3.	Traffic Filtering (observation) Details	81
11.3.4.	Sampling Distribution	81
11.3.5.	Run-time Parameters and Data Format	81
11.3.6.	Roles	81
11.4.	Output	82
11.4.1.	Type	82

11.4.2.	Reference Definition	82
11.4.3.	Metric Units	82
11.4.4.	Calibration	82
11.5.	Administrative items	82
11.5.1.	Status	82
11.5.2.	Requestor	82
11.5.3.	Revision	82
11.5.4.	Revision Date	82

11.6.	Comments and Remarks	82
12.	Example RTCP-XR Registry Entry	83
12.1.	Registry Indexes	83
12.1.1.	Identifier	83
12.1.2.	Name	83
12.1.3.	URI	83
12.1.4.	Status	83
12.1.5.	Requestor	83
12.1.6.	Revision	83
12.1.7.	Revision Date	83
12.1.8.	Description	83
12.1.9.	Reference Specification(s)	84
12.2.	Metric Definition	84
12.2.1.	Reference Definition	84
12.2.2.	Fixed Parameters	84
12.3.	Method of Measurement	85
12.3.1.	Reference Method	85
12.3.2.	Stream Type and Stream Parameters	85
12.3.3.	Output Type and Data Format	85
12.3.4.	Metric Units	85
12.3.5.	Run-time Parameters and Data Format	86
12.4.	Comments and Remarks	87
13.	Revision History	87
14.	Security Considerations	88
15.	IANA Considerations	88
16.	Acknowledgements	88
17.	References	89
17.1.	Normative References	89
17.2.	Informative References	91
	Authors' Addresses	93

1. Introduction

Note: Efforts to synchronize structure and terminology with [[I-D.ietf-ippm-metric-registry](#)] will likely be incomplete until both drafts are stable.

This memo proposes an initial set of entries for the Performance Metric Registry. It uses terms and definitions from the IPPM literature, primarily [[RFC2330](#)].

Although there are several standard templates for organizing

specifications of performance metrics (see [[RFC2679](#)] for an example of the traditional IPPM template, based to large extent on the Benchmarking Methodology Working Group's traditional template in [[RFC1242](#)], and see [[RFC6390](#)] for a similar template), none of these templates were intended to become the basis for the columns of an IETF-wide registry of metrics. While examining aspects of metric specifications which need to be registered, it became clear that none of the existing metric templates fully satisfies the particular needs of a registry.

Therefore, [[I-D.ietf-ippm-metric-registry](#)] defines the overall format for a Performance Metric Registry. Section 5 of [[I-D.ietf-ippm-metric-registry](#)] also gives guidelines for those requesting registration of a Metric, that is the creation of entry(s) in the Performance Metric Registry: "In essence, there needs to be evidence that a candidate Registered Performance Metric has significant industry interest, or has seen deployment, and there is agreement that the candidate Registered Performance Metric serves its intended purpose." The process in [[I-D.ietf-ippm-metric-registry](#)] also requires that new entries are administered by IANA through Expert Review, which will ensure that the metrics are tightly defined.

[2.](#) Scope

This document defines the initial set of Performance Metrics Registry entries, for which IETF approval (following development in the IP Performance Metrics (IPPM) Working Group) will satisfy the requirement for Expert Review. Most are Active Performance Metrics, which are based on RFCs prepared in the IPPM working group of the IETF, according to their framework [[RFC2330](#)] and its updates.

[3.](#) Registry Categories and Columns

This section provides the categories and columns of the registry, for easy reference. An entry (row) therefore gives a complete description of a Registered Metric.

Registry Categories and Columns, shown as

```

Category
-----
Column | Column |

```

Summary

```

-----
Identifier | Name | URIs | Desc. | Reference | Change Controller | Ver |

```

Metric Definition

```

-----
Reference Definition | Fixed Parameters |

```

Method of Measurement

```

-----
Reference | Packet      | Traffic | Sampling | Run-time | Role |
Method    | Stream      | Filter  | Distribution | Parameters |      |
          | Generation  |         |              |           |      |

```

Output

```

-----
Type | Reference | Units | Calibration |
     | Definition |      |              |

```

Administrative Information

```

-----
Status | Request | Rev | Rev.Date |

```

Comments and Remarks

4. UDP Round-trip Latency and Loss Registry Entries

This section specifies an initial registry entry for the UDP Round-trip Latency, and another entry for UDP Round-trip Loss Ratio.

Note: Each Registry entry only produces a "raw" output or a statistical summary. To describe both "raw" and one or more statistics efficiently, the Identifier, Name, and Output Categories can be split and a single section can specify two or more closely-related metrics. This section specifies two Registry entries with many common columns. See [Section 7](#) for an example specifying multiple Registry entries with many common columns.

All column entries beside the ID, Name, Description, and Output Reference Method categories are the same, thus this section proposes

two closely-related registry entries. As a result, IANA is also asked to assign corresponding URNs and URLs to each Named Metric.

[4.1.](#) Summary

This category includes multiple indexes to the registry entry: the element ID and metric name.

[4.1.1.](#) ID (Identifier)

<insert a numeric identifier, an integer, TBD>

IANA is asked to assign different numeric identifiers to each of the two Named Metrics.

[4.1.2.](#) Name

<insert name according to metric naming convention>

RTDelay_Active_IP-UDP-Periodic_RFCXXXXsecY_Seconds_95Percentile

RTLoss_Active_IP-UDP-Periodic_RFCXXXXsecY_Percent_LossRatio

[4.1.3.](#) URIs

URN: Prefix urn:ietf:metrics:perf:<name>

URL: http://<TBD by IANA>/<name>

[4.1.4.](#) Description

RTDelay: This metric assesses the delay of a stream of packets exchanged between two hosts (which are the two measurement points), and the Output is the Round-trip delay for all successfully exchanged packets expressed as the 95th percentile of their conditional delay distribution.

RTLoss: This metric assesses the loss ratio of a stream of packets exchanged between two hosts (which are the two measurement points), and the Output is the Round-trip loss ratio for all successfully exchanged packets expressed as a percentage.

[4.1.5.](#) Change Controller

IETF

[4.1.6.](#) Version (of Registry Format)

1.0

Morton, et al.

Expires April 25, 2019

[Page 10]

Internet-Draft

Initial Registry

October 2018

[4.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[4.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

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

[RFC2681]

<specific section reference and additional clarifications, if needed>

[Section 2.4 of \[RFC2681\]](#) provides the reference definition of the singleton (single value) Round-trip delay metric. [Section 3.4 of \[RFC2681\]](#) provides the reference definition expanded to cover a multi-singleton sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

Note that although the [\[RFC2681\]](#) definition of "Round-trip-Delay between Src and Dst" is directionally ambiguous in the text, this metric tightens the definition further to recognize that the host in the "Src" role will send the first packet to "Dst", and ultimately receive the corresponding return packet from "Dst" (when neither are lost).

Finally, note that the variable "dT" is used in [\[RFC2681\]](#) to refer to the value of Round-trip delay in metric definitions and methods. The

variable "dT" has been re-used in other IPPM literature to refer to different quantities, and cannot be used as a global variable name.

Morton, A., "Round-trip Packet Loss Metrics", [RFC 6673](#), August 2012.

[RFC6673]

Both delay and loss metrics employ a maximum waiting time for received packets, so the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 6.1 of \[RFC6673\]](#).

Morton, et al.

Expires April 25, 2019

[Page 11]

Internet-Draft

Initial Registry

October 2018

[4.2.2.](#) Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Type-P as defined in [Section 13 of \[RFC2330\]](#):

- o IPv4 header values:
 - * DSCP: set to 0
 - * TTL: set to 255
 - * Protocol: Set to 17 (UDP)
- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Count: set to 255
 - * Protocol: Set to 17 (UDP)
- o UDP header values:

- * Checksum: the checksum MUST be calculated and included in the header
- o UDP Payload
 - * total of 100 bytes

Other measurement parameters:

- o Tmax: a loss threshold waiting time
 - * 3.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

[4.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[4.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

The methodology for this metric is defined as Type-P-Round-trip-Delay-Poisson-Stream in [section 2.6 of RFC 2681 \[RFC2681\]](#) and [section 3.6 of RFC 2681 \[RFC2681\]](#) using the Type-P and Tmax defined under Fixed Parameters. However, the Periodic stream will be generated according to [\[RFC3432\]](#).

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a packet lost. Lost packets SHALL be designated as having undefined delay, and counted for the RTLoss metric.

The calculations on the delay (RTT) SHALL be performed on the

conditional distribution, conditioned on successful packet arrival within T_{max} . Also, when all packet delays are stored, the process which calculates the RTT value MAY enforce the T_{max} threshold on stored values before calculations. See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving packet. Sequence numbers or other send-order identification MUST be retained at the Src or included with each packet to disambiguate packet reordering if it occurs.

If a standard measurement protocol is employed, then the measurement process will determine the sequence numbers or timestamps applied to test packets after the Fixed and Runtime parameters are passed to that process. The chosen measurement protocol will dictate the format of sequence numbers and time-stamps, if they are conveyed in the packet payload.

Refer to [Section 4.4 of \[RFC6673\]](#) for expanded discussion of the instruction to "send a Type-P packet back to the Src as quickly as possible" in [Section 2.6 of RFC 2681 \[RFC2681\]](#). [Section 8 of \[RFC6673\]](#) presents additional requirements which MUST be included in the method of measurement for this metric.

[4.3.2.](#) Packet Stream Generation

<list of generation parameters and section/spec references if needed>

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

[Section 3 of \[RFC3432\]](#) prescribes the method for generating Periodic streams using associated parameters.

incT the nominal duration of inter-packet interval, first bit to first bit, with value 0.0200, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms).

dT the duration of the interval for allowed sample start times, with value 1.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms).

T0 the actual start time of the periodic stream, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)).

NOTE: an initiation process with a number of control exchanges resulting in unpredictable start times (within a time interval) may be sufficient to avoid synchronization of periodic streams, and therefore a valid replacement for selecting a start time at random from a fixed interval.

The T0 parameter will be reported as a measured parameter. Parameters incT and dT are Fixed Parameters.

[4.3.3. Traffic Filtering \(observation\) Details](#)

The measured results based on a filtered version of the packets observed, and this section provides the filter details (when present).

<section reference>.

NA

[4.3.4. Sampling Distribution](#)

<insert time distribution details, or how this is diff from the filter>

NA

[4.3.5.](#) Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Tf a time, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a end time date is ignored and Tf is interpreted as the Duration of the measurement interval.

[4.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

Src launches each packet and waits for return transmissions from Dst.

Dst waits for each packet from Src and sends a return packet to Src.

[4.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

[4.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

Percentile -- for the conditional distribution of all packets with a valid value of Round-trip delay (undefined delays are excluded), a single value corresponding to the 95th percentile, as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The percentile = 95, meaning that the reported delay, "95Percentile", is the smallest value of Round-trip delay for which the Empirical Distribution Function (EDF), $F(95\text{Percentile}) \geq 95\%$ of the singleton Round-trip delay values in the conditional distribution. See [section 11.3 of \[RFC2330\]](#) for the definition of the percentile statistic using the EDF.

LossRatio -- the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 6.1 of \[RFC6673\]](#).

[4.4.2.](#) Reference Definition

<describe the reference data format for each type of result>

For all outputs ---

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

TotalPkts the count of packets sent by the Src to Dst during the measurement interval.

For

Internet-Draft

Initial Registry

October 2018

RTDelay_Active_IP-UDP-Periodic_RFCXXXsecY_Seconds_95Percentile:

95Percentile The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as

For

RTLoss_Active_IP-UDP-Periodic_RFCXXXsecY_Percent_LossRatio:

Percentile The numeric value of the result is expressed in units of lost packets to total packets times 100%, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.0000000001.

[4.4.3.](#) Metric Units

<insert units for the measured results, and the reference specification>.

The 95th Percentile of Round-trip Delay is expressed in seconds.

The Round-trip Loss Ratio is expressed as a percentage of lost packets to total packets sent.

[4.4.4.](#) Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback at the Source host that includes as much of the measurement system as possible, performs address manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a

normal measurement with additional indication that it is a calibration result.

Both internal loopback calibration and clock synchronization can be used to estimate the *available accuracy* of the Output Metric Units. For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

Morton, et al.

Expires April 25, 2019

[Page 17]

Internet-Draft

Initial Registry

October 2018

[4.5.](#) Administrative items

[4.5.1.](#) Status

<current or deprecated>

[4.5.2.](#) Requestor (keep?)

name or RFC, etc.

[4.5.3.](#) Revision

1.0

[4.5.4.](#) Revision Date

YYYY-MM-DD

[4.6.](#) Comments and Remarks

Additional (Informational) details for this entry

[5.](#) Packet Delay Variation Registry Entry

This section gives an initial registry entry for a Packet Delay Variation metric.

Note: If each Registry entry should only produce a "raw" output or a statistical summary, then the "Output" Category can be split and this section can become two closely-related metrics.

[5.1.](#) Summary

This category includes multiple indexes to the registry entries, the element ID and metric name.

<skipping some Summary columns for now>

[5.1.1.](#) ID (Identifier)

<insert numeric identifier, an integer>

[5.1.2.](#) Name

<insert name according to metric naming convention>

OWPDV_Active_IP-UDP-Periodic_RFCXXXsecY_Seconds_95Percentile

Morton, et al.

Expires April 25, 2019

[Page 18]

Internet-Draft

Initial Registry

October 2018

[5.1.3.](#) URIs

URI: Prefix urn:ietf:metrics:perf:<name>

URL: http://<TBD by IANA>/<name>

[5.1.4.](#) Description

An assessment of packet delay variation with respect to the minimum delay observed on the periodic stream, and the Output is expressed as the 95th percentile of the packet delay variation distribution.

[5.1.5.](#) Change Controller

<org or person >

IETF

[5.1.6.](#) Version (of Registry Format)

1.0

[5.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference

and values of input factors, called fixed parameters.

5.2.1. Reference Definition

<Full bibliographic reference to an immutable doc.>

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

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

Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", [RFC 5481](#), March 2009. [[RFC5481](#)]

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

<specific section reference and additional clarifications, if needed>

See sections [2.4](#) and [3.4](#) of [[RFC3393](#)]. Singleton delay differences measured are referred to by the variable name "ddT" (applicable to all forms of delay variation). However, this metric entry specifies the PDV form defined in [section 4.2 of \[\[RFC5481\]\(#\)\]](#), where the singleton PDV for packet *i* is referred to by the variable name "PDV(*i*)".

5.2.2. Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

o IPv4 header values:

- * DSCP: set to 0
- * TTL: set to 255
- * Protocol: Set to 17 (UDP)

- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Count: set to 255
 - * Protocol: Set to 17 (UDP)
- o UDP header values:
 - * Checksum: the checksum MUST be calculated and included in the header
- o UDP Payload
 - * total of 200 bytes

Other measurement parameters:

Tmax: a loss threshold waiting time with value 3.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

F a selection function unambiguously defining the packets from the stream selected for the metric. See [section 4.2 of \[RFC5481\]](#) for the PDV form.

See the Packet Stream generation category for two additional Fixed Parameters.

[5.3](#). Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[5.3.1](#). Reference Method

<for metric, insert relevant section references and supplemental

info>

See [section 2.6](#) and 3.6 of [\[RFC3393\]](#) for general singleton element calculations. This metric entry requires implementation of the PDV form defined in [section 4.2 of \[RFC5481\]](#). Also see measurement considerations in [section 8 of \[RFC5481\]](#).

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a packet lost. Lost packets SHALL be designated as having undefined delay.

The calculations on the one-way delay SHALL be performed on the conditional distribution, conditioned on successful packet arrival within Tmax. Also, when all packet delays are stored, the process which calculates the one-way delay value MAY enforce the Tmax threshold on stored values before calculations. See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving packet. Sequence numbers or other send-order identification MUST be retained at the Src or included with each packet to disambiguate packet reordering if it occurs.

If a standard measurement protocol is employed, then the measurement process will determine the sequence numbers or timestamps applied to test packets after the Fixed and Runtime parameters are passed to that process. The chosen measurement protocol will dictate the format of sequence numbers and time-stamps, if they are conveyed in the packet payload.

[5.3.2](#). Packet Stream Generation

<list of generation parameters and section/spec references if needed>

This section gives the details of the packet traffic which is the

basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

[Section 3 of \[RFC3432\]](#) prescribes the method for generating Periodic streams using associated parameters.

incT the nominal duration of inter-packet interval, first bit to first bit, with value 0.0200, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms).

dT the duration of the interval for allowed sample start times, with value 1.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms).

T0 the actual start time of the periodic stream, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)).

NOTE: an initiation process with a number of control exchanges resulting in unpredictable start times (within a time interval) may be sufficient to avoid synchronization of periodic streams, and therefore a valid replacement for selecting a start time at random from a fixed interval.

The T0 parameter will be reported as a measured parameter. Parameters incT and dT are Fixed Parameters.

[5.3.3](#). Traffic Filtering (observation) Details

<insert the measured results based on a filtered version of the packets observed, and this section provides the filter details (when present), and section reference>.

NA

[5.3.4.](#) Sampling Distribution

<insert time distribution details, or how this is diff from the filter>

NA

[5.3.5.](#) Run-time Parameters and Data Format

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Tf a time, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a end time date is ignored and Tf is interpreted as the Duration of the measurement interval.

[5.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

Src launches each packet to Dst.

Dst waits for each packet from Src.

[5.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

Internet-Draft

Initial Registry

October 2018

[5.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

Percentile -- for the conditional distribution of all packets with a valid value of one-way delay (undefined delays are excluded), a single value corresponding to the 95th percentile, as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The percentile = 95, meaning that the reported delay, "95Percentile", is the smallest value of one-way PDV for which the Empirical Distribution Function (EDF), $F(95\text{Percentile}) \geq 95\%$ of the singleton one-way PDV values in the conditional distribution. See [section 11.3 of \[RFC2330\]](#) for the definition of the percentile statistic using the EDF.

[5.4.2.](#) Reference Definition

<the output type and data format for each type of result>

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

95Percentile The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[5.4.3.](#) Metric Units

<insert units for the measured results, and the reference

specification>.

The 95th Percentile of one-way PDV is expressed in seconds.

[5.4.4.](#) Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback that includes as much of the measurement system as possible, performs address manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

For one-way delay measurements, the error calibration must include an assessment of the internal clock synchronization with its external reference (this internal clock is supplying timestamps for measurement). In practice, the time offsets of clocks at both the source and destination are needed to estimate the systematic error due to imperfect clock synchronization (the time offsets are smoothed, thus the random variation is not usually represented in the results).

`time_offset` The time value of the result is expressed in units of seconds, as a signed value of type `decimal64` with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a normal measurement with additional indication that it is a calibration result. In any measurement, the measurement function SHOULD report its current estimate of time offset as an indicator of the degree of synchronization.

Both internal loopback calibration and clock synchronization can be used to estimate the **available accuracy** of the Output Metric Units.

For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

[5.5.](#) Administrative items

[5.5.1.](#) Status

<current or deprecated>

Morton, et al.

Expires April 25, 2019

[Page 25]

Internet-Draft

Initial Registry

October 2018

[5.5.2.](#) Requestor (keep?)

<name of individual or RFC, etc.>

[5.5.3.](#) Revision

1.0

[5.5.4.](#) Revision Date

YYYY-MM-DD

[5.6.](#) Comments and Remarks

<Additional (Informational) details for this entry>

Lost packets represent a challenge for delay variation metrics. See [section 4.1 of \[RFC3393\]](#) and the delay variation applicability statement[RFC5481] for extensive analysis and comparison of PDV and an alternate metric, IPDV.

[6.](#) DNS Response Latency and Loss Registry Entries

@@@ comment from Brian: there is an interesting method for DNS measurement by encoding information in the query itself. It is a question of what exactly we are trying to measure: specific RR, or the infrastructure itself. (at this time we measure a specific RR).

This section gives initial registry entries for DNS Response Latency and Loss from a network user's perspective, for a specific named resource. The metric can be measured repeatedly using different names. [RFC 2681](#) [[RFC2681](#)] defines a Round-trip delay metric. We build on that metric by specifying several of the input parameters to precisely define two metrics for measuring DNS latency and loss.

Note to IANA: Each Registry "Name" below specifies a single registry entry, whose output format varies in accordance with the name.

All column entries beside the ID, Name, Description, and Output Reference Method categories are the same, thus this section proposes two closely-related registry entries. As a result, IANA is also asked to assign corresponding URNs and URLs to each Named Metric.

[6.1.](#) Summary

This category includes multiple indexes to the registry entries, the element ID and metric name.

Morton, et al.

Expires April 25, 2019

[Page 26]

Internet-Draft

Initial Registry

October 2018

<skipping some admin columns>

[6.1.1.](#) ID (Identifier)

<insert numeric identifier, an integer>

IANA is asked to assign different numeric identifiers to each of the two Named Metrics.

[6.1.2.](#) Name

<insert name according to metric naming convention>

RTDNS_Active_IP-UDP-Poisson_RFCXXXXsecY_Seconds_Raw

RLDNS_Active_IP-UDP-Poisson_RFCXXXXsecY_Logical_Raw

[6.1.3.](#) URI

URI: Prefix urn:ietf:metrics:perf:<name>

URL: <http://<TBD by IANA>/<name>>

[6.1.4.](#) Description

This is a metric for DNS Response performance from a network user's perspective, for a specific named resource. The metric can be measured repeatedly using different resource names.

RTDNS: This metric assesses the response time, the interval from the query transmission to the response.

RLDNS: This metric indicates that the response was deemed lost. In other words, the response time exceeded the maximum waiting time.

[6.1.5.](#) Change Controller

IETF

[6.1.6.](#) Version (of Registry Format)

1.0

[6.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

Morton, et al.

Expires April 25, 2019

[Page 27]

Internet-Draft

Initial Registry

October 2018

[6.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987. (and updates)

[RFC1035]

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

[RFC2681]

<specific section reference and additional clarifications, if needed>

[Section 2.4 of \[RFC2681\]](#) provides the reference definition of the singleton (single value) Round-trip delay metric. [Section 3.4 of \[RFC2681\]](#) provides the reference definition expanded to cover a multi-singleton sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

For DNS Response Latency, the entities in [\[RFC1035\]](#) must be mapped to [\[RFC2681\]](#). The Local Host with its User Program and Resolver take the role of "Src", and the Foreign Name Server takes the role of "Dst".

Note that although the [\[RFC2681\]](#) definition of "Round-trip-Delay between Src and Dst at T" is directionally ambiguous in the text, this metric tightens the definition further to recognize that the host in the "Src" role will send the first packet to "Dst", and ultimately receive the corresponding return packet from "Dst" (when neither are lost).

Morton, A., "Round-trip Packet Loss Metrics", [RFC 6673](#), August 2012.

[RFC6673]

Both response time and loss metrics employ a maximum waiting time for received responses, so the count of lost packets to total packets sent is the basis for the loss determination as per [Section 4.3 of \[RFC6673\]](#).

[6.2.2](#). Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Morton, et al.

Expires April 25, 2019

[Page 28]

Internet-Draft

Initial Registry

October 2018

Type-P as defined in [Section 13 of \[RFC2330\]](#):

o IPv4 header values:

* DSCP: set to 0

* TTL set to 255

- * Protocol: Set to 17 (UDP)
- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Count: set to 255
 - * Protocol: Set to 17 (UDP)
- o UDP header values:
 - * Source port: 53
 - * Destination port: 53
 - * Checksum: the checksum must be calculated and included in the header
- o Payload: The payload contains a DNS message as defined in [RFC 1035](#) [[RFC1035](#)] with the following values:
 - * The DNS header section contains:
 - + Identification (see the Run-time column)
 - + QR: set to 0 (Query)
 - + OPCODE: set to 0 (standard query)
 - + AA: not set
 - + TC: not set
 - + RD: set to one (recursion desired)
 - + RA: not set
 - + RCODE: not set

- + QDCOUNT: set to one (only one entry)
- + ANCOUNT: not set
- + NSCOUNT: not set
- + ARCOUNT: not set
- * The Question section contains:
 - + QNAME: the Fully Qualified Domain Name (FQDN) provided as input for the test, see the Run-time column
 - + QTYPE: the query type provided as input for the test, see the Run-time column
 - + QCLASS: set to 1 for IN
- * The other sections do not contain any Resource Records.

Other measurement parameters:

- o Tmax: a loss threshold waiting time (and to help disambiguate queries)
 - * 5.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

Observation: reply packets will contain a DNS response and may contain RRs.

[6.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[6.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

The methodology for this metric is defined as Type-P-Round-trip-Delay-Poisson-Stream in [section 2.6 of RFC 2681 \[RFC2681\]](#) and section

Internet-Draft

Initial Registry

October 2018

3.6 of [RFC 2681](#) [[RFC2681](#)] using the Type-P and Timeout defined under Fixed Parameters.

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a response packet lost. Lost packets SHALL be designated as having undefined delay and counted for the RLDNS metric.

The calculations on the delay (RTT) SHALL be performed on the conditional distribution, conditioned on successful packet arrival within Tmax. Also, when all packet delays are stored, the process which calculates the RTT value MAY enforce the Tmax threshold on stored values before calculations. See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving reply.

DNS Messages bearing Queries provide for random ID Numbers in the Identification header field, so more than one query may be launched while a previous request is outstanding when the ID Number is used. Therefore, the ID Number MUST be retained at the Src or included with each response packet to disambiguate packet reordering if it occurs.

IF a DNS response does not arrive within Tmax, the response time RTDNS is undefined, and RLDNS = 1. The Message ID SHALL be used to disambiguate the successive queries that are otherwise identical.

Since the ID Number field is only 16 bits in length, it places a limit on the number of simultaneous outstanding DNS queries during a stress test from a single Src address.

Refer to [Section 4.4 of \[RFC6673\]](#) for expanded discussion of the instruction to "send a Type-P packet back to the Src as quickly as possible" in [Section 2.6 of RFC 2681](#) [[RFC2681](#)]. However, the DNS Server is expected to perform all required functions to prepare and send a response, so the response time will include processing time and network delay. [Section 8 of \[RFC6673\]](#) presents additional

requirements which SHALL be included in the method of measurement for this metric.

In addition to operations described in [[RFC2681](#)], the Src MUST parse the DNS headers of the reply and prepare the information for

subsequent reporting as a measured result, along with the Round-Trip Delay.

[6.3.2.](#) Packet Stream Generation

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

<list of generation parameters and section/spec references if needed>

[Section 11.1.3 of RFC 2681](#) [[RFC2330](#)] provides three methods to generate Poisson sampling intervals. The reciprocal of lambda is the average packet rate, thus the Run-time Parameter is Reciprocal_lambda = 1/lambda, in seconds.

Method 3 is used, where given a start time (Run-time Parameter), the subsequent send times are all computed prior to measurement by computing the pseudo-random distribution of inter-packet send times, (truncating the distribution as specified in the Run-time Parameters), and the Src sends each packet at the computed times.

Note that Trunc is the upper limit on inter-packet times in the Poisson distribution. A random value greater than Trunc is set equal to Trunc instead.

[6.3.3.](#) Traffic Filtering (observation) Details

The measured results based on a filtered version of the packets observed, and this section provides the filter details (when present).

<section reference>.

[6.3.4. Sampling Distribution](#)

<insert time distribution details, or how this is diff from the filter>

NA

Morton, et al.

Expires April 25, 2019

[Page 32]

Internet-Draft

Initial Registry

October 2018

[6.3.5. Run-time Parameters and Data Format](#)

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Tf a time, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a end time date is ignored and Tf is interpreted as the Duration of the measurement interval.

Reciprocal_lambda average packet interval for Poisson Streams expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.0001 seconds (0.1 ms), and with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

Trunc Upper limit on Poisson distribution expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.0001 seconds (0.1 ms), and with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#) (values above this limit will be clipped and set to the limit value). (if fixed, Trunc = 30.0000 seconds.)

ID The 16-bit identifier assigned by the program that generates the query, and which must vary in successive queries, see [Section 4.1.1 of \[RFC1035\]](#). This identifier is copied into the

corresponding reply and can be used by the requester (Src) to match-up replies to outstanding queries.

QNAME The domain name of the Query, formatted as specified in [section 4 of \[RFC6991\]](#).

QTYPE The Query Type, which will correspond to the IP address family of the query (decimal 1 for IPv4 or 28 for IPv6, formatted as a uint16, as per [section 9.2 of \[RFC6020\]](#)).

[6.3.6](#). Roles

<lists the names of the different roles from the measurement method>

Src launches each packet and waits for return transmissions from Dst.

Dst waits for each packet from Src and sends a return packet to Src.

[6.4](#). Output

This category specifies all details of the Output of measurements

using the metric.

6.4.1. Type

<insert name of the output type, raw or a selected summary statistic>

Raw -- for each DNS Query packet sent, sets of values as defined in the next column, including the status of the response, only assigning delay values to successful query-response pairs.

6.4.2. Reference Definition

<describe the data format for each type of result>

For all outputs:

T the time the DNS Query was sent during the measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

dT The time value of the round-trip delay to receive the DNS response, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per

[section 6](#) of RFC [\[RFC5905\]](#). This value is undefined when the response packet is not received at Src within waiting time Tmax seconds.

Rcode The value of the Rcode field in the DNS response header, expressed as a uint64 as specified in [section 9.2 of \[RFC6020\]](#). Non-zero values convey errors in the response, and such replies must be analyzed separately from successful requests.

6.4.3. Metric Units

<insert units for the measured results, and the reference specification>.

RTDNS: Round-trip Delay, dT, is expressed in seconds.

RTLDNS: the Logical value, where 1 = Lost and 0 = Received.

[6.4.4.](#) Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback at the Source host that includes as much of the measurement system as possible, performs address and payload manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a normal measurement with additional indication that it is a calibration result.

Both internal loopback calibration and clock synchronization can be used to estimate the *available accuracy* of the Output Metric Units. For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

[6.5.](#) Administrative items

[6.5.1.](#) Status

<current or deprecated>

Morton, et al.

Expires April 25, 2019

[Page 35]

Internet-Draft

Initial Registry

October 2018

[6.5.2.](#) Requestor

name or RFC, etc.

[6.5.3.](#) Revision

1.0

[6.5.4.](#) Revision Date

YYYY-MM-DD

[6.6.](#) Comments and Remarks

Additional (Informational) details for this entry

[7.](#) UDP Poisson One-way Delay and Loss Registry Entries

This section specifies five initial registry entries for the UDP Poisson One-way Delay, and one for UDP Poisson One-way Loss.

IANA Note: Registry "Name" below specifies a single registry entry, whose output format varies according to the <statistic> element of the name that specifies one form of statistical summary. There is an additional metric name for the Loss metric.

All column entries beside the ID, Name, Description, and Output Reference Method categories are the same, thus this section proposes six closely-related registry entries. As a result, IANA is also asked to assign corresponding URNs and URLs to each Named Metric.

[7.1.](#) Summary

This category includes multiple indexes to the registry entries, the element ID and metric name.

[7.1.1.](#) ID (Identifier)

<insert numeric identifier, an integer, one corresponding to each name below>

IANA is asked to assign different numeric identifiers to each of the six Metrics.

[7.1.2.](#) Name

<insert name according to metric naming convention>

OWDelay_Active_IP-UDP-Poisson-
Payload250B_RFCXXXXsecY_Seconds_<statistic>

where <statistic> is one of:

- o 95Percentile
- o Mean
- o Min
- o Max
- o StdDev

OWLoss_Active_IP-UDP-Poisson-
Payload250B_RFCXXXXsecY_Percent_LossRatio

[7.1.3.](#) URI and URL

URI: Prefix urn:iETF:metrics:perf:<name>

URL: http://www.iana.org/ ... <name>

[7.1.4.](#) Description

OWDelay: This metric assesses the delay of a stream of packets exchanged between two hosts (or measurement points), and reports the <statistic> One-way delay for all successfully exchanged packets based on their conditional delay distribution.

where <statistic> is one of:

- o 95Percentile
- o Mean
- o Min
- o Max
- o StdDev

OWLoss: This metric assesses the loss ratio of a stream of packets exchanged between two hosts (which are the two measurement points), and the Output is the One-way loss ratio for all successfully received packets expressed as a percentage.

[7.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[7.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

For Delay:

Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, [RFC 7679](#), DOI 10.17487/RFC7679, January 2016, <<http://www.rfc-editor.org/info/rfc7679>>.

[RFC7679]

Morton, A., and Stephan, E., "Spatial Composition of Metrics", [RFC 6049](#), January 2011.

[RFC6049]

<specific section reference and additional clarifications, if needed>

[Section 3.4 of \[RFC7679\]](#) provides the reference definition of the singleton (single value) One-way delay metric. [Section 4.4 of \[RFC7679\]](#) provides the reference definition expanded to cover a multi-value sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

Only successful packet transfers with finite delay are included in the sample, as prescribed in [section 4.1.2 of \[RFC6049\]](#).

For loss:

Almes, G., Kalidini, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", [RFC 7680](#), DOI 10.17487/RFC7680, January 2016, <<http://www.rfc-editor.org/info/rfc7680>>.

Internet-Draft

Initial Registry

October 2018

[Section 2.4 of \[RFC7680\]](#) provides the reference definition of the singleton (single value) one-way loss metric. [Section 3.4 of \[RFC7680\]](#) provides the reference definition expanded to cover a multi-singleton sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

[7.2.2.](#) Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Type-P:

- o IPv4 header values:
 - * DSCP: set to 0
 - * TTL: set to 255
 - * Protocol: Set to 17 (UDP)
- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Count: set to 255
 - * Protocol: Set to 17 (UDP)
- o UDP header values:
 - * Checksum: the checksum MUST be calculated and included in the header
- o UDP Payload: TWAMP Test Packet Formats, [Section 4.1.2 of \[RFC5357\]](#)
 - * Security features in use influence the number of Padding octets.

* 250 octets total, including the TWAMP format

Other measurement parameters:

Tmax: a loss threshold waiting time with value 3.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with

Morton, et al.

Expires April 25, 2019

[Page 39]

Internet-Draft

Initial Registry

October 2018

resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

See the Packet Stream generation category for two additional Fixed Parameters.

[7.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[7.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

The methodology for this metric is defined as Type-P-One-way-Delay-Poisson-Stream in [section 3.6 of \[RFC7679\]](#) and [section 4.6 of \[RFC7679\]](#) using the Type-P and Tmax defined under Fixed Parameters.

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a packet lost. Lost packets SHALL be designated as having undefined delay, and counted for the OWLoss metric.

The calculations on the one-way delay SHALL be performed on the conditional distribution, conditioned on successful packet arrival within Tmax. Also, when all packet delays are stored, the process which calculates the one-way delay value MAY enforce the Tmax threshold on stored values before calculations. See [section 4.1 of](#)

[\[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving packet. Sequence numbers or other send-order identification MUST be retained at the Src or included with each packet to disambiguate packet reordering if it occurs.

Since a standard measurement protocol is employed [\[RFC5357\]](#), then the measurement process will determine the sequence numbers or timestamps applied to test packets after the Fixed and Runtime parameters are passed to that process. The measurement protocol dictates the format

of sequence numbers and time-stamps conveyed in the TWAMP-Test packet payload.

[7.3.2.](#) Packet Stream Generation

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

<list of generation parameters and section/spec references if needed>

[Section 11.1.3 of RFC 2681 \[RFC2330\]](#) provides three methods to generate Poisson sampling intervals. The reciprocal of lambda is the average packet spacing, thus the Run-time Parameter is `Reciprocal_lambda = 1/lambda`, in seconds.

Method 3 SHALL be used, where given a start time (Run-time Parameter), the subsequent send times are all computed prior to measurement by computing the pseudo-random distribution of inter-packet send times, (truncating the distribution as specified in the Parameter Trunc), and the Src sends each packet at the computed times.

Note that Trunc is the upper limit on inter-packet times in the

Poisson distribution. A random value greater than Trunc is set equal to Trunc instead.

Reciprocal_lambda average packet interval for Poisson Streams expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.0001 seconds (0.1 ms), and with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#). Reciprocal_lambda = 1 packet per second.

Trunc Upper limit on Poisson distribution expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.0001 seconds (0.1 ms), and with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#) (values above this limit will be clipped and set to the limit value). Trunc = 30.0000 seconds.

[7.3.3.](#) Traffic Filtering (observation) Details

NA

[7.3.4.](#) Sampling Distribution

NA

[7.3.5.](#) Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6,

see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Tf a time, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a end time date is ignored and Tf is interpreted as the Duration of the measurement interval.

[7.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

Src launches each packet and waits for return transmissions from Dst. This is the TWAMP Session-Sender.

Dst waits for each packet from Src and sends a return packet to Src. This is the TWAMP Session-Reflector.

[7.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

[7.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

See subsection titles below for Types.

[7.4.2.](#) Reference Definition

<describe the data format for each type of result>

For all output types ---

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

For LossRatio -- the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 4.1 of \[RFC7680\]](#).

For each <statistic>, one of the following sub-sections apply:

[7.4.2.1](#). Percentile95

The 95th percentile SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3 of \[RFC3393\]](#) for details on the percentile statistic (where Round-trip delay should be substituted for "ipdv").

The percentile = 95, meaning that the reported delay, "95Percentile", is the smallest value of one-way delay for which the Empirical Distribution Function (EDF), $F(95\text{Percentile}) \geq 95\%$ of the singleton one-way delay values in the conditional distribution. See [section 11.3 of \[RFC2330\]](#) for the definition of the percentile statistic using the EDF.

95Percentile The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from

the 64-bit NTP timestamp as per [section 6](#) of RFC [[RFC5905](#)]

[7.4.2.2](#). Mean

The mean SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.2.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.2.3 of [[RFC6049](#)].

Mean The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [[RFC5905](#)]

[7.4.2.3](#). Min

The minimum SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.3.3 of [[RFC6049](#)].

Min The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [[RFC5905](#)]

[7.4.2.4](#). Max

The maximum SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is as follows:

$$\text{Max} = (\text{FiniteDelay } [j])$$

such that for some index, j , where $1 \leq j \leq N$
 $\text{FiniteDelay}[j] \geq \text{FiniteDelay}[n]$ for all n

Max The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[7.4.2.5](#). Std_Dev

The Std_Dev SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is the classic calculation for standard deviation of a population.

Std_Dev The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[7.4.3.](#) Metric Units

<insert units for the measured results, and the reference specification>.

The <statistic> of One-way Delay is expressed in seconds.

The One-way Loss Ratio is expressed as a percentage of lost packets to total packets sent.

[7.4.4.](#) Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback that includes as much of the measurement system as possible, performs address manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

For one-way delay measurements, the error calibration must include an assessment of the internal clock synchronization with its external reference (this internal clock is supplying timestamps for measurement). In practice, the time offsets of clocks at both the source and destination are needed to estimate the systematic error due to imperfect clock synchronization (the time offsets are smoothed, thus the random variation is not usually represented in the results).

`time_offset` The time value of the result is expressed in units of seconds, as a signed value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a normal measurement with additional indication that it is a calibration result. In any measurement, the measurement function

SHOULD report its current estimate of time offset as an indicator of the degree of synchronization.

Both internal loopback calibration and clock synchronization can be used to estimate the *available accuracy* of the Output Metric Units. For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

Morton, et al.

Expires April 25, 2019

[Page 46]

Internet-Draft

Initial Registry

October 2018

[7.5.](#) Administrative items

[7.5.1.](#) Status

<current or deprecated>

[7.5.2.](#) Requestor (keep?)

name or RFC, etc.

[7.5.3.](#) Revision

1.0

[7.5.4.](#) Revision Date

YYYY-MM-DD

[7.6.](#) Comments and Remarks

Additional (Informational) details for this entry

[8.](#) UDP Periodic One-way Delay and Loss Registry Entries

This section specifies five initial registry entries for the UDP Periodic One-way Delay, and one for UDP Periodic One-way Loss.

IANA Note: Registry "Name" below specifies a single registry entry, whose output format varies according to the <statistic> element of the name that specifies one form of statistical summary. There is an additional metric name for the Loss metric.

All column entries beside the ID, Name, Description, and Output

Reference Method categories are the same, thus this section proposes six closely-related registry entries. As a result, IANA is also asked to assign corresponding URNs and URLs to each Named Metric.

[8.1.](#) Summary

This category includes multiple indexes to the registry entries, the element ID and metric name.

[8.1.1.](#) ID (Identifier)

<insert numeric identifier, an integer, one corresponding to each name below>

Morton, et al.

Expires April 25, 2019

[Page 47]

Internet-Draft

Initial Registry

October 2018

IANA is asked to assign a different numeric identifiers to each of the six Metrics.

[8.1.2.](#) Name

<insert name according to metric naming convention>

OWDelay_Active_IP-UDP-Periodic20m-
Payload142B_RFCXXXXsecY_Seconds_<statistic>

where <statistic> is one of:

- o 95Percentile
- o Mean
- o Min
- o Max
- o StdDev

OWLoss_Active_IP-UDP-Periodic-
Payload142B_RFCXXXXsecY_Percent_LossRatio

[8.1.3.](#) URIs

URI: Prefix urn:ietf:metrics:perf:<name>

URL: http:\\www.iana.org\ ... <name>

[8.1.4.](#) Description

OWDelay: This metric assesses the delay of a stream of packets exchanged between two hosts (or measurement points), and reports the <statistic> One-way delay for all successfully exchanged packets based on their conditional delay distribution.

where <statistic> is one of:

- o 95Percentile
- o Mean
- o Min
- o Max

- o StdDev

OWLoss: This metric assesses the loss ratio of a stream of packets exchanged between two hosts (which are the two measurement points), and the Output is the One-way loss ratio for all successfully received packets expressed as a percentage.

[8.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[8.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

For Delay:

Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, [RFC 7679](#), DOI 10.17487/RFC7679, January 2016, <<http://www.rfc-editor.org/info/rfc7679>>.

[RFC7679]

Morton, A., and Stephan, E., "Spatial Composition of Metrics", [RFC 6049](#), January 2011.

[RFC6049]

<specific section reference and additional clarifications, if needed>

[Section 3.4 of \[RFC7679\]](#) provides the reference definition of the singleton (single value) One-way delay metric. [Section 4.4 of \[RFC7679\]](#) provides the reference definition expanded to cover a multi-value sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

Only successful packet transfers with finite delay are included in the sample, as prescribed in [section 4.1.2 of \[RFC6049\]](#).

For loss:

Almes, G., Kalidini, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", [RFC 7680](#), DOI 10.17487/RFC7680, January 2016, <<http://www.rfc-editor.org/info/rfc7680>>.

Morton, et al.

Expires April 25, 2019

[Page 49]

Internet-Draft

Initial Registry

October 2018

[Section 2.4 of \[RFC7680\]](#) provides the reference definition of the singleton (single value) one-way loss metric. [Section 3.4 of \[RFC7680\]](#) provides the reference definition expanded to cover a multi-singleton sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

[8.2.2](#). Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Type-P:

- o IPv4 header values:
 - * DSCP: set to 0
 - * TTL: set to 255
 - * Protocol: Set to 17 (UDP)
- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Count: set to 255
 - * Protocol: Set to 17 (UDP)
- o UDP header values:
 - * Checksum: the checksum MUST be calculated and included in the header
- o UDP Payload: TWAMP Test Packet Formats, [Section 4.1.2 of \[RFC5357\]](#)
 - * Security features in use influence the number of Padding octets.
 - * 142 octets total, including the TWAMP format (if used)

Other measurement parameters:

Tmax: a loss threshold waiting time with value 3.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with

resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

See the Packet Stream generation category for two additional Fixed Parameters.

[8.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[8.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

The methodology for this metric is defined as Type-P-One-way-Delay-Poisson-Stream in [section 3.6 of \[RFC7679\]](#) and [section 4.6 of \[RFC7679\]](#) using the Type-P and Tmax defined under Fixed Parameters. However, a Periodic stream is used, as defined in [\[RFC3432\]](#).

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a packet lost. Lost packets SHALL be designated as having undefined delay, and counted for the OWLoss metric.

The calculations on the one-way delay SHALL be performed on the conditional distribution, conditioned on successful packet arrival within Tmax. Also, when all packet delays are stored, the process which calculates the one-way delay value MAY enforce the Tmax threshold on stored values before calculations. See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving packet. Sequence numbers or other send-order identification MUST be retained at the Src or included with each packet to disambiguate packet reordering if it occurs.

Since a standard measurement protocol is employed [\[RFC5357\]](#), then the measurement process will determine the sequence numbers or timestamps applied to test packets after the Fixed and Runtime parameters are passed to that process. The measurement protocol dictates the format

of sequence numbers and time-stamps conveyed in the TWAMP-Test packet payload.

8.3.2. Packet Stream Generation

<list of generation parameters and section/spec references if needed>

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

Section 3 of [RFC3432] prescribes the method for generating Periodic streams using associated parameters.

incT the nominal duration of inter-packet interval, first bit to first bit, with value 0.0200 expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see section 9.3 of [RFC6020]) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per section 6 of [RFC5905].

dT the duration of the interval for allowed sample start times, with value 1.0000, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see section 9.3 of [RFC6020]) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per section 6 of [RFC5905].

T0 the actual start time of the periodic stream, determined from T0 and dT.

NOTE: an initiation process with a number of control exchanges resulting in unpredictable start times (within a time interval) may be sufficient to avoid synchronization of periodic streams, and therefore a valid replacement for selecting a start time at random from a fixed interval.

These stream parameters will be specified as Run-time parameters.

8.3.3. Traffic Filtering (observation) Details

NA

Internet-Draft

Initial Registry

October 2018

[8.3.4.](#) Sampling Distribution

NA

[8.3.5.](#) Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Tf a time, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a end time date is ignored and Tf is interpreted as the Duration of the measurement interval.

[8.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

Src launches each packet and waits for return transmissions from Dst. This is the TWAMP Session-Sender.

Dst waits for each packet from Src and sends a return packet to Src. This is the TWAMP Session-Reflector.

[8.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

Morton, et al.

Expires April 25, 2019

[Page 53]

Internet-Draft

Initial Registry

October 2018

[8.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

See subsection titles in Reference Definition for Latency Types.

[8.4.2.](#) Reference Definition

<describe the data format for each type of result>

For all output types ---

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

For LossRatio -- the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 4.1 of \[RFC7680\]](#).

For each <statistic>, one of the following sub-sections apply:

[8.4.2.1.](#) Percentile95

The 95th percentile SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional

distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3 of \[RFC3393\]](#) for details on the percentile statistic (where Round-trip delay should be substituted for "ipdv").

The percentile = 95, meaning that the reported delay, "95Percentile", is the smallest value of one-way delay for which the Empirical Distribution Function (EDF), $F(95\text{Percentile}) \geq 95\%$ of the singleton one-way delay values in the conditional distribution. See [section 11.3 of \[RFC2330\]](#) for the definition of the percentile statistic using the EDF.

95Percentile The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[8.4.2.2](#). Mean

The mean SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.2.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.2.3 of [\[RFC6049\]](#).

Mean The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[8.4.2.3](#). Min

The minimum SHALL be calculated using the conditional distribution of

all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.3.3 of [\[RFC6049\]](#).

Min The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[8.4.2.4](#). Max

The maximum SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is as follows:

$$\text{Max} = (\text{FiniteDelay } [j])$$

such that for some index, j , where $1 \leq j \leq N$
 $\text{FiniteDelay}[j] \geq \text{FiniteDelay}[n]$ for all n

Max The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9

(see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[8.4.2.5](#). Std_Dev

The Std_Dev SHALL be calculated using the conditional distribution of all packets with a finite value of One-way delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is the classic calculation for standard deviation of a population.

Std_Dev The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[8.4.3](#). Metric Units

<insert units for the measured results, and the reference specification>.

The <statistic> of One-way Delay is expressed in seconds, where <statistic> is one of:

- o 95Percentile
- o Mean
- o Min

- o Max
- o StdDev

The One-way Loss Ratio is expressed as a percentage of lost packets to total packets sent.

8.4.4. Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback that includes as much of the measurement system as possible, performs address manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

For one-way delay measurements, the error calibration must include an assessment of the internal clock synchronization with its external reference (this internal clock is supplying timestamps for measurement). In practice, the time offsets of clocks at both the source and destination are needed to estimate the systematic error due to imperfect clock synchronization (the time offsets are smoothed, thus the random variation is not usually represented in the results).

`time_offset` The time value of the result is expressed in units of seconds, as a signed value of type `decimal64` with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a normal measurement with additional indication that it is a calibration result. In any measurement, the measurement function SHOULD report its current estimate of time offset as an indicator of the degree of synchronization.

Both internal loopback calibration and clock synchronization can be used to estimate the *available accuracy* of the Output Metric Units. For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

[8.5.](#) Administrative items

[8.5.1.](#) Status

<current or deprecated>

[8.5.2.](#) Requestor (keep?)

name or RFC, etc.

[8.5.3.](#) Revision

1.0

[8.5.4.](#) Revision Date

YYYY-MM-DD

[8.6.](#) Comments and Remarks

Additional (Informational) details for this entry

[9.](#) ICMP Round-trip Latency and Loss Registry Entries

This section specifies three initial registry entries for the ICMP Round-trip Latency, and another entry for ICMP Round-trip Loss Ratio.

This section specifies four Registry entries with many common columns.

All column entries beside the ID, Name, Description, and Output Reference Method categories are the same, thus this section proposes two closely-related registry entries. As a result, IANA is also asked to assign four corresponding URNs and URLs to each Named Metric.

[9.1.](#) Summary

This category includes multiple indexes to the registry entry: the element ID and metric name.

[9.1.1.](#) ID (Identifier)

<insert a numeric identifier, an integer, TBD>

IANA is asked to assign different numeric identifiers to each of the four Named Metrics.

[9.1.2.](#) Name

<insert name according to metric naming convention>

RTDelay_Active_IP-ICMP-SendOnRcv_RFCXXXXsecY_Seconds_<statistic>

where <statistic> is one of:

- o Mean
- o Min
- o Max

RTLoss_Active_IP-ICMP-SendOnRcv_RFCXXXXsecY_Percent_LossRatio

[9.1.3.](#) URIs

URN: Prefix urn:ietf:metrics:perf:<name>

URL: http://<TBD by IANA>/<name>

[9.1.4.](#) Description

RTDelay: This metric assesses the delay of a stream of ICMP packets exchanged between two hosts (which are the two measurement points), and the Output is the Round-trip delay for all successfully exchanged packets expressed as the <statistic> of their conditional delay distribution, where <statistic> is one of:

- o Mean
- o Min
- o Max

Internet-Draft

Initial Registry

October 2018

RTLoss: This metric assesses the loss ratio of a stream of ICMP packets exchanged between two hosts (which are the two measurement points), and the Output is the Round-trip loss ratio for all successfully exchanged packets expressed as a percentage.

[9.1.5.](#) Change Controller

IETF

[9.1.6.](#) Version (of Registry Format)

1.0

[9.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[9.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

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

[RFC2681]

<specific section reference and additional clarifications, if needed>

[Section 2.4 of \[RFC2681\]](#) provides the reference definition of the singleton (single value) Round-trip delay metric. [Section 3.4 of \[RFC2681\]](#) provides the reference definition expanded to cover a multi-singleton sample. Note that terms such as singleton and sample are defined in [Section 11 of \[RFC2330\]](#).

Note that although the [\[RFC2681\]](#) definition of "Round-trip-Delay between Src and Dst" is directionally ambiguous in the text, this metric tightens the definition further to recognize that the host in the "Src" role will send the first packet to "Dst", and ultimately receive the corresponding return packet from "Dst" (when neither are lost).

Finally, note that the variable "dT" is used in [[RFC2681](#)] to refer to the value of Round-trip delay in metric definitions and methods. The variable "dT" has been re-used in other IPPM literature to refer to different quantities, and cannot be used as a global variable name.

Morton, A., "Round-trip Packet Loss Metrics", [RFC 6673](#), August 2012.

[RFC6673]

Both delay and loss metrics employ a maximum waiting time for received packets, so the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 6.1 of \[RFC6673\]](#).

[9.2.2](#). Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Type-P as defined in [Section 13 of \[RFC2330\]](#):

- o IPv4 header values:
 - * DSCP: set to 0
 - * TTL: set to 255
 - * Protocol: Set to 01 (ICMP)
- o IPv6 header values:
 - * DSCP: set to 0
 - * Hop Limit: set to 255
 - * Protocol: Set to 01 (ICMP)
- o ICMP header values:
 - * Type: 8 (Echo Request)

- * Code: 0
- * Checksum: the checksum MUST be calculated and included in the header
- * (Identifier and Sequence Number set at Run-Time)
- o ICMP Payload
 - * total of 32 bytes of random info

Other measurement parameters:

- o Tmax: a loss threshold waiting time
 - * 3.0, expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 4 (see [section 9.3 of \[RFC6020\]](#)) and with resolution of 0.0001 seconds (0.1 ms), with lossless conversion to/from the 32-bit NTP timestamp as per [section 6 of \[RFC5905\]](#).

[9.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[9.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

The methodology for this metric is defined as Type-P-Round-trip-Delay-Poisson-Stream in [section 2.6 of RFC 2681 \[RFC2681\]](#) and [section 3.6 of RFC 2681 \[RFC2681\]](#) using the Type-P and Tmax defined under Fixed Parameters.

The reference method distinguishes between long-delayed packets and lost packets by implementing a maximum waiting time for packet arrival. Tmax is the waiting time used as the threshold to declare a

packet lost. Lost packets SHALL be designated as having undefined delay, and counted for the RTLoss metric.

The calculations on the delay (RTD) SHALL be performed on the conditional distribution, conditioned on successful packet arrival within Tmax. Also, when all packet delays are stored, the process which calculates the RTD value MAY enforce the Tmax threshold on stored values before calculations. See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

The reference method requires some way to distinguish between different packets in a stream to establish correspondence between sending times and receiving times for each successfully-arriving packet. Sequence numbers or other send-order identification MUST be retained at the Src or included with each packet to disambiguate packet reordering if it occurs.

The measurement process will determine the sequence numbers applied to test packets after the Fixed and Runtime parameters are passed to that process. The ICMP measurement process and protocol will dictate the format of sequence numbers and other identifiers.

Refer to [Section 4.4 of \[RFC6673\]](#) for expanded discussion of the instruction to "send a Type-P packet back to the Src as quickly as possible" in [Section 2.6 of RFC 2681 \[RFC2681\]](#). [Section 8 of \[RFC6673\]](#) presents additional requirements which MUST be included in the method of measurement for this metric.

[9.3.2](#). Packet Stream Generation

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

The ICMP metrics use a sending discipline called "SendOnRcv" or Send On Receive. This is a modification of [Section 3 of \[RFC3432\]](#), which prescribes the method for generating Periodic streams using associated parameters:

incT the nominal duration of inter-packet interval, first bit to first bit

dT the duration of the interval for allowed sample start times

T0 the actual start time of the periodic stream

The incT and T0 stream parameters will be specified as Run-time parameters, dT is not used in SendOnRcv.

A SendOnRcv sender behaves exactly like a Periodic stream generator while all reply packets arrive with $RTD < incT$, and the inter-packet interval will be constant.

If a reply packet arrives with $RTD \geq incT$, then the inter-packet interval for the next sending time is nominally RTD.

If a reply packet fails to arrive within Tmax, then the inter-packet interval for the next sending time is nominally Tmax.

If an immediate send on reply arrival is desired, then set $incT=0$.

[9.3.3.](#) Traffic Filtering (observation) Details

The measured results based on a filtered version of the packets observed, and this section provides the filter details (when present).

<section reference>.

NA

[9.3.4.](#) Sampling Distribution

<insert time distribution details, or how this is diff from the filter>

NA

9.3.5. Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters, and their data formats>

Src the IP address of the host in the Src Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the Dst Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Count The total count of ICMP Echo Requests to send, formatted as a uint16, as per [section 9.2 of \[RFC6020\]](#).

(see the Packet Stream Generation section for additional Run-time parameters)

9.3.6. Roles

<lists the names of the different roles from the measurement method>

Src launches each packet and waits for return transmissions from Dst.

Dst waits for each packet from Src and sends a return packet to Src.

9.4. Output

This category specifies all details of the Output of measurements using the metric.

9.4.1. Type

<insert name of the output type, raw or a selected summary statistic>

See subsection titles in Reference Definition for Latency Types.

LossRatio -- the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 6.1 of \[RFC6673\]](#).

9.4.2. Reference Definition

<describe the data format for each type of result>

For all output types ---

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

TotalCount the count of packets actually sent by the Src to Dst during the measurement interval.

For LossRatio -- the count of lost packets to total packets sent is the basis for the loss ratio calculation as per [Section 4.1 of \[RFC7680\]](#).

For each <statistic>, one of the following sub-sections apply:

9.4.2.1. Mean

The mean SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.2.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.2.3 of [\[RFC6049\]](#).

Mean The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[9.4.2.2](#). Min

The minimum SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.3.3 of [\[RFC6049\]](#).

Min The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[9.4.2.3](#). Max

The maximum SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is as follows:

$$\text{Max} = (\text{FiniteDelay } [j])$$

such that for some index, j , where $1 \leq j \leq N$
 $\text{FiniteDelay}[j] \geq \text{FiniteDelay}[n]$ for all n

Max The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[9.4.3.](#) Metric Units

<insert units for the measured results, and the reference specification>.

The <statistic> of Round-trip Delay is expressed in seconds, where <statistic> is one of:

- o Mean
- o Min
- o Max

The Round-trip Loss Ratio is expressed as a percentage of lost packets to total packets sent.

[9.4.4.](#) Calibration

[Section 3.7.3 of \[RFC7679\]](#) provides a means to quantify the systematic and random errors of a time measurement. In-situ calibration could be enabled with an internal loopback at the Source host that includes as much of the measurement system as possible, performs address manipulation as needed, and provides some form of isolation (e.g., deterministic delay) to avoid send-receive interface contention. Some portion of the random and systematic error can be characterized this way.

When a measurement controller requests a calibration measurement, the loopback is applied and the result is output in the same format as a normal measurement with additional indication that it is a

calibration result.

Internet-Draft

Initial Registry

October 2018

Both internal loopback calibration and clock synchronization can be used to estimate the *available accuracy* of the Output Metric Units. For example, repeated loopback delay measurements will reveal the portion of the Output result resolution which is the result of system noise, and thus inaccurate.

[9.5.](#) Administrative items

[9.5.1.](#) Status

<current or deprecated>

[9.5.2.](#) Requestor (keep?)

name or RFC, etc.

[9.5.3.](#) Revision

1.0

[9.5.4.](#) Revision Date

YYYY-MM-DD

[9.6.](#) Comments and Remarks

Additional (Informational) details for this entry

[10.](#) TCP Round-Trip Delay and Loss Registry Entries

This section specifies three initial registry entries for the Passive assessment of TCP Round-Trip Delay (RTD) and another entry for TCP Round-trip Loss Count.

This section specifies four Registry entries with many common columns.

All column entries beside the ID, Name, Description, and Output Reference Method categories are the same, thus this section proposes

four closely-related registry entries. As a result, IANA is also asked to assign four corresponding URNs and URLs to each Named Metric.

[10.1.](#) Summary

This category includes multiple indexes to the registry entry: the element ID and metric name.

Morton, et al.

Expires April 25, 2019

[Page 68]

Internet-Draft

Initial Registry

October 2018

[10.1.1.](#) ID (Identifier)

<insert a numeric identifier, an integer, TBD>

IANA is asked to assign different numeric identifiers to each of the four Named Metrics.

[10.1.2.](#) Name

<insert name according to metric naming convention>

RTDelay_Passive_IP-TCP_RFCXXXsecY_Seconds_<statistic>

where <statistic> is one of:

- o Mean
- o Min
- o Max

RTDelay_Passive_IP-TCP-HS_RFCXXXsecY_Seconds_Singleton

Note that a mid-point observer only has the opportunity to compose a single RTDelay on the TCP Hand Shake.

RTLoss_Passive_IP-TCP_RFCXXXsecY_Packet_Count

[10.1.3.](#) URIs

URN: Prefix urn:ietf:metrics:perf:<name>

URL: `http://<TBD by IANA>/<name>`

[10.1.4.](#) Description

RTDelay: This metric assesses the round-trip delay of TCP packets constituting a single connection, exchanged between two hosts. We consider the measurement of round-trip delay based on a single Observation Point [[RFC7011](#)] somewhere in the network. The Output is the Round-trip delay for all successfully exchanged packets expressed as the <statistic> of their conditional delay distribution, where <statistic> is one of:

- o Mean
- o Min

Morton, et al.

Expires April 25, 2019

[Page 69]

Internet-Draft

Initial Registry

October 2018

- o Max

RTLoss: This metric assesses the estimated loss count for TCP packets constituting a single connection, exchanged between two hosts. We consider the measurement of round-trip delay based on a single Observation Point [[RFC7011](#)] somewhere in the network. The Output is the estimated Loss Count for the measurement interval.

[10.1.5.](#) Change Controller

IETF

[10.1.6.](#) Version (of Registry Format)

1.0

[10.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[10.2.1.](#) Reference Definitions

<Full bibliographic reference to an immutable doc.>

Although there is no RFC that describes passive measurement of Round-Trip Delay, the parallel definition for Active measurement is:

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

[RFC2681]

<specific section reference and additional clarifications, if needed>

This metric definition uses the terms singleton and sample as defined in [Section 11 of \[RFC2330\]](#). ([Section 2.4 of \[RFC2681\]](#) provides the reference definition of the singleton (single value) Round-trip delay metric. [Section 3.4 of \[RFC2681\]](#) provides the reference definition expanded to cover a multi-singleton sample.)

With the Observation Point [\[RFC7011\]](#) (OP) typically located between the hosts participating in the TCP connection, the Round-trip Delay metric requires two individual measurements between the OP and each host, such that the Spatial Composition [\[RFC6049\]](#) of the measurements yields a Round-trip Delay singleton (we are extending the composition of one-way subpath delays to subpath round-trip delay).

Using the direction of TCP SYN transmission to anchor the nomenclature, host A sends the SYN and host B replies with SYN-ACK during connection establishment. The direction of SYN transfer is considered the Forward direction of transmission, from A through OP to B (Reverse is B through OP to A).

Traffic filters reduce the packet stream at the OP to a Qualified bidirectional flow packets.

In the definitions below, Corresponding Packets are transferred in different directions and convey a common value in a TCP header field that establishes correspondence (to the extent possible). Examples may be found in the TCP timestamp fields.

For a real number, RTD_{fwd} , \gg the Round-trip Delay in the Forward direction from OP to host B at time T' is $RTD_{fwd} \ll$ REQUIRES that OP observed a Qualified Packet to host B at wire-time T' , that host B received that packet and sent a Corresponding Packet back to host A,

and OP observed the Corresponding Packet at wire-time $T' + \text{RTD_fwd}$.

For a real number, RTD_rev , \gg the Round-trip Delay in the Reverse direction from OP to host A at time T'' is $\text{RTD_rev} \ll$ REQUIRES that OP observed a Qualified Packet to host A at wire-time T'' , that host A received that packet and sent a Corresponding Packet back to host B, and that OP observed the Corresponding Packet at wire-time $T'' + \text{RTD_rev}$.

Ideally, the packet sent from host B to host A in both definitions above SHOULD be the same packet (or, when measuring RTD_rev first, the packet from host A to host B in both definitions should be the same).

The REQUIRED Composition Function for a singleton of Round-trip Delay at time T (where T is the earliest of T' and T'' above) is:

$$\text{RTDelay} = \text{RTD_fwd} + \text{RTD_rev}$$

Note that when OP is located at host A or host B, one of the terms composing RTDelay will be zero or negligible.

When the Qualified and Corresponding Packets are a TCP-SYN and a TCP-SYN-ACK, then $\text{RTD_fwd} == \text{RTD_HS_fwd}$.

When the Qualified and Corresponding Packets are a TCP-SYN-ACK and a TCP-ACK, then $\text{RTD_rev} == \text{RTD_HS_rev}$.

The REQUIRED Composition Function for a singleton of Round-trip Delay for the connection Hand Shake:

$$\text{RTDelay_HS} = \text{RTD_HS_fwd} + \text{RTD_HS_rev}$$

The definition of Round-trip Loss Count uses the nomenclature developed above, based on observation of the TCP header sequence numbers and storing the sequence number gaps observed. Packet Losses can be inferred from:

- o Out-of-order segments: TCP segments are transmitted with monotonically increasing sequence numbers, but these segments may be received out of order. [Section 3 of \[RFC4737\]](#) describes the notion of "next expected" sequence numbers which can be adapted to

TCP segments (for the purpose of detecting reordered packets). Observation of out-of-order segments indicates loss on the path prior to the OP, and creates a gap.

- o Duplicate segments: [Section 2 of \[RFC5560\]](#) defines identical packets and is suitable for evaluation of TCP packets to detect duplication. Observation of duplicate segments *without a corresponding gap* indicates loss on the path following the OP (because they overlap part of the delivered sequence numbers already observed at OP).

Each observation of an out-of-order or duplicate infers a singleton of loss, but composition of Round-trip Loss Counts will be conducted over a measurement interval which is synonymous with a single TCP connection.

With the above observations in the Forward direction over a measurement interval, the count of out-of-order and duplicate segments is defined as RTL_fwd. Comparable observations in the Reverse direction are defined as RTL_rev.

For a measurement interval (corresponding to a single TCP connection), T_0 to T_f , the REQUIRED Composition Function for a the two single-direction counts of inferred loss is:

$$RTL_{Loss} = RTL_{fwd} + RTL_{rev}$$

[10.2.2.](#) Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

Traffic Filters:

- o IPv4 header values:

* DSCP: set to 0

* Protocol: Set to 06 (TCP)

- o IPv6 header values:
 - * DSCP: set to 0
 - * Protocol: Set to 06 (TCP)
- o TCP header values:
 - * Flags: ACK, SYN, FIN, @@@@ others??
 - * Timestamp Option (TSopt): Set
 - + Kind: 8
 - + Length: 10 bytes
- o

10.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

10.3.1. Reference Methods

<for metric, insert relevant section references and supplemental info>

The foundation methodology for this metric is defined in [Section 4 of \[RFC7323\]](#) using the Timestamp Option with modifications that allow application at a mid-path Observation Point (OP) [\[RFC7011\]](#). Further details and applicable heuristics were derived from [\[Strowes\]](#) and [\[Trammell-14\]](#).

The Traffic Filter at the OP is configured to observe a single TCP connection. When the SYN, SYN-ACK, ACK handshake occurs, it offers the first opportunity to measure both RTD_fwd (on the SYN to SYN-ACK pair) and RTD_rev (on the SYN-ACK to ACK pair). Label this singleton of RTDelay as RTDelay_HS (composed using the forward and reverse measurement pair). RTDelay_HS SHALL be treated separately from other RTDelays on data-bearing packets and their ACKs. The RTDelay_HS value MAY be used as a sanity check on other Composed values of RTDelay.

For payload bearing packets, the OP measures the time interval between observation of a packet with Sequence Number s , and the corresponding ACK with same Sequence number. When the payload is transferred from host A to host B, the observed interval is `RTD_fwd`.

Because many data transfers are unidirectional (say, in the Forward direction from host A to host B), it is necessary to use pure ACK packets with Timestamp (`TSval`) and their Timestamp value echo to perform a `RTD_rev` measurement. The time interval between observation of the ACK from B to A, and the corresponding packet with Timestamp echo (`TSecr`) is the `RTD_rev`.

Delay Measurement Filtering Heuristics:

If Data payloads were transferred in both Forward and Reverse directions, then the Round-Trip Time Measurement Rule in [Section 4.1 of \[RFC7323\]](#) could be applied. This rule essentially excludes any measurement using a packet unless it makes progress in the transfer (advances the left edge of the send window, consistent with [Strowes]).

A different heuristic from [\[Trammell-14\]](#) is to exclude any `RTD_rev` that is larger than previously observed values. This would tend to exclude Reverse measurements taken when the Application has no data ready to send, because considerable time could be added to `RTD_rev` from this source of error.

@@@ Note that the above Heuristic assumes that host A is sending data. Host A expecting a download would mean that this heuristic should be applied to `RTD_fwd`.

The statistic calculations to summarize the delay (`RTDelay`) SHALL be performed on the conditional distribution, conditioned on successful Forward and Reverse measurements which follow the Heuristics.

Method for Inferring Loss:

The OP tracks sequence numbers and stores gaps for each direction of transmission, as well as the next-expected sequence number as in [\[Trammell-14\]](#) and [\[RFC4737\]](#). Loss is inferred from Out-of-order segments and Duplicate segments.

Loss Measurement Filtering Heuristics:

[\[Trammell-14\]](#) adds a window of evaluation based on the `RTDelay`.

Distinguish Re-ordered from OOO due to loss, because sequence number

gap is filled during the same RTDelay window. Segments detected as

re-ordered according to [[RFC4737](#)] MUST reduce the Loss Count inferred from Out-of-order segments.

Spurious (unneeded) retransmissions (observed as duplicates) can also be reduced this way, as described in [[Trammell-14](#)].

Sources of Error:

The principal source of RTDelay error is the host processing time to return a packet that defines the termination of a time interval. The heuristics above intend to mitigate these errors by excluding measurements where host processing time is a significant part of RTD_fwd or RTD_rev.

A key source of RTLoss error is observation loss, described in section 3 of [[Trammell-14](#)].

[10.3.2.](#) Packet Stream Generation

This section gives the details of the packet traffic which is the basis for measurement. In IPPM metrics, this is called the Stream, and can easily be described by providing the list of stream parameters.

NA

[10.3.3.](#) Traffic Filtering (observation) Details

The measured results based on a filtered version of the packets observed, and this section provides the filter details (when present).

The Fixed Parameters above give a portion of the Traffic Filter. Other aspects will be supplied as Run-time Parameters (below).

[10.3.4.](#) Sampling Distribution

<insert time distribution details, or how this is diff from the filter>

This metric requires a complete sample of all packets that qualify according to the Traffic Filter criteria.

[10.3.5.](#) Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

Morton, et al.

Expires April 25, 2019

[Page 75]

Internet-Draft

Initial Registry

October 2018

<list of run-time parameters, and their data formats>

Src the IP address of the host in the host A Role (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [Section 4 of \[RFC6991\]](#))

Dst the IP address of the host in the host B (format ipv4-address-no-zone value for IPv4, or ipv6-address-no-zone value for IPv6, see [section 4 of \[RFC6991\]](#))

T0 a time, the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). When T0 is "all-zeros", a start time is unspecified and Td is to be interpreted as the Duration of the measurement interval. The start time is controlled through other means.

Td Optionally, the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)), or the duration (see T0). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). Alternatively, the end of the measurement interval MAY be controlled by the measured connection, where the second pair of FIN and ACK packets exchanged between host A and B effectively ends the interval.

TTL or Hop Limit Set at desired value.

[10.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

host A launches the SYN packet to open the connection, and synonymous with an IP address.

host B replies with the SYN-ACK packet to open the connection, and synonymous with an IP address.

[10.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

[10.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

See subsection titles in Reference Definition for RTDelay Types.

Morton, et al.

Expires April 25, 2019

[Page 76]

Internet-Draft

Initial Registry

October 2018

For RTLoss -- the count of lost packets.

[10.4.2.](#) Reference Definition

<describe the data format for each type of result>

For all output types ---

T0 the start of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#).

Tf the end of a measurement interval, (format "date-and-time" as specified in [Section 5.6 of \[RFC3339\]](#), see also [Section 3 of \[RFC6991\]](#)). The UTC Time Zone is required by [Section 6.1 of \[RFC2330\]](#). The end of the measurement interval MAY be controlled by the measured connection, where the second pair of FIN and ACK packets exchanged between host A and B effectively ends the interval.

... ..

For RTDelay_HS -- the Round trip delay of the Handshake.

For RTLoss -- the count of lost packets.

For each <statistic>, one of the following sub-sections apply:

[10.4.2.1.](#) Mean

The mean SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.2.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.2.3 of [\[RFC6049\]](#).

Mean The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[10.4.2.2.](#) Min

The minimum SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for details on calculating this statistic, and 4.3.3 of [\[RFC6049\]](#).

Min The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[10.4.2.3.](#) Max

The maximum SHALL be calculated using the conditional distribution of all packets with a finite value of Round-trip delay (undefined delays are excluded), a single value as follows:

See [section 4.1 of \[RFC3393\]](#) for details on the conditional distribution to exclude undefined values of delay, and [Section 5 of \[RFC6703\]](#) for background on this analysis choice.

See [section 4.3.2 of \[RFC6049\]](#) for a closely related method for calculating this statistic, and 4.3.3 of [\[RFC6049\]](#). The formula is as follows:

$$\text{Max} = (\text{FiniteDelay } [j])$$

such that for some index, j , where $1 \leq j \leq N$
 $\text{FiniteDelay}[j] \geq \text{FiniteDelay}[n]$ for all n

Max The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see [section 9.3 of \[RFC6020\]](#)) with resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per [section 6](#) of RFC [\[RFC5905\]](#)

[10.4.3.](#) Metric Units

<insert units for the measured results, and the reference specification>.

The <statistic> of Round-trip Delay is expressed in seconds, where <statistic> is one of:

- o Mean
- o Min

- o Max

The Round-trip Delay of the Hand Shake is expressed in seconds.

The Round-trip Loss Count is expressed as a number of packets.

[10.4.4.](#) Calibration

Passive measurements at an OP could be calibrated against an active measurement (with loss emulation) at host A or B, where the active measurement represents the ground-truth.

[10.5.](#) Administrative items

[10.5.1.](#) Status

<current or deprecated>

[10.5.2.](#) Requestor (keep?)

name or RFC, etc.

[10.5.3.](#) Revision

1.0

[10.5.4.](#) Revision Date

YYYY-MM-DD

[10.6.](#) Comments and Remarks

Additional (Informational) details for this entry

Morton, et al.

Expires April 25, 2019

[Page 79]

Internet-Draft

Initial Registry

October 2018

[11.](#) ver08 BLANK Registry Entry

This section gives an initial registry entry for

[11.1.](#) Summary

This category includes multiple indexes to the registry entries, the element ID and metric name.

[11.1.1.](#) ID (Identifier)

<insert numeric identifier, an integer>

[11.1.2.](#) Name

<insert name according to metric naming convention>

[11.1.3.](#) URIs

URI: Prefix urn:ietf:metrics:perf:<name>

URL:

[11.1.4.](#) Description

TBD.

[11.1.5.](#) Reference

<reference to the RFC of spec where the registry entry is defined>

[11.1.6.](#) Change Controller

<org or person >

[11.1.7.](#) Version (of Registry Format)

<currently 1.0>

[11.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

[11.2.1.](#) Reference Definition

<Full bibliographic reference to an immutable doc.>

<specific section reference and additional clarifications, if needed>

[11.2.2.](#) Fixed Parameters

<list and specify Fixed Parameters, input factors that must be determined and embedded in the measurement system for use when needed>

[11.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

[11.3.1.](#) Reference Method

<for metric, insert relevant section references and supplemental info>

[11.3.2.](#) Packet Stream Generation

<list of generation parameters and section/spec references if needed>

[11.3.3.](#) Traffic Filtering (observation) Details

<insert the measured results based on a filtered version of the packets observed, and this section provides the filter details (when present), and section reference>.

[11.3.4.](#) Sampling Distribution

<insert time distribution details, or how this is diff from the filter>

[11.3.5.](#) Run-time Parameters and Data Format

<list of run-time parameters, and any reference(s)>.

[11.3.6.](#) Roles

<lists the names of the different roles from the measurement method>

Internet-Draft

Initial Registry

October 2018

[11.4.](#) Output

This category specifies all details of the Output of measurements using the metric.

[11.4.1.](#) Type

<insert name of the output type, raw or a selected summary statistic>

[11.4.2.](#) Reference Definition

<pointer to section/spec where output type/format is defined>

[11.4.3.](#) Metric Units

<insert units for the measured results, and the reference specification>.

[11.4.4.](#) Calibration

<describe the error calibration, a way to indicate that the results were collected in a calibration mode of operation, and a way to report internal status metrics related to calibration, such as time offset>

[11.5.](#) Administrative items

[11.5.1.](#) Status

<current or deprecated>

[11.5.2.](#) Requestor

<name of individual or Internet Draft, etc.>

[11.5.3.](#) Revision

1.0

[11.5.4.](#) Revision Date

YYYY-MM-DD

[11.6.](#) Comments and Remarks

Additional (Informational) details for this entry

Morton, et al.

Expires April 25, 2019

[Page 82]

Internet-Draft

Initial Registry

October 2018

[12.](#) Example RTCP-XR Registry Entry

This section is MAY BE DELETED or adapted before submission.

This section gives an example registry entry for the end-point metric described in [RFC 7003](#) [[RFC7003](#)], for RTCP-XR Burst/Gap Discard Metric reporting.

[12.1.](#) Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

[12.1.1.](#) Identifier

An integer having enough digits to uniquely identify each entry in the Registry.

[12.1.2.](#) Name

A metric naming convention is TBD.

[12.1.3.](#) URI

Prefix urn:ietf:metrics:param:<name>

[12.1.4.](#) Status

current

[12.1.5.](#) Requestor

Alcelip Mornuley

[12.1.6.](#) Revision

1.0

[12.1.7.](#) Revision Date

2014-07-04

[12.1.8.](#) Description

TBD.

Morton, et al.

Expires April 25, 2019

[Page 83]

Internet-Draft

Initial Registry

October 2018

[12.1.9.](#) Reference Specification(s)

[[RFC3611](#)] [[RFC4566](#)] [[RFC6776](#)] [[RFC6792](#)] [[RFC7003](#)]

[12.2.](#) Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters. [Section 3.2 of \[RFC7003\]\(#\)](#) provides the reference information for this category.

[12.2.1.](#) Reference Definition

Packets Discarded in Bursts:

The total number of packets discarded during discard bursts. The measured value is unsigned value. If the measured value exceeds 0xFFFFFD, the value 0xFFFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFFFFFF MUST be reported.

[12.2.2.](#) Fixed Parameters

Fixed Parameters are input factors that must be determined and embedded in the measurement system for use when needed. The values of these parameters is specified in the Registry.

Threshold: 8 bits, set to value = 3 packets.

The Threshold is equivalent to Gmin in [\[RFC3611\]](#), i.e., the number of successive packets that must not be discarded prior to and following a discard packet in order for this discarded packet to be regarded as part of a gap. Note that the Threshold is set in accordance with the Gmin calculation defined in [Section 4.7.2 of \[RFC3611\]](#).

Interval Metric flag: 2 bits, set to value 11=Cumulative Duration

This field is used to indicate whether the burst/gap discard metrics are Sampled, Interval, or Cumulative metrics [\[RFC6792\]](#):

I=10: Interval Duration - the reported value applies to the most recent measurement interval duration between successive metrics reports.

I=11: Cumulative Duration - the reported value applies to the accumulation period characteristic of cumulative measurements.

Senders MUST NOT use the values I=00 or I=01.

[12.3.](#) Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure unambiguous methods for implementations. For the Burst/Gap Discard Metric, it appears that the only guidance on methods of measurement is in [Section 3.0 of \[RFC7003\]](#) and its supporting references. Relevant information is repeated below, although there appears to be no section titled "Method of Measurement" in [\[RFC7003\]](#).

[12.3.1.](#) Reference Method

Metrics in this block report on burst/gap discard in the stream arriving at the RTP system. Measurements of these metrics are made at the receiving end of the RTP stream. Instances of this metrics block use the synchronization source (SSRC) to refer to the separate auxiliary Measurement Information Block [\[RFC6776\]](#), which describes measurement periods in use (see [\[RFC6776\]](#), [Section 4.2](#)).

This metrics block relies on the measurement period in the Measurement Information Block indicating the span of the report.

Senders MUST send this block in the same compound RTCP packet as the Measurement Information Block. Receivers MUST verify that the measurement period is received in the same compound RTCP packet as this metrics block. If not, this metrics block MUST be discarded.

12.3.2. Stream Type and Stream Parameters

Since RTCP-XR Measurements are conducted on live RTP traffic, the complete description of the stream is contained in SDP messages that proceed the establishment of a compatible stream between two or more communicating hosts. See Run-time Parameters, below.

12.3.3. Output Type and Data Format

The output type defines the type of result that the metric produces.

- o Value: Packets Discarded in Bursts
- o Data Format: 24 bits
- o Reference: [Section 3.2 of \[RFC7003\]](#)

12.3.4. Metric Units

The measured results are apparently expressed in packets, although there is no section of [\[RFC7003\]](#) titled "Metric Units".

12.3.5. Run-time Parameters and Data Format

Run-Time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete. However, the values of these parameters is not specified in the Registry, rather these parameters are listed as an aid to the measurement system implementor or user (they must be left as variables, and supplied on execution).

The Data Format of each Run-time Parameter SHALL be specified in this column, to simplify the control and implementation of measurement devices.

SSRC of Source: 32 bits As defined in [Section 4.1 of \[RFC3611\]](#).

SDP Parameters: As defined in [[RFC4566](#)]

Session description v= (protocol version number, currently only 0)

o= (originator and session identifier : username, id, version number, network address)

s= (session name : mandatory with at least one UTF-8-encoded character)

i=* (session title or short information) u=* (URI of description)

e=* (zero or more email address with optional name of contacts)

p=* (zero or more phone number with optional name of contacts)

c=* (connection information--not required if included in all media)

b=* (zero or more bandwidth information lines) One or more Time descriptions ("t=" and "r=" lines; see below)

z=* (time zone adjustments)

k=* (encryption key)

a=* (zero or more session attribute lines)

Zero or more Media descriptions (each one starting by an "m=" line; see below)

m= (media name and transport address)

i=* (media title or information field)

c=* (connection information -- optional if included at session level)

b=* (zero or more bandwidth information lines)

k=* (encryption key)

a=* (zero or more media attribute lines -- overriding the Session

attribute lines)

An example Run-time SDP description follows:

```
v=0

o=jdoe 2890844526 2890842807 IN IP4 192.0.2.5

s=SDP Seminar i=A Seminar on the session description protocol

u=http://www.example.com/seminars/sdp.pdf e=j.doe@example.com (Jane Doe)

c=IN IP4 233.252.0.12/127

t=2873397496 2873404696

a=recvonly

m=audio 49170 RTP/AVP 0

m=video 51372 RTP/AVP 99

a=rtpmap:99 h263-1998/90000
```

[12.4.](#) Comments and Remarks

TBD.

[13.](#) Revision History

This section may be removed for publication. It contains overview information on updates.

This draft replaced [draft-mornuley-ippm-initial-registry](#).

In version 02, [Section 4](#) has been edited to reflect recent discussion on the ippm-list: * Removed the combination of "Raw" and left 95th percentile. * Hanging Indent on Run-time parameters (Fixed parameters use bullet lists and other indenting formats. * Payload format for

measurement has been removed. * Explanation of Conditional delay distribution.

Version 03 addressed Phil Eardley's comments and suggestions in sections [1-4](#). and resolved the definition of Percentiles.

Version 04 * All [section 4](#) parameters reference YANG types for alternate data formats. * Discussion has concluded that usecase(s) for machine parse-able registry columns are not needed.

Version 05 * Revised several Poisson streams to Periodic, sections [4](#) & 5. * Addition of ICMP (ping) metrics in [section 9](#). * First implementation of Passive TCP RTT metrics in [section 10](#).

[14](#). Security Considerations

These registry entries represent no known security implications for Internet Security. Each referenced Metric contains a Security Considerations section.

[15](#). IANA Considerations

IANA is requested to populate The Performance Metric Registry defined in [[I-D.ietf-ippm-metric-registry](#)] with the values defined above.

See the IANA Considerations section of [[I-D.ietf-ippm-metric-registry](#)] for additional requests and considerations.

[16](#). Acknowledgements

The authors thank Brian Trammell for suggesting the term "Run-time Parameters", which led to the distinction between run-time and fixed parameters implemented in this memo, for identifying the IPFIX metric with Flow Key as an example, for suggesting the Passive TCP RTD metric and supporting references, and for many other productive suggestions. Thanks to Peter Koch, who provided several useful suggestions for disambiguating successive DNS Queries in the DNS Response time metric.

The authors also acknowledge the constructive reviews and helpful suggestions from Barbara Stark, Juergen Schoenwaelder, Tim Carey, and participants in the LMAP working group. Thanks to Michelle Cotton for her early IANA review, and to Amanda Barber for answering questions related to the presentation of the registry and accessibility of the complete template via URL.

Internet-Draft

Initial Registry

October 2018

17. References

17.1. Normative References

- [I-D.ietf-ippm-metric-registry]
Bagnulo, M., Claise, B., Eardley, P., and A. Morton,
"Registry for Performance Metrics", Internet Draft (work
in progress) [draft-ietf-ippm-metric-registry](#), 2014.
- [RFC1035] Mockapetris, P., "Domain names - implementation and
specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035,
November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#),
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis,
"Framework for IP Performance Metrics", [RFC 2330](#),
DOI 10.17487/RFC2330, May 1998,
<<https://www.rfc-editor.org/info/rfc2330>>.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way
Delay Metric for IPPM", [RFC 2679](#), DOI 10.17487/RFC2679,
September 1999, <<https://www.rfc-editor.org/info/rfc2679>>.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way
Packet Loss Metric for IPPM", [RFC 2680](#),
DOI 10.17487/RFC2680, September 1999,
<<https://www.rfc-editor.org/info/rfc2680>>.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip
Delay Metric for IPPM", [RFC 2681](#), DOI 10.17487/RFC2681,
September 1999, <<https://www.rfc-editor.org/info/rfc2681>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet:
Timestamps", [RFC 3339](#), DOI 10.17487/RFC3339, July 2002,
<<https://www.rfc-editor.org/info/rfc3339>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation
Metric for IP Performance Metrics (IPPM)", [RFC 3393](#),
DOI 10.17487/RFC3393, November 2002,

Internet-Draft

Initial Registry

October 2018

- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", [RFC 3432](#), DOI 10.17487/RFC3432, November 2002, <<https://www.rfc-editor.org/info/rfc3432>>.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", [RFC 4737](#), DOI 10.17487/RFC4737, November 2006, <<https://www.rfc-editor.org/info/rfc4737>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), DOI 10.17487/RFC5357, October 2008, <<https://www.rfc-editor.org/info/rfc5357>>.
- [RFC5560] Uijterwaal, H., "A One-Way Packet Duplication Metric", [RFC 5560](#), DOI 10.17487/RFC5560, May 2009, <<https://www.rfc-editor.org/info/rfc5560>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", [RFC 5905](#), DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#), DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6049] Morton, A. and E. Stephan, "Spatial Composition of Metrics", [RFC 6049](#), DOI 10.17487/RFC6049, January 2011, <<https://www.rfc-editor.org/info/rfc6049>>.
- [RFC6673] Morton, A., "Round-Trip Packet Loss Metrics", [RFC 6673](#), DOI 10.17487/RFC6673, August 2012, <<https://www.rfc-editor.org/info/rfc6673>>.

[RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", [RFC 6991](#), DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, [RFC 7011](#), DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.

Morton, et al.

Expires April 25, 2019

[Page 90]

Internet-Draft

Initial Registry

October 2018

[RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", [RFC 7323](#), DOI 10.17487/RFC7323, September 2014, <<https://www.rfc-editor.org/info/rfc7323>>.

[RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, [RFC 7679](#), DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.

[RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, [RFC 7680](#), DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

17.2. Informative References

[RFC1242] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", [RFC 1242](#), DOI 10.17487/RFC1242, July 1991, <<https://www.rfc-editor.org/info/rfc1242>>.

[RFC3611] Friedman, T., Ed., Caceres, R., Ed., and A. Clark, Ed., "RTP Control Protocol Extended Reports (RTCP XR)", [RFC 3611](#), DOI 10.17487/RFC3611, November 2003, <<https://www.rfc-editor.org/info/rfc3611>>.

- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", [BCP 108](#), [RFC 4148](#), DOI 10.17487/RFC4148, August 2005, <<https://www.rfc-editor.org/info/rfc4148>>.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), DOI 10.17487/RFC4566, July 2006, <<https://www.rfc-editor.org/info/rfc4566>>.
- [RFC5472] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IP Flow Information Export (IPFIX) Applicability", [RFC 5472](#), DOI 10.17487/RFC5472, March 2009, <<https://www.rfc-editor.org/info/rfc5472>>.
- [RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", [RFC 5477](#), DOI 10.17487/RFC5477, March 2009, <<https://www.rfc-editor.org/info/rfc5477>>.

Morton, et al.

Expires April 25, 2019

[Page 91]

Internet-Draft

Initial Registry

October 2018

- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", [RFC 5481](#), DOI 10.17487/RFC5481, March 2009, <<https://www.rfc-editor.org/info/rfc5481>>.
- [RFC6248] Morton, A., "[RFC 4148](#) and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", [RFC 6248](#), DOI 10.17487/RFC6248, April 2011, <<https://www.rfc-editor.org/info/rfc6248>>.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", [BCP 170](#), [RFC 6390](#), DOI 10.17487/RFC6390, October 2011, <<https://www.rfc-editor.org/info/rfc6390>>.
- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of View", [RFC 6703](#), DOI 10.17487/RFC6703, August 2012, <<https://www.rfc-editor.org/info/rfc6703>>.
- [RFC6776] Clark, A. and Q. Wu, "Measurement Identity and Information Reporting Using a Source Description (SDS) Item and an RTCP Extended Report (XR) Block", [RFC 6776](#),

DOI 10.17487/RFC6776, October 2012,
<<https://www.rfc-editor.org/info/rfc6776>>.

- [RFC6792] Wu, Q., Ed., Hunt, G., and P. Arden, "Guidelines for Use of the RTP Monitoring Framework", [RFC 6792](#), DOI 10.17487/RFC6792, November 2012, <<https://www.rfc-editor.org/info/rfc6792>>.
- [RFC7003] Clark, A., Huang, R., and Q. Wu, Ed., "RTP Control Protocol (RTCP) Extended Report (XR) Block for Burst/Gap Discard Metric Reporting", [RFC 7003](#), DOI 10.17487/RFC7003, September 2013, <<https://www.rfc-editor.org/info/rfc7003>>.
- [RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", [RFC 7594](#), DOI 10.17487/RFC7594, September 2015, <<https://www.rfc-editor.org/info/rfc7594>>.
- [Strowes] Strowes, S., "Passively Measuring TCP Round Trip Times, Communications of the ACM, Vol. 56 No. 10, Pages 57-64", September 2013.

Morton, et al.

Expires April 25, 2019

[Page 92]

Internet-Draft

Initial Registry

October 2018

[Trammell-14]

Trammell, B., "Inline Data Integrity Signals for Passive Measurement, TMA 2014
<https://trammell.ch/pdf/qof-tma14.pdf>", March 2014.

Authors' Addresses

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192

Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
Email: marcelo@it.uc3m.es
URI: <http://www.it.uc3m.es>

Philip Eardley
BT
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

Kevin D'Souza
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 xxxx
Email: kld@att.com