

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

M. Bagnulo
UC3M
A. Morton
AT&T Labs
P. Eardley
BT
October 21, 2013

A(nother) Registry for Performance Metrics
draft-mornulo-ippm-registry-columns-01

Abstract

This memo investigates a scheme to organize registry entries, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, and applicable to all IETF metrics. Three aspects make IPPM metric registration difficult: (1) Use of the Type-P notion to allow users to specify their own packet types. (2) Use of flexible input variables, called Parameters in IPPM definitions, some which determine the quantity measured and others which should not be specified until execution of the measurement. (3) Allowing flexibility in choice of statistics to summarize the results on a stream of measurement packets. Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

This version contains an example registry entry for a passive endpoint metric based on RFC7003, an example active metric entry based on RFC3393 and RFC5481, and an example pure passive metric based on RFC5472. Also, this version *continues* to allow blank entries in columns which have no applicability to a specific metric, or class of metrics. This is preferred to more general registry organization because each column serves as a check-list item and helps to avoid omissions during registration and expert review.

Requirements Language

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

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 <http://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 24, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1.	Introduction	5
1.1.	Background and Motivation	5
2.	Scope	7
3.	Registry Categories and Columns	7
3.1.	Registry Indexes	8
3.1.1.	Element ID	8
3.1.2.	Metric Name	8
3.2.	Metric Definition	9
3.2.1.	Reference Definition	9
3.2.2.	Fixed Parameters	9
3.2.3.	Metric Units	9
3.3.	Method of Measurement	10
3.3.1.	Reference Method	10
3.3.2.	Stream Type and Stream Parameters	10
3.3.3.	Output Type and Data Format	10
3.3.4.	Run-time Parameters and Data Format	11
3.4.	Comments and Remarks	11
4.	Example RTCP-XR Registry Entry	12
4.1.	Registry Indexes	12
4.1.1.	Element ID	12
4.1.2.	Metric Name	12
4.2.	Metric Definition	12
4.2.1.	Reference Definition	12
4.2.2.	Fixed Parameters	12
4.2.3.	Metric Units	13
4.3.	Method of Measurement	13
4.3.1.	Reference Method	13
4.3.2.	Stream Type and Stream Parameters	14
4.3.3.	Output Type and Data Format	14
4.3.4.	Run-time Parameters and Data Format	14
4.4.	Comments and Remarks	16
5.	Example IPPM Active Registry Entry	16
5.1.	Registry Indexes	16
5.1.1.	Element ID	16
5.1.2.	Metric Name	16
5.2.	Metric Definition	16
5.2.1.	Reference Definition	16
5.2.2.	Fixed Parameters	17
5.2.3.	Metric Units	17
5.3.	Method of Measurement	17
5.3.1.	Reference Method	17
5.3.2.	Stream Type and Stream Parameters	17
5.3.3.	Output Type and Data Format	18
5.3.4.	Run-time Parameters and Data Format	18
5.4.	Comments and Remarks	19
6.	Example IPFIX RTT Pair Matching Registry Entry	19

6.1. Registry Indexes	19
6.1.1. Element ID	19
6.1.2. Metric Name	19
6.2. Metric Definition	19
6.2.1. Reference Definition	20
6.2.2. Fixed Parameters	20
6.2.3. Metric Units	20
6.3. Method of Measurement	20
6.3.1. Reference Method	20
6.3.2. Stream Type and Stream Parameters	21
6.3.3. Output Type and Data Format	21
6.3.4. Run-time Parameters and Data Format	21
6.4. Comments and Remarks	21
7. Security Considerations	21
8. IANA Considerations	22
9. Acknowledgements	22
10. References	22
10.1. Normative References	22
10.2. Informative References	23
Authors' Addresses	23

1. Introduction

This memo investigates a scheme to organize registry entries, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, according to their framework [RFC2330]. Three aspects make IPPM metric registration difficult: (1) Use of the Type-P notion to allow users to specify their own packet types. (2) Use of Flexible input variables, called Parameters in IPPM definitions, some which determine the quantity measured and others which should not be specified until execution of the measurement. (3) Allowing flexibility in choice of statistics to summarize the results on a stream of measurement packets. This memo uses terms and definitions from the IPPM literature, primarily [RFC2330], and the reader is assumed familiar with them or may refer questions there as necessary.

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 was intended to become the basis for the columns of an IETF-wide registry of metrics. As we examine the aspects of metric specifications which need to be registered, we will see that none of the existing metric templates fully satisfies the needs of a registry.

The authors of [draft-bagnulo-ippm-new-registry] and [draft-bagnulo-ippm-new-registry-independent] made important contributions to this memo in the registry column structure, and the problem of registry development in general. We also acknowledge input from the authors of [draft-claise-ippm-perf-metric-registry], especially the value of an Element ID and the need for naming conventions.

1.1. Background and Motivation

The motivation for having such registry is to allow a controller to request a measurement agent to execute a measurement using a specific metric. Such request can be performed using any control protocol that refers to the value assigned to the specific metric in the registry. Similarly, the measurement agent can report the results of the measurement and by referring to the metric value it can unequivocally identify the metric that the results correspond to.

There was a previous attempt to define a metric registry RFC 4148 [RFC4148]. However, it was obsoleted by RFC 6248 [RFC6248] because it was "found to be insufficiently detailed to uniquely identify IPPM

metrics... [there was too much] variability possible when characterizing a metric exactly" which led to the RFC4148 registry having "very few users, if any".

Our approach learns from this, by tightly defining each entry in the registry with only a few parameters open for each. The idea is that the entries in the registry represent different measurement tests, whilst the run-time parameters set things like source and destination addresses that don't change the fundamental nature of the test. The downside of this approach is that it could result in an explosion in the number of entries in the registry. We believe that less is more in this context - it is better to have a reduced set of useful metrics rather than a large set of metrics with questionable usefulness. Therefore this document defines that the registry only includes commonly used metrics that are well defined; hence we require both reference specification required AND expert review policies for the assignment of values in the registry.

There are several side benefits of having such a registry. First the registry could serve as an inventory of useful and used metrics, that are normally supported by different implementations of measurement agents. Second, the results of the metrics would be comparable even if they are performed by different implementations and in different networks, as the metric is properly defined.

The registry forms part of a Characterization Plan. It describes various factors that need to be set by the party controlling the measurements, for example: specific values for the parameters associated with the selected registry entry (for instance, source and destination addresses); and how often the measurement is made. The Characterization Plan determines the individual Measurement Instructions that will be communicated to measurement agents, whose task is then to execute the Instruction autonomously.

Measurement Instructions might look something like: "Dear measurement agent: Please start test DNS(example.com) and RTT(server.com,150) every day at 2000 GMT. Run the DNS test 5 times and the RTT test 50 times. Do that when the network is idle. Generate both raw results and 99th percentile mean. Send measurement results to collector.com in IPFIX format". The Characterization Plan depends on the requirements of the controlling party. For instance the broadband consumer might want a one-off measurement made immediately to one specific server; a regulator might want the same measurement made once a day until further notice to the 'top 10' servers; whilst an operator might want a varying series of tests (some of which will be beyond those defined in the registry) as determined from time to time by their operational support system. While the registries defined in this document help to define the Characterization Plan, its full

specification falls outside the scope of this document, and other IETF work as currently chartered.

2. Scope

Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

In this memo, we attempt a combinatoric registry, where all factors that can be reasonably specified ARE specified, and changing even one factor would require a new registry entry (row). It is believed that this exercise can also be instructive for a registry based on independent factors, [draft-bagnulo-ippm-new-registry-independent] but that topic is beyond the scope of this effort.

Entries in the registry must reference an existing RFC or other recognized standard, and are subject to expert review. The expert review must make sure that the proposed metric is operationally useful. This means that the metric has proven to be useful in operational/real scenarios.

3. Registry Categories and Columns

This section briefly describes the categories and columns proposed for the registry, as this is likely to be a topic for discussion and revision. Below, categories are described at the 3.x heading level, and columns are at the 3.x.y heading level. The Figure below illustrates this organization.

Taken as a whole, the entries in the columns give a registered instance of a metric with sufficient specificity to promote comparable results across independent implementations. In other words, a *complete description* of a Metric Instance. Some instances may not require entries in all columns, but this is preferred to more general organization because each column serves as a check-list item and helps to avoid omissions during registration and expert review.

Registry Categories and Columns, shown as

Registry Indexes		Category	
-----		-----	
Element ID	Metric Name	Column	Column
Metric Definition			

Reference Definition	Fixed Parameters	Metric Units	
Method of Measurement			

Reference Method	Stream Type and Param	Output Type	Run-time Param
Comments and Remarks			

3.1. Registry Indexes

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

3.1.1. Element ID

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

3.1.2. Metric Name

A metric naming convention is TBD.

The current guidance from Section 13 of [RFC2330], where Type-P is a feature of all IPPM metric names, is:

"... we introduce the generic notion of a "packet of type P", where in some contexts P will be explicitly defined (i.e., exactly what type of packet we mean), partially defined (e.g., "with a payload of B octets"), or left generic. Thus we may talk about generic IP-type-P-connectivity or more specific IP-port-HTTP-connectivity. Some metrics and methodologies may be fruitfully defined using generic type P definitions which are then made specific when performing

actual measurements. Whenever a metric's value depends on the type of the packets involved in the metric, the metric's name will include either a specific type or a phrase such as "type-P". ..."

Registry entries are a context where Type-P must be defined.

IPPM Metric names have also typically included the stream type, to distinguish between singleton and sample metrics (see [RFC2330] for the definition of these terms).

3.2. 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.

3.2.1. Reference Definition

This entry provides references to relevant sections of the RFC(s) defining the metric, as well as any supplemental information needed to ensure an unambiguous definition for implementations.

3.2.2. 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.

Where referenced metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Fixed Parameters. For example, Fixed Parameters determine most or all of the IPPM Framework convention "packets of Type-P" as described in [RFC2330], such as transport protocol, payload length, TTL, etc.

A Parameter which is Fixed for one Registry entry may be designated as a Run-time Parameter for another Registry entry.

3.2.3. Metric Units

The measured results of a metric must be expressed using some standard dimension or units of measure. This column provides the units (and if possible, the data format, whose specification will simplify both measurement implementation and collection/storage tasks, see the Output Type column below).

When a sample of singletons (see [RFC2330] for definitions of these terms) is collected, this entry will specify the units for each measured value.

3.3. 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.

3.3.1. Reference Method

This entry provides references to relevant sections of the RFC(s) describing the method of measurement, as well as any supplemental information needed to ensure unambiguous interpretation for implementations referring to the RFC text.

3.3.2. Stream Type and Stream Parameters

Principally, two different streams are used in IPPM metrics, Poisson distributed as described in [RFC2330] and Periodic as described in [RFC3432]. Both Poisson and Periodic have their own unique parameters, and the relevant set of values is specified in this column.

Some metrics, such as those intended for passive monitoring or RTCP and RTCP-XR metrics, will not specify an entry for this column.

Each entry for this column contains the following information:

- o Value: The name of the packet stream scheduling discipline
- o Stream Parameters: The values and formats of input factors for each type of stream. For example, the average packet rate and distribution truncation value for streams with Poisson-distributed inter-packet sending times.
- o Reference: the specification where the stream is defined

The simplest example of stream specification is Singleton scheduling, where a single atomic measurement is conducted. Each atomic measurement could consist of sending a single packet (such as a DNS request) or sending several packets (for example, to request a webpage). Other streams support a series of atomic measurements in a "sample", with a schedule defining the timing between each transmitted packet and subsequent measurement.

3.3.3. Output Type and Data Format

For entries which involve a stream and many singleton measurements, a statistic may be specified in this column to summarize the results to a single value. If the complete set of measured singletons is

output, this will be specified here.

Some metrics embed one specific statistic in the reference metric definition, while others allow several output types or statistics.

Each entry in the output type column contains the following information:

- o Value: The name of the output type
- o Data Format: provided to simplify the communication with collection systems and implementation of measurement devices.
- o Reference: the specification where the output type is defined

The output type defines the type of result that the metric produces. It can be the raw results or it can be some form of statistic. The specification of the output type must define the format of the output. Note that if two different statistics are required from a single measurement (for example, both "Xth percentile mean" and "Raw"), then a new output type must be defined ("Xth percentile mean AND Raw").

3.3.4. 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).

Where metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Run-Time Parameters.

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

Examples of Run-time Parameters include IP addresses, measurement point designations, start times and end times for measurement, and other measurement-specific information.

3.4. Comments and Remarks

Besides providing additional details which do not appear in other categories, this open Category (single column) allows for unforeseen

issues to be addressed by simply updating this Informational entry.

4. Example RTCP-XR Registry Entry

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

4.1. Registry Indexes

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

4.1.1. Element ID

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

4.1.2. Metric Name

A metric naming convention is TBD.

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. Section 3.2 of [RFC7003] provides the reference information for this category.

4.2.1. 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 0xFFFFF MUST be reported.

4.2.2. 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.

4.2.3. Metric Units

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

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. 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].

4.3.1. 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.

4.3.2. 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.

4.3.3. 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]

4.3.4. 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 10.47.16.5
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 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
m=audio 49170 RTP/AVP 0
```

```
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
```

4.4. Comments and Remarks

Besides providing additional details which do not appear in other categories, this open Category (single column) allows for unforeseen issues to be addressed by simply updating this Informational entry.

5. Example IPPM Active Registry Entry

This section gives an example registry entry for the active metric described in [RFC3393], on Packet Delay Variation.

5.1. Registry Indexes

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

5.1.1. Element ID

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

5.1.2. Metric Name

A metric naming convention is TBD.

One possibility based on IPPM's framework is:

```
IP-UDP-One-way-pdv-95th-percentile-Poisson
```

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

See sections 2.4 and 3.4 of [RFC3393]. Singleton delay differences measured are referred to by the variable name "ddT".

5.2.2. Fixed Parameters

Where metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Fixed Parameters.

- o F, a selection function defining unambiguously the packets from the stream selected for the metric. See section 4.2 of [RFC5481] for the PDV form.
- o L, a packet length in bits. L = 200 bits.
- o Tmax, a maximum waiting time for packets to arrive at Dst, set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost). Tmax = 3 seconds.
- o Type-P, as defined in [RFC2330], which includes any field that may affect a packet's treatment as it traverses the network. The packets are IP/UDP, with DSCP = 0 (BE).

5.2.3. Metric Units

See section 3.3 of [RFC3393] for singleton elements.

[RFC2330] recommends that when a time is given, it will be expressed in UTC.

The timestamp format (for T, Tf, etc.) is the same as in [RFC5905] (64 bits) and is as follows: the first 32 bits represent the unsigned integer number of seconds elapsed since 0h on 1 January 1900; the next 32 bits represent the fractional part of a second that has elapsed since then.

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

See section 2.6 and 3.6 of [RFC3393] for singleton elements.

5.3.2. Stream Type and Stream Parameters

Poisson distributed as described in [RFC2330], with the following Parameters.

- o λ , a rate in reciprocal seconds (for Poisson Streams). $\lambda = 1$ packet per second
- o Upper limit on Poisson distribution (values above this limit will be clipped and set to the limit value). Upper limit = 30 seconds.

5.3.3. Output Type and Data Format

See section 4.3 of [RFC3393] for details on the percentile statistic.

The percentile = 95.

Data format is a 32-bit unsigned floating point value.

Individual results (singletons) should be represented by the following triple

- o T_1 and T_2 , times as described below in the Run-time parameters section.
- o ddT as defined in section 2.4 of [RFC3393]

if needed. The result format for ddT is *similar to* the short format in [RFC5905] (32 bits) and is as follows: the first 16 bits represent the *signed* integer number of seconds; the next 16 bits represent the fractional part of a second.

5.3.4. Run-time Parameters and Data Format

Where metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Run-Time Parameters. In related registry entries, some of the parameters below may be designated as Fixed Parameters instead.

- o Src , the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)
- o Dst , the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)
- o T , a time (start of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])
- o T_f , a time (end of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])
- o T_l , the wire time of the first packet in a pair, measured at $MP(Src)$ as it leaves for Dst (64-bit NTP Timestamp Format, see

section 6 of [RFC5905]).

- o T2, the wire time of the second packet in a pair, measured at MP(Src) as it leaves for Dst (64-bit NTP Timestamp Format, see section 6 of [RFC5905]).
- o I(i),I(i+1), i >=0, pairs of times which mark the beginning and ending of the intervals in which the packet stream from which the measurement is taken occurs. Here, I(0) = T0 and assuming that n is the largest index, I(n) = Tf (pairs of 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).

5.4. Comments and Remarks

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 IPDV.

6. Example IPFIX RTT Pair Matching Registry Entry

This section gives an example registry entry for the passive metric described in section 2.5.2.1 of [RFC5472], for Round-Trip Time (RTT) Measurements with Packet Pair Matching (Single-Point).

6.1. Registry Indexes

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

6.1.1. Element ID

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

6.1.2. Metric Name

A metric naming convention is TBD.

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. Section 2.5.2.1 of [RFC5472] provides the reference information for this category.

6.2.1. Reference Definition

Observations of both directions are required to correlate request/response packet pairs.

Pair matching techniques are described in [Brow00].

6.2.2. 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.

Protocol (Pair Type): TCP (SYN/SYN_ACK)

Note: other possibilities are DNS, ICMP, SNMP or TCP (DATA/ACK), discussed in [Brow00].

6.2.3. Metric Units

The measured results are expressed in microseconds, which follows the format of Information Elements per observed packet, see section 8.4.3 of[RFC5477] titled "observationTimeMicroseconds".

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 the TCP(SYN/SYN_ACK) RTT metric, the guidance on methods of measurement is in slides 12 and 15 of [Brow00].

Recognition of request response pairs is a REQUIRED function, as is the correlation of data from both directions of transmission, see section 2.5.2.1 of [RFC5472].

The method requires the collection of the following Information Elements per packet:

- o Packet arrival time: observationTimeMicroseconds, see section 8.4.3 of[RFC5477]
- o TCP header: ipPayloadPacketSection, see section 8.5.2 of[RFC5477]

6.3.2. Stream Type and Stream Parameters

Since IPFIX passive Measurements are conducted on live/production network traffic, the measurement methods rely on user-generated packet flows. Such flows are not described in this column.

6.3.3. Output Type and Data Format

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

- o Value: RTT in microseconds
- o Data Format: (There may be some precedent to follow here, but otherwise use 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).
- o Reference: Section 2.5.2.1 of [RFC5472]

6.3.4. 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 list of Run-time parameters is not specified for purely passive metrics, as there are infinite possibilities.

A likely Run-time parameter is the Destination host, which may be given as a Fully-Qualified Domain Name as done in [Brow00], or an IP address of the host (32-bit value for IPv4, 128-bit value for IPv6).

6.4. Comments and Remarks

Additional (Informational) details for this entry, from [Brow00]:

Can't get RTT for every packet, only those which are ACKed.

Overlapping packets (resent) are counted as lost, but not queued. This means the first copy of resent packets are used for RTTs, giving a high RTT estimate.

7. Security Considerations

This registry has no known implications on Internet Security.

8. IANA Considerations

Metrics previously defined in IETF were registered in the IANA IPPM METRICS REGISTRY, however this process was discontinued when the registry structure was found to be inadequate, and the registry was declared Obsolete [RFC6248].

The form of metric registration will be finalized in the future, and no IANA Action is requested at this time.

9. Acknowledgements

The author thanks Brian Trammell for suggesting the term "Run-time Parameters", which led to the distinction between run-time and fixed parameters implemented in this memo.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737,

November 2006.

- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarez, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

10.2. Informative References

- [Brow00] Brownlee, N., "Packet Matching for NeTraMet Distributions", March 2000.
- [RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.
- [RFC5472] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IP Flow Information Export (IPFIX) Applicability", RFC 5472, March 2009.
- [RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, April 2011.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.
- [RFC7003] Clark, A., Huang, R., and Q. Wu, "RTP Control Protocol (RTCP) Extended Report (XR) Block for Burst/Gap Discard Metric Reporting", RFC 7003, September 2013.

Authors' Addresses

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>

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/>

Philip Eardley
British Telecom
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

