

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 30, 2014

M. Bagnulo
UC3M
A. Morton
AT&T Labs
P. Eardley
BT
September 30, 2013

A Registry for Performance Metrics
draft-mornulo-ippm-registry-00

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.

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 March 22, 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	4
1.1.	Background and Motivation	4
2.	Scope	5
3.	Registry Columns and Sub-Columns	5
3.1.	Metric ID	6
3.2.	Metric Name	6
3.3.	Metric Description	7
3.4.	Method of Measurement	7
3.4.1.	Reference Method	7
3.4.2.	Fixed Parameters	7
3.4.3.	Schedule	8
3.4.4.	Output Type	8
3.4.5.	Run-time Parameters	9
3.5.	Metric Units	10
3.6.	Measurement Point	10
3.7.	Timing	10
3.8.	Other	10
4.	Example of allocation	10
4.1.	UDP latency metric	10
5.	Security Considerations	13
6.	IANA Considerations	13
7.	Acknowledgements	13
8.	References	13
8.1.	Normative References	13
8.2.	Informative References	14
	Authors' Addresses	14

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.

This registry is based on the template defined in [RFC6390] expanded with further details to fully cover 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 a couple of side benefits of having such 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.

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 Columns and Sub-Columns

This section describes the columns and sub-columns proposed for the registry. Below, columns are described at the 3.x heading level, and sub-columns are at the 3.x.y heading level. The Figure below illustrates this organization.

Taken as a whole, each entries (row) in the registry gives 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 sub-columns, but this is preferred to more general organization because each sub-column serves as a check-list item and helps to avoid omissions during registration and expert review. The columns are extracted directly from [RFC6390] while the sub-columns provide additional detailed required for each column.

ID	Name	Description	Method	Units	Measurement Points	Timing	Other

Figure 1: Registry columns

We describe the content of each of the columns next, some of them contain sub-columns.

3.1. Metric ID

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

3.2. Metric Name

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

Based on this, the metric name is composed in the following way:

P_type-Descriptive_name-Schedule-Output-Other, where:

- o P-type is a text describing the P-type selected
- o Descriptive_name describes the nature of the metric
- o The schedule describes the so-called stream type
- o The output describes the expected output of the metric, in particular, the type of statistic which is outputted, if it is one.
- o Other, describes other consideration that affects the nature of the metric, for example the presence or absence of cross traffic

3.3. Metric Description

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.

3.4. Method of Measurement

This column is composed by the following sub-columns:

Method				
Reference Method	Fixed Parameters	Schedule	Output Type	Run-time Param

3.4.1. Reference Method

This sub-column provides references to relevant sections of the specifications or 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.

3.4.2. Fixed Parameters

In the case that the metric is defined as a more specific instance of a broader metric defined in the specification pointed in the "Reference method" column, this is done by defining as fixed some of the open parameters defined in the broader metric. If this should be the case of the specific entry in the registry, this sub-column specifies the values of these parameters in the Registry.

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

3.4.3. Schedule

Principally, two different schedules 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 sub-column contains the following information:

- o Value: The name of the packet stream scheduling discipline
- o Schedule 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

Schedule		
Value	Schedule Parameters	Reference

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.

3.4.4. Output Type

For some entries, 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

Output type		
Value	Data format	Reference

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").

3.4.5. Run-time Parameters

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 actual 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.

3.5. 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.

3.6. Measurement Point

Measurement Point(s) with potential Measurement Domain: A pointer to the specification that defines whether the metric is specific to a given measurement point or measurement domain. A canonical reference path is defined in [I-D.ietf-ippm-lmap-path].

3.7. Timing

A pointer to the specification where the acceptable range of timing intervals or sampling intervals are defined, if any.

3.8. Other

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 of allocation

In this section we provide a few example of allocations.

4.1. UDP latency metric

The registry entry for for the Xth percentile mean of the UDP latency using a Poisson stream of packets would look like this:

ID: 344 (for example, typically assigned by IANA)

Name: UDP-Latency-Poisson-Xth_percentile_mean.

Description: This metric is a specific instance of the Round trip metric defined in RFC2681 and it measures the Xth percentile mean of the UDP latency of a Poisson stream of packets.

Method:

Reference Method: The methodology for this metric is defined as Type-P-Round-trip-Delay-Poisson-Stream in RFC 2681.

Fixed Parameters:

P-Type:

IPv4 header values:

DSCP: set to 0

TTL set to 255

Protocol: Set to 17 (UDP)

UDP header values: Checksum: the checksum must be calculated

Payload

Sequence number: 8-byte integer

Timestamp: 8 byte integer. Expressed as 64-bit NTP timestamp as per section 6 of RFC 5905

No padding

Timeout: 3 seconds

Schedule:

Value: Poisson

Schedule Parameters:

lambda: the parameter defining the Poisson distribution. Lambda is the mean number of distinct measurements per second in the sample.

T0: time to begin a test

Tf: time to end a test

T0 and Tf are both in seconds and use the date (yyyy-mm-dd) and NTP 64 bit timestamp. T0 includes any control handshaking before the test stream or singleton. Tf is the time the last test data is sent. As a result, we have that the time when test devices may close the test socket is Tf + Waiting Time (the time to wait before declaring a packet lost is fixed for each metric) and the Total duration of the test: Tf - T0 + Waiting Time

Reference: The Poisson scheduling is defined in section 11.1.1 of RFC 2330

Output Type

Value: Xth percentile mean

Data format:

Reference:

Run-time Param

Source IP Address

Destination IP Address

Source UDP port

Destination UDP port

Initial time T0

end time Tf

Rate lambda

X

Units: milliseconds

Measurement Points: The metric is not specific to any particular measurement point.

Timing: between microseconds and seconds

Other

5. Security Considerations

This registry has no known implications on Internet Security.

6. 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.

7. 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.

8. References

8.1. 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.

8.2. Informative References

- [I-D.ietf-ippm-lmap-path] Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and A. Morton, "A Reference Path and Measurement Points for LMAP", draft-ietf-ippm-lmap-path-00 (work in progress), July 2013.
- [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.
- [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.

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