

Registry for Performance Metrics

draft-ietf-ippm-metric-registry-08

M. Bagnulo, B. Claise, P. Eardley,

A. Morton, A. Akhter

Overall Registry Concept

- Problem: How can we specify with Precision the Metrics and Methods to Implement and Use?
 - Many Standardized Metrics with similar names
 - Registry enables all parties to be sure they're talking about the same Metric
 - Flexibility and customization of Generic Metrics seen as an advantage in standards development
 - Methods allow variables, system issues out-of-scope
- Provide Unique ID and detailed exposition
 - Raise the bar from Standard to Registered Metrics
 - (How do we do that? Read on...)

Full Registry Concept & Format

- draft-ietf-ippm-metric-registry-08
- Each entry in the registry is a row
 - Series of columns
 - Typically ~1 column may be Not Applicable
 - Clustered in categories
- Each row is indexed by ID
 - 16 bit flat identifier
 - With associated name (i-d defines naming convention)
 - Auto-generate URI (pre-pend urn:ietf:metric: to name)
 - Auto-generate URL (location of text file with registry entry)
- Control & report protocols use URI
- Next slide shows category /column headings
 - Layout is purely presentational (slide not wide enough, neither is anyone's screen, which is why the text file presentation is available)

Revised Categories & Columns

Category	Column
----------	--------	-------

Summary	ID	Name	URIs	Description	Ref	Change Ctrl	Ver
---------	----	------	------	-------------	-----	-------------	-----

Metric definition	Reference	Fixed parameters
-------------------	-----------	------------------

Method of measurement	Ref. Meth. (eg Section 3 of RFC XXXX)	Packet stream generation (active tests)	Traffic filter (passive tests)	Sampling distribution (for traffic filter)	Run-Time Parameter(s) (eg.MPaddress)	Role(s) (eg sender)
-----------------------	--	--	-----------------------------------	---	---	------------------------

Maybe a lot of info (~sub-columns)

Don't change nature of Method

Output	Type	Reference Method	Units	Calibration
--------	------	------------------	-------	-------------

Admin info	Status	Requestor	Revision #	Date
------------	--------	-----------	------------	------

Comments	Full history
----------	--------------	-------

Registry Draft Updates

- Broader set of Name components in the SubTypeMethod element (and elsewhere)
 - ICMP
 - IP
 - DSCPxx (where xx is replaced by a decimal code point)
 - UDP
 - TCP
 - Poisson
 - Periodic
 - PayloadxxxxB (where xxxx is replaced by an integer)
 - SustainedBurst
 - StandingQueue
 - @@@@<add others from MBM draft?
- Hyphen as Separation Character for SubTypeMethods

More Updates

- URI
 - urn:ietf:metric: (was urn:ietf:param:ippm:metric)
- Expanded IANA Section
 - 7 Columns and URL to HTML-ized Template
 - 2 forms of extension: New Categories/Columns, Name Element sets and SubTypeMethods
- 3 New Columns
 - Calibration supported for some metrics
 - Approach involves Loop-Back test and Clk offset

Naming: Buenos Aires Scheme

MetricType_Method_SubTypeMethod_ ...
Spec_Units_Output

MetricType: RTDelay, OwDelay, OwPDV

Method: Active, Passive, HybridType1

SubTypeMethod: ICMP, IP, UDP, Poisson, Periodic (one or more sub-types)

Spec: RFCXXXsecY (containing the approved entry)

Units: Seconds, RatioPercent,

Output: Singleton, Minimum, Maximum, Mean, Median, StdDev, 95Percentile, 99Percentile

Note: The examples above are candidates for a sub-registry, if IANA prefers that approach.

Revised Elements of BA names

- RTDelay_Active_IP-UDP-Poisson_RFCXXXXsecY_Seconds_95Percentile
 ^^^^^^^^^^^^^^
- OWPDV_Active_IP-UDP-Poisson_RFCXXXXsecY_Seconds_95Percentile
- RTDNS_Active_IP-UDP-Poisson_RFCXXXXsecY_Seconds_95Percentile
- OWDelay_Active_IP-UDP-Poisson-Payload250B_RFCXXXXsecY_Seconds_<statistic>
- OWDelay_Active_IP-UDP-Periodic-Payload142B_RFCXXXXsecY_Seconds_<statistic>
where <statistic> is currently one of
 - Mean, Minimum, Maximum, StdDev, 95Percentile

To Do (?)

- URN agreement?
- Do the Name Element Sets cover Passive well-enough?
- Suggest to maintain the Name element sets as a part of the draft till done, then request IANA make it a sub-registry.

Initial Performance Metric Registry Entries Part 2: MBM

draft-morton-ippm-mbm-registry-00

A. Morton, M. Mathis

A Test of the Registry Design

- Brian/Chair: How can Speed be specified as a Registry Entry, using Model-Based Metrics?
 - Should be part of the initial registry...
 - Name:
- `OWMBM_Active_IP-TCP-SustainedBurst_RFCXXXXsecY_Enumerated_PFI`
- Sustained Burst is described in 8.5.1 of [MBM]
- This test describes stream conditions to evaluate a `target_rate`, at target RTT and MTU
- Loss and RTT are the primary measurements

Was the Registry Design efficient?

- Packet Stream Generation Parameters:
 - **paced_bursts** Send bursts on a timer. Specify any 3 of: average data rate, packet size, burst size (number of packets) and burst headway (burst start to start).
- Many Run-time Parameters:
 - Targets(rate,RTT,MTU), derating, SPRT params
- Additional parameters calculated in the model
- Primary Output: Pass/Fail/Inconclusive, or PFI
- Also Output Loss ratio and Mean RTT results
- More to do (data formats), but short answer: YES

Next Steps

- Combine with Part 1
- Add remaining data formats, etc.

Initial Performance Metric Registry Entries

~~draft-mornuley-ippm-initial-registry-01~~
~~draft-morton-ippm-initial-registry-0,1,2,3,4~~
draft-ietf-ippm-initial-registry-02

A. Morton, M. Bagnulo, P. Eardley, K. D'Souza

Feedback on the Registry Contents

- Seeking feedback on the current contents, and what else the WG (and guests today) want
- Sections 4 and 5 of version 02 has been updated to changing format (IANA comments)
 - New Sections
 - New Name conventions

Discussion and comments

- Loss (obviously, but a simple way to add it?)
- Traceroute
 - Many methods of measurement
- ICMP Echo Request/Reply (wanted?)
 - Yes

Next Steps

- Discuss and close Open Issues (Any?)
- Update sections (>5) with registry entries.

BACKUP

How do I get a registry entry?

- Submit request to IANA, with columns filled in
 - Likely prior review in WG
- Review by performance metric experts
 - If necessary, work on improvements with requester
 - Does the proposed registry entry clearly define the metric & method of measurement?
 - Is it different from existing registry entries?
 - Is it operationally useful (significant industry interest or been deployed)?
- IANA adds to registry
- Similar process for revisions
 - Must be backwards compatible (eg editorial)
 - Otherwise create a new metric (& maybe deprecate old one)

Names, identifiers and URIs

- We keep identifiers, names and we automatically generate URIs
 - Identifiers are flat 16-bit integers
 - Names are unique within the registered metrics
 - URIs are generated by prepending `urn:ietf:params:performance:metric` to the name
- Also, a URL to a text file containing the Registry Entry

End Review, now some Entries

4. UDP Round-trip Latency Registry Entry

4.1. Summary

4.1.1. ID (Identifier)

4.1.2. Name

4.1.3. URI

4.1.4. Description

4.2. Metric Definition

4.2.1. Reference Definition

4.2.2. Fixed Parameters .

4.3. Method of Measurement

4.3.1. Reference Method

4.3.2. Packet Generation Stream

4.3.3. Traffic Filtering (observation) Details

4.3.4. Sampling Distribution

4.3.5. Run-time Parameters and Data Format

4.3.6. Roles

4.4. Output

4.4.1. Type

4.4.2. Data Format

4.4.3. Reference

4.4.4. Metric Units

4.5. Administrative items

4.5.1. Status

4.5.2. Requestor (keep?)

4.5.3. Revision

4.5.4. Revision Date

4.6. Comments and Remarks

passive ex: <https://tools.ietf.org/html/draft-mornulo-ippm-registry-columns-01#section-6>

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-value sample. Note that terms such as singleton and sample are defined in Section 11 of [RFC2330].

Note that although the 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).

4.2.2 Fixed Parameters

Type-P:

- o IPv4 header values:

- * DSCP: set to 0

- * TTL set to 255

- * Protocol: Set to 17 (UDP)

- o UDP header values:

- * Checksum: the checksum must be calculated

- o Payload

- * Sequence number: 8-byte integer

- * Timestamp: 8 byte integer. Expressed as 64-bit NTP timestamp as per section 6 of RFC 5905 [RFC5905]

- * No padding (total of 9 bytes)

Timeout, Tmax: 3 seconds

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 Timeout defined under Fixed Parameters.

The method requires sequence numbers or other send-order information to be retained at the Src or included with each packet to disambiguate packet reordering if it occurs. Sequence number is part of the payload described under Fixed Parameters.

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 shall be included in the method of measurement for this metric.

4.3.5 Run-time Parameters and Data Format

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

- 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 T0, a time (start of measurement interval, 128-bit NTP Date Format, see section 6 of [RFC5905]). When T0 is "all-zeros", a start time is unspecified and Tf is to be interpreted as the Duration of the measurement interval.
- o Tf, a time (end of measurement interval, 128-bit NTP Date Format, see section 6 of [RFC5905]), interpreted as the Duration of the measurement interval.
- o $1/\lambda$, average packet rate (for Poisson Streams). ($1/\lambda = 1$ packet per second, if fixed)
- o Upper limit on Poisson distribution (values above this limit will be clipped and set to the limit value). (if fixed, Upper limit = 30 seconds.)

4.3.5 Run-time Parameters and Data Format

(continued)

The format for $1/\lambda$ and Upper limit of Poisson Dist. are the short format in [RFC5905] (32 bits) and is as follows: the first 16 bits represent the integer number of seconds; the next 16 bits represent the fractional part of a second.

>>> should Poisson run-time params be fixed instead? probably yes if modeling a specific version of MBA tests.

MORE QUESTIONS -----

>>> Should we require that each Registry entry have a SINGLE output Format and Statistic ?

(now, the answer is yes)

>>> Should we require that each Registry entry specify the Test Protocol used to collect the metric ?

(seems impractical, MUCH duplication)

>>> Current Entries are Detailed. A kind of roadmap to IPPM Literature. Should we retain this practice (at the risk of non-equivalent metrics)? If you were implementing, would you find this detail helpful?

Section

Example Registry Entry Names: