ALTO Performance Metrics

draft-ietf-alto-performance-metrics-12

Luis Miguel Contreras Murillo
Sabine Randriamasy
Y. Richard Yang
D. Dhody
Young Lee
Qin Wu

July 27, 2020

IETF 108, Virtual Meeting

Outline

- Updates from v10-v12: address discussions and reviews at interim meeting in April 2020
  - How to choose types of metrics
  - Conforming to RFC6390 (how much details to specify)
  - How to handle different statistics of the same metric
  - How to convey freshness of metric values

Recall of Key WG Decision

- ALTO provides *guidance*, not measurement framework
- There can be ***multiple*** types of guidance; we chose 4 types, but how to specify the type

RFC8571

```
2.  Link Attribute TLVs for TE Metric Extensions

   The following new Link Attribute TLVs are defined:

oint                    Value
-------------------------------------------------
                        Unidirectional Link Delay

                        Min/Max Unidirectional Link Delay

                        Unidirectional Delay Variation

                        Unidirectional Link Loss

                        Unidirectional Residual Bandwidth

                        Unidirectional Available Bandwidth

                        Unidirectional Utilized Bandwidth
```

https://www.sprint.net/sla_performance.php

SLA Performance

To view current IP network performance, visit IP Network Performance Map

Choose a Network: SprintLink | Global MPLS

Name	Metric	Committed Value	Jun 2018	Jul 2018	Aug 2018	Sep 2018	Oct 2018	Nov 2018	Dec 2018	Jan 2019	Feb 2019	Mar 2019	Apr 2019	May 2019	Jun 2019
INTRA-REGION															
North America	Backbone Delay	55.00 ms	34.18 ms	34.16 ms	34.20 ms	34.25 ms	34.26 ms	34.22 ms	34.16 ms	34.00 ms	34.14 ms	33.82 ms	34.09 ms	34.64 ms	34.55 ms
	Packet Loss	0.30 %	0.0050 %	0.0219 %	0.0052 %	0.0060 %	0.0010 %	0.0044 %	0.0007 %	0.0021 %	0.0010 %	0.0086 %	0.0132 %	0.0077 %	0.0104 %
	Data Delivery Rate	99.70 %	99.9950 %	99.9781 %	99.9948 %	99.9940 %	99.9990 %	99.9956 %	99.9993 %	99.9979 %	99.9990 %	99.9914 %	99.9868 %	99.9923 %	99.9896 %
	Jitter	2 ms	0.0031 ms	0.0005 ms	0.0002 ms	0.0001 ms	0.0002 ms	0.0003 ms	0.0002 ms	0.0001 ms	0.0004 ms	0.0017 ms	0.0064 ms	0.0067 ms	0.0098 ms
Europe	Backbone Delay	45.00 ms	17.35 ms	17.46 ms	17.50 ms	17.26 ms	17.45 ms	17.34 ms	17.53 ms	17.57 ms	17.23 ms	17.13 ms	17.15 ms	17.18 ms	17.36 ms
	Packet Loss	0.30													
	Data Delivery Rate	99.7(													
	Jitter	2 ms													
Asia	Backbone Delay	105.(													
	Packet Loss	0.30													
	Data Delivery Rate	99.7(													
	Jitter	2 ms													
Europe to North America	Backbone Delay	95.(													
	Packet Loss	0.30													
	Data Delivery Rate	99.7(													
	Jitter	2 ms													
Japan to North America	Backbone Delay	130.(													
	Packet Loss	0.30													
	Data Delivery Rate	99.7(													
	Jitter	2 ms													
Hong Kong to North America	Backbone Delay	190.(													
	Packet Loss	0.30													
	Data Delivery Rate	99.7(													
	Jitter	2 ms													
Korea to North America	Backbone Delay	190.(													
	Packet Loss	0.30													
	Data Delivery Rate	99.70 %	99.9975 %	99.9980 %	99.9970 %	99.9995 %	99.9927 %	99.9855 %	99.9945 %	99.9993 %	99.9989 %	99.9992 %	99.9989 %	99.9968 %	99.9984 %
	Jitter	2 ms	0.0019 ms	0.0023 ms	0.0003 ms	0.0000 ms	0.0004 ms	0.0012 ms	0.0003 ms	0.0006 ms	0.0002 ms	0.0000 ms	0.0016 ms	0.0006 ms	0.0005 ms
	Backbone Delay	250.00 ms	187.77 ms	191.62 ms	191.94 ms	192.00 ms	197.97 ms	191.82 ms	194.65 ms	191.96 ms	192.00 ms	191.97 ms	191.99 ms	185.13 ms	191.91 ms

NETWORK AVERAGES

Averages

Monthly Network Averages	Target Values	Observed Values: June	May	April
U.S. Network Averages				
Roundtrip Latency	< 37 ms	30.9	30.9	30.6
Roundtrip Loss*	< 0.05%	0.00%	0.01%	0.00%
Network Reliability	> 99.95%	99.9980%	N/A%	100.0000%
Network Jitter	< 1 ms	0.57	0.56	0.56

*Loss% is (100 - Data Delivery%)

http://ipnetwork.bgtmo.ip.att.net/pws/network_delay.html

Import vs estimation

- Section 2.1:
  - The "estimation" category indicates that the value of the metric is computed through an estimation process. An ALTO server may compute "estimation" values by retrieving and/or aggregating information from routing protocols (e.g., [RFC8571]) and traffic measurement management tools (e.g., TWAMP [RFC5357]), with corresponding operational issues.
  - **A particular type of "estimation is direct "import"**, which indicates that the value of the metric is imported directly from a specific existing protocol or system. Specifying "import" as source instead of the more generic "estimation" may allow better tracing of information flow. For an "import" metric, it is RECOMMENDED that the “parameters" field provides details to the system from which raw data is imported. In particular, one may notice that the set of end-to-end metrics defined in Table 1 has large overlap with the set defined in [RFC8571], in the setting of IGP traffic engineering performance metrics for each link (i.e., unidirectional link delay, min/max unidirectional link delay, unidirectional delay variation, unidirectional link loss, unidirectional residual bandwidth, unidirectional available bandwidth, unidirectional utilized bandwidth). Hence, an ALTO server may use "import" to indicate that its end-to-end metrics are computed from link metrics imported from [RFC8571].
  - There can be overlap in deciding the cost-source category. **It is the operator of an ALTO server who chooses the category. If a metric does not include a "cost-source" value, the application MUST assume that the value of "cost-source" is the most generic "estimation".**

Conforming to [RFC6390]

- Section 2.
  - "When defining the metrics in Table 1, this document **considers** the guidelines specified in [RFC6390], which requires fine-grained specification of (i) Metric Name, (ii) Metric Description, (iii) Method of Measurement or Calculation, (iv) Units of Measurement, (v) Measurement Points, and (vi) Measurement Timing. In particular, for each metric, this document defines (i) Metric Name, (ii) Metric Description, and (iv) Units of Measurement. The Measurement Points are always specified by the specific ALTO services; for example, endpoint cost service is between the two end points.

    On the other hand, to be able to use coarse-grained information such as routing system information (e.g., [RFC8571]), which may not provide fine-grained information such as (iii) Method of Measurement or Calculation and (vi) Measurement Timing, this document provides context information to indicate the source of information and hence available metric details."

How to Handle Statistics of Same Metric

- "Q2: How to handle statistics of the same metric
  - statistics: min, max, x-percentile, avg, …
  - related to Martin Duke comment"
- Updated Sec. 2.2:
  - "The measurement of a performance metric often yields a set of samples from an observation distribution ([Prometheus]), instead of a single value. This document considers that the samples are aggregated as a statistic when reported. Hence, each performance metric's identifier should indicate the statistic (i.e., an aggregation operation), to become <metric-base-identifier>-<stat>"
  - Compared with other common statistics such as those in Prometheus [1]
    - They use quantile but we decide to keep percentile
  - Make complete statistics: add stddev, stdvar
    => percentile, min, max, median, mean, stddev, stdvar

[1] https://prometheus.io/docs/concepts/data_model/

Freshness of Metrics

- “Despite the introduction of the additional cost-context information, there is not a built-in field to indicate the timestamps of the data used to compute a metric. To indicate this attribute, the ALTO server SHOULD return HTTP "Last-Modified", to indicate the freshness of the data used to compute the performance metrics. If the ALTO client obtains updates through an incremental update mechanism (e.g., RFC editor: Fix the RFC number when available. [ALTO SSE]), the client SHOULD assume that the metric is computed using a snapshot at the time that is approximated by the receiving time.”

Next Step

- The authors are quite happy about the documents, with some very minor edits needed, e.g.,
  - Adding Content-Length in the examples after no more changes