# ALTO Performance Metrics

## draft-ietf-alto-performance-metrics-10

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

April 21, 2020

ALTO Interim Meeting

# Outline

- Updates from v08-v10
  - Systematic refinement of the acquisition context
  - Refinement of detail of individual performance metrics
  - General text edit
- Remaining issues to be discussed

# Recall of Key WG Decision

- ALTO provides *guidance*, not measurement framework
- There can be ***multiple*** types of guidance

RFC8571

2. Link Attribute TLVs for TE Metric Extensions

The following new Link Attribute TLVs are defined:

| 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

# v08-v10 Change: Cost-Context Structure to Specify Context the Metric Value is Acquired

- Consist of two members, “cost-source”, and “parameters”
  - cost-source: high-level category of type of guidance
  - parameters: structure to lower-level, additional information

```
object {
        CostMetric   cost-metric;
        CostMode     cost-mode;
        [CostContext cost-context;]
        [JSONString  description;]
       } CostType;
```

```
object {
 JSONString    cost-source;
 [JSONValue    parameters;]
} CostContext;
```

# v08-v10 Change: Cost-Context Structure to Specify Context the Metric Value is Acquired

- Finalized 4 types of “cost-source”

> The "cost-source" field of the "cost-context" field MUST be one of four category values: "nominal", "sla", "import", and "estimation". 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 "estimation". "cost-context" will not be used as a key to distinguish among performance metrics. Hence, an ALTO information resource SHOULD NOT announce multiple CostType with the same "cost-metric" and "cost-mode". They can be placed into different information resources.

# v08-v10 Change: Cost-Context Structure to Specify Context the Metric Value is Acquired

- Give descriptions to the 4 types of “cost-source”

> The "nominal" category indicates that the value of the metric is statically configured by the underlying devices. Not all metrics have reasonable "nominal" values. For example, throughput can have a nominal value, which indicates the configured transmission rate of the devices; latency typically do not have a nominal value.
>
> The "sla" category indicates that the value of the metric is derived from some commitment which this document refers to as service-level agreement (SLA). Some operators also used terms such as "target" or "committed" values. For a "sla" metric, it is RECOMMENDED that the "parameters" field provides a link to the SLA definition.
>
> The "import" category indicates that the value of the metric is derived from importing from a specific existing protocol or system. For an "import" metric, it is RECOMMENDED that the "parameters" field provides details to the system from which raw data is imported.
>
> 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), with corresponding operational issues. A potential architecture on estimating these metrics is shown in Figure 1 below. Section 5 will discuss in more detail the operational issues and how a network may address them.

# Measurement Considerations -> Cost-Context Specification Considerations

- For each performance metric, instead of specifying measurements details, specify cost-context specification considerations:

| 2.1.3. Measurement Considerations | 3.1.4. Cost-Context Specification Considerations |
|---|---|
| Method of Measurement or Calculation: | |
| See section 8.3 of [I-D.ietf-ippm-initial-registry] for potential measurement method. | |
| Measurement Point(s) with Potential Measurement Domain: | "nominal": Typically network delay does not have a nominal value. |
| See Section 4.1, Data sources for potential data sources. | "sla": Many networks provide delay in their application-level service level agreements. It is RECOMMENDED that the "parameters" field of an "sla" one-way-delay metric provides a link ("link") to the SLA definition. |
| Measurement Timing: | "import": There can be multiple sources to import one-way delay. If the estimation is based on [RFC8571], it is RECOMMENDED that "parameters" provides "protocol" as a field and "RFC8571" as the value. If the estimation is computed from the IPPM framework, it is recommended that "parameters" provides "protocol" as a field and "ippm" as the value; see Section 4 of [I-D.ietf-ippm-initial-registry] for additional fields which can be specified for "ippm" in "parameters". |
| See section 8.3.5 of [I-D.ietf-ippm-initial-registry] for potential measurement timing considerations. | "estimation": The exact estimation method is out of the scope of this document. It is RECOMMENDED that the "parameters" field of an "estimation" one-way-delay metric provides a link ("link") to a description of the "estimation" method. |

# Other Changes

- Slight restructure to be consistent w/ RFC7285

```
2.1.  Cost Metric: One Way Delay (owdelay)

   Metric name:

      One Way Delay

   Metric Identifier:

      owdelay
```

```
3.1.  Cost Metric: One Way Delay (owdelay)

3.1.1.  Identifier

   The identifier for this performance metric is "owdelay".

3.1.2.  Value Representation

   The metric value type is a single 'JSONNumber' type value conforming
   to the number specification of [RFC8259] Section 6.  Hence, the
   number can be a floating point number.  The number MUST be non-
   negative.  The unit is expressed in milliseconds.
```

- Adjustment of normative references
  - Refer to related metric documents, but do not make them normative.

| Metric | Definition | Origin |
|---|---|---|
| One Way Delay | Section 2.1 | [RFC2679] Section 3.6 |
| Round Trip Delay | Section 2.2 | [RFC2681] Section 2.6 |
| Packet Delay Variation | Section 2.3 | [RFC3393] Section 2.6 |
| Hop Count | Section 2.4 | [RFC7285] |
| Packet Loss | Section 2.5 | [RFC7680] Section 2.6 |
| Throughput | Section 2.6 | [RFC6349] Section 3.3 |
| Max Reservable Bandwidth | Section 3.1 | [RFC5305] Section 3.5 |
| Residue Bandwidth | Section 3.2 | [RFC7810] Section 4.5 |

Table 1. Cost Metrics Defined in this Document

# Remaining Issues/Suggestions

- How much detail to give guidance on specifying “parameters”
  - Current recommendation
    - “nominal”: NO
    - “sla”: ”link”: <uri>
    - “import”: “protocol”: <value>
    - “estimation”: “link”: <uri>

# Remaining Issues/Suggestions: Metrics Finalization

- Q1: Consistency w/ others (e.g., RFC8571 always say unidirectional?)
- Q2: How to handle statistics of the same metric
  - statistics: min, max, x-percentile, avg, …
  - related to Martin Duke comment
  - approach one (concatenation naming):
    - <metric>[-<stat>]
    - OK for now, but gen cross product in metrics

| Identifier | Intended Semantics |
|---|---|
| owdelay | See Section 2.1 |
| rtt | See Section 2.2 |
| pdv | See Section 2.3 |
| hopcount | See Section 2.4 |
| pktloss | See Section 2.5 |
| throughput | See Section 2.6 |
| maxresbw | See Section 3.1 |
| residuebw | See Section 3.2 |