

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: 22 September 2022

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
Samsung
D. Dhody
Huawei
S. Randriamasy
Nokia Bell Labs
L. Contreras
Telefonica
21 March 2022

ALTO Performance Cost Metrics
draft-ietf-alto-performance-metrics-28

Abstract

The cost metric is a basic concept in Application-Layer Traffic Optimization (ALTO), and different applications may use different types of cost metrics. Since the ALTO base protocol (RFC 7285) defines only a single cost metric (namely, the generic "routingcost" metric), if an application wants to issue a cost map or an endpoint cost request in order to identify a resource provider that offers better performance metrics (e.g., lower delay or loss rate), the base protocol does not define the cost metric to be used.

This document addresses this issue by extending the specification to provide a variety of network performance metrics, including network delay, delay variation (a.k.a, jitter), packet loss rate, hop count, and bandwidth.

There are multiple sources (e.g., estimation based on measurements or service-level agreement) to derive a performance metric. This document introduces an additional "cost-context" field to the ALTO "cost-type" field to convey the source of a performance metric.

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 <https://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 22 September 2022.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Requirements Language	6
3. Performance Metric Attributes	6
3.1. Performance Metric Context: "cost-context"	7
3.2. Performance Metric Statistics	9
4. Packet Performance Metrics	11
4.1. Cost Metric: One-Way Delay (delay-ow)	11
4.1.1. Base Identifier	11
4.1.2. Value Representation	12
4.1.3. Intended Semantics and Use	12
4.1.4. Cost-Context Specification Considerations	14
4.2. Cost Metric: Round-trip Delay (delay-rt)	16
4.2.1. Base Identifier	16
4.2.2. Value Representation	16
4.2.3. Intended Semantics and Use	16
4.2.4. Cost-Context Specification Considerations	17
4.3. Cost Metric: Delay Variation (delay-variation)	18
4.3.1. Base Identifier	18
4.3.2. Value Representation	18
4.3.3. Intended Semantics and Use	18
4.3.4. Cost-Context Specification Considerations	19
4.4. Cost Metric: Loss Rate (lossrate)	20
4.4.1. Base Identifier	20
4.4.2. Value Representation	20
4.4.3. Intended Semantics and Use	20

4.4.4.	Cost-Context Specification Considerations	21
4.5.	Cost Metric: Hop Count (hopcount)	22
4.5.1.	Base Identifier	22
4.5.2.	Value Representation	22
4.5.3.	Intended Semantics and Use	22
4.5.4.	Cost-Context Specification Considerations	23
5.	Throughput/Bandwidth Performance Metrics	24
5.1.	Cost Metric: TCP Throughput (tput)	24
5.1.1.	Base Identifier	24
5.1.2.	Value Representation	24
5.1.3.	Intended Semantics and Use	24
5.1.4.	Cost-Context Specification Considerations	25
5.2.	Cost Metric: Residual Bandwidth (bw-residual)	26
5.2.1.	Base Identifier	26
5.2.2.	Value Representation	26
5.2.3.	Intended Semantics and Use	26
5.2.4.	Cost-Context Specification Considerations	28
5.3.	Cost Metric: Available Bandwidth (bw-available)	28
5.3.1.	Base Identifier	28
5.3.2.	Value Representation	28
5.3.3.	Intended Semantics and Use	29
5.3.4.	Cost-Context Specification Considerations	30
6.	Operational Considerations	30
6.1.	Source Considerations	31
6.2.	Metric Timestamp Consideration	31
6.3.	Backward Compatibility Considerations	31
6.4.	Computation Considerations	32
6.4.1.	Configuration Parameters Considerations	32
6.4.2.	Aggregation Computation Considerations	32
7.	Security Considerations	32
8.	IANA Considerations	33
9.	Acknowledgments	35
10.	References	35
10.1.	Normative References	35
10.2.	Informative References	37
	Authors' Addresses	38

1. Introduction

Application-Layer Traffic Optimization (ALTO) provides a means for network applications to obtain network information so that the applications can identify efficient application-layer traffic patterns using the networks. Cost metrics are used in both the ALTO cost map service and the ALTO endpoint cost service in the ALTO base protocol [RFC7285].

Since different applications may use different cost metrics, the ALTO base protocol introduces an ALTO Cost Metric Registry (Section 14.2 of [RFC7285]) as a systematic mechanism to allow different metrics to be specified. For example, a delay-sensitive application may want to use latency related metrics, and a bandwidth-sensitive application may want to use bandwidth related metrics. However, the ALTO base protocol has registered only a single cost metric, i.e., the generic "routingcost" metric (Section 14.2 of [RFC7285]); no latency or bandwidth related metrics are defined in the base protocol.

This document registers a set of new cost metrics (Table 1) to allow applications to determine "where" to connect based on network performance criteria including delay and bandwidth related metrics.

Metric	Definition in this doc	Semantics Based On
One-way Delay	Section 4.1	Base: [RFC7471,8570,8571] sum Unidirectional Delay
Round-trip Delay	Section 4.2	Base: Sum of two directions from above
Delay Variation	Section 4.3	Base: [RFC7471,8570,8571] sum of Unidirectional Delay Variation
Loss Rate	Section 4.4	Base: [RFC7471,8570,8571] aggr Unidirectional Link Loss
Residual Bandwidth	Section 5.2	Base: [RFC7471,8570,8571] min Unidirectional Residual BW
Available Bandwidth	Section 5.3	Base: [RFC7471,8570,8571] min Unidirectional Avail. BW
TCP Throughput	Section 5.1	[I-D.ietf-tcpm-rfc8312bis]
Hop Count	Section 4.5	[RFC7285]

Table 1. Cost Metrics Defined in this Document.

The first 6 metrics listed in Table 1 (i.e., One-way Delay, Round-trip Delay, Delay Variation, Loss Rate, Residual Bandwidth, and Available Bandwidth) are derived from the set of traffic engineering performance metrics commonly defined in OSPF [RFC3630], [RFC7471]; IS-IS [RFC5305], [RFC8570]; and BGP-LS [RFC8571]. Deriving ALTO cost performance metrics from existing network-layer traffic engineering performance metrics, to expose to application-layer traffic optimization, can be a typical mechanism by network operators to deploy ALTO [RFC7971], [FlowDirector]. This document defines the base semantics of these metrics by extending them from link metrics

to end-to-end metrics for ALTO. The "Semantics Based On" column specifies at a high level how the end-to-end metric is computed from link metrics; the details will be specified in the following sections.

The common metrics Min/Max Unidirectional Delay defined in [RFC8570][RFC8571] and Max Link Bandwidth defined in [RFC3630,RFC5305] are not listed in Table 1 because they can be handled by applying the statistical operators defined in this document. The metrics related with utilized bandwidth and reservable bandwidth (i.e., Max Reservable BW and Unreserved BW defined in [RFC3630,RFC5305]) are outside the scope of this document.

The 7th metric (the estimated TCP-flow throughput metric) provides an estimation of the bandwidth of a TCP flow, using TCP throughput modeling, to support use cases of adaptive applications [Prophet], [G2]. Note that other transport-specific metrics can be defined in the future. For example, QUIC-related metrics [RFC9000] can be considered when the methodology to measure such metrics is more mature (e.g., [I-D.corre-quick-throughput-testing]).

The 8th metric (the hop count metric) in Table 1 is mentioned in the ALTO base protocol [RFC7285], but not defined, and this document defines it to be complete.

These 8 performance metrics can be classified into two categories: those derived from the performance of individual packets (i.e., One-way Delay, Round-trip Delay, Delay Variation, Loss Rate, and Hop Count), and those related to bandwidth/throughput (Residual bandwidth, and Available Bandwidth, and TCP throughput). These two categories are defined in Sections 4 and 5 respectively. Note that all metrics except Round-trip Delay are unidirectional. An ALTO client will need to query both directions if needed.

The purpose of this document is to ensure proper usage of these 8 performance metrics in the context of ALTO. This document follows the guideline defined in Section 14.2 of the ALTO base protocol [RFC7285] on registering ALTO cost metrics. Hence, it specifies the identifier, the intended semantics, and the security considerations of each one of the metrics specified in Table 1.

The definitions of the intended semantics of the metrics tend to be coarse-grained, for guidance only, and they may work well for ALTO. On the other hand, a performance measurement framework, such as the IP Performance Measurement (IPPM) framework, may provide more details in defining a performance metric. This document introduces a mechanism called "cost-context" to provide additional details, when they are available; see Section 3.

Following the ALTO base protocol, this document uses JSON to specify the value type of each defined metric. See [RFC8259] for JSON data type specification. In particular, [RFC7285] specifies that cost values should be assumed by default as JSONNumber. When defining the value representation of each metric in Table 1, this document conforms to [RFC7285], but specifies additional, generic constraints on valid JSONNumbers for each metric. For example, each new metric in Table 1 will be specified as non-negative (≥ 0); Hop Count is specified to be an integer.

An ALTO server may provide only a subset of the metrics described in this document. For example, those that are subject to privacy concerns should not be provided to unauthorized ALTO clients. Hence, all cost metrics defined in this document are optional; not all of them need to be exposed to a given application. When an ALTO server supports a cost metric defined in this document, it announces the metric in its information resource directory (IRD) as defined in Section 9.2 of [RFC7285].

An ALTO server introducing these metrics should consider related security issues. As a generic security consideration on the reliability and trust in the exposed metric values, applications SHOULD rapidly give up using ALTO-based guidance if they detect that the exposed information does not preserve their performance level or even degrades it. Section 7 discusses security considerations in more detail.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Performance Metric Attributes

The definitions of the metrics in this document are coarse-grained, based on network-layer traffic engineering performance metrics, for guidance only. A fine-grained framework specified in [RFC6390] requires that the fine-grained specification of a network performance metric include 6 components: (i) Metric Name, (ii) Metric Description, (iii) Method of Measurement or Calculation, (iv) Units of Measurement, (v) Measurement Points, and (vi) Measurement Timing. Requiring that an ALTO server provides precise, fine-grained values for all 6 components for each metric that it exposes may not be feasible or necessary for all ALTO use cases. For example, an ALTO server computing its metrics from network-layer traffic-engineering

performance metrics may not have information about the method of measurement or calculation (e.g., measured traffic patterns).

To address the issue and realize ALTO use cases, for metrics in Table 1, this document defines performance metric identifiers which can be used in the ALTO protocol with well-defined (i) Metric Name, (ii) Metric Description, (iv) Units of Measurement, and (v) Measurement Points, which are always specified by the specific ALTO services; for example, endpoint cost service is between the two endpoints. Hence, the ALTO performance metric identifiers provide basic metric attributes.

To allow the flexibility of allowing an ALTO server to provide fine-grained information such as Method of Measurement or Calculation, according to its policy and use cases, this document introduces context information so that the server can provide these additional details.

3.1. Performance Metric Context: "cost-context"

The core additional details of a performance metric specify "how" the metric is obtained. This is referred to as the source of the metric. Specifically, this document defines three types of coarse-grained metric information sources: "nominal", and "sla" (service level agreement), and "estimation".

For a given type of source, precise interpretation of a performance metric value can depend on specific measurement and computation parameters.

To make it possible to specify the source and the aforementioned parameters, this document introduces an optional "cost-context" field to the "cost-type" field defined by the ALTO base protocol (Section 10.7 of [RFC7285]) as the following:

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

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

"cost-context" will not be used as a key to distinguish among performance metrics. Hence, an ALTO information resource MUST NOT announce multiple CostType with the same "cost-metric", "cost-mode" and "cost-context". They must be placed into different information resources.

The "cost-source" field of the "cost-context" field is defined as a string consisting of only US-ASCII alphanumeric characters (U+0030-U+0039, U+0041-U+005A, and U+0061-U+007A). The cost-source is used in this document to indicate a string of this format.

As mentioned above, this document defines three values for "cost-source": "nominal", "sla", and "estimation". The "cost-source" field of the "cost-context" field MUST be one registered in "ALTO Cost Source" registry (Section 8).

The "nominal" category indicates that the metric value 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 involved devices; latency typically does not have a nominal value.

The "sla" category indicates that the metric value is derived from some commitment which this document refers to as service-level agreement (SLA). Some operators also use terms such as "target" or "committed" values. For an "sla" metric, it is RECOMMENDED that the "parameters" field provide a link to the SLA definition.

The "estimation" category indicates that the metric value is computed through an estimation process. An ALTO server may compute "estimation" values by retrieving and/or aggregating information from routing protocols (e.g., [RFC7471], [RFC8570], [RFC8571]), traffic measurement management tools (e.g., TWAMP [RFC5357]), and measurement frameworks (e.g., IPPM), with corresponding operational issues. An illustration of potential information flows used for estimating these metrics is shown in Figure 1. Section 6 discusses in more detail the operational issues and how a network may address them.




Figure 1. A framework to compute estimation to performance metrics

There can be multiple choices 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 source, i.e., "estimation".

3.2. Performance Metric Statistics

The measurement of a performance metric often yields a set of samples from an observation distribution ([Prometheus]), instead of a single value. A statistical operator is applied to the samples to obtain a value to be reported to the client. Multiple statistical operators (e.g., min, median, and max) are commonly being used.

Hence, this document extends the general US-ASCII alphanumeric cost metric strings, formally specified as the CostMetric type defined in Section 10.6 of [RFC7285], as follows:

A cost metric string consists of a base metric identifier (or base identifier for short) string, followed by an optional statistical operator string, connected by the ASCII character colon (':', U+003A), if the statistical operator string exists. The total length of the cost metric string MUST NOT exceed 32, as required by [RFC7285].

The statistical operator string MUST be one of the following:

cur:

the instantaneous observation value of the metric from the most recent sample (i.e., the current value).

percentile, with letter 'p' followed by a number:

gives the percentile specified by the number following the letter 'p'. The number MUST be a non-negative JSON number in the range [0, 100] (i.e., greater than or equal to 0 and less than or equal to 100), followed by an optional decimal part, if a higher precision is needed. The decimal part should start with the '.' separator (U+002E), and followed by a sequence of one or more ASCII numbers between '0' and '9'. Assume this number is y and consider the samples coming from a random variable X. Then the metric returns x, such that the probability of X is less than or equal to x, i.e., $\text{Prob}(X \leq x) = y/100$. For example, delay-ow:p99 gives the 99% percentile of observed one-way delay; delay-ow:p99.9 gives the 99.9% percentile. Note that some systems use quantile, which is in the range [0, 1]. When there is a more common form for a given percentile, it is RECOMMENDED that the common form be used; that is, instead of p0, use min; instead of p50, use median; instead of p100, use max.

min:

the minimal value of the observations.

max:

the maximal value of the observations.

median:

the mid-point (i.e., p50) of the observations.

mean:

the arithmetic mean value of the observations.

stddev:

the standard deviation of the observations.

stdvar:

the standard variance of the observations.

Examples of cost metric strings then include "delay-ow", "delay-ow:min", "delay-ow:p99", where "delay-ow" is the base metric identifier string; "min" and "p99" are example statistical operator strings.

If a cost metric string does not have the optional statistical operator string, the statistical operator SHOULD be interpreted as the default statistical operator in the definition of the base metric. If the definition of the base metric does not provide a definition for the default statistical operator, the metric MUST be considered as the median value.

Note that RFC 7258 limits the overall cost metric identifier to 32 characters. The cost metric variants with statistical operator suffixes defined by this document are also subject to the same overall 32-character limit, so certain combinations of (long) base metric identifier and statistical operator will not be representable. If such a situation arises, it could be addressed by defining a new base metric identifier that is an "alias" of the desired base metric, with identical semantics and just a shorter name.

4. Packet Performance Metrics

This section introduces ALTO network performance metrics on one way delay, round-trip delay, delay variation, packet loss rate, and hop count. They measure the "quality of experience" of the stream of packets sent from a resource provider to a resource consumer. The measures of each individual packet (pkt) can include the delay from the time when the packet enters the network to the time when the packet leaves the network (pkt.delay); whether the packet is dropped before reaching the destination (pkt.dropped); the number of network hops that the packet traverses (pkt.hopcount). The semantics of the performance metrics defined in this section are that they are statistics computed from these measures; for example, the x-percentile of the one-way delay is the x-percentile of the set of delays {pkt.delay} for the packets in the stream.

4.1. Cost Metric: One-Way Delay (delay-ow)

4.1.1. Base Identifier

The base identifier for this performance metric is "delay-ow".

4.1.2. Value Representation

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

4.1.3. Intended Semantics and Use

Intended Semantics: To specify the temporal and spatial aggregated delay of a stream of packets from the specified source to the specified destination. The base semantics of the metric is the Unidirectional Delay metric defined in [RFC8571,RFC8570,RFC7471], but instead of specifying the delay for a link, it is the (temporal) aggregation of the link delays from the source to the destination. A non-normative reference definition of end-to-end one-way delay is [RFC7679]. The spatial aggregation level is specified in the query context, e.g., provider-defined identifier (PID) to PID, or endpoint to endpoint, where PID is defined in Section 5.1 of [RFC7285].

Use: This metric could be used as a cost metric constraint attribute or as a returned cost metric in the response.

Example 1: Delay value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 239
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json

{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "delay-ow"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 247
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "delay-ow"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 10,
      "ipv4:198.51.100.34": 20
    }
  }
}
```

Note that since the "cost-type" does not include the "cost-source" field, the values are based on "estimation". Since the identifier does not include the statistical operator string component, the values will represent median values.

Example 1a below shows an example that is similar to Example 1, but for IPv6.

Example 1a: Delay value on source-destination endpoint pairs for IPv6

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 252
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "delay-ow"
  },
  "endpoints": {
    "srcs": [
      "ipv6:2001:db8:100::1"
    ],
    "dsts": [
      "ipv6:2001:db8:100::2",
      "ipv6:2001:db8:100::3"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 257
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "delay-ow"
    }
  },
  "endpoint-cost-map": {
    "ipv6:2001:db8:100::1": {
      "ipv6:2001:db8:100::2": 10,
      "ipv6:2001:db8:100::3": 20
    }
  }
}
```

4.1.4. Cost-Context Specification Considerations

"nominal": Typically network one-way delay does not have a nominal value.

"sla": Many networks provide delay-related parameters in their application-level SLAs. It is RECOMMENDED that the "parameters" field of an "sla" one-way delay metric include a link (i.e., a field named "link") providing a URI to the specification of SLA details, if available. Such a specification can be either free text for possible presentation to the user, or a formal specification. The format of the specification is out of the scope of this document.

"estimation": The exact estimation method is out of the scope of this document. There can be multiple sources to estimate one-way delay. For example, the ALTO server may estimate the end-to-end delay by aggregation of routing protocol link metrics; the server may also estimate the delay using active, end-to-end measurements, for example, using the IPPM framework [RFC2330].

If the estimation is computed by aggregation of routing protocol link metrics (e.g., OSPF [RFC7471], IS-IS [RFC8570], or BGP-LS [RFC8571]) Unidirectional Delay link metrics, it is RECOMMENDED that the "parameters" field of an "estimation" one-way delay metric include the following information: (1) the RFC defining the routing protocol metrics (e.g., <https://www.rfc-editor.org/info/rfc7471> for RFC7471 derived metrics); (2) configurations of the routing link metrics such as configured intervals; and (3) the aggregation method from link metrics to end-to-end metrics. During aggregation from link metrics to the end-to-end metric, the server should be cognizant of potential issues when computing an end-to-end summary statistic from link statistics. The default end-to-end average one-way delay is the sum of average link one-way delays. If an ALTO server provides the min and max statistical operators for the one-way delay metric, the values can be computed directly from the routing link metrics, as [RFC7471,RFC8570,RFC8571] provide Min/Max Unidirectional Link Delay.

If the estimation is from the IPPM measurement framework, it is RECOMMENDED that the "parameters" field of an "estimation" one-way delay metric includes the following information: the URI to the URI field of the IPPM metric defined in the IPPM performance metric [IANA-IPPM] registry (e.g., https://www.iana.org/assignments/performance-metrics/OWDelay_Active_IP-UDP-Poisson-Payload250B_RFC8912sec7_Seconds_95Percentile). The IPPM metric MUST be one-way delay (i.e., IPPM OWDelay* metrics). The statistical operator of the ALTO metric MUST be consistent with the IPPM statistical property (e.g., 95-th percentile).

4.2. Cost Metric: Round-trip Delay (delay-rt)

4.2.1. Base Identifier

The base identifier for this performance metric is "delay-rt".

4.2.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of Section 6 of [RFC8259]. The number MUST be non-negative. The unit is expressed in microseconds.

4.2.3. Intended Semantics and Use

Intended Semantics: To specify temporal and spatial aggregated round-trip delay between the specified source and specified destination. The base semantics is that it is the sum of one-way delay from the source to the destination and the one-way delay from the destination back to the source, where the one-way delay is defined in Section 4.1. A non-normative reference definition of end-to-end round-trip delay is [RFC2681]. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Note that it is possible for a client to query two one-way delays (delay-ow) and then compute the round-trip delay. The server should be cognizant of the consistency of values.

Use: This metric could be used either as a cost metric constraint attribute or as a returned cost metric in the response.

Example 2: Round-trip Delay of source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 238
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "delay-rt"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 245
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "delay-rt"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 4,
      "ipv4:198.51.100.34": 3
    }
  }
}
```

4.2.4. Cost-Context Specification Considerations

"nominal": Typically network round-trip delay does not have a nominal value.

"sla": See the "sla" entry in Section 4.1.4.

"estimation": See the "estimation" entry in Section 4.1.4. For estimation by aggregation of routing protocol link metrics, the aggregation should include all links from the source to the destination and then back to the source; for estimation using IPPM, the IPPM metric MUST be round-trip delay (i.e., IPPM RTDelay* metrics). The statistical operator of the ALTO metric MUST be consistent with the IPPM statistical property (e.g., 95-th percentile).

4.3. Cost Metric: Delay Variation (delay-variation)

4.3.1. Base Identifier

The base identifier for this performance metric is "delay-variation".

4.3.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of Section 6 of [RFC8259]. The number MUST be non-negative. The unit is expressed in microseconds.

4.3.3. Intended Semantics and Use

Intended Semantics: To specify temporal and spatial aggregated delay variation (also called delay jitter) with respect to the minimum delay observed on the stream over the one-way delay from the specified source and destination, where the one-way delay is defined in Section 4.1. A non-normative reference definition of end-to-end one-way delay variation is [RFC3393]. Note that [RFC3393] allows the specification of a generic selection function F to unambiguously define the two packets selected to compute delay variations. This document defines the specific case that F selects as the "first" packet the one with the smallest one-way delay. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Note that in statistics, variations are typically evaluated by the distance from samples relative to the mean. In networking context, it is more commonly defined from samples relative to the min. This definition follows the networking convention.

Use: This metric could be used either as a cost metric constraint attribute or as a returned cost metric in the response.

Example 3: Delay variation value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 245
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "delay-variation"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 252
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "delay-variation"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 0,
      "ipv4:198.51.100.34": 1
    }
  }
}
```

4.3.4. Cost-Context Specification Considerations

"nominal": Typically network delay variation does not have a nominal value.

"sla": See the "sla" entry in Section 4.1.4.

"estimation": See the "estimation" entry in Section 4.1.4. For estimation by aggregation of routing protocol link metrics, the default aggregation of the average of delay variations is the sum of the link delay variations; for estimation using IPPM, the IPPM metric MUST be delay variation (i.e., IPPM OWPDV* metrics). The statistical operator of the ALTO metric MUST be consistent with the IPPM statistical property (e.g., 95-th percentile).

4.4. Cost Metric: Loss Rate (lossrate)

4.4.1. Base Identifier

The base identifier for this performance metric is "lossrate".

4.4.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of Section 6 of [RFC8259]. The number MUST be non-negative. The value represents the percentage of packet losses.

4.4.3. Intended Semantics and Use

Intended Semantics: To specify temporal and spatial aggregated one-way packet loss rate from the specified source and the specified destination. The base semantics of the metric is the Unidirectional Link Loss metric defined in [RFC8571,RFC8570,RFC7471], but instead of specifying the loss for a link, it is the aggregated loss of all links from the source to the destination. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Use: This metric could be used as a cost metric constraint attribute or as a returned cost metric in the response.

Example 5: Loss rate value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 238
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "lossrate"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 248
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "lossrate"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 0,
      "ipv4:198.51.100.34": 0.01
    }
  }
}
```

4.4.4. Cost-Context Specification Considerations

"nominal": Typically packet loss rate does not have a nominal value, although some networks may specify zero losses.

"sla": See the "sla" entry in Section 4.1.4..

"estimation": See the "estimation" entry in Section 4.1.4. For estimation by aggregation of routing protocol link metrics, the default aggregation of the average of loss rate is the sum of the link link loss rates. But this default aggregation is valid only if two conditions are met: (1) it is valid only when link loss rates are low, and (2) it assumes that each link's loss events are uncorrelated with every other link's loss events. When loss rates at the links are high but independent, the general formula for aggregating loss assuming each link is independent is to compute end-to-end loss as one minus the product of the success rate for each link. Aggregation when losses at links are correlated can be more complex and the ALTO server should be cognizant of correlated loss rates. For estimation using IPPM, the IPPM metric MUST be packet loss (i.e., IPPM OWLoss* metrics). The statistical operator of the ALTO metric MUST be consistent with the IPPM statistical property (e.g., 95-th percentile).

4.5. Cost Metric: Hop Count (hopcount)

The hopcount metric is mentioned in Section 9.2.3 of [RFC7285] as an example. This section further clarifies its properties.

4.5.1. Base Identifier

The base identifier for this performance metric is "hopcount".

4.5.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of Section 6 of [RFC8259]. The number MUST be a non-negative integer (greater than or equal to 0). The value represents the number of hops.

4.5.3. Intended Semantics and Use

Intended Semantics: To specify the number of hops in the path from the specified source to the specified destination. The hop count is a basic measurement of distance in a network and can be exposed as the number of router hops computed from the routing protocols originating this information. A hop, however, may represent other units. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Use: This metric could be used as a cost metric constraint attribute or as a returned cost metric in the response.

Example 4: hopcount value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 238
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "hopcount"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 245
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "hopcount"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 5,
      "ipv4:198.51.100.34": 3
    }
  }
}
```

4.5.4. Cost-Context Specification Considerations

"nominal": Typically hop count does not have a nominal value.

"sla": Typically hop count does not have an SLA value.

"estimation": The exact estimation method is out of the scope of this document. An example of estimating hopcounts is by importing from IGP routing protocols. It is RECOMMENDED that the "parameters" field of an "estimation" hop count define the meaning of a hop.

5. Throughput/Bandwidth Performance Metrics

This section introduces four throughput/bandwidth related metrics. Given a specified source to a specified destination, these metrics reflect the volume of traffic that the network can carry from the source to the destination.

5.1. Cost Metric: TCP Throughput (tput)

5.1.1. Base Identifier

The base identifier for this performance metric is "tput".

5.1.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of Section 6 of [RFC8259]. The number MUST be non-negative. The unit is bytes per second.

5.1.3. Intended Semantics and Use

Intended Semantics: To give the throughput of a TCP congestion-control conforming flow from the specified source to the specified destination. The throughput SHOULD be interpreted as only an estimation, and the estimation is designed only for bulk flows.

Use: This metric could be used as a cost metric constraint attribute or as a returned cost metric in the response.

Example 5: TCP throughput value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 234
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json
```

```
{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "tput"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 251
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "tput"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 256000,
      "ipv4:198.51.100.34": 128000
    }
  }
}
```

5.1.4. Cost-Context Specification Considerations

"nominal": Typically TCP throughput does not have a nominal value, and SHOULD NOT be generated.

"sla": Typically TCP throughput does not have an SLA value, and SHOULD NOT be generated.

"estimation": The exact estimation method is out of the scope of this document. It is RECOMMENDED that the "parameters" field of an "estimation" TCP throughput metric include the following information: (1) the congestion-control algorithm; and (2) the estimation methodology. To specify (1), it is RECOMMENDED that the "parameters" field (object) include a field named "congestion-control-algorithm", which provides a URI for the specification of the algorithm; for example, for an ALTO server to provide estimation to the throughput of a Cubic Congestion control flow, its "parameters" includes a field "congestion-control-algorithm", with value being set to [I-D.ietf-tcpm-rfc8312bis]; for an ongoing congestion control algorithm such as BBR, a link to its specification. To specify (2), the "parameters" includes as many details as possible; for example, for TCP Cubic throughput estimation, the "parameters" field specifies that the throughput is estimated by setting `_C_` to 0.4, and the Equation in Figure 8 of [I-D.ietf-tcpm-rfc8312bis] is applied; as an alternative, the methodology may be based on the NUM model [Prophet], or the G2 model [G2]. The exact specification of the parameters field is out of the scope of this document.

5.2. Cost Metric: Residual Bandwidth (bw-residual)

5.2.1. Base Identifier

The base identifier for this performance metric is "bw-residual".

5.2.2. Value Representation

The metric value type is a single 'JSONNumber' type value that is non-negative. The unit of measurement is bytes per second.

5.2.3. Intended Semantics and Use

Intended Semantics: To specify temporal and spatial residual bandwidth from the specified source and the specified destination. The base semantics of the metric is the Unidirectional Residual Bandwidth metric defined in [RFC8571,RFC8570,RFC7471], but instead of specifying the residual bandwidth for a link, it is the residual bandwidth of the path from the source to the destination. Hence, it is the minimal residual bandwidth among all links from the source to the destination. When the max statistical operator is defined for the metric, it typically provides the minimum of the link capacities along the path, as the default value of the residual bandwidth of a link is its link capacity [RFC8571,8570,7471]. The spatial aggregation unit is specified in the query context (e.g., PID to PID,

or endpoint to endpoint).

The default statistical operator for residual bandwidth is the current instantaneous sample; that is, the default is assumed to be "cur".

Use: This metric could be used either as a cost metric constraint attribute or as a returned cost metric in the response.

Example 7: bw-residual value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 241
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json

{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "bw-residual"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 255
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "bw-residual"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 0,
      "ipv4:198.51.100.34": 2000
    }
  }
}
```

5.2.4. Cost-Context Specification Considerations

"nominal": Typically residual bandwidth does not have a nominal value.

"sla": Typically residual bandwidth does not have an "sla" value.

"estimation": See the "estimation" entry in Section 4.1.4 on aggregation of routing protocol link metrics. The current ("cur") residual bandwidth of a path is the minimal of the residual bandwidth of all links on the path.

5.3. Cost Metric: Available Bandwidth (bw-available)

5.3.1. Base Identifier

The base identifier for this performance metric is "bw-available".

5.3.2. Value Representation

The metric value type is a single 'JSONNumber' type value that is non-negative. The unit of measurement is bytes per second.

5.3.3. Intended Semantics and Use

Intended Semantics: To specify temporal and spatial available bandwidth from the specified source to the specified destination. The base semantics of the metric is the Unidirectional Available Bandwidth metric defined in [RFC8571,RFC8570,RFC7471], but instead of specifying the available bandwidth for a link, it is the available bandwidth of the path from the source to the destination. Hence, it is the minimal available bandwidth among all links from the source to the destination. The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

The default statistical operator for available bandwidth is the current instantaneous sample; that is, the default is assumed to be "cur".

Use: This metric could be used either as a cost metric constraint attribute or as a returned cost metric in the response.

Example 8: bw-available value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
Host: alto.example.com
Content-Length: 244
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json

{
  "cost-type": {
    "cost-mode": "numerical",
    "cost-metric": "bw-available"
  },
  "endpoints": {
    "srcs": [
      "ipv4:192.0.2.2"
    ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34"
    ]
  }
}
```

```
HTTP/1.1 200 OK
Content-Length: 255
Content-Type: application/alto-endpointcost+json
```

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "bw-available"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89": 0,
      "ipv4:198.51.100.34": 2000
    }
  }
}
```

5.3.4. Cost-Context Specification Considerations

"nominal": Typically available bandwidth does not have a nominal value.

"sla": Typically available bandwidth does not have an "sla" value.

"estimation": See the "estimation" entry in Section 4.1.4 on aggregation of routing protocol link metrics. The current ("cur") available bandwidth of a path is the minimum of the available bandwidth of all links on the path.

6. Operational Considerations

The exact measurement infrastructure, measurement condition, and computation algorithms can vary from different networks, and are outside the scope of this document. Both the ALTO server and the ALTO clients, however, need to be cognizant of the operational issues discussed in the following sub-sections.

Also, the performance metrics specified in this document are similar, in that they may use similar data sources and have similar issues in their calculation. Hence, this document specifies common issues unless one metric has its unique challenges.

6.1. Source Considerations

The addition of the "cost-source" field is to solve a key issue: An ALTO server needs data sources to compute the cost metrics described in this document, and an ALTO client needs to know the data sources to better interpret the values.

To avoid too fine-grained information, this document introduces "cost-source" to indicate only the high-level type of data sources: "estimation", "nominal" or "lsa", where "estimation" is a type of measurement data source, "nominal" is a type of static configuration, and "sla" is a type that is more based on policy.

For estimation, for example, the ALTO server may use log servers or the OAM system as its data source as recommended by [RFC7971]. In particular, the cost metrics defined in this document can be computed using routing systems as the data sources.

6.2. Metric Timestamp Consideration

Despite the introduction of the additional cost-context information, the metrics do not have a field to indicate the timestamps of the data used to compute the metrics. 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 [RFC8895], the client SHOULD assume that the metric is computed using a snapshot at the time that is approximated by the receiving time.

6.3. Backward Compatibility Considerations

One potential issue introduced by the optional "cost-source" field is backward compatibility. Consider that an IRD which defines two cost-types with the same "cost-mode" and "cost-metric", but one with "cost-source" being "estimation" and the other being "sla". Then an ALTO client that is not aware of the extension will not be able to distinguish between these two types. A similar issue can arise even with a single cost-type, whose "cost-source" is "sla": an ALTO client that is not aware of this extension will ignore this field and consider the metric estimation.

To address the backward-compatibility issue, if a "cost-metric" is "routingcost" and the metric contains a "cost-context" field, then it MUST be "estimation"; if it is not, the client SHOULD reject the information as invalid.

6.4. Computation Considerations

The metric values exposed by an ALTO server may result from additional processing on measurements from data sources to compute exposed metrics. This may involve data processing tasks such as aggregating the results across multiple systems, removing outliers, and creating additional statistics. There are two challenges on the computation of ALTO performance metrics.

6.4.1. Configuration Parameters Considerations

Performance metrics often depend on configuration parameters, and exposing such configuration parameters can help an ALTO client to better understand the exposed metrics. In particular, an ALTO server may be configured to compute a TE metric (e.g., packet loss rate) in fixed intervals, say every T seconds. To expose this information, the ALTO server may provide the client with two pieces of additional information: (1) when the metrics are last computed, and (2) when the metrics will be updated (i.e., the validity period of the exposed metric values). The ALTO server can expose these two pieces of information by using the HTTP response headers Last-Modified and Expires.

6.4.2. Aggregation Computation Considerations

An ALTO server may not be able to measure the performance metrics to be exposed. The basic issue is that the "source" information can often be link level. For example, routing protocols often measure and report only per link loss, not end-to-end loss; similarly, routing protocols report link level available bandwidth, not end-to-end available bandwidth. The ALTO server then needs to aggregate these data to provide an abstract and unified view that can be more useful to applications. The server should consider that different metrics may use different aggregation computation. For example, the end-to-end latency of a path is the sum of the latency of the links on the path; the end-to-end available bandwidth of a path is the minimum of the available bandwidth of the links on the path; in contrast, aggregating loss values is complicated by the potential for correlated loss events on different links in the path

7. Security Considerations

The properties defined in this document present no security considerations beyond those in Section 15 of the base ALTO specification [RFC7285].

However, concerns addressed in Sections 15.1, 15.2, and 15.3 of [RFC7285] remain of utmost importance. Indeed, Traffic Engineering (TE) performance is highly sensitive ISP information; therefore, sharing TE metric values in numerical mode requires full mutual confidence between the entities managing the ALTO server and the ALTO client. ALTO servers will most likely distribute numerical TE performance to ALTO clients under strict and formal mutual trust agreements. On the other hand, ALTO clients must be cognizant on the risks attached to such information that they would have acquired outside formal conditions of mutual trust.

To mitigate confidentiality risks during information transport of TE performance metrics, the operator should address the risk of ALTO information being leaked to malicious Clients or third parties, through attacks such as the person-in-the-middle (PITM) attacks. As specified in "Protection Strategies" (Section 15.3.2 of [RFC7285]), the ALTO Server should authenticate ALTO Clients when transmitting an ALTO information resource containing sensitive TE performance metrics. "Authentication and Encryption" (Section 8.3.5 of [RFC7285]) specifies that "ALTO Server implementations as well as ALTO Client implementations MUST support the "https" URI scheme of [RFC7230] and Transport Layer Security (TLS) of [RFC8446]".

8. IANA Considerations

IANA has created and now maintains the "ALTO Cost Metric" registry, listed in Section 14.2, Table 3 of [RFC7285]. This registry is located at <https://www.iana.org/assignments/alto-protocol/alto-protocol.xhtml#cost-metrics>. This document requests to add the following entries to the "ALTO Cost Metric" registry.

Identifier	Intended Semantics
delay-ow	Section 4.1 of [RFCXXX]
delay-rt	Section 4.2 of [RFCXXX]
delay-variation	Section 4.3 of [RFCXXX]
lossrate	Section 4.4 of [RFCXXX]
hopcount	Section 4.5 of [RFCXXX]
tput	Section 5.1 of [RFCXXX]
bw-residual	Section 5.2 of [RFCXXX]
bw-available	Section 5.3 of [RFCXXX]

* [Note to the RFC Editor]: Please replace RFCXXX with the RFC number assigned to this document.

This document requests the creation of the "ALTO Cost Source" registry. This registry serves two purposes. First, it ensures uniqueness of identifiers referring to ALTO cost source types. Second, it provides references to particular semantics of allocated cost source types to be applied by both ALTO servers and applications utilizing ALTO clients.

A new ALTO cost source can be added after IETF Review [RFC8126], to ensure that proper documentation regarding the new ALTO cost source and its security considerations have been provided. The RFC(s) documenting the new cost source should be detailed enough to provide guidance to both ALTO service providers and applications utilizing ALTO clients as to how values of the registered ALTO cost source should be interpreted. Updates and deletions of ALTO cost source follow the same procedure.

Registered ALTO address type identifiers MUST conform to the syntactical requirements specified in Section 3.1. Identifiers are to be recorded and displayed as strings.

Requests to add a new value to the registry MUST include the following information:

- * Identifier: The name of the desired ALTO cost source type.
- * Intended Semantics: ALTO cost source type carry with them semantics to guide their usage by ALTO clients. Hence, a document defining a new type should provide guidance to both ALTO service providers and applications utilizing ALTO clients as to how values of the registered ALTO endpoint property should be interpreted.
- * Security Considerations: ALTO cost source types expose information to ALTO clients. ALTO service providers should be made aware of the security ramifications related to the exposure of a cost source type.

This specification requests registration of the identifiers "nominal", "sla", and "estimation" listed in the table below. Semantics for these are documented in Section 3.1, and security considerations are documented in Section 7.

Identifier	Intended Semantics	Security Considerations
nominal	Values in nominal cases; Section 3.1 of [RFCXXX]	Section 7 of [RFCXXX]
sla	Values reflecting service level agreement; Section 3.1 of [RFCXXXX]	Section 7 of [RFCXXX]
estimation	Values by estimation; Section 3.1 of [RFCXXX]	Section 7 of [RFCXXX]

9. Acknowledgments

The authors of this document would also like to thank Martin Duke for the highly informative, thorough AD reviews and comments. We thank Christian Amsuess, Elwyn Davies, Haizhou Du, Kai Gao, Geng Li, Lili Liu, Danny Alex Lachos Perez, and Brian Trammell for the reviews and comments. We thank Benjamin Kaduk, Eric Kline, Francesca Palombini, Lars Eggert, Martin Vigoureux, Murraray Kucherawy, Roman Danyliw, Zaheduzzaman Sarker, Eric Vyncke for discussions and comments that improve this document.

10. References

10.1. Normative References

[I-D.ietf-tcpm-rfc8312bis]

Xu, L., Ha, S., Rhee, I., Goel, V., and L. Eggert, "CUBIC for Fast and Long-Distance Networks", Work in Progress, Internet-Draft, draft-ietf-tcpm-rfc8312bis-07, 4 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-tcpm-rfc8312bis-07.txt>>.

[IANA-IPPM]

IANA, "Performance Metrics Registry, <https://www.iana.org/assignments/performance-metrics/performance-metrics.xhtml>".

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, DOI 10.17487/RFC6390, October 2011, <<https://www.rfc-editor.org/info/rfc6390>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, DOI 10.17487/RFC7285, September 2014, <<https://www.rfc-editor.org/info/rfc7285>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.
- [RFC8895] Roome, W. and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Incremental Updates Using Server-Sent Events (SSE)", RFC 8895, DOI 10.17487/RFC8895, November 2020, <<https://www.rfc-editor.org/info/rfc8895>>.

10.2. Informative References

- [FlowDirector]
Pujol, E., Poese, I., Zerwas, J., Smaragdakis, G., and A. Feldmann, "Steering Hyper-Giants' Traffic at Scale", ACM CoNEXT 2020, 2020.
- [G2] Ros-Giralt, J., Bohara, A., Yellamraju, S., and et. al., "On the Bottleneck Structure of Congestion-Controlled Networks", ACM SIGMETRICS 2019, 2020.
- [I-D.corre-quick-throughput-testing]
Corre, K., "Framework for QUIC Throughput Testing", Work in Progress, Internet-Draft, draft-corre-quick-throughput-testing-00, 17 September 2021, <<https://www.ietf.org/archive/id/draft-corre-quick-throughput-testing-00.txt>>.
- [Prometheus]
Volz, J. and B. Rabenstein, "Prometheus: A Next-Generation Monitoring System", 2015.
- [Prophet] Gao, K., Zhang, J., and YR. Yang, "Prophet: Fast, Accurate Throughput Prediction with Reactive Flows", ACM/IEEE Transactions on Networking July, 2020.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <<https://www.rfc-editor.org/info/rfc2330>>.

- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, DOI 10.17487/RFC2681, September 1999, <<https://www.rfc-editor.org/info/rfc2681>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008, <<https://www.rfc-editor.org/info/rfc5357>>.
- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [RFC7971] Stiemerling, M., Kiesel, S., Scharf, M., Seidel, H., and S. Previdi, "Application-Layer Traffic Optimization (ALTO) Deployment Considerations", RFC 7971, DOI 10.17487/RFC7971, October 2016, <<https://www.rfc-editor.org/info/rfc7971>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.

Authors' Addresses

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing
Jiangsu, 210012
China
Email: bill.wu@huawei.com

Y. Richard Yang
Yale University
51 Prospect St
New Haven, CT 06520
United States of America
Email: yry@cs.yale.edu

Young Lee
Samsung
Email: young.lee@gmail.com

Dhruv Dhody
Huawei
Leela Palace
Bangalore 560008
Karnataka
India
Email: dhruv.ietf@gmail.com

Sabine Randriamasy
Nokia Bell Labs
Route de Villejust
91460 Nozay
France
Email: sabine.randriamasy@nokia-bell-labs.com

Luis Miguel Contreras Murillo
Telefonica
Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com