

Internet Engineering Task Force (IETF)
Request for Comments: 9439
Category: Standards Track
ISSN: 2070-1721

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
Samsung
D. Dhody
Huawei
S. Randriamasy
Nokia Networks France
L. Contreras
Telefonica
August 2023

Application-Layer Traffic Optimization (ALTO) Performance Cost Metrics

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., estimations based on measurements or a Service Level Agreement) available for deriving 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 is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9439>.

Copyright Notice

Copyright (c) 2023 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
2. Requirements Language
3. Performance Metric Attributes
 - 3.1. Performance Metric Context: "cost-context"
 - 3.2. Performance Metric Statistics
4. Packet Performance Metrics
 - 4.1. Cost Metric: One-Way Delay (delay-ow)
 - 4.1.1. Base Identifier
 - 4.1.2. Value Representation
 - 4.1.3. Intended Semantics and Use
 - 4.1.4. Cost-Context Specification Considerations
 - 4.2. Cost Metric: Round-Trip Delay (delay-rt)
 - 4.2.1. Base Identifier
 - 4.2.2. Value Representation
 - 4.2.3. Intended Semantics and Use
 - 4.2.4. Cost-Context Specification Considerations
 - 4.3. Cost Metric: Delay Variation (delay-variation)
 - 4.3.1. Base Identifier
 - 4.3.2. Value Representation
 - 4.3.3. Intended Semantics and Use
 - 4.3.4. Cost-Context Specification Considerations
 - 4.4. Cost Metric: Loss Rate (lossrate)
 - 4.4.1. Base Identifier
 - 4.4.2. Value Representation
 - 4.4.3. Intended Semantics and Use
 - 4.4.4. Cost-Context Specification Considerations
 - 4.5. Cost Metric: Hop Count (hopcount)
 - 4.5.1. Base Identifier
 - 4.5.2. Value Representation
 - 4.5.3. Intended Semantics and Use
 - 4.5.4. Cost-Context Specification Considerations
5. Throughput/Bandwidth Performance Metrics
 - 5.1. Cost Metric: TCP Throughput (tput)
 - 5.1.1. Base Identifier
 - 5.1.2. Value Representation
 - 5.1.3. Intended Semantics and Use
 - 5.1.4. Cost-Context Specification Considerations
 - 5.2. Cost Metric: Residual Bandwidth (bw-residual)
 - 5.2.1. Base Identifier
 - 5.2.2. Value Representation
 - 5.2.3. Intended Semantics and Use
 - 5.2.4. Cost-Context Specification Considerations
 - 5.3. Cost Metric: Available Bandwidth (bw-available)
 - 5.3.1. Base Identifier
 - 5.3.2. Value Representation
 - 5.3.3. Intended Semantics and Use
 - 5.3.4. Cost-Context Specification Considerations
6. Operational Considerations
 - 6.1. Source Considerations
 - 6.2. Metric Timestamp Considerations
 - 6.3. Backward-Compatibility Considerations
 - 6.4. Computation Considerations
 - 6.4.1. Configuration Parameter Considerations
 - 6.4.2. Aggregation Computation Considerations
7. Security Considerations
8. IANA Considerations
 - 8.1. ALTO Cost Metrics Registry
 - 8.2. ALTO Cost Source Types Registry

9. References

9.1. Normative References

9.2. Informative References

Acknowledgments

Authors' Addresses

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 introduced the "ALTO Cost Metrics" 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 Document	Semantics Based On
One-Way Delay	Section 4.1	Base: [RFC7471] [RFC8570] [RFC8571] sum of Unidirectional Delay of links along the path
Round-Trip Delay	Section 4.2	Base: Sum of two directions of Unidirectional Delay
Delay Variation	Section 4.3	Base: [RFC7471] [RFC8570] [RFC8571] Sum of Unidirectional Delay Variation of links along the path
Loss Rate	Section 4.4	Base: [RFC7471] [RFC8570] [RFC8571] aggr Unidirectional Link Loss
Residual Bandwidth	Section 5.2	Base: [RFC7471] [RFC8570] [RFC8571] min Unidirectional Residual BW
Available Bandwidth	Section 5.3	Base: [RFC7471] [RFC8570] [RFC8571] min Unidirectional Available BW
TCP Throughput	Section 5.1	[RFC9438]
Hop Count	Section 4.5	[RFC7285]

Table 1: Cost Metrics Defined in This Document

The first six 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 (TE) performance metrics commonly defined in OSPF [RFC3630]

[RFC7471], IS-IS [RFC5305] [RFC8570], and BGP - Link State (BGP-LS) [RFC8571]. Deriving ALTO cost performance metrics from existing network-layer TE performance metrics, and making it exposed to ALTO, can be a typical mechanism used 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 metrics are computed from link metrics; details will be specified in the following sections.

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

The seventh metric in Table 1 (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 for measuring such metrics is more mature (e.g., see [QUIC-THROUGHPUT-TESTING]).

The eighth metric in Table 1 (the hop count metric) is mentioned, but not defined, in the ALTO base protocol [RFC7285]; this document provides a definition for it.

These eight 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, 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 eight performance metrics in the context of ALTO. This document follows the guidelines defined in Section 14.2 of [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 and are for guidance only, and they may work well for ALTO. On the other hand, a performance measurement framework, such as the IP Performance Metrics (IPPM) framework, may provide more details for 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 specifications. In particular, [RFC7285] specifies that cost values should be assumed by default to be '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 regarding reliability and trust in the exposed metric values, applications SHOULD promptly stop 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 TE performance metrics, and for guidance only. A fine-grained framework as specified in [RFC6390] requires that the fine-grained specification of a network performance metric include six components: (1) Metric Name, (2) Metric Description, (3) Method of Measurement or Calculation, (4) Units of Measurement, (5) Measurement Points, and (6) Measurement Timing. Requiring that an ALTO server provide precise, fine-grained values for all six 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 TE 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 the metrics listed in Table 1, this document defines performance metric identifiers that can be used in the ALTO Protocol with the following well-defined items: (1) Metric Name, (2) Metric Description, (3) Units of Measurement, and (4) Measurement Points, which are always specified by the specific ALTO services; for example, the 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", "sla", 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 follows:

```
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 entries 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 ASCII alphanumeric characters (U+0030-U+0039, U+0041-U+005A, and U+0061-U+007A). The "cost-source" field 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 that is registered in the "ALTO Cost Source Types" 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 a 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., see [RFC7471], [RFC8570], and [RFC8571]), traffic measurement management tools (e.g., the Two-Way Active Measurement Protocol (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 of Performance Metrics

There can be multiple options available when choosing the "cost-source" category; the operator of an ALTO server will make that choice. 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 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 colon character (':', 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 the 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 higher precision is needed. The decimal part should start with the '.' separator (U+002E) and be followed by a sequence of one or more ASCII numbers between '0' and '9'. Assume that this number is y , and consider the case where the samples are coming from a random variable X . The metric then 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 99th percentile of observed one-way delay; delay-ow:p99.9 gives the 99.9th 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 midpoint (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", and "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 the median value.

Note that [RFC7285] 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 identifiers and statistical operators 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 measurements 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), and 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 measurements; 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 specifications provided in 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 as defined in [RFC8571], [RFC8570],

and [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 the end-to-end one-way delay metric is provided in [RFC7679]. The spatial aggregation level is specified in the query context, e.g., provider-defined identifier (PID) to PID, or endpoint to endpoint, where the PID is as 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.

```
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
    }
  }
}
```

Figure 2: Delay Value on Source-Destination Endpoint Pairs
(Example 1)

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.

Figure 3 shows an example that is similar to Example 1 (Figure 2), but 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
    }
  }
}
```

Figure 3: Delay Value on Source-Destination Endpoint Pairs for IPv6 (Example 1a)

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 for the specification of SLA details, if available. Such a specification can be either (1) free text for possible presentation to the user or (2) a formal specification. The format of the specification is outside the scope of this document.

"estimation": The exact estimation method is outside the scope of this document. There can be multiple sources for estimating 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., Unidirectional Link Delay metrics for OSPF [RFC7471],

IS-IS [RFC8570], or BGP-LS [RFC8571]), 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., see [RFC7471] for 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 end-to-end metrics, 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], and [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 include the URI in the "URI" field of the IPPM metric defined in the IPPM "Performance Metrics" registry [IANA-IPPM] (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., 95th 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 specifications provided in 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 the 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 as defined in Section 4.1. A non-normative reference definition of the end-to-end round-trip delay metric is provided in [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 delay (delay-ow) items and then compute the round-trip delay. The server should be cognizant of the consistency of values.

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

```
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
    }
  }
}

```

Figure 4: Round-Trip Delay of Source-Destination Endpoint Pairs
(Example 2)

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., 95th 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 specifications provided in 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 as defined in Section 4.1. A non-normative reference definition of the end-to-end one-way delay variation metric is provided in [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 where F selects the packet with the smallest one-way delay as the "first" packet. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

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

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

```
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
    }
  }
}
```

Figure 5: Delay Variation Value on Source-Destination Endpoint Pairs (Example 3)

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., 95th 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 specifications provided in 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 the 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 as defined in [RFC8571], [RFC8570], and [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.

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
    }
  }
}
```

Figure 6: Loss Rate Value on Source-Destination Endpoint Pairs
(Example 4)

4.4.4. Cost-Context Specification Considerations

"nominal": Typically, the 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 loss rate is the sum of the link loss rates. But this default aggregation is valid only if two conditions are met: (1) link loss rates are low and (2) one 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 that 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., 95th percentile).

4.5. Cost Metric: Hop Count (hopcount)

The hop count (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 specifications provided in 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.

```
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
    }
  }
}
```

Figure 7: Hop Count Value on Source-Destination Endpoint Pairs
(Example 5)

4.5.4. Cost-Context Specification Considerations

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

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

"estimation": The exact estimation method is outside the scope of this document. An example of estimating hop count values 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 three metrics related to throughput and bandwidth. Given a specified source and 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 specifications provided in 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 congestion control conforming TCP 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.

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
  }
}

```

Figure 8: TCP Throughput Value on Source-Destination Endpoint Pairs (Example 6)

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 outside 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 of the throughput of a CUBIC congestion control flow, its "parameters" field includes the "congestion-control-algorithm" field, with value being set to the URI for [RFC9438]; for an ongoing congestion control algorithm such as BBR, a link to its specification can be added. To specify (2), the "parameters" field includes as many details as possible; for example, for the TCP Cubic throughput estimation, the "parameters" field specifies that the throughput is estimated by setting `_C_` to 0.4, and the equation in [RFC9438], Section 5.1, Figure 8 is applied; as an alternative, the methodology may be based on the NUM model [Prophet] or the model described in [G2]. The exact specification of the "parameters" field is outside 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 to the specified destination. The base semantics of the metric is the Unidirectional Residual Bandwidth metric as defined in [RFC8571], [RFC8570], and [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] [RFC8570] [RFC7471]. 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 as a cost metric constraint attribute or as a returned cost metric in the response.

```
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
    }
  }
}
```

Figure 9: Residual Bandwidth Value on Source-Destination Endpoint Pairs (Example 7)

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. The current ("cur") residual bandwidth of a path is the minimal 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 as defined in [RFC8571], [RFC8570], and [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 as a cost metric constraint attribute or as a returned cost metric in the response.

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
    }
  }
}
```

```

    }
  }
}

```

Figure 10: Available Bandwidth Value on Source-Destination
Endpoint Pairs (Example 8)

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. 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 conditions, and computation algorithms can vary between 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 subsections.

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 issues that the performance metrics might have in common and also discusses challenges regarding the computation of ALTO performance metrics (Section 6.4).

6.1. Source Considerations

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

For example, for "estimation", the ALTO server may use log servers or the Operations, Administration, and Maintenance (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 Considerations

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 an HTTP Last-Modified value 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 the case where an IRD defines two "cost-type" entries with the same "cost-mode" and "cost-metric", but one with "cost-source" being "estimation" and the other being "sla". In such a case, 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 instead 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 of 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. The computation of ALTO performance metrics can present two challenges.

6.4.1. Configuration Parameter 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) at 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 were 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 information. For example, routing protocols often measure and report only per-link loss and not end-to-end loss; similarly, routing protocols report link-level available bandwidth and 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 be aware that different metrics may use different aggregation computations. For example, the end-to-end latency of a path is the sum of the latencies 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, 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 of 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 such attacks as person-in-the-middle (PITM) attacks. As specified in Section 15.3.2 ("Protection Strategies") of [RFC7285], the ALTO server should authenticate ALTO clients when transmitting an ALTO information resource containing sensitive TE performance metrics. Section 8.3.5 ("Authentication and Encryption") of [RFC7285] specifies that ALTO server implementations as well as ALTO client implementations MUST support the "https" URI scheme [RFC9110] and Transport Layer Security (TLS) [RFC8446].

8. IANA Considerations

8.1. ALTO Cost Metrics Registry

IANA created and now maintains the "ALTO Cost Metrics" registry, as listed in [RFC7285], Section 14.2, Table 3. This registry is located at <<https://www.iana.org/assignments/alto-protocol/>>. IANA has added the following entries to the "ALTO Cost Metrics" registry.

Identifier	Intended Semantics	Reference
delay-ow	See Section 4.1	RFC 9439
delay-rt	See Section 4.2	RFC 9439
delay-variation	See Section 4.3	RFC 9439
lossrate	See Section 4.4	RFC 9439
hopcount	See Section 4.5	RFC 9439
tput	See Section 5.1	RFC 9439
bw-residual	See Section 5.2	RFC 9439
bw-available	See Section 5.3	RFC 9439

Table 2: ALTO Cost Metrics Registry

8.2. ALTO Cost Source Types Registry

IANA has created the "ALTO Cost Source Types" registry. This registry serves two purposes. First, it ensures the 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 type can be added after IETF Review [RFC8126], to ensure that proper documentation regarding the new ALTO cost source type and its security considerations has been provided. The RFC(s) documenting the new cost source type 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 type should be interpreted. Updates and deletions of ALTO cost source types 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 types 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.

IANA has registered the identifiers "nominal", "sla", and "estimation" as listed in the table below.

Identifier	Intended Semantics	Security Considerations	Reference
nominal	Values in nominal cases (Section 3.1)	Section 7	RFC 9439
sla	Values reflecting Service Level Agreement (Section 3.1)	Section 7	RFC 9439
estimation	Values by estimation (Section 3.1)	Section 7	RFC 9439

Table 3: ALTO Cost Source Types Registry

9. References

9.1. Normative References

[IANA-IPPM]

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

[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>>.
- [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>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/info/rfc9110>>.
- [RFC9438] Xu, L., Ha, S., Rhee, I., Goel, V., and L. Eggert, Ed., "CUBIC for Fast and Long-Distance Networks", RFC 9438, DOI 10.17487/RFC9438, August 2023, <<https://www.rfc-editor.org/info/rfc9438>>.

9.2. Informative References

- [FlowDirector]
Pujol, E., Poese, I., Zerwas, J., Smaragdakis, G., and A.

Feldmann, "Steering Hyper-Giants' Traffic at Scale", ACM CoNEXT '19, December 2019.

- [G2] Ros-Giralt, J., Bohara, A., Yellamraju, S., Harper Langston, M., Lethin, R., Jiang, Y., Tassiulas, L., Li, J., Tan, Y., and M. Veeraraghavan, "On the Bottleneck Structure of Congestion-Controlled Networks", Proceedings of the ACM on Measurement and Analysis of Computing Systems, Vol. 3, No. 3, Article No. 59, pp. 1-31, DOI 10.1145/3366707, December 2019, <<https://dl.acm.org/doi/10.1145/3366707>>.
- [Prometheus] Volz, J. and B. Rabenstein, "Prometheus: A Next-Generation Monitoring System (Talk)", SREcon15 Europe, May 2015.
- [Prophet] Zhang, J., Gao, K., Yang, YR., and J. Bi, "Prophet: Toward Fast, Error-Tolerant Model-Based Throughput Prediction for Reactive Flows in DC Networks", IEEE/ACM Transactions on Networking, Volume 28, Issue 601, pp. 2475-2488, December 2020, <<https://dl.acm.org/doi/10.1109/TNET.2020.3016838>>.
- [QUIC-THROUGHPUT-TESTING] Corre, K., "Framework for QUIC Throughput Testing", Work in Progress, Internet-Draft, draft-corre-quic-throughput-testing-00, 17 September 2021, <<https://datatracker.ietf.org/doc/html/draft-corre-quic-throughput-testing-00>>.
- [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] Stiernerling, 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>>.

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

Authors' Addresses

Qin Wu
Huawei
Yuhua District
101 Software Avenue
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: younglee.tx@gmail.com

Dhruv Dhody
Huawei
India
Email: dhruv.ietf@gmail.com

Sabine Randriamasy
Nokia Networks France
France
Email: sabine.randriamasy@nokia-bell-labs.com

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