ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2020

Q. Wu
Huawei
Y. Yang
Yale University
D. Dhody
Huawei
S. Randriamasy
Nokia Bell Labs
L. Contreras
Telefonica
March 09, 2020

ALTO Performance Cost Metrics
draft-ietf-alto-performance-metrics-09

Abstract

Cost metric is a basic concept in Application-Layer Traffic Optimization (ALTO), and is used in basic ALTO services including both the cost map service and the endpoint cost service.

Different applications may use different cost metrics, but the ALTO base protocol [RFC7285] documents only one single cost metric, i.e., the generic "routingcost" metric; see Sec. 14.2 of [RFC7285]. Hence, if the resource consumer of an application prefers a resource provider that offers low-delay delivery to the resource consumer, the base protocol does not define the cost metric to be used.

ALTO cost metrics can be generic metrics and this document focuses on network performance metrics, including network delay, jitter, packet loss, hop count, and bandwidth.

When using an ALTO performance metric, an application may need additional contextual information beyond the metric value. For example, whether the metric is an estimation based on measurements or a service-level agreement (SLA) can define the meaning of a performance metric. Hence, this document introduces an additional "cost-context" field to the ALTO "cost-type" field to convey such information. To report an estimated value of a performance metric, the ALTO server may derive and aggregate from routing protocols with different granularity and scope, such as BGP-LS, OSPF-TE and ISIS-TE, or from end-to-end traffic management tools. These metrics may then be exposed by an ALTO Server to allow applications to determine "where" to connect based on network performance criteria.

Requirements Language The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of This Memo

This Internet-Draft will expire on September 10, 2020.

Copyright Notice

Table of Contents

1. Introduction

Cost Metric is a basic concept in Application-Layer Traffic Optimization (ALTO). It is used in both the ALTO cost map service and the ALTO endpoint cost service, to allow applications to request network cost metrics.

Different applications may use different cost metrics. Hence, the ALTO base protocol [RFC7285] 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. The ALTO base protocol [RFC7285], however, has registered only one single cost metric, i.e., the generic "routingcost" metric; no latency or bandwidth related metrics are defined.

This document registers a set of new cost metrics specified in Table 1, to support the aforementioned applications, to allow them to better determine "where" to connect based on network performance criteria. This document follows the guideline (Section 14.2 of [RFC7285]) of the ALTO base protocol on registering ALTO cost metrics. Hence it specifies the identifier, the intended semantics, and the security considerations of each one of the metrics defined in Table 1.

Metric	Definition	Origin
One Way Delay	Section 2.1	[RFC7679]
Round Trip Delay	Section 2.2	[RFC2681]
Packet Delay Variation	Section 2.3	[RFC3393]
Hop Count	Section 2.4	[RFC7285]
Packet Loss	Section 2.5	[RFC7680]
Throughput	Section 2.6	[RFC6349]
Max Reservable Bandwidth	Section 3.1	[RFC5305]
Residue Bandwidth	Section 3.2	[RFC7810]

Table 1. Cost Metrics Defined in this Document.

The purpose of this document is to ensure proper usage of the metrics by ALTO clients. It does not claim novelty of the metrics; see Table 1 for the source definition of each metric.

An ALTO server may provide only a subset of the cost metrics described in this document. Hence, all cost metrics defined in this document are optional and not all of them need to be exposed to a given application. For example, those that are subject to privacy concerns should not be provided to unauthorized ALTO clients.

When an ALTO server supports a cost metric defined in this document, it MUST announce this metric in its information resource directory (IRD).

An ALTO server introducing these metrics should also consider 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. We discuss security considerations in more details in Section 6.

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.

2. Performance Metric Context: cost-context

The semantics of a performance metric depends on the context. Specifically, this document defines three sources when defining performance metrics: "estimation", "nominal", and "sla".

Even given the source, precise interpretation of a performance metric value, if needed, depends on an additional set of measurement and computation parameters. For example, see Section 3.8 of [RFC7679] on items which a more complete measurement-based report should include.

To make it possible to specify both the source and the additional 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;
```

The "cost-source" MUST be one of only three values: "estimation", "nominal", and "sla". If a "cost-type" does not include the optional "cost-context" field which includes the "cost-source" field, the application MUST assume that the value of "cost-source" is "estimation".

An ALTO server may compute "estimation" values by retrieving and/or aggregating information from routing protocols or other traffic measurement management tools, with corresponding operational issues. A potential architecture on estimating these metrics is shown in Figure 1 below. In Section 5, we discuss in more detail the operational issues and how a network may address them.

```
     +--------+   +--------+  +--------+
     | Client |   | Client |  | Client |
     +----^---+   +---^----+  +---^----+
          |           |           |
          +-----------|-----------+
                NBI   |ALTO protocol
                      |
                      |
                   +--+-----+  retrieval      +---------+
                   |  ALTO  |<----------------| Routing |
                   | Server |  and aggregation|         |
                   |        |<-------------+  | Protocol|
                   +--------+              |  +---------+
                                           |
                                           |  +---------+
                                           |  |Management
                                           ---|         |
                                              |  Tool   |
                                              +---------+
```

Figure 1. Potential framework to compute performance cost metrics

3. Network Performance Cost Metrics

This section introduces generic ALTO network performance metrics such as one way delay, round trip delay, hop count, packet loss, and throughput.

3.1. Cost Metric: One Way Delay (owdelay)

3.1.1. Identifier

The identifier for this performance metric is "owdelay".

3.1.2. Value Representation

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

3.1.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal aggregated delay of a stream of packets exchanged between the specified source and destination or the time that the packet spends to travel from source to 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 used either together with cost metric attribute 'routingcost' or on its own 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: TBA
Content-Type: application/alto-endpointcostparams+json
Accept:
  application/alto-endpointcost+json,application/alto-error+json

{
  "cost-type": {"cost-mode" : "numerical",
                "cost-metric" : "owdelay"},
  "endpoints" : {
    "srcs": [ "ipv4:192.0.2.2" ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv6:2000::1:2345:6789:abcd"
    ]
  }
}

HTTP/1.1 200 OK
Content-Length: TBA
Content-Type: application/alto-endpointcost+json
{
  "meta" :{
    "cost-type": {"cost-mode" : "numerical",
                  "cost-metric" : "owdelay"
     }
   },
    "endpoint-cost-map" : {
      "ipv4:192.0.2.2": {
        "ipv4:192.0.2.89"    : 10,
        "ipv4:198.51.100.34" : 20,
        "ipv6:2000::1:2345:6789:abcd"  : 30,
    }
  }
}
```

Comment: Since the "cost-type" does not include the "cost-source" field, the values are based on "estimation".

3.1.4. Measurement Considerations and Parameters

See Section 4 of [I-D.ietf-ippm-initial-registry] for measurement considerations and parameters which may be specified in "parameters".

Note that the "parameters" field is an optional field providing non-normative information.

3.2. Cost Metric: RoundTrip Time (rtt)

3.2.1. Identifier

The identifier for this performance metric is "rtt".

3.2.2. Value Representation

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

3.2.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal aggregated round trip delay between the specified source and destination or the time that the packet spends to travel from source to destination and then from destination to source. 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 2: Roundtrip Delay value on source-destination endpoint pairs

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

{
 "cost-type": {"cost-mode" : "numerical",
               "cost-metric" : "rtt"},
  "endpoints" : {
     "srcs": [ "ipv4:192.0.2.2" ],
     "dsts": [
       "ipv4:192.0.2.89",
       "ipv4:198.51.100.34",
       "ipv6:2000::1:2345:6789:abcd"
     ]
   }
}

 HTTP/1.1 200 OK
 Content-Length: TBA
 Content-Type: application/alto-endpointcost+json
 {
   "meta" :{
     "cost-type": {"cost-mode" : "numerical",
                   "cost-metric" : "rtt"
      }
    },
     "endpoint-cost-map" : {
       "ipv4:192.0.2.2": {
         "ipv4:192.0.2.89"    : 4,
         "ipv4:198.51.100.34" : 3,
         "ipv6:2000::1:2345:6789:abcd"  : 2,
     }
   }
 }
```

3.2.4. Measurement Considerations and Parameters

See Section 4 of [I-D.ietf-ippm-initial-registry] for measurement considerations and parameters which may be specified in "parameters". Note that the "parameters" field is an optional field providing non-normative information.

3.3. Cost Metric: Packet Delay Variation (pdv)

3.3.1. Identifier

The identifier for this performance metric is "pdv".

3.3.2. Value Representation

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

3.3.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal aggregated jitter (packet delay variation) with respect to the minimum delay observed on the stream over the specified source and 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 3: PDV value on source-destination endpoint pairs

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

{
  "cost-type": {"cost-mode" : "numerical",
   "cost-metric" : "pdv"},
  "endpoints" : {
    "srcs": [ "ipv4:192.0.2.2" ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv6:2000::1:2345:6789:abcd"
    ]
  }
}
HTTP/1.1 200 OK
 Content-Length: TBA
 Content-Type: application/alto-endpointcost+json
{
  "meta": {
           "cost type": {
           "cost-mode": "numerical",
           "cost-metric":"delayjitter"
    }
   },
  "endpoint-cost-map": {
           "ipv4:192.0.2.2": {
           "ipv4:192.0.2.89"    : 0
           "ipv4:198.51.100.34" : 1
           "ipv6:2000::1:2345:6789:abcd"  : 5
         }
      }
   }
```

3.3.4. Measurement Considerations and Parameters

See Section 5 of [I-D.ietf-ippm-initial-registry] for measurement considerations and parameters which may be specified in "parameters". Note that the "parameters" field is an optional field providing non-normative information.

3.4. Cost Metric: Hop Count

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

3.4.1. Identifier

The identifier for this performance metric is "hopcount".

3.4.2. Value Representation

The metric value type is a single 'JSONNumber' type value conforming to the number specification of [RFC8259] Section 6. The number MUST be an integer and non-negative. The value represents the number of hops.

3.4.3. Intended Semantics and Use

Intended Semantics: To specify the number of hops in the path between the source endpoint and the destination endpoint. The hop count is a basic measurement of distance in a network and can be exposed as Router Hops, in direct relation to the routing protocols originating this information.

Use: This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own 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: TBA
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",
        "ipv6:2000::1:2345:6789:abcd"
      ]
    }
  }

HTTP/1.1 200 OK
Content-Length: TBA
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,
           "ipv6:2000::1:2345:6789:abcd" : 2,
                             }
             }
 }
```

3.4.4. Measurement Considerations and Parameters

The hop count can be calculated based on the number of routers from the source endpoint through which data must pass to reach the destination endpoint. This count can be measured at the source endpoint by traceroute.

Upon need, the traceroute can use UDP probe message or other implementations that use ICMP and TCP to discover the hop counts along the path from source endpoint to destination endpoint.

3.5. Cost Metric: Packet Loss

3.5.1. Identifier

The identifier for this performance metric is "pktloss".

3.5.2. Value Representation

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

3.5.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal aggregated packet loss over the specified source and 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 5: pktloss value on source-destination endpoint pairs

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

  {
    "cost-type": {"cost-mode" : "numerical",
     "cost-metric" : "pktloss"},
    "endpoints" : {
      "srcs": [ "ipv4:192.0.2.2" ],
      "dsts": [
        "ipv4:192.0.2.89",
        "ipv4:198.51.100.34",
        "ipv6:2000::1:2345:6789:abcd"
      ]
    }
  }
```

```
HTTP/1.1 200 OK
Content-Length: TBA
Content-Type: application/alto-endpointcost+json
{
    "meta": {
               "cost type": {
             "cost-mode": "numerical",
             "cost-metric":"pktloss"}
       }
    },
   "endpoint-cost-map": {
           "ipv4:192.0.2.2": {
           "ipv4:192.0.2.89"   : 0,
           "ipv4:198.51.100.34": 0,
           "ipv6:2000::1:2345:6789:abcd" : 0,
                             }
             }
 }
```

3.5.4. Measurement Considerations and Parameters

See Section 4 of [I-D.ietf-ippm-initial-registry] for measurement considerations and parameters which may be specified in "parameters". Note that the "parameters" field is an optional field providing non-normative information.

3.6. Cost Metric: Throughput

3.6.1. Identifier

The identifier for this performance metric is "throughput".

3.6.2. Value Representation

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

3.6.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal throughput over the specified source and 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 5: throughtput value on source-destination endpoint pairs

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


```
{
  "cost-type": {"cost-mode" : "numerical",
                "cost-metric" : "throughput"},
  "endpoints" : {
    "srcs": [ "ipv4:192.0.2.2" ],
    "dsts": [
       "ipv4:192.0.2.89",
       "ipv4:198.51.100.34",
       "ipv6:2000::1:2345:6789:abcd"
    ]
  }
}
```


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


```
{
  "meta": {
     "cost type": {
        "cost-mode": "numerical",
        "cost-metric":"throughput"
    }
  }
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89"   : 25.6,
      "ipv4:198.51.100.34": 12.8,
      "ipv6:2000::1:2345:6789:abcd" : 42.8,
  }
}
```


3.6.4. Measurement Considerations and Parameters

See Section 3.3 of [RFC6349] for measurement method and parameters which may be specified in "parameters". Note that the "parameters" field is an optional field providing non-normative information.

4. Traffic Engineering Performance Cost Metrics

This section introduces ALTO network performance metrics that may be aggregated from network metrics measured on links and specified in other documents. In particular, the bandwidth related metrics specified in this section are only available through link level measurements. For some of these metrics, the ALTO Server may further expose aggregated values while specifying the aggregation laws.

4.1. Cost Metric: Link Maximum Reservable Bandwidth

4.1.1. Identifier

The identifier for this performance metric is "maxresbw".

4.1.2. Value Representation

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

4.1.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal maximum reservable bandwidth over the specified source and destination. The value is corresponding to the maximum bandwidth that can be reserved (motivated from RFC 3630 Sec. 2.5.7.). The spatial aggregation unit 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 6: maxresbw value on source-destination endpoint pairs

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

  {
    "cost-type" { "cost-mode":  "numerical",
                  "cost-metric":  "maxresbw"},
    "endpoints":  {
      "srcs": [ "ipv4 : 192.0.2.2" ],
      "dsts": [
        "ipv4:192.0.2.89",
        "ipv4:198.51.100.34",
        "ipv6:2000::1:2345:6789:abcd"
      ]
    }
  }

HTTP/1.1 200 OK
Content-Length: TBA
Content-Type: application/alto-endpointcost+json
{
    "meta": {
           "cost-type": {
           "cost-mode": "numerical",
           "cost-metric": "maxresbw"
           }
    },
  " endpoint-cost-map": {
          "ipv4:192.0.2.2" {
          "ipv4:192.0.2.89" :    0,
          "ipv4:198.51.100.34": 2000,
          "ipv6:2000::1:2345:6789:abcd":  5000,
                            }
           }
}
```

4.1.4. Measurement Considerations and Parameters

Method of Measurement or Calculation:

Maximum Reservable Bandwidth is the bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors. See Section 3.5 of [RFC5305] for Measurement Method.

Measurement Point(s) with Potential Measurement Domain:

See Section 4.1 this document for discussions.

Measurement Timing:

See Section 3.5 of [RFC5305] and Section 5 of [RFC7810] for Measurement Timing.

4.2. Cost Metric: Link Residue Bandwidth

4.2.1. Identifier

The identifier for this performance metric is "residuebw".

4.2.2. Value Representation

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

4.2.3. Intended Semantics and Use

Intended Semantics: To specify spatial and temporal residual bandwidth over the specified source and destination. The value is calculated by subtracting tunnel reservations from Maximum Bandwidth (motivated from [RFC7810], Section 4.5.). The spatial aggregation unit 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 used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 7: residuebw value on source-destination endpoint pairs

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

  {
   "cost-type": { "cost-mode":  "numerical",
                  "cost-metric":  "residuebw"},
   "endpoints":  {
     "srcs": [ "ipv4 : 192.0.2.2" ],
     "dsts": [
       "ipv4:192.0.2.89",
       "ipv4:198.51.100.34",
       "ipv6:2000::1:2345:6789:abcd"
     ]
   }
  }

HTTP/1.1 200 OK
Content-Length: TBA
Content-Type: application/alto-endpointcost+json
{
   "meta": {
          "cost-type" {
          "cost-mode": "numerical",
          "cost-metric": "residuebw"
        }
    },
"endpoint-cost-map" {
         "ipv4:192.0.2.2" {
         "ipv4:192.0.2.89" :    0,
         "ipv4:198.51.100.34": 2000,
         "ipv6:2000::1:2345:6789:abcd":  5000,
                       }
        }
}
```

4.2.4. Measurement Considerations and Parameters

Method of Measurement or Calculation:

Residue Bandwidth is the Unidirectional Residue bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors. See Section 4.5 of [RFC7810] for Measurement Method.

Measurement Point(s) with Potential Measurement Domain:

See Section 4.1 of this document.

Measurement Timing:

See Section 5 of [RFC7810] for Measurement Timing.

5. 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 below.

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, we specify common issues unless one metric has its unique challenges.

5.1. 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" or "sla", where "estimation" is a type of measurement data source 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 [RFC7971]. In particular, the cost metrics defined in this document can be computed using routing systems as the data sources. Mechanisms defined in [RFC2681], [RFC3393], [RFC7679], [RFC7680], [RFC3630], [RFC3784], [RFC7471], [RFC7810], [RFC7752] and [I-D.ietf-idr-te-pm-bgp] that allow an ALTO Server to retrieve and derive the necessary information to compute the metrics that we describe in this document.

5.2. 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 which has "cost-source" being "sla", but the backward client will ignore this field and consider the metric estimation.

To address this issue, the only defined "routingcost" metric can be ONLY "estimation".

5.3. 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.

5.3.1. Configuration Parameters Considerations

Performance metrics often depend on configuration parameters. For example, the value of packet loss rate depends on the measurement interval and varies over time. To handle this issue, an ALTO server may collect data on time periods covering the previous and current time or only collect data on present time. The ALTO server may further aggregate these data to provide an abstract and unified view that can be more useful to applications. To make the ALTO client better understand how to use these performance data, the ALTO server may provide the client with the validity period of the exposed metric values.

5.3.2. Availability Considerations

Applications value information relating to bandwidth availability whereas bandwidth related metrics can often be only measured at the link level. This document specifies a set of link-level bandwidth related values that may be exposed as such by an ALTO server. The server may also expose other metrics derived from their aggregation and having different levels of endpoint granularity, e.g., link endpoints or session endpoints. The metric specifications may also expose the utilized aggregation laws.

6. 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 Authenticity and Integrity of ALTO Information", "15.2 Potential Undesirable Guidance from Authenticated ALTO Information" and "15.3 Confidentiality of ALTO Information" remain of utmost importance. Indeed, TE performance is a highly sensitive ISP information, therefore, sharing TE metric values in numerical mode requires full mutual confidence between the entities managing the ALTO Server and Client. Numerical TE performance information will most likely be distributed by ALTO Servers to 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.

7. 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 <http://www.iana.org/assignments/alto-protocol/alto-protocol.xhtml#cost-metrics>. This document requests to add the following entries to "ALTO Cost Metric Registry".

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

This document requests the creation of the "ALTO Cost Source Registry" with the following currently defined values:

Identifier	Intended Semantics
estimation	Values by estimation
nominal	Values in nominal cases
sla	Values reflecting service level agreement

8. Acknowledgments

The authors of this document would also like to thank Brian Trammell, Haizhou Du, Kai Gao, Lili Liu, Li, Geng, Danny Alex Lachos Perez for the reviews and comments. Young Lee is an author of an earlier version of the document.

9. References

9.1. Normative References

[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>.

[RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, DOI 10.17487/RFC2679, September 1999, <https://www.rfc-editor.org/info/rfc2679>.

[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>.

[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>.

[RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, DOI 10.17487/RFC6349, August 2011, <https://www.rfc-editor.org/info/rfc6349>.

[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>.

[RFC7810] Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 7810, DOI 10.17487/RFC7810, May 2016, <https://www.rfc-editor.org/info/rfc7810>.

[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>.

9.2. Informative References

[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>.

[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>.

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
USA

Email: yry@cs.yale.edu

Dhruv Dhody
Huawei
Leela Palace
Bangalore, Karnataka 560008
INDIA

Email: dhruv.ietf@gmail.com

Sabine Randriamasy
Nokia Bell Labs
Route de Villejust
Nozay 91460
FRANCE

Email: sabine.randriamasy@nokia-bell-labs.com

Luis Miguel Contreras Murillo
Telefonica

Email: luismiguel.contrerasmurillo@telefonica.com