

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 25, 2018

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
D. Dhody
Huawei
S. Randriamasy
Nokia Bell Labs
December 22, 2017

ALTO Performance Cost Metrics
draft-ietf-alto-performance-metrics-03

Abstract

Cost Metric is a basic concept in Application-Layer Traffic Optimization (ALTO). It is used in both the Cost Map Service and the Endpoint Cost Service.

Different applications may benefit from different Cost Metrics. For example, a Resource Consumer may prefer Resource Providers that offer a low delay delivery to the Resource Consumer. However the base ALTO protocol [ALTO] has documented only one single cost metric, i.e., the generic "routingcost" metric (Sec. 14.2 of ALTO base specification [ALTO]).

This document, proposes a set of Cost Metrics, derived and aggregated 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. It currently documents Network Performance Cost Metrics reporting on network delay, jitter, packet loss, hop count, and bandwidth. These metrics may be exposed by an ALTO Server to allow applications to determine "where" to connect based on network performance criteria. Additional Cost Metrics involving ISP specific considerations or other network technologies may be documented in further versions of this draft.

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 [RFC 2119](#) [[RFC2119](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 25, 2018.

Copyright Notice

Copyright (c) 2017 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Challenges on data sources and computation of ALTO performance metrics	5
2.1.	Data sources Challenge	5
2.2.	ALTO performance metrics Computation Challenges	5
2.2.1.	Configuration Parameters Challenge	5
2.2.2.	Availability of end to end path values Challenge	6
3.	Cost Metric: POWDelay	6
4.	Cost Metric: RTT	8
5.	Cost Metric: PDV	10
6.	Cost Metric: Hop Count	12
7.	Cost Metric: Packet Loss	14
8.	Traffic Engineering Performance Cost Metrics	16
8.1.	Cost Metric: Link Maximum Reservable Bandwidth	17

8.2.	Cost Metric: Link Residue Bandwidth	18
8.3.	Cost Metric: Link Available Bandwidth	20
8.4.	Cost Metric: Link Utilized Bandwidth	23
9.	Security Considerations	25
10.	IANA Considerations	25
11.	References	26
11.1.	Normative References	26
11.2.	Informative References	27
Appendix A.	Open Issue List	28
	Authors' Addresses	28

[1.](#) Introduction

Cost Metric is a basic concept in Application-Layer Traffic Optimization (ALTO). It is used in both the Cost Map Service and the Endpoint Cost Service. In particular, applications may benefit from knowing network performance measured on several Cost Metrics. For example, a more delay-sensitive application may focus on latency, and a more bandwidth-sensitive application may focus on available bandwidth.

This document introduces a set of new cost metrics, listed in Table 1, to support the aforementioned applications and allow them to determine "where" to connect based on network performance criteria. Hence, this document extends the base ALTO protocol [ALTO], which defines only a single cost metric, i.e., the generic "routingcost" metric (Sec. 14.2 of ALTO base specification [ALTO]).

Namespace	Property	Reference
	owdelay	See Section 3 , [RFC2679] Section 3.6
	rtt	See Section 4 , [RFC2681] Section 2.6
	pdv	See Section 5 , [RFC3393] Section 2.6
	hopcount	See Section 6 , [RFC7285]
	pktloss	See Section 7 , [RFC7680] Section 2.6
	maxresbw	See Section 8.1 , [RFC5305] Section 3.5
	residbw	See Section 8.2 , [RFC7810] Section 4.5
	availbw	See Section 8.3 , [RFC7810] Section 4.6
	utilbw	See Section 8.4 , [RFC7810] Section 4.7

Table 1.

The purpose of this draft is to list the metrics likely to be exposed to ALTO Clients, including those already specified in other standardization groups and as such it does not claim novelty on all the specified metrics. Some metrics may have values produced by standard measurement methods such as those specified in IPPM, some

may be ISP dependent such as those registered in ISIS or OSPF-TE. In this case, this document will refer to the relevant specifications.

An ALTO server may provide a subset of the cost metrics described in this document. These cost metrics can be retrieved and aggregated from routing protocols or other traffic measurement management tools (See Figure 1). Note that these cost metrics are optional and not all them need to be exposed to applications. For example, those that are subject to privacy concerns should not be provided to unauthorized ALTO clients.

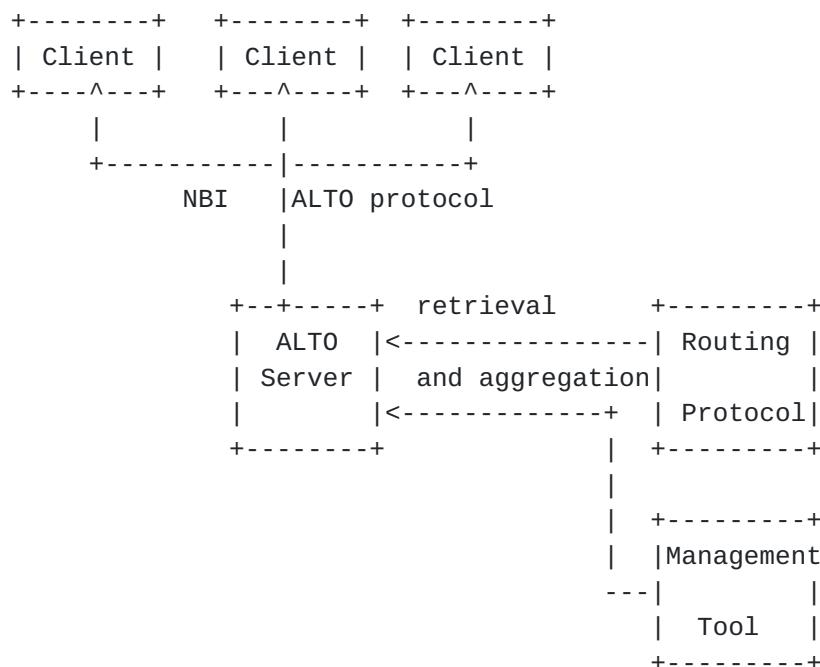


Figure 1. End-to-End Path Cost Metrics Exposing

When an ALTO server supports a cost metric defined in this document, it MUST announce this metric in its IRD.

Additionally, future versions of this document may define network metric values that stem from both measurements and provider policies such as many metrics related to end-to-end path bandwidth.

As for the reliability and trust in the exposed metric values, applications SHOULD rapidly give up using ALTO-based guidance if they feel the exposed information does not preserve their performance level or even degrades it.

Following the ALTO base protocol, this document uses JSON to specify the value type of each defined metric. See [[RFC4627](#)] for JSON data type specification.

2. Challenges on data sources and computation of ALTO performance metrics

2.1. Data sources Challenge

An ALTO server needs data sources to compute the cost metrics described in this document. This document does not define the exact data sources. For example, the ALTO server may use log servers or the OAM system as its data source [ALTO-DEPLOYMENT]. In particular, the cost metrics defined in this document can be computed using routing systems as the data sources. Mechanisms defined in [[RFC3630](#)], [[RFC3784](#)], [OSPF-TE], [ISIS-TE], [BGP-LS] and [BGP-PM] that allow an ALTO Server to retrieve and derive the necessary information to compute the metrics that we describe in this document.

One challenge lies in the data sources originating the ALTO metric values. The very important purpose of ALTO is to guide application traffic with provider network centric information that may be exposed to ALTO Clients in the form of network performance metric values. Not all of these metrics have values produced by standardized measurement methods or routing protocols. Some of them involve provider-centric policy considerations. Some of them may describe wireless or cellular networks. To reliably guide users and applications while preserving provider privacy, ALTO performance metric values may also add abstraction to measurements or provide unitless performance scores.

2.2. ALTO performance metrics Computation Challenges

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 computation of ALTO performance metrics.

2.2.1. Configuration Parameters Challenge

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.

2.2.2. Availability of end to end path values Challenge

Applications value information relating to bandwidth availability where as 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.

3. Cost Metric: POWDelay

Metric name:

Periodic One Way Delay

Metric Description:

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).

Method of Measurement or Calculation:

See section 8.3 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Method.

Units of Measurement:

See section 8.4.3 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Unit. The unit is expressed in seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See section 8.3.5 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Timing.

Use and Applications:

The Metric value Type is a single 'JSONNumber' type value containing a non-negative integer component that may be followed by an exponent part. The Cost Mode is encoded as a US-ASCII string.

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" : "powdelay"},
  "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" : "powdelay"
    }
  },
  "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,
    }
  }
}
```

4. Cost Metric: RTT

Metric name:

Round Trip Delay

Metric Description:

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).

Method of Measurement or Calculation:

See section 4.3 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Method.

Units of Measurement:

See section 4.4.3 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Unit. The unit is expressed in seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See section 4.3.5 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

Example 2: Round Trip 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,
    }
  }
}
```

5. Cost Metric: PDV

Metric name:

Packet Delay Variation

Metric Description:

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).

Method of Measurement or Calculation:

See section 5.3 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Method.

Units of Measurement:

See section 5.4.4 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Unit. The unit is expressed in seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See section 5.3.5 of [[I-D.ietf-ippm-initial-registry](#)] for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

Example 3: Delay jitter 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" : "delayjitter"},
  "endpoints" : {
    "srcs": [ "ipv4:192.0.2.2" ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv6:2000::1:2345:6789:abcd"
    ]
  }
}
```


Example 3: Delay jitter 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" : "delayjitter"},
  "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
    }
  }
}
```

6. Cost Metric: Hop Count

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

Metric name:

Hop count

Metric Description:

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, IP hops or other hops in direct relation to the routing protocols originating this information. It might also result from the aggregation of such information.

Method of Measurement or Calculation:

See [section 2.2](#), Computation of metrics.

Units of Measurement:

The unit is integer number.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See [section 2.1](#), second paragraph for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

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,
    }
  }
}
```

7. Cost Metric: Packet Loss

Metric name:

Packet loss

Metric Description:

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).

Method of Measurement or Calculation:

See [section 2.6 of \[RFC7680\]](#) for Measurement Method.

Units of Measurement:

The unit is percentile.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See [section 2](#) and section3 of [\[RFC7680\]](#) for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

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": 1,
      "ipv6:2000::1:2345:6789:abcd" : 2,
    }
  }
}
```

8. 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.

8.1. Cost Metric: Link Maximum Reservable Bandwidth

Metric name:

Maximum Reservable Bandwidth

Metric Description:

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).

Method of Measurement or Calculation:

Maximum Reserveable Bandwidth is the bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors, See [section 3.5 of \[RFC5305\]](#) for Measurement Method.

Units of Measurement:

The unit of measurement is byte per seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See [section 3.5 of \[RFC5305\]](#) and [section 5 of \[RFC7810\]](#) for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

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,
    }
  }
}
```

8.2. Cost Metric: Link Residue Bandwidth

Metric name:

Residue Bandwidth

Metric Description:

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 [\[RFC7810\]](#), Sec.4.5.). The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

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.

Units of Measurement:

The unit of measurement is byte per seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See [section 5 of \[RFC7810\]](#) for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

Example 7: residubw 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": "residubw"},
  "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": "residubw"
    }
  },
  "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,
    }
  }
}
```

8.3. Cost Metric: Link Available Bandwidth

Metric name:

Available Bandwidth

Metric Description:

To specify spatial and temporal available bandwidth over the specified source and destination. The value is calculated by subtracting the measured bandwidth used for the actual forwarding of best effort traffic from Residue Bandwidth (motivated from [\[RFC7810\]](#), Sec.4.6.). The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Method of Measurement or Calculation:

Available bandwidth is the Unidirectional Available bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors, See [section 4.6 of \[RFC7810\]](#) for Measurement Method.

Units of Measurement:

The unit of measurement is byte per seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

See [section 5 of \[RFC7810\]](#) for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application. Besides, knowledge about available bandwidth is essential for applications to distribute or schedule their transmissions. The example below illustrates how this metric is provided in the form of an ALTO calendar, as specified in [XXXX] to help deciding "where" and "when" to transmit.

Example 8: availbw value on source-destination endpoint pairs

This example assumes that the ALTO Server provides the values for metric "availbw" in the form of an ALTO calendar and declares it in its IRD.

```
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": "availbw"},
  "calendared" : [true],
  "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": "availbw"
    }
    "calendar-response-attributes" : [
      "calendar-start-time" : Tue, 1 Mar 2017 13:00:00 GMT,
      "time-interval-size" : "1 hour",
      "numb-intervals" : 8
    ]
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2" {
      "ipv4:192.0.2.89" : [6,5,7,8,4,10,7,6],
      "ipv4:198.51.100.34" : [7,4,6,8,5,9,6,7],
      "ipv6:2000::1:2345:6789:abcd" : [7,6,8,5,7,9,6,8],
    }
  }
}
```

8.4. Cost Metric: Link Utilized Bandwidth

Metric name:

Utilized Bandwidth

Metric Description:

To specify spatial and temporal utilized bandwidth over the specified source and destination. The value is corresponding to the actual measured bandwidth used for all traffic (motivated from [\[RFC7810\]](#), Sec.4.7.). The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Method of Measurement or Calculation:

Link Utilized bandwidth is Unidirectional utilization bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors, See [section 4.7 of \[RFC7810\]](#) for Measurement Method.

Units of Measurement:

The unit of measurement is byte per seconds.

Measurement Point(s) with Potential Measurement Domain:

See [section 2.1](#), Data sources.

Measurement Timing:

Link Utilized bandwidth is Unidirectional utilization bandwidth measured between two directly connected IS-IS neighbors or OSPF neighbors, See [section 5 of \[RFC7810\]](#) for Measurement Timing.

Use and Applications:

See [section 3](#) for use and application.

Example 9: utilbw 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" : "utilbw"},
  "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": "utilbw"
    }
  },
  "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,
    }
  }
}
```

9. Security Considerations

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

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.

10. 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 Meric Registry".

Namespace	Property	Reference
	owdelay	[thisdraft] Section 3 , [RFC2679] Section 3.6
	rtt	[thisdraft] Section 4 , [RFC2681] , Section 2.6
	pdv	[thisdraft] Section 5 , [RFC3393] , Section 2.6
	hopcount	[thisdraft] Section 6 , [RFC7285]
	pktloss	[thisdraft] Section 7 , [RFC7680] , Section 2.6
	maxresbw	[thisdraft] Section 8.1 , [RFC5305] , Section 3.5
	residbw	[thisdraft] Section 8.2 , [RFC7810] , Section 4.5
	availbw	[thisdraft] Section 8.3 , [RFC7810] , Section 4.6
	utilbw	[thisdraft] Section 8.4 , [RFC7810] , Section 4.7

11. References

11.1. Normative References

[I-D.ietf-idr-te-pm-bgp]

Ginsberg, L., Previdi, S., Wu, Q., Gredler, H., Ray, S., Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions", [draft-ietf-idr-te-pm-bgp-08](#) (work in progress), August 2017.

[I-D.ietf-ippm-initial-registry]

Morton, A., Bagnulo, M., Eardley, P., and K. D'Souza, "Initial Performance Metric Registry Entries", [draft-ietf-ippm-initial-registry-05](#) (work in progress), October 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

[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>>.

- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", [RFC 4627](#), DOI 10.17487/RFC4627, July 2006, <<https://www.rfc-editor.org/info/rfc4627>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [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>>.
- [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>>.
- [RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, [RFC 7680](#), DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [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>>.

11.2. Informative References

- [I-D.ietf-alto-deployments]
Stiemerling, M., Kiesel, S., Scharf, M., Seidel, H., and S. Previdi, "ALTO Deployment Considerations", [draft-ietf-alto-deployments-16](#) (work in progress), July 2016.

[RFC6390] Clark, A. and B. Claise, "Framework for Performance Metric Development", [RFC 6390](#), July 2011.

[Appendix A](#). Open Issue List

We need to consider to add Cellular endpoint format support in the example, the Cellular endpoint format is specified in [draft-randriamasy-alto-cellular-adresses](#).

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

Young Lee
Huawei
1700 Alma Drive, Suite 500
Plano, TX 75075
USA

Email: leeyoung@huawei.com

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