

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 8, 2017

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
D. Dhody
Huawei
S. Randriamasy
Nokia Bell Labs
July 7, 2016

ALTO Performance Cost Metrics
draft-wu-alto-te-metrics-08

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. Future extensions to ALTO may also use Cost Metric.

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

In this document, we propose 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 tool. We currently define 11 new Performance Metric to measure network delay, jitter, packet loss, hop count, and bandwidth. The metrics defined in this document provide a relatively comprehensive set of Cost Metrics for ALTO and allow applications to determine "where" to connect based on end to end network performance criteria. Additional Cost Metrics such as financial cost metrics may be defined in other documents.

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 <http://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 January 8, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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. | Data sources, computation of defined cost metrics | 5 |
| 2.1. | Data sources | 5 |
| 2.2. | Computation of metrics | 5 |
| 3. | Cost Metric: Delay | 6 |
| 4. | Cost Metric: Delay Jitter | 7 |
| 5. | Cost Metric: Packet Loss | 9 |
| 6. | Cost Metric: Hop Count | 11 |
| 7. | Cost Metric: Bandwidth | 13 |
| 8. | Cost Metric: Maximum Bandwidth | 14 |
| 9. | Cost Metric: Maximum Reservable Bandwidth | 16 |
| 10. | Cost Metric: Unreserved Bandwidth | 18 |
| 11. | Cost Metric: Residue Bandwidth | 21 |
| 12. | Cost Metric: Available Bandwidth | 22 |

| | | |
|-----------------------|---|--------------------|
| 13. | Cost Metric: Utilized Bandwidth | 24 |
| 14. | Security Considerations | 26 |
| 15. | IANA Considerations | 27 |
| 16. | References | 27 |
| 16.1. | Normative References | 27 |
| 16.2. | Informative References | 28 |
| | Authors' Addresses | 29 |

[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.

The objective of this document is to introduce 11 new performance cost metrics, listed in Table 1, to support the aforementioned applications and allow applications to determine "where" to connect based on end to end 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 |
|-----------|-------------|---------------------------------------|
| | delay | [RFCxxxx], Section 3 |
| | delayjitter | [RFCxxxx], Section 4 |
| | pktloss | [RFCxxxx], Section 5 |
| | hopcount | [RFCxxxx], Section 6 |
| | bandwidth | [RFCxxxx], Section 7 |
| | maxbw | [RFCxxxx], Section 8 |
| | maxresbw | [RFCxxxx], Section 9 |
| | unresdbw | [RFCxxxx], Section 10 |
| | residbw | [RFCxxxx], Section 11 |
| | availbw | [RFCxxxx], Section 12 |
| | utilbw | [RFCxxxx], Section 13 |

Table 1.

An ALTO server may provide a subset of the cost metrics defined in this document. These cost metrics can be retrieved and aggregated from routing protocol or other traffic measurement management tool (See Figure 1).

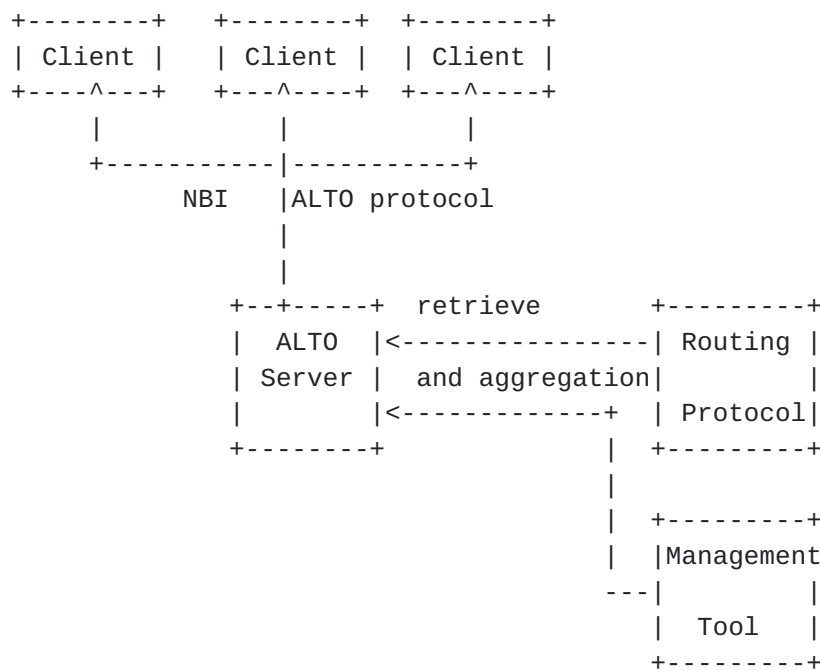


Figure 1. End to End Path Cost Metrics Exposing

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

The definitions of a set of cost metrics can allow us to extend the ALTO base protocol (e.g., allowing output and constraints use different cost metrics), but such extensions are not in the scope of this document.

One challenge in describing the metrics is that 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, ALTO server may collect data on time periods covering the past, present or only collect data on present time. The ALTO may further aggregate these data to provide an abstract and unified view that can be more useful to applications. To make the ALTO client understand whether the performance data is past data or present data, the ALTO server needs to expose to the client the validity period of each performance metric.

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. Data sources, computation of defined cost metrics

The cost metrics described in this document are similar, in that they may use similar data sources and have similar issues in their calculation. Hence, instead of specifying such issues for each metric individually, we specify the common issue in this section.

2.1. Data sources

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 described in this document.

2.2. Computation of metrics

An ALTO server processes measurements from data sources to compute exposed metrics. It may need performance data processing tasks such as aggregating the results across multiple systems, removing outliers, and creating additional statistics.

One specific challenge in deriving the metrics in this document is that these performance metrics depend on some configuration parameters. For example, the value of packet loss rate depends on the measurement interval and varies over time. If the ALTO server uses aforementioned routing protocol based mechanisms as data sources, then the measurement interval may be preconfigured by the routing protocol. For example, [Section 5](#) of [ISIS-TE] defines a default measurement interval of 30 seconds. This document uses the term Measurement Interval to refer to the measurement interval used by the data sources. In the [ISIS-TE] case, it is a measurement interval set by routing protocol. The Measurement Interval(s) of the data sources can be different from the interval that this document derives the metric, e.g., the interval used by this document is multiple of measurement interval of the data sources. Hence, an ALTO server needs to resolve the mismatch, when it happens.

Another issue of converting from data source measurements to ALTO exposed metric values is that the measurement results that the ALTO Server retrieves may be defined for only links, and hence, the server will need to compose the link metrics to obtain path metrics used in services such as the Cost Map Service. In this definition, we define

the metrics to be independent of link or path, considering that future ALTO extensions may define link-based services, and hence the defined metrics should still be usable.

3. Cost Metric: Delay

Metric name:

Delay

Metric Description:

To specify spatial and temporal aggregated delay between the specified source and destination or the time that the packet spends to travel from source to destination. The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

The unit is microsecond.

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:

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

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" : "delay"},
  "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" : "delay"}
  },
  "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: Delay Jitter

Metric name:

Delay jitter

Metric Description:

To specify spatial and temporal aggregated jitter (latency variation) over the specified source and destination. The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

The unit is microsecond.

Measurement Point(s) with Potential Measurement Domain:

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

Measurement Timing:

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

Measurement Timing:Use and Applications:

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

Example 2: Delayjitter 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
    }
  }
}
```

5. 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 unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

The unit is percentile.

Measurement Point(s) with Potential Measurement Domain:

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

Measurement Timing:

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

Use and Applications:

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

Example 3: 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,
    }
  }
}
```

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: Bandwidth

Metric name:

Bandwidth

Metric Description:

To specify spatial and temporal aggregated bandwidth over the specified source and destination. The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endhost to endhost); and the temporal unit is specified as the measurement interval in the query context.

This is just a definition of a class of cost metric 'bandwidth'. The use of this cost metric is always in conjunction with what it represents, which could be Max Bandwidth (maxbw), Residual Bandwidth (residuebw) etc.

Method of Measurement or Calculation:

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

Units of Measurement:

The units are bytes per second.

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.

8. Cost Metric: Maximum Bandwidth

Metric name:

Maximum Bandwidth

Metric Description:

To specify spatial and temporal maximum bandwidth over the specified source and destination. The values correspond to the

maximum bandwidth that can be used (motivated from [RFC 3630](#) Sec. 2.5.6.). The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endhost to endhost); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition for the Bandwidth Cost Metric.

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 5: maxbw 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": "maxbw"},
  "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": "maxbw"
    }
  },
  "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. Cost Metric: 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); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition of the Bandwidth Cost Metric.

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 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,
    }
  }
}
```

10. Cost Metric: Unreserved Bandwidth

Metric name:

Unreserved Bandwidth

Metric Description:

To specify spatial and temporal unreserved bandwidth over the specified source and destination. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7. Therefore this metric is encoded as an array of 8 values. The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition for the bandwidth Cost Metric.

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 7: unresbw value on source-destination endpoint pairs

In this example, the Collection method specifies that the 'unresbw' values are defined as the 'unavailable bandwidth' specified in [section 2.5.8 of RFC3630](#): 8 unavailable bandwidth value are reported in the same OSPF message using the same TLV. Each value is corresponding to the bandwidth that can be reserved with a setup priority of 0 through 7.

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": "unresbw[1,8]" },
  "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": "unresbw[1,8]"
    }
  },
  "endpoint-cost-map" {
    "ipv4:192.0.2.2" {
      "ipv4:192.0.2.89" : [0,0,0,0,0,0,0,0],
      "ipv4:198.51.100.34": [0,0,0,0,0,0,0,2000],
      "ipv6:2000::1:2345:6789:abcd": [0,0,0,0,0,0,0,5000],
    }
  }
}
```


11. Cost Metric: 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 [I-D. ietf-isis-te-metric-extensions], Sec.4.5.). The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition of the general Bandwidth.

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 8: 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,
    }
  }
}
```

12. Cost Metric: 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 [I-D. ietf-isis-te-metric-extensions], Sec.4.6.). The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition of the general Bandwidth.

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 9: availbw 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": "numeric",
  "cost-metric": "availbw"},
  "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": "numeric",
      "cost-metric": "availbw"
    }
  },
  "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],
    }
  }
}
```

13. Cost Metric: 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 [I-D. ietf-isis-te-metric-extensions], Sec.4.7.). The spatial aggregation unit is specified in the query context (e.g., PID to PID, or endpoint to endpoint); and the temporal unit is specified as the measurement interval in the query context.

Method of Measurement or Calculation:

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

Units of Measurement:

See definition of the general Bandwidth.

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 10: 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,
    }
  }
}
```

14. 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 and 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. One 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.

15. IANA Considerations

IANA has added the following entries to the ALTO cost map Properties registry, defined in [Section 3](#) of [RFCXXX].

| Namespace | Property | Reference |
|-----------|-------------|---------------------------------------|
| | delay | [RFCxxxx], Section 3 |
| | delayjitter | [RFCxxxx], Section 4 |
| | pktloss | [RFCxxxx], Section 5 |
| | hopcount | [RFCxxxx], Section 6 |
| | bandwidth | [RFCxxxx], Section 7 |
| | maxbw | [RFCxxxx], Section 8 |
| | maxresbw | [RFCxxxx], Section 9 |
| | unresdbw | [RFCxxxx], Section 10 |
| | residbw | [RFCxxxx], Section 11 |
| | availbw | [RFCxxxx], Section 12 |
| | utilbw | [RFCxxxx], Section 13 |

16. References

16.1. Normative References

[I-D.ietf-idr-te-pm-bgp]
 Previdi, S., Wu, Q., Gredler, H., Ray, S.,
 jeffrant@gmail.com, j., Filsfils, C., and L. Ginsberg,
 "BGP-LS Advertisement of IGP Traffic Engineering
 Performance Metric Extensions", [draft-ietf-idr-te-pm-bgp-03](#) (work in progress), May 2016.

- [I-D.ietf-isis-te-metric-extensions]
Previdi, S., Giacalone, S., Ward, D., Drake, J., and W. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", [draft-ietf-isis-te-metric-extensions-11](#) (work in progress), February 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", [RFC 4627](#), DOI 10.17487/RFC4627, July 2006, <<http://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, <<http://www.rfc-editor.org/info/rfc5234>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc7471>>.
- [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, <<http://www.rfc-editor.org/info/rfc7752>>.

16.2. 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-15](#) (work in progress), May 2016.
- [RFC6390] Clark, A. and B. Claise, "Framework for Performance Metric Development", [RFC 6390](#), July 2011.

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

