

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 18, 2014

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
D. Dhody
Huawei
S. Randriamasy
Alcatel-Lucent
February 14, 2014

ALTO Traffic Engineering Cost Metrics
draft-wu-alto-te-metrics-01

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 define eleven Cost Metrics, derived from OSPF-TE and ISIS-TE, 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 focusing on traffic engineering (TE). Additional Cost Metrics such as financial cost metrics may be defined in other documents.

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 August 18, 2014.

Copyright Notice

Copyright (c) 2014 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	4
2.	Conventions used in this document	5
3.	Metric: Delay	6
4.	Metric: Delayjitter	9
5.	Metric: Pktloss	12
6.	Metric: Hopcount	15
7.	Metric: Bandwidth	16
8.	Metric: Maximum Bandwidth	18
9.	Metric: Maximum Reserved Bandwdith	21
10.	Metric: Unavailable Reserved Bandwidth	24
11.	Metric: Residue Bandwidth	27
12.	Metric: Available Bandwidth	30
13.	Metric: Utilized Bandwidth	33
14.	Security Considerations	36
15.	IANA Considerations	37
16.	References	38
16.1.	Normative References	38
16.2.	Informative References	38
	Authors' Addresses	39

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.

In this document, we define eleven Cost Metrics, extending the base ALTO protocol [[ALTO](#)], which has defined only a single Cost Metric, i.e., the generic "routingcost" metric (Sec. 14.2 of ALTO base specification [[ALTO](#)]).

The Cost Metrics that we define in this document focus on traffic engineering. Additional metrics may be defined in other documents. In particular, the Cost Metrics that we define in this document can be gathered from routing systems; [[RFC3630](#)], [[RFC3784](#)], [[OSPF-TE](#)], [[ISIS-TE](#)], [[BGP-LS](#)] and [[BGP-PM](#)] define mechanisms that allow an ALTO Server to retrieve and derive the necessary information to provide the metrics that we define in this document.

Note that the metrics 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.

One challenge in defining 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 definitions of a set of Cost Metrics can allow us to extend the base protocol (e.g., allowing output and constraints use different Cost Metrics), but such extensions are not in the scope of this document.

2. Conventions used in this document

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](#) [[RFC2119](#)].

Syntax specifications shown here use the augmented Backus-Naur Form (ABNF) as described in [[RFC5234](#)], and are specified as in the base JSON specification [[RFC4627](#)].

3. Metric: Delay

Cost Metric name:

Delay

Metric Description:

To specify spatial and temporal aggregated delay between 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. The delay is the time that the packet spends to travel from source to destination.

Metric Unit:

The unit is microsecond.

Metric Value Type:

A single 'JSONNumber' type value containing a non-negative integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

Collection Method:

The delay metric is part of ALTO information provided by ALTO server. ALTO server may collect delay metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the delay metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as

aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, Either scheduling can be used so that the past and present results for the delay metrics over a period of time are recorded, e.g., measured every hours and aggregated, or only the latest update is recorded.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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,
      "ipv4:203.0.113.45" : 30,
    }
  }
}
```


4. Metric: Delayjitter

Cost Metric name:

Delayjitter

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 endhost to endhost); and the temporal unit is specified as the measurement interval in the query context.

Metric Unit:

The unit is microsecond.

Metric Value Type:

A single 'JSONNumber' type value containing an integer component that may be followed by exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string **MUST** either have the value 'numerical' or 'ordinal'.

Collection Method:

The delayjitter metric is part of ALTO information provided by ALTO server. ALTO server may collect delayjitter metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the delayjitter metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers,

create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for the delayjitter metrics over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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
      "ipv4:203.0.113.45" : 5
    }
  }
}
```


5. Metric: Pktloss

Cost Metric name:

pktloss

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 endhost to endhost); and the temporal unit is specified as the measurement interval in the query context.

Metric Unit:

The unit is percentile.

Metric Value Type:

A single number value containing an integer component that may be followed by a fraction part and/or an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

Collection Method:

The pktloss metric is part of ALTO information provided by ALTO server. ALTO server may collect pktloss metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the pktloss metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers,

create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
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,
      "ipv4:203.0.113.45" : 2,
    }
  }
}
```


6. Metric: Hopcount

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

Cost Metric name:

Hopcount

Metric Description:

To specify the number of hops in the path between the source endpoint and the destination endpoint.

Editor Note: Need to specify which level (AS, IP perhaps), details TBD for multiple-layer aspect.

Metric Unit:

The unit is integer number.

Metric Value Type:

A single 'JSONNumber' type value containing an integer component.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. string MUST either have the value 'numerical' or 'ordinal'.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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.

7. Metric: Bandwidth

Cost 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONNumber' type value containing an integer component that may be prefixed with an optional minus sign, which may be followed by a fraction part and/or an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. string MUST either have the value 'numerical' or 'ordinal'.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

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.

This is intended to be a constraint attribute value. It could be used as a cost metric constraint attribute used together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

8. Metric: Maximum Bandwidth

Cost Metric name:

Maximum Bandwidth

Metric Description:

To specify spatial and temporal maximum 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 context interval.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONNumber' type value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

Collection Method:

The Maximum Bandwidth metric is gathered using [[RFC3630](#)], [[RFC3784](#)] or [[BGP-LS](#)]. It is extended from Bandwidth cost metric and is part of ALTO information provided by ALTO server. ALTO server may collect Maximum Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the Maximum Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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,
      "ipv4:203.0.113.45": 5000,
    }
  }
}
```


9. Metric: Maximum Reserved Bandwidth

Cost Metric name:

Maximum Reserved Bandwidth

Metric Description:

To specify spatial and temporal maximum reserved 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONNumber' type value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string **MUST** either have the value 'numerical' or 'ordinal'.

Collection Method:

The Maximum Reserved Bandwidth is gathered using [[RFC3630](#)], [[RFC3784](#)] or [[BGP-LS](#)]. It is extended from Bandwidth Cost metric defined as part of ALTO information provided by ALTO server. ALTO server may collect Maximum Reserved Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the Maximum Reserved Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute value. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
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,
      "ipv4:203.0.113.45": 5000,
    }
  }
}
```


10. Metric: Unavailable Reserved Bandwidth

Cost Metric name:

Unavailable Reserved Bandwidth

Metric Description:

To specify spatial and temporal unavailable reserved 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONArray' type with each value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

Collection Method:

The Unavailable Reserved Bandwidth metric is gathered using [[RFC3630](#)], [[RFC3784](#)] or [[BGP-LS](#)]. It is extended from Bandwidth Cost metric and defined as part of ALTO information provided by ALTO server. ALTO server may collect Unavailable Reserved Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the Unavailable Reserved Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute values. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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],
      "ipv4:203.0.113.45": [0,0,0,0,0,0,0,5000],
    }
  }
}
```


11. Metric: Residue Bandwidth

Cost Metric name:

Residue Bandwidth

Metric Description:

To specify spatial and temporal residual 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONArray' type with each value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string **MUST** either have the value 'numerical' or 'ordinal'.

Collection Method:

The Residue Bandwidth metric is gathered using [[OSPF-TE](#)], [[ISIS-TE](#)] or [[BGP-PM](#)]. It is extended from Bandwidth Cost metric and is part of ALTO information provided by ALTO server. ALTO server may collect Residue Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the Residue Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute values. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7. It 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: 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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,
      "ipv4:203.0.113.45": 5000,
    }
  }
}
```


12. Metric: Available Bandwidth

Cost Metric name:

Available Bandwidth

Metric Description:

To specify spatial and temporal available 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONArray' type with each value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

Collection Method:

The Available Bandwidth metric is gathered using [[OSPF-TE](#)], [[ISIS-TE](#)] or [BGP- PM]. It is extended from Bandwidth Cost metric and is part of ALTO information provided by ALTO server. ALTO server may collect Available Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When the Available Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute values. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7. It 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 8: 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": "numerical",
  "cost-metric": "availbw"},
  "endpoints": {
    "srcs": [ "ipv4 : 192.0.2.2" ],
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv4:203.0.113.45"
    ]
  }
}
```

HTTP/1.1 200 OK

Content-Length: TBA

Content-Type: application/alto-endpointcost+json

```
{
  "meta": {
    "cost-type": {
      "cost-mode": "numerical",
      "cost-metric": "availbw"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2" {
      "ipv4:192.0.2.89" : 0,
      "ipv4:198.51.100.34": 2000,
      "ipv4:203.0.113.45": 5000,
    }
  }
}
```


13. Metric: Utilized Bandwidth

Cost Metric name:

Utilized Bandwidth

Metric Description:

To specify spatial and temporal utilized 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.

Metric Unit:

The units are bytes per second.

Metric Value Type:

A single 'JSONArray' type with each value containing an integer component that may be followed by an exponent part.

Cost Mode:

A Cost Mode is encoded as a US-ASCII string. The string **MUST** either have the value 'numerical' or 'ordinal'.

Collection Method:

The Utilized Bandwidth metric is gathered using [[OSPF-TE](#)], [[ISIS-TE](#)] or [[BGP-PM](#)]. It is extended from bandwidth cost metric and is part of ALTO information provided by ALTO server. ALTO server may collect Utilized Bandwidth metric from routing protocol or deploy data source to collect data to compute the delay metrics [[ALTO-DEPLOYMENT](#)]. The data source can be either log server or OAM system or P2P client, etc.

When The Utilized Bandwidth metric is collected from routing protocol, it is measured over a measurement interval preconfigured by routing protocol. As described in Section 5 of [[ISIS-TE](#)], the default measurement interval is set to 30 seconds.

ALTO server may also have additional data processing such as aggregating the results across multiple systems, remove outliers, create additional statistics.

When the data is collected from data source or collected from routing protocol, either scheduling can be used so that the past and present results for this metric over a period of time can be recorded, e.g., measured every hours, or only measure the latest update.

ALTO Sample interval:

The ALTO Server may collect values from the sources measurements done e.g. every second over a period of 30 seconds. The ALTO Server may then aggregate these values over ALTO Sample intervals of at least 30 seconds and provide updates every hour.

Use and Applications:

This is intended to be a constraint attribute values. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7. It 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 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",
      "ipv4:203.0.113.45"
    ]
  }
}
```

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,
      "ipv4:203.0.113.45" : 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. 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.

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
	jitter	[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

- [ALTO] Alimi, R., "ALTO Protocol", ID [draft-ietf-alto-protocol-16](#), May 2013.
- [BGP-LS] Gredler, H., "North-Bound Distribution of Link-State and TE Information using BGP", ID [draft-ietf-idr-ls-distribution-03](#), May 2013.
- [BGP-PM] Wu, Q., "BGP attribute for North-Bound Distribution of Traffic Engineering (TE) performance Metrics", ID [draft-wu-idr-te-pm-bgp-02](#), October 2013.
- [ISIS-TE] Giacalone, S., "ISIS Traffic Engineering (TE) Metric Extensions", ID [draft-ietf-isis-te-metric-extensions-01](#), October 2013.
- [OSPF-TE] Giacalone, S., "OSPF Traffic Engineering (TE) Metric Extensions", ID [draft-ietf-ospf-te-metric-extensions-04](#), June 2013.
- [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](#), July 2006.
- [RFC5234] Crocker, D., "Augmented BNF for Syntax Specifications: ABNF", [RFC 5234](#), January 2008.

16.2. Informative References

- [ALTO-DEPLOYMENT] Stiemerling, M., Kiesel, S., Previdi, S., and M. Scharf, "ALTO Deployment Considerations", ID [draft-ietf-alto-deployments-08](#), October 2013.
- [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: sunseawq@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
Alcatel-Lucent
Route de Villejust
Nozay, 91460
FRANCE

Email: Sabine.Randriamasy@alcatel-lucent.com

