

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

Q. Wu
Huawei
Y. Yang
Yale University
Y. Lee
D. Dhody
Huawei
S. Randriamasy
Nokia Bell Labs
July 08, 2019

ALTO Performance Cost Metrics
draft-ietf-alto-performance-metrics-07

Abstract

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

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

ALTO cost metrics can be generic metrics and this document focuses on network performance metrics, including network delay, jitter, packet loss, hop count, and bandwidth. These metrics can be 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. These metrics may then be exposed by an ALTO Server to allow applications to determine "where" to connect based on network performance criteria. Additional cost metrics may be documented 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 9, 2020.

Copyright Notice

Copyright (c) 2019 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.	Network Performance Cost Metrics	5
2.1.	Cost Metric: One Way Delay (owdelay)	5
2.1.1.	Intended Semantics	6
2.1.2.	Use and Example	6
2.1.3.	Measurement Considerations	7
2.2.	Cost Metric: RoundTrip Time (rtt)	7
2.2.1.	Intended Semantics	8
2.2.2.	Use and Example	8
2.2.3.	Measurement Considerations	9
2.3.	Cost Metric: Packet Delay Variation (pdv)	9
2.3.1.	Intended Semantics	10
2.3.2.	Use and Example	10
2.3.3.	Measurement Considerations	11

2.4.	Cost Metric: Hop Count	12
2.4.1.	Intended Semantics	12
2.4.2.	Use and Example	13
2.4.3.	Measurement Considerations	14
2.5.	Cost Metric: Packet Loss	14
2.5.1.	Intended Semantics	14
2.5.2.	Use and Example	15
2.5.3.	Measurement Considerations	16
2.6.	Cost Metric: Throughput	16
2.6.1.	Intended Semantics	17
2.6.2.	Use and Example	17
2.6.3.	Measurement Considerations	18
3.	Traffic Engineering Performance Cost Metrics	18
3.1.	Cost Metric: Link Maximum Reservable Bandwidth	19
3.1.1.	Intended Semantics	19
3.1.2.	Use and Example	19
3.1.3.	Measurement Considerations	20
3.2.	Cost Metric: Link Residue Bandwidth	21
3.2.1.	Intended Semantics	21
3.2.2.	Use and Example	21
3.2.3.	Measurement Considerations	22
4.	Operational Considerations	23
4.1.	Data Source Considerations	23
4.2.	Computation Considerations	24
4.2.1.	Configuration Parameters Considerations	24
4.2.2.	Availability Considerations	24
5.	Security Considerations	24
6.	IANA Considerations	25
7.	Acknowledgments	25
8.	References	25
8.1.	Normative References	25
8.2.	Informative References	27
	Authors' Addresses	27

1. Introduction

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

Different applications may use different cost metrics. Hence, the ALTO base protocol [RFC7285] introduces an ALTO Cost Metric Registry (Section 14.2 of [RFC7285]) as a systematic mechanism to allow different metrics to be specified. For example, a more delay-sensitive application may want to use latency related metrics, and a more bandwidth-sensitive application may want to use bandwidth related metrics. The ALTO base protocol [RFC7285], however, has

registered only one single cost metric, i.e., the generic "routingcost" metric; no latency or bandwidth related metrics are defined.

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

Metric	Definition	Origin
One Way Delay	Section 2.1	[RFC2679] Section 3.6
Round Trip Delay	Section 2.2	[RFC2681] Section 2.6
Packet Delay Variation	Section 2.3	[RFC3393] Section 2.6
Hop Count	Section 2.4	[RFC7285]
Packet Loss	Section 2.5	[RFC7680] Section 2.6
Throughput	Section 2.6	[RFC6349] Section 3.3
Max Reservable Bandwidth	Section 3.1	[RFC5305] Section 3.5
Residue Bandwidth	Section 3.2	[RFC7810] Section 4.5

Table 1. Cost Metrics Defined in this Document

The purpose of this document is to ensure proper usage of the metrics by ALTO clients. It does not claim novelty of the metrics. Some of these metrics are already specified by standards such as IPPM; some are ISP dependent such as those registered in ISIS or OSPF-TE. This document will refer to the relevant specifications.

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

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

The cost metrics defined in this document can be retrieved and aggregated from routing protocols or other traffic measurement management tools, with corresponding operational issues. A potential architecture on computing these metrics is shown in Figure 1 below. In [Section 4](#), we discuss in more detail the operations issues and how to address them.

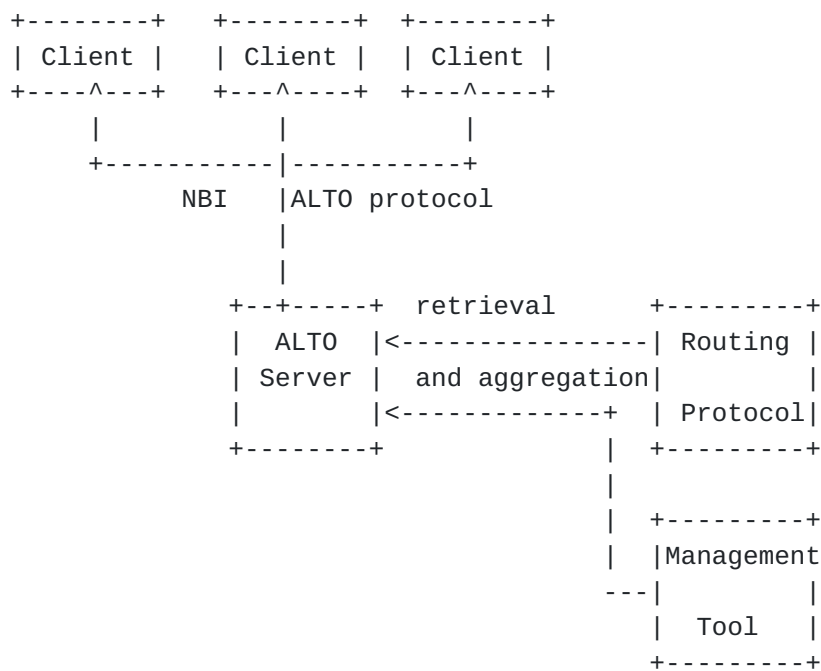


Figure 1. Potential framework to compute performance cost metrics

An ALTO server introducing these metrics should also consider security issues. As a generic security consideration on the reliability and trust in the exposed metric values, applications SHOULD rapidly give up using ALTO-based guidance if they feel the exposed information does not preserve their performance level or even degrades it. We discuss security considerations in more details in [Section 5](#).

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. Network Performance Cost Metrics

This section introduces generic ALTO network performance metrics such as one way delay, round trip delay, hop count, packet loss, throughput derived and aggregated from routing protocols or from end to end traffic management tools.

2.1. Cost Metric: One Way Delay (owdelay)

Metric name:

One Way Delay

Metric Identifier:

owdelay

2.1.1. Intended Semantics

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).

Metric Representation: The metric value type is a single 'JSONNumber' type value containing a non-negative integer component that may be followed by an exponent part. See [section 8.4.3](#) of [I-D.ietf-ippm-initial-registry] for metric unit. The unit is expressed in milliseconds in this document.

2.1.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 1: Delay value on source-destination endpoint pairs

POST /endpointcost/lookup HTTP/1.1

Host: alto.example.com

Content-Length: TBA

Content-Type: application/alto-endpointcostparams+json

Accept:

application/alto-endpointcost+json,application/alto-error+json

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



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

2.1.3. Measurement Considerations

Method of Measurement or Calculation:

See section 8.3 of [[I-D.ietf-ippm-initial-registry](#)] for potential measurement method.

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#), Data sources for potential data sources.

Measurement Timing:

See section 8.3.5 of [[I-D.ietf-ippm-initial-registry](#)] for potential measurement timing considerations.

2.2. Cost Metric: RoundTrip Time (rtt)

Metric name:

Round Trip Time

Metric Identifier:

rtt

2.2.1. Intended Semantics

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).

Metric Representation: The metric value type is a single 'JSONNumber' type value containing a non-negative integer component that may be followed by an exponent part. See [section 4.4.3](#) of [I-D.ietf-ippm-initial-registry] for Measurement Unit. The unit is expressed in milliseconds in this document.

2.2.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

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

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

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



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

2.2.3. Measurement Considerations

Method of Measurement or Calculation:

See section 4.3 of [[I-D.ietf-ippm-initial-registry](#)] for potential measurement method.

Measurement Point(s) with Potential Measurement Domain:

See [section 4.1](#), Data sources.

Measurement Timing:

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

2.3. Cost Metric: Packet Delay Variation (pdv)

Metric name:

Packet Delay Variation

Metric Identifier:

pdv

2.3.1. Intended Semantics

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).

Metric Representation: The metric value type is a single 'JSONNumber' type value containing a non-negative integer component that may be followed by an exponent part. See [section 5.4.4](#) of [I-D.ietf-ippm-initial-registry] for Measurement Unit. The unit is expressed in milliseconds in this document.

2.3.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 3: PDV value on source-destination endpoint pairs

POST /endpointcost/lookup HTTP/1.1

Host: alto.example.com

Content-Length: TBA

Content-Type: application/alto-endpointcostparams+json

Accept:

application/alto-endpointcost+json,application/alto-error+json

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

2.3.3. Measurement Considerations

Method of Measurement or Calculation:

See Section 5.3 of [[I-D.ietf-ippm-initial-registry](#)] for potential measurement method.

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#), Data sources for potential data sources.

Measurement Timing:

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

[2.4.](#) Cost Metric: Hop Count

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

Metric name:

Hop count

Metric Identifier:

hopcount

[2.4.1.](#) Intended Semantics

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, in direct relation to the routing protocols originating this information.

Metric Representation:

The metric value type is a single 'JSONNumber' type value containing a non-negative integer component. The unit is integer number.

[2.4.2.](#) Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 4: hopcount value on source-destination endpoint pairs

POST /endpointcost/lookup HTTP/1.1

Host: alto.example.com

Content-Length: TBA

Content-Type: application/alto-endpointcostparams+json

Accept:

application/alto-endpointcost+json,application/alto-error+json

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

HTTP/1.1 200 OK

Content-Length: TBA

Content-Type: application/alto-endpointcost+json

```
{
  "meta": {
    "cost type": {
      "cost-mode": "numerical",
      "cost-metric": "hopcount"
    }
  },
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89" : 5,
      "ipv4:198.51.100.34": 3,
      "ipv6:2000::1:2345:6789:abcd" : 2,
    }
  }
}
```


2.4.3. Measurement Considerations

Method of Measurement or Calculation:

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

Measurement Point(s) with Potential Measurement Domain:

The hop count can be measured at the source endpoint by traceroute.

Measurement Timing:

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

2.5. Cost Metric: Packet Loss

Metric name:

Packet loss

Metric Identifier:

pktloss

2.5.1. Intended Semantics

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).

Metric Representation:

The metric value type is a single 'JSONNumber' type value which be be non-negative integer. The unit is percentile.

2.5.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 5: pktloss value on source-destination endpoint pairs

```
POST /endpointcost/lookup HTTP/1.1
```

```
Host: alto.example.com
```

```
Content-Length: TBA
```

```
Content-Type: application/alto-endpointcostparams+json
```

```
Accept:
```

```
application/alto-endpointcost+json,application/alto-error+json
```

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



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

2.5.3. Measurement Considerations

Method of Measurement or Calculation:

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

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#) this document, Data sources.

Measurement Timing:

See [Section 2](#) and [Section 3 of \[RFC7680\]](#) for Measurement Timing.

2.6. Cost Metric: Throughput

Metric name:

Throughput

Metric Identifier:

throughput

2.6.1. Intended Semantics

Metric Description:

To specify spatial and temporal throughput over the specified source and destination. The spatial aggregation level is specified in the query context (e.g., PID to PID, or endpoint to endpoint).

Metric Representation:

The unit is Mbps.

2.6.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 5: throughput value on source-destination endpoint pairs

POST /endpointcost/lookup HTTP/1.1

Host: alto.example.com

Content-Length: TBA

Content-Type: application/alto-endpointcostparams+json

Accept:

application/alto-endpointcost+json,application/alto-error+json

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



```
HTTP/1.1 200 OK
Content-Length: TBA
Content-Type: application/alto-endpointcost+json
{
  "meta": {
    "cost type": {
      "cost-mode": "numerical",
      "cost-metric": "throughput"
    }
  }
  "endpoint-cost-map": {
    "ipv4:192.0.2.2": {
      "ipv4:192.0.2.89" : 25.6,
      "ipv4:198.51.100.34": 12.8,
      "ipv6:2000::1:2345:6789:abcd" : 42.8,
    }
  }
}
```

2.6.3. Measurement Considerations

Method of Measurement or Calculation:

See [Section 3.3 of \[RFC6349\]](#) for Measurement Method.

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#) of this document.

Measurement Timing:

Similar to RTT. See [Section 4.3.5](#) of [I-D.ietf-ippm-initial-registry] for Measurement Timing.

3. 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.

3.1. Cost Metric: Link Maximum Reservable Bandwidth

Metric name:

Maximum Reservable Bandwidth

Metric Identifier:

maxresbw

3.1.1. Intended Semantics

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).

Metric Representation:

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

3.1.2. Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 6: maxresbw value on source-destination endpoint pairs

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

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

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

3.1.3. Measurement Considerations

Method of Measurement or Calculation:

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

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#) this document for discussions.

Measurement Timing:

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

[3.2.](#) Cost Metric: Link Residue Bandwidth

Metric name:

Residue Bandwidth

Metric Identifier:

residuebw

[3.2.1.](#) Intended Semantics

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

Metric Representation:

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

[3.2.2.](#) Use and Example

This metric could be used as a cost metric constraint attribute used either together with cost metric attribute 'routingcost' or on its own or as a returned cost metric in the response.

Example 7: residuebw value on source-destination endpoint pairs

POST/ endpointcost/lookup HTTP/1.1

Host: alto.example.com

Content-Length: TBA

Content-Type: application/alto-endpointcostparams+json

Accept:

application/alto-endpointcost+json,application/alto-error+json

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

HTTP/1.1 200 OK

Content-Length: TBA

Content-Type: application/alto-endpointcost+json

```
{
  "meta": {
    "cost-type" {
      "cost-mode": "numerical",
      "cost-metric": "residuebw"
    }
  },
  "endpoint-cost-map" {
    "ipv4:192.0.2.2" {
      "ipv4:192.0.2.89" : 0,
      "ipv4:198.51.100.34": 2000,
      "ipv6:2000::1:2345:6789:abcd": 5000,
    }
  }
}
```

3.2.3. Measurement Considerations

Method of Measurement or Calculation:

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

Measurement Point(s) with Potential Measurement Domain:

See [Section 4.1](#) of this document.

Measurement Timing:

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

4. Operational Considerations

It can be non-trivial for an ALTO server to derive the metrics. Also, the exact infrastructure and algorithms can vary from different networks, and are outside the scope of this document. However, since they present challenges, we discuss these common challenges.

Also, the performance metrics specified in this document are similar, in that they may use similar data sources and have similar issues in their calculation. Hence, we specify common issues unless one metric has its unique challenges.

4.1. Data Source Considerations

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 [\[RFC7971\]](#). In particular, the cost metrics defined in this document can be computed using routing systems as the data sources. Mechanisms defined in [\[RFC2681\]](#), [\[RFC3393\]](#), [\[RFC7679\]](#), [\[RFC7680\]](#), [\[RFC3630\]](#), [\[RFC3784\]](#), [\[RFC7471\]](#), [\[RFC7810\]](#), [\[RFC7752\]](#) and [\[I-D.ietf-idr-te-pm-bgp\]](#) that allow an ALTO Server to retrieve and derive the necessary information to compute the metrics that we describe in this document.

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.

4.2. Computation Considerations

The metric values exposed by an ALTO server may result from additional processing on measurements from data sources to compute exposed metrics. This may involve data processing tasks such as aggregating the results across multiple systems, removing outliers, and creating additional statistics. There are two challenges on the computation of ALTO performance metrics.

4.2.1. Configuration Parameters Considerations

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

4.2.2. Availability Considerations

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

5. Security Considerations

The properties defined in this document present no security considerations beyond those in [Section 15](#) of the base ALTO specification [[RFC7285](#)].

However concerns addressed in Sections "15.1 Authenticity and Integrity of ALTO Information", "15.2 Potential Undesirable Guidance from Authenticated ALTO Information" and "15.3 Confidentiality of ALTO Information" remain of utmost importance. Indeed, TE performance is a highly sensitive ISP information, therefore, sharing TE metric values in numerical mode requires full mutual confidence between the entities managing the ALTO Server and Client. Numerical TE performance information will most likely be distributed by ALTO Servers to Clients under strict and formal mutual trust agreements.

On the other hand, ALTO Clients must be cognizant on the risks attached to such information that they would have acquired outside formal conditions of mutual trust.

6. IANA Considerations

IANA has created and now maintains the "ALTO Cost Metric Registry", listed in [Section 14.2](#), Table 3 of [RFC7285]. This registry is located at <<http://www.iana.org/assignments/alto-protocol/alto-protocol.xhtml#cost-metrics>>. This document requests to add the following entries to "ALTO Cost Metric Registry".

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

7. Acknowledgments

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

8. References

8.1. Normative References

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

Ginsberg, L., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions", [draft-ietf-idr-te-pm-bgp-18](#) (work in progress), December 2018.

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

Morton, A., Bagnulo, M., Eardley, P., and K. D'Souza, "Initial Performance Metrics Registry Entries", [draft-ietf-ippm-initial-registry-11](#) (work in progress), March 2019.

- [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>>.
- [RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", [RFC 6349](#), DOI 10.17487/RFC6349, August 2011, <<https://www.rfc-editor.org/info/rfc6349>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", [RFC 7285](#), DOI 10.17487/RFC7285, September 2014, <<https://www.rfc-editor.org/info/rfc7285>>.
- [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>>.

- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, [RFC 7679](#), DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [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>>.

8.2. Informative References

- [RFC6390] Clark, A. and B. Claise, "Framework for Performance Metric Development", [RFC 6390](#), July 2011.
- [RFC7971] Stiemerling, M., Kiesel, S., Scharf, M., Seidel, H., and S. Previdi, "Application-Layer Traffic Optimization (ALTO) Deployment Considerations", [RFC 7971](#), DOI 10.17487/[RFC7971](#), October 2016, <<https://www.rfc-editor.org/info/rfc7971>>.

Authors' Addresses

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: bill.wu@huawei.com

Y. Richard Yang
Yale University
51 Prospect St
New Haven, CT 06520
USA

Email: yry@cs.yale.edu

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

