ALTO Performance Metrics

draft-ietf-alto-performance-metrics-08

Qin Wu
D. Dhody
Young Lee
Sabine Randriamasy
Y. Richard Yang

IETF 106
November 21, 2019
Singapore

Outline

- Updates from v07-v08
- Remaining issue
- Plan for next step

Recall: Remaining Issue at 105: Metric Definition Consistency and Reusability

- A basic issue is whether to define ALTO metrics precisely on existing measurement metrics such as those defined by IPPM

ALTO performance metrics [this document]

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

IPPM metrics [2][1]

- UDP
  - RTDelay_Active_IP-UDP-Periodic_RFCXXXXsecY_Seconds_95Percentile
  - RTLoss_Active_IP-UDP-Periodic_RFCXXXXsecY_Percent_LossRatio
  - OWPDV_Active_IP-UDP-Periodic_RFCXXXXsecY_Seconds_95Percentile
  - OWDelay_Active_IP-UDP-Poisson-Payload250B_RFCXXXXsecY_Seconds_<statistic>
  - OWDelay_Active_IP-UDP-Periodic20m-Payload142B_RFCXXXXsecY_Seconds_<statistic>
- TCP
  - RTDelay_Passive_IP-TCP_RFCXXXXsecY_Seconds_<statistic>
- DNS
  - RTDNS_Active_IP-UDP-Poisson_RFCXXXXsecY_Seconds_Raw RLDNS_Active_IP-UDP-Poisson_RFCXXXXsecY_Logical_Raw

- 105 WG feedback: clarify that ALTO metrics are guidance

[1] https://datatracker.ietf.org/doc/draft-ietf-ippm-metric-registry/
[2] https://datatracker.ietf.org/doc/draft-ietf-ippm-initial-registry/

IETF105 WG Feedback and Issue

- Issue: There can be multiple sources/types of guidance, even for the same metric (e.g., delay)
  - SLA, estimation

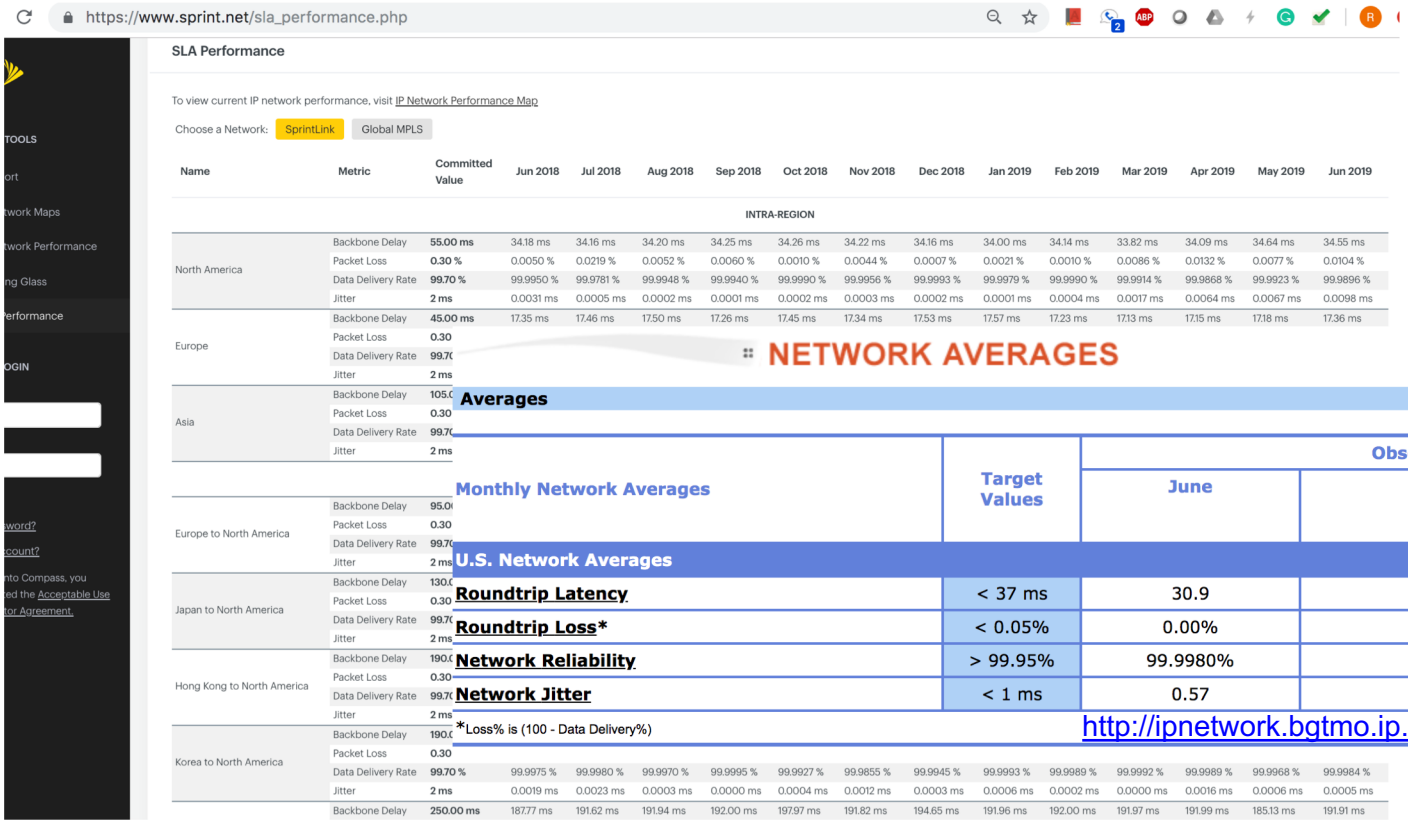

Key Change v07-v08

- To allow multiple types of guidance, specify guidance source by “cost-source”

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.

4.1. Source Considerations

The additiona of the "cost-source" field is to solve a key issue: An ALTO server needs data sources to compute the cost metrics described in this document and an ALTO client needs to know the data sources to better interpret the values.

To avoid too fine-grained information, this document introduces "cost-source" to indicate only the high-level type of data sources: "estimation" or "sla", where "estimation" is a type of measurement data source and "sla" is a type that is more based on policy.

For estimation, for example, the ALTO server may use log servers or the OAM system as its data source [RFC7971]. In particular, the cost metrics defined in this document can be computed using routing systems as the data sources. Mechanisms defined in [RFC2681], [RFC3393], [RFC7679], [RFC7680], [RFC3630], [RFC3784], [RFC7471], [RFC7810], [RFC7752] and [I-D.ietf-idr-te-pm-bgp] that allow an ALTO Server to retrieve and derive the necessary information to compute the metrics that we describe in this document.

“cost-source” Definition Details

- “cost-source” is an optional extension field of ”cost-type”

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.

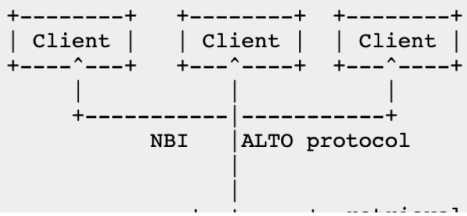

To make clear how the value of an ALTO performance metric is derived this document defines an optional field named "cost-source" to extend the ALTO "cost-type". The "cost-source" indicates how the metric is derived, and currently it can be either "estimation" or "sla". If a "cost-type" does not include the optional "cost-source" field, the application MUST assume that the value of "cost-source" is "estimation".

An ALTO server may compute "estimation" values by retrieving and/or aggregating information from routing protocols or other traffic measurement management tools, with corresponding operational issues. A potential architecture on estimating these metrics is shown in Figure 1 below. In Section 4, we discuss in more detail the operations issues and how to address them.

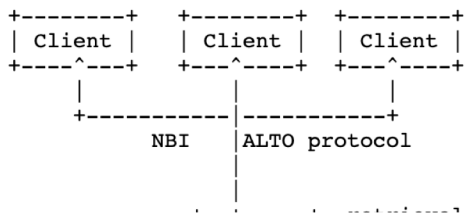

- “cost-source” should be registered to allow extension

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

Identifier	Intended Semantics
estimation	Values by estimation
sla	Values reflect service level agreement

Optional “cost-source” and Examples

- Minimal updated examples:

```
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
```

skipping to change at *page 7, line 22*

```
    },
     "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,
     }
   }
}
```

```
2.1.2.  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
```

skipping to change at *page 7, line 45*

```
    },
     "endpoint-cost-map" : {
       "ipv4:192.0.2.2": {
         "ipv4:192.0.2.89"     : 10,
         "ipv4:198.51.100.34" : 20,
         "ipv6:2000::1:2345:6789:abcd"   : 30,
     }
   }
}

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

Change in Representation Specification after Decoupling from IPPM

```
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.  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 conforming to the number specification of [RFC8259]
   Section 6.  The number MUST be non-negative.  The unit is expressed
   in milliseconds in this document.
```

- Same for other metrics defined in the doc

Optional “cost-source” and Backward Compatibility

One potential issue introduced by the optional "cost-source" field is backward compatibility. Consider that an IRD which defines two cost-types with the same "cost-mode" and "cost-metric", but one with "cost-source" being "estimation" and the other being "sla". Then an ALTO client that is not aware of the extension will not be able to distinguish between these two types. A similar issue can arise even with a single cost-type which has "cost-source" being "sla", but the backward client will ignore this field and consider the metric estimation. [RFC7285]

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

Remaining Issue Needing WG Feedback

- 3 types instead of 2 types of “cost-source”
  - sla, estimation, nominal (e.g., nominal link capacity being 10G)

Next Step Plan

- Wait for WG feedback and discuss with IPPM
- Submit an update, if any, by mid December, 2019

Backup Slides

Updated Abstract

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.

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. Additional cost metrics may be documented in other documents.

When using an ALTO performance metric, applications need additional information beyond the metric. In particular, its "cost-source", such as it being an estimation or an SLA, is key to define the meaning of a performance metric. Hence, each ALTO performance metric should include the "cost-source" of the metric. To report an estimated value of a performance metric, the ALTO server may derive and aggregate from routing protocols with different granularity and scope, such as BGP-LS, OSPF-TE and ISIS-TE, or from end-to-end traffic management tools. These metrics may then be exposed by an ALTO Server to allow applications to determine "where" to connect based on network performance criteria.