

ALTO Traffic Engineering Cost Metrics

draft-wu-alto-te-metrics-08

Qin Wu (bill.wu@huawei.com)

Y. Richard Yang (yry@cs.yale.edu)

Young Lee (leeyoung@huawei.com)

Dhruv Dhody (dhruv.ietf@gmail.com)

Sabine Randriamasy (sabine.randriamasy@nokia-bell-labs.com)

Document Status

- Started with draft-wu-alto-json-te, and was first presented at IETF 87 Berlin, during ALTO re-chartering
 - A lot of interests from vendors and operators in this work.
- Consolidation with draft-lee-alto-app-net-info-exchange and draft-ietf-alto-multi-cost, renamed as draft-wu-alto-te-metrics-00, second presented at IETF90 Toronto
- Accept as Chartered item together with ALTO Cost Calendar to fulfill the Cost Property Extension milestone (May 2015)

Update since last version

- The changes in the latest version (draft-wu-alto-te-metrics-08):
 - Revised -08 to use template of RFC 6390 (Guidelines for Considering New Performance Metric Development) to describe performance cost metric.
 - Change the title as ALTO performance Cost Metric
 - Add IPv6 example to address Fred Baker's comment on the list.
 - Highlight the metrics are end to end metrics. Hop Count property is clarified.
 - Clarified the validity period of each metric needs to be negotiated between the server and client
 - The validity period of a metric can be one or multiple of measurement interval on the metric.
 - Clear confusion on whether we redefine the metric and how is related to BGP-LS in routing area
 - Metric redefine Clarifications:
 - We are not redefine metric and we feed derive and aggregate TE metrics from routing protocol and feed them to ALTO server and provide them to ALTO client.
 - ALTO server may aggregate metric data gathered using various routing protocols or other management tool and provide them in an e2e sense.
 - Relationship with BGP-LS
 - This draft is complementary to BGP-LS since BGP-LS is used to collect TE metric and send it to ALTO server and then ALTO server aggregate these data and expose it to the ALTO client in the northbound interface.

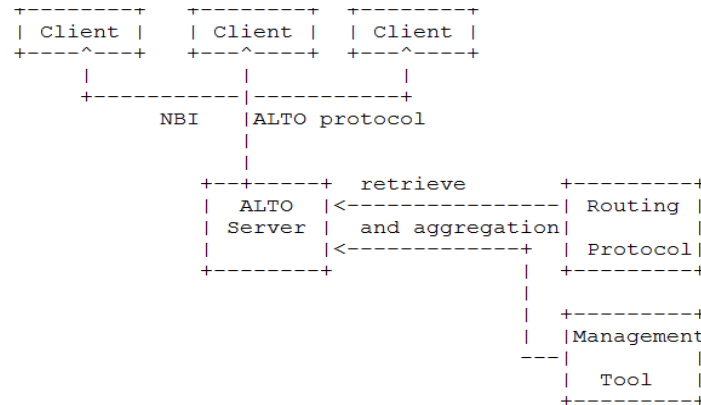
The Cost Metrics

- Delay (delay)
- Delay Jitter (jitter) -> delayjitter
- Packet Loss (pktloss)
- Hop Count (hopcount)
- Bandwidth

Metric name	Metric Description	Relation to other metric
Maximum Bandwidth (maxbw):	the maximum bandwidth that can be used; motivated from RFC 3630 Sec. 2.5.6.	
Maximum Reservable Bandwidth (maxresbw):	the maximum bandwidth that can be reserved; motivated from RFC 3630 Sec. 2.5.7.	Maxresbw can be larger than maxbw if the link is oversubscribed
UnReserved Bandwidth (unresbw[x,y]):	the amount of bandwidth not yet reserved at each of the eight priority levels in IEEE floating point format, return an array, motivated from RFC 3630 Sec.2.5.8	The initial value of unresbw for each priority can be maxbw.
Residue Bandwidth (residbw):	subtracts tunnel reservations from Maximum Bandwidth, motivated from [I-D. ietf-isis-te-metric-extensions], Sec.4.5.	Residbw = maxbw - tunnel reservation bw
Available Bandwidth (availbw):	subtracts the tunnel reservation and the measured bandwidth used for the actual forwarding of best effort traffic from Maximum Bandwidth, motivated from [I-D. ietf-isis-te-metric-extensions], Sec.4.6.	Availbw = residbw- measured bw for best effort traffic
Utilized Bandwidth (utilbw):	Actual measured bandwidth used for all traffic, motivated from [I-D. ietf-isis-te-metric-extensions], Sec.4.7.	Utilbw = measured bw for best effort traffic + bw for TE traffic

Why Propose new Cost Metric?

- Fulfill the Cost Property Extension milestone together with ALTO Cost Calendar
- Allow applications to determine "where" to connect based on end to end network performance criteria
- Fill the gap on Alto Server Network API Use case proposed in RFC7752, section 2.2 and draft-ietf-idr-te-pm-bgp-02, section3.2



- ALTO base protocol defines only a single cost metric, i.e., the generic "routingcost" metric
 - (Sec. 14.2 of ALTO base specification [ALTO])

Next Step

- WG adoption?