ALTO Performance Metrics

draft-ietf-alto-performance-metrics-07

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

IETF 105
July 25, 2019
Montreal

Outline

- Updates from v06-v07
- Remaining issues requiring WG discussions
- Plan for next step

Updates Overview (v06-v07)

- Structure changes
  - Move challenges to back; restructure metric definition
- Many small text changes as well

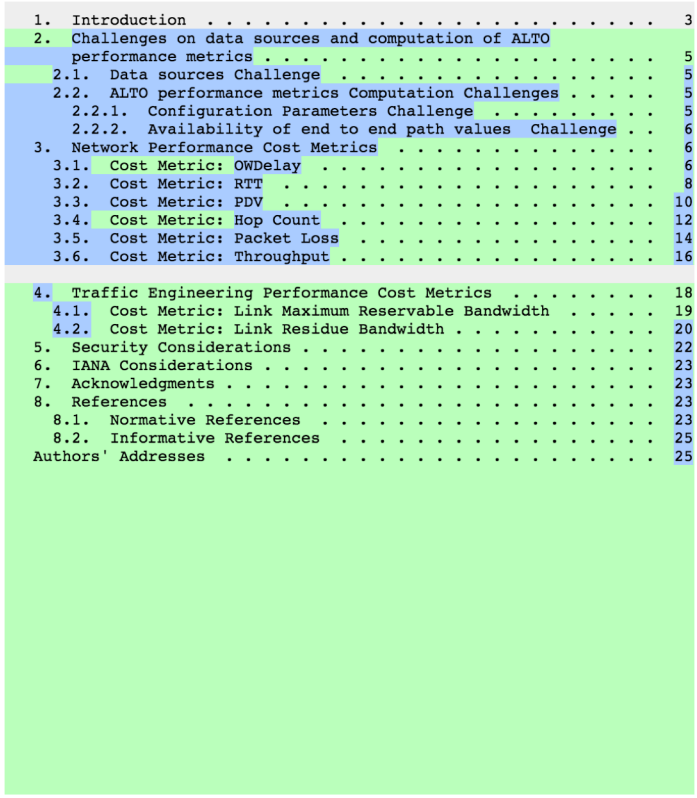

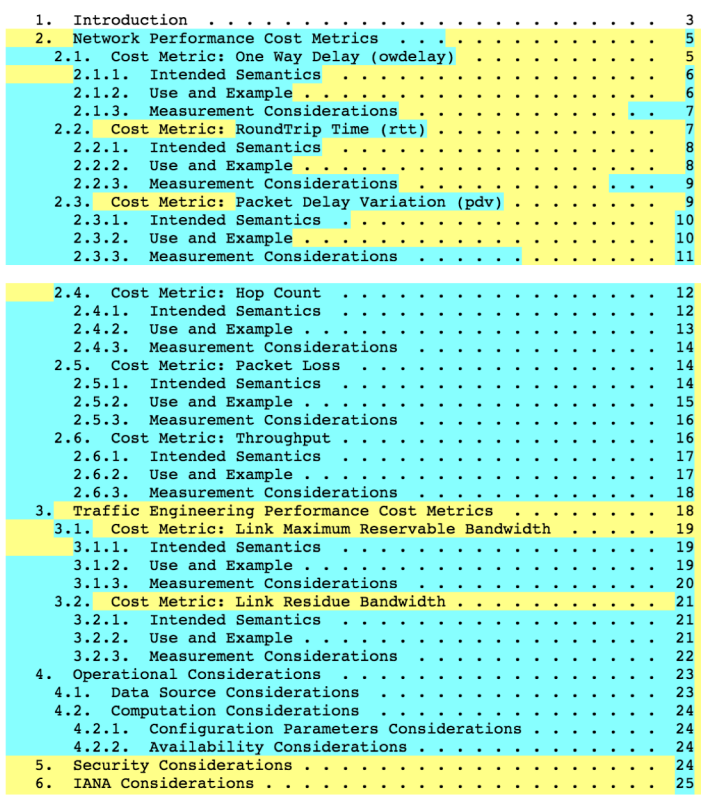

Main Update (v06-v07): Metric Definition

- Restructure the definition of each metric to be consistent with ALTO base protocol (RFC 7285): why
  - v06, structure, for each metric,
    - Metric Name
    - Metric Description
    - Method of Measurement or Calculation
    - Units of Measurement
    - Measurement Point(s) with Potential Measurement Domain
    - Measurement Timing
    - Use and Applications

- RFC7285

14.2. ALTO Cost Metric Registry

IANA has created and now maintains the "ALTO Cost Metric Registry", listed in Table 3.

Identifier	Intended Semantics
routingcost	See Section 6.1.1.1
priv:	Private use

Table 3: ALTO Cost Metrics

This registry serves two purposes. First, it ensures uniqueness of identifiers referring to ALTO cost metrics. Second, it provides references to particular semantics of allocated cost metrics to be applied by both ALTO servers and applications utilizing ALTO clients.

Requests to add a new value to the registry MUST include the following information:

- Identifier: The name of the desired ALTO cost metric.
- Intended Semantics: ALTO costs carry with them semantics to guide their usage by ALTO clients. For example, if a value refers to a measurement, the measurement units must be documented. For proper implementation of the ordinal cost mode (e.g., by a third-party service), it should be documented whether higher or lower values of the cost are more preferred.
- Security Considerations: ALTO costs expose information to ALTO clients. As such, proper usage of a particular cost metric may require certain information to be exposed by an ALTO service provider. Since network information is frequently regarded as proprietary or confidential, ALTO service providers should be made aware of the security ramifications related to usage of a cost metric.

Main Update (v06-v07): Metric Definition

- Restructure the definition of each metric to be consistent with ALTO base protocol (RFC 7285): change

  - v06, structure, for each metric, defines
    - Metric Name
    - Metric Description
    - Method of Measurement or Calculation
    - Units of Measurement
    - Measurement Point(s) with Potential Measurement Domain
    - Measurement Timing
    - Use and Applications

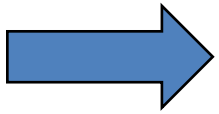

  - v07, structure, for each metric, defines
    - Metric Name
    - Metric identifier
    - Intended semantics
      - Metric Description
      - Metric Representation
    - Use and Example
    - Measurement Considerations
      - Method of Measurement or Calculation
      - Measurement Point(s) with Potential Measurement Domain
      - Measurement Timing

Metric Details

Metric	Representation
One-Way Delay, Round-trip Time, Packet Delay Variation	A single JSONNumber conforming to Sec. 6 [RFC8259] (int [frac] [exp]); Must be non-negative; unit is ms;
Hop Count	The metric value type is a single 'JSONNumber' type value conforming to the number specification (Section 6, [RFC8259]). The number MUST be an integer and non-negative.
Packet Loss	The metric value type is a single 'JSONNumber' type value conforming to the number specification (Section 6, [RFC8259]). The number MUST be non-negative. The value represents the percentage of packet loss.
Throughput, Max Reservable BW, Residue BW	The metric value type is a single 'JSONNumber' type value conforming to the number specification (Section 6, [RFC8259]). The number MUST be non-negative. The unit is Mbps.

- Suggestions but not adopted: (1) add *infinity*; (2) allow units such as ms, s, Mbps/Kbps/Gbps, …

Update: Operations Considerations

- Substantially extended the section on operations considerations, to emphasize,
  - Some performance metric can be a complex function of multiple factors:
    - Traffic type (e.g., UDP, TCP; video)
    - Client behavior (e.g., arrival patterns such as Poisson, periodical…)
    - Network settings (e.g., scheduling policies, cross traffic interference, …)
    - Time
  - A network may adopt different measurement approaches
    - Active (e.g., probe measured, packet pair measured, …)
    - Passive (e.g., derivation from existing data such as logs)
  - Computing some performance metrics can involve non-trivial computation, which has implications on timeliness, denial-of-service , …
    - Data cleaning, aggregation, inference, …

Remaining Issue (1): Metric Definition Consistency and Reusability

- A basic issue is consistency and reusability in IETF

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

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

Discussion

- IPPM metrics are more for infrastructure management
  - Underlying network technology ***aware*** -- implementation focus

- ALTO metrics are more for applications
  - Underlying network technology ***transparent*** (e.g., do not care if the transport uses IP/MPLS, …) – interface focus

Example High-Level Network Metrics Exposure

https://www.sprint.net/sla_performance.php

TOOLS
ort
twork Maps
twork Performance
ng Glass
erformance
OGIN
sword?
count?
nto Compass, you
ed the Acceptable Use
tor Agreement.

SLA Performance

To view current IP network performance, visit IP Network Performance Map

Choose a Network: SprintLink | Global MPLS

Name	Metric	Committed Value	Jun 2018	Jul 2018	Aug 2018	Sep 2018	Oct 2018	Nov 2018	Dec 2018	Jan 2019	Feb 2019	Mar 2019	Apr 2019	May 2019	Jun 2019
INTRA-REGION															
North America	Backbone Delay	**55.00 ms**	34.18 ms	34.16 ms	34.20 ms	34.25 ms	34.26 ms	34.22 ms	34.16 ms	34.00 ms	34.14 ms	33.82 ms	34.09 ms	34.64 ms	34.55 ms
	Packet Loss	**0.30 %**	0.0050 %	0.0219 %	0.0052 %	0.0060 %	0.0010 %	0.0044 %	0.0007 %	0.0021 %	0.0010 %	0.0086 %	0.0132 %	0.0077 %	0.0104 %
	Data Delivery Rate	**99.70 %**	99.9950 %	99.9781 %	99.9948 %	99.9940 %	99.9990 %	99.9956 %	99.9993 %	99.9979 %	99.9990 %	99.9914 %	99.9868 %	99.9923 %	99.9896 %
	Jitter	**2 ms**	0.0031 ms	0.0005 ms	0.0002 ms	0.0001 ms	0.0002 ms	0.0003 ms	0.0002 ms	0.0001 ms	0.0004 ms	0.0017 ms	0.0064 ms	0.0067 ms	0.0098 ms
Europe	Backbone Delay	**45.00 ms**	17.35 ms	17.46 ms	17.50 ms	17.26 ms	17.45 ms	17.34 ms	17.53 ms	17.57 ms	17.23 ms	17.13 ms	17.15 ms	17.18 ms	17.36 ms
	Packet Loss	**0.30 %**	0.0007 %	0.0001 %	0.0118 %	0.0024 %	0.0077 %	0.0075 %	0.0121 %	0.0126 %	0.0012 %	0.0001 %	0.0152 %	0.0051 %	0.0033 %
	Data Delivery Rate	**99.70 %**	99.9993 %	99.9999 %	99.9882 %	99.9976 %	99.9923 %	99.9925 %	99.9879 %	99.9874 %	99.9988 %	99.9999 %	99.9848 %	99.9949 %	99.9967 %
	Jitter	**2 ms**	0.0006 ms	0.0001 ms	0.0014 ms	0.0000 ms	0.0047 ms	0.0065 ms	0.0048 ms	0.0006 ms	0.0005 ms	0.0000 ms	0.0006 ms	0.0001 ms	0.0003 ms
Asia	Backbone Delay	**105.00 ms**	65.07 ms	65.41 ms	68.15 ms	70.15 ms	68.23 ms	68.08 ms	68.97 ms	69.12 ms	78.11 ms	70.75 ms	68.80 ms	74.91 ms	70.24 ms
	Packet Loss	**0.30 %**	0.0245 %	0.0030 %	0.0038 %	0.0127 %	0.0027 %	0.0044 %	0.0034 %	0.0019 %	0.0031 %	0.0021 %	0.0064 %	0.0116 %	0.0060 %
	Data Delivery Rate	**99.70 %**	99.9755 %	99.9970 %	99.9962 %	99.9873 %	99.9973 %	99.9956 %	99.9966 %	99.9981 %	99.9969 %	99.9979 %	99.9936 %	99.9884 %	99.9940 %
	Jitter	**2 ms**	0.0030 ms	0.0041 ms	0.0143 ms	0.0118 ms	0.0185 ms	0.0145 ms	0.0167 ms	0.0114 ms	0.0101 ms	0.0091 ms	0.0075 ms	0.0255 ms	0.0076 ms
INTER-REGION															
Europe to North America	Backbone Delay	**95.00 ms**	78.46 ms	78.50 ms	78.02 ms	79.36 ms	77.17 ms	76.02 ms	76.09 ms	76.02 ms	76.02 ms	76.05 ms	75.28 ms	73.04 ms	72.95 ms
	Packet Loss	**0.30 %**	0.0049 %	0.0007 %	0.0000 %	0.0001 %	0.0018 %	0.0044 %	0.0015 %	0.0035 %	0.0021 %	0.0052 %	0.0110 %	0.0145 %	0.0045 %
	Data Delivery Rate	**99.70 %**	99.9950 %	99.9993 %	100.0000 %	99.9999 %	99.9982 %	99.9956 %	99.9985 %	99.9965 %	99.9979 %	99.9948 %	99.9890 %	99.9855 %	99.9955 %
	Jitter	**2 ms**	0.0099 ms	0.0001 ms	0.0000 ms	0.0000 ms	0.0076 ms	0.0043 ms	0.0078 ms	0.0008 ms	0.0058 ms	0.0109 ms	0.0005 ms	0.0001 ms	0.0000 ms
Japan to North America	Backbone Delay	**130.00 ms**	98.56 ms	98.33 ms	98.29 ms	98.33 ms	98.44 ms	98.81 ms	98.81 ms	98.82 ms	98.82 ms	98.80 ms	98.86 ms	98.81 ms	98.82 ms
	Packet Loss	**0.30 %**	0.0006 %	0.0000 %	0.0592 %	0.0001 %	0.0001 %	0.0000 %	0.0002 %	0.0001 %	0.0001 %	0.0000 %	0.0032 %	0.0000 %	0.0000 %
	Data Delivery Rate	**99.70 %**	99.9994 %	100.0000 %	99.9408 %	99.9999 %	99.9999 %	100.0000 %	99.9998 %	99.9999 %	99.9999 %	100.0000 %	99.9968 %	100.0000 %	100.0000 %
	Jitter	**2 ms**	0.0003 ms	0.0006 ms	0.0016 ms	0.0015 ms	0.0009 ms	0.0019 ms	0.0003 ms	0.0007 ms	0.0004 ms	0.0001 ms	0.0001 ms	0.0006 ms	0.0003 ms
Hong Kong to North America	Backbone Delay	**190.00 ms**	150.80 ms	151.01 ms	150.75 ms	151.00 ms	150.87 ms	150.47 ms	150.70 ms	147.61 ms	146.84 ms	146.47 ms	154.38 ms	154.88 ms	149.31 ms
	Packet Loss	**0.30 %**	0.0009 %	0.0000 %	0.0569 %	0.0092 %	0.0003 %	0.0003 %	0.0055 %	0.0003 %	0.0014 %	0.0003 %	0.0004 %	0.0194 %	0.0013 %
	Data Delivery Rate	**99.70 %**	99.9991 %	100.0000 %	99.9431 %	99.9908 %	99.9997 %	99.9997 %	99.9945 %	99.9997 %	99.9986 %	99.9997 %	99.9996 %	99.9806 %	99.9987 %
	Jitter	**2 ms**	0.0021 ms	0.0010 ms	0.0007 ms	0.0001 ms	0.0017 ms	0.0004 ms	0.0337 ms	0.0009 ms	0.0011 ms	0.0012 ms	0.0001 ms	0.0094 ms	0.0003 ms
Korea to North America	Backbone Delay	**190.00 ms**	130.97 ms	132.00 ms	133.82 ms	132.99 ms	133.32 ms	132.81 ms	132.04 ms	131.71 ms	131.58 ms	133.48 ms	134.38 ms	133.54 ms	131.85 ms
	Packet Loss	**0.30 %**	0.0025 %	0.0020 %	0.0030 %	0.0005 %	0.0073 %	0.0145 %	0.0055 %	0.0007 %	0.0011 %	0.0008 %	0.0011 %	0.0032 %	0.0016 %
	Data Delivery Rate	**99.70 %**	99.9975 %	99.9980 %	99.9970 %	99.9995 %	99.9927 %	99.9855 %	99.9945 %	99.9993 %	99.9989 %	99.9992 %	99.9989 %	99.9968 %	99.9984 %
	Jitter	**2 ms**	0.0019 ms	0.0023 ms	0.0003 ms	0.0000 ms	0.0004 ms	0.0012 ms	0.0003 ms	0.0006 ms	0.0002 ms	0.0000 ms	0.0016 ms	0.0006 ms	0.0005 ms
	Backbone Delay	**250.00 ms**	187.77 ms	191.62 ms	191.94 ms	192.00 ms	197.97 ms	191.82 ms	194.65 ms	191.96 ms	192.00 ms	191.97 ms	191.99 ms	185.13 ms	191.91 ms

Example High-Level Network Metrics Exposure

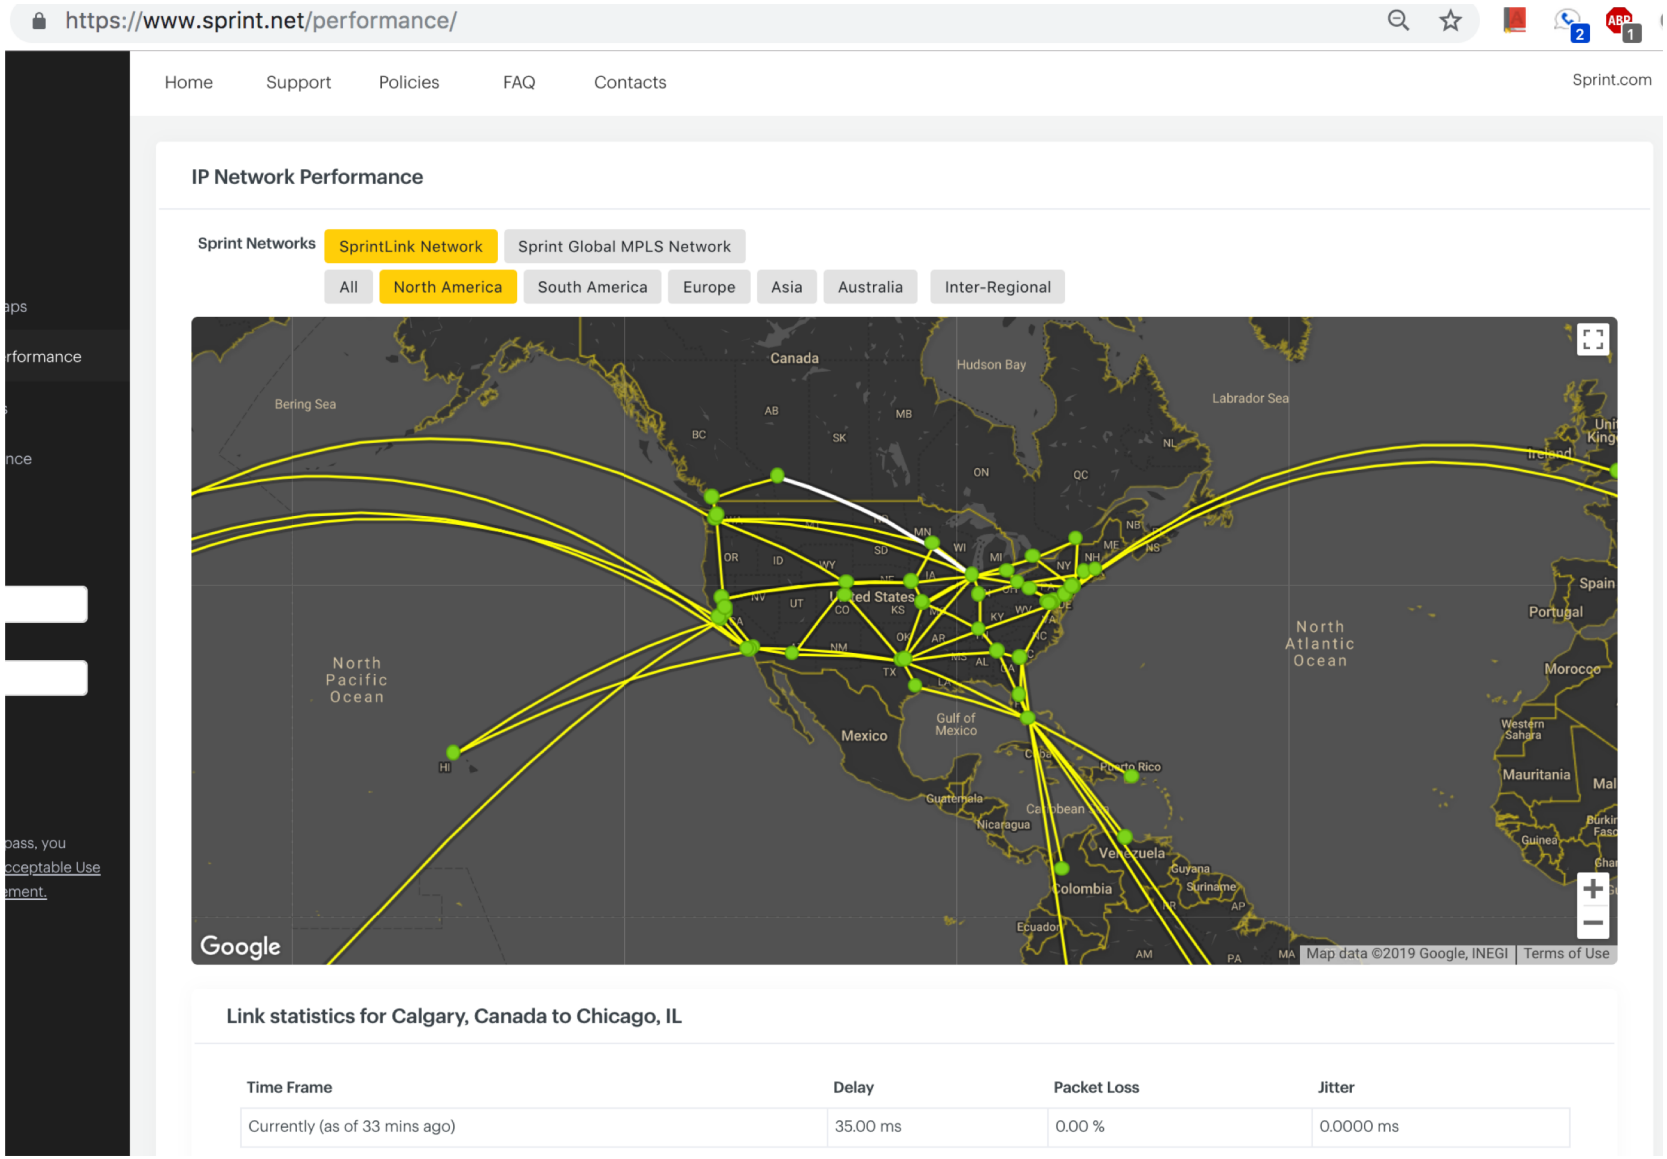

Example High-Level Network Metrics Exposure

NETWORK AVERAGES

Averages

Monthly Network Averages	Target Values	Observed Values		
		June	May	April
U.S. Network Averages				
Roundtrip Latency	< 37 ms	30.9	30.9	30.6
Roundtrip Loss*	< 0.05%	0.00%	0.01%	0.00%
Network Reliability	> 99.95%	99.9980%	N/A%	100.0000%
Network Jitter	< 1 ms	0.57	0.56	0.56

*Loss% is (100 - Data Delivery%)

NETWORK LATENCY

U.S. Network Latency

Figures are in **ms**. Thresholds are distance sensitive.

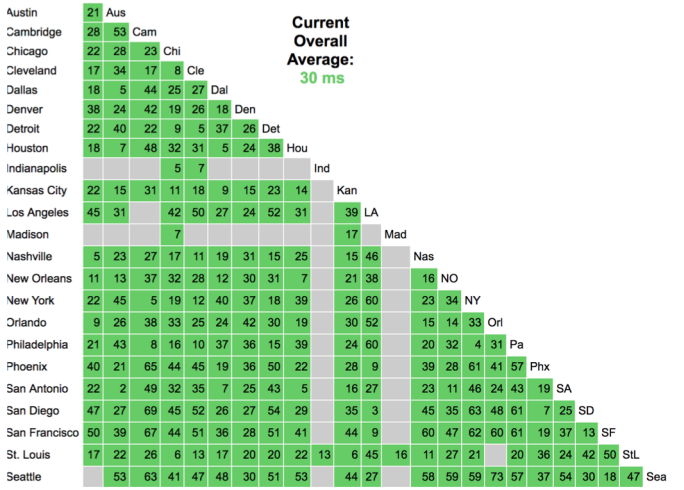

http://ipnetwork.bgtmo.ip.att.net/pws/network_delay.html

Author Discussion: Metric Definition Consistency and Reusability

- Many levels of reusability and consistency
  - Reusability:
    - base ALTO metrics on IPPM metric registry, or
    - not
      - Consistency
        - » Same ID
        - » Same metric unit (e.g., ippm latency unit is second, current document is ms)

Author Discussion

- IPPM metrics are more for infrastructure management
  - Underlying network technology ***aware*** -- implementation focus
- ALTO metrics are more for applications
  - Underlying network technology ***transparent*** (e.g., do not care if the transport uses IP/MPLS, …) – interface focus

- UDP
  - RTDelay_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>

Proposed Moving Forward

- Application-layer performance depends on
  - Network equivalent classes (e.g., categories)
  - Application behaviors

- Distinguish performance metrics
  - reflecting categories
  - dependency on application behaviors or not
    - propagation delay vs
    - traffic pattern

Remaining Issue (2): Operations and Security Considerations

- How much to update
  - Operations considerations
  - Security considerations

Next Step Plan

- Discuss with IPPM to finalize updates

- Finalize updates and submit an update by end of August