v6 Health Metric with NOMA (guest-starring RIPE Atlas)

Network Operator Measurements Activity

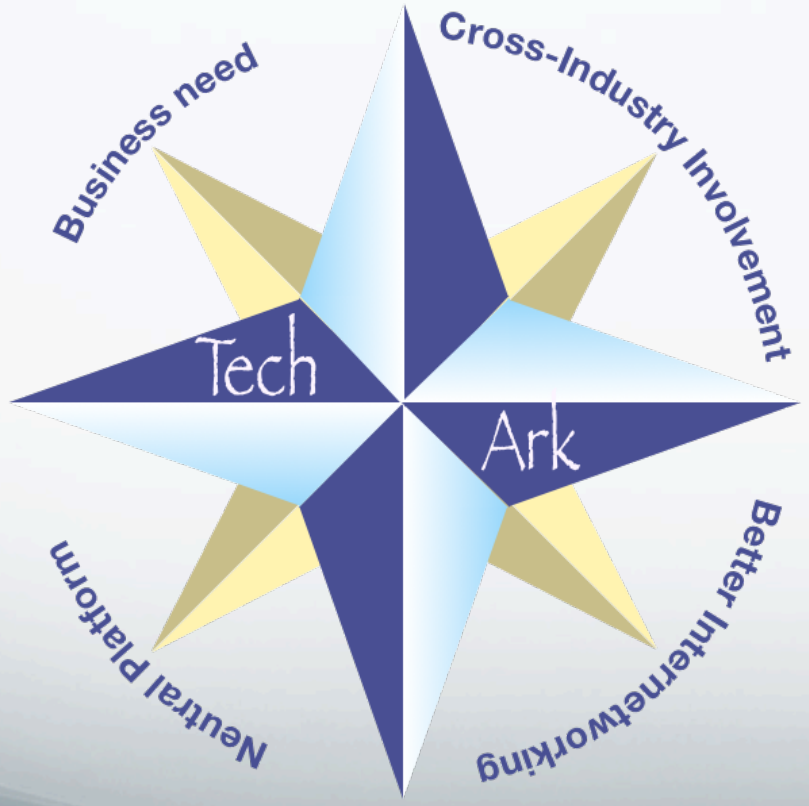

Leslie Daigle

May 2017

Summary

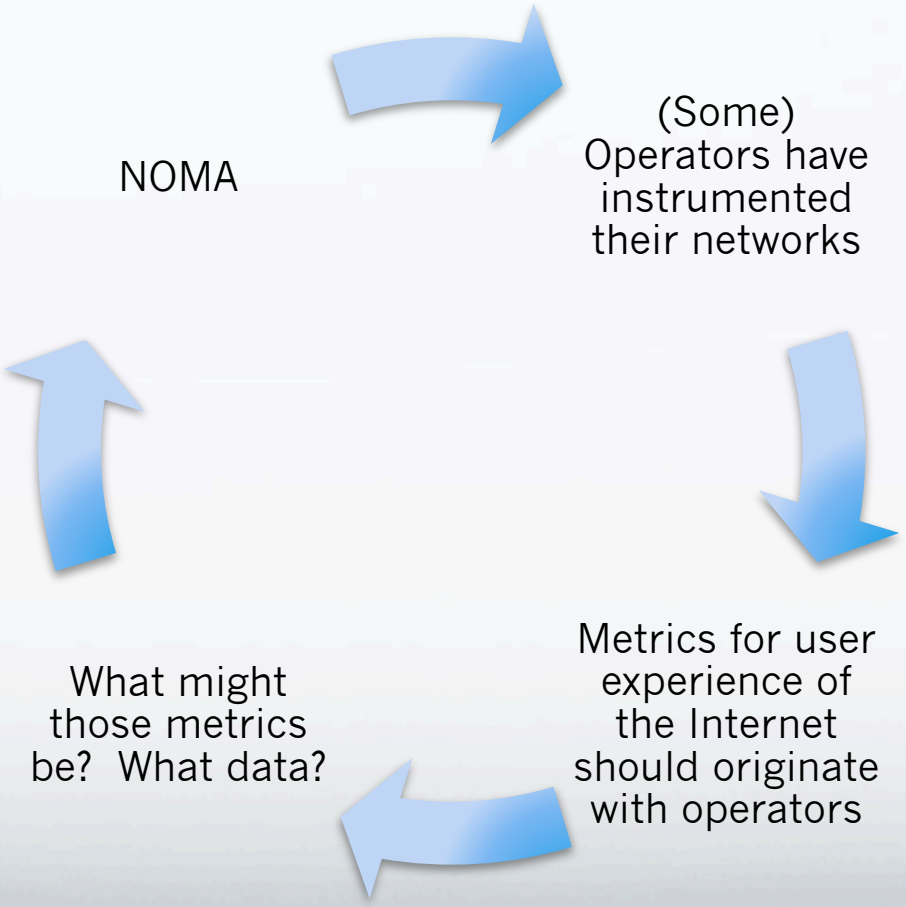
NOMA
(Some) Operators have instrumented their networks
Metrics for user experience of the Internet should originate with operators
What might those metrics be? What data?

Measure, Learn, Change, Measure

- Comcast has previously presented about instrumenting their network
  - Simple measurements from within their network, to put numbers to describe customers' experience
- What could we do if more operators self-instrumented that way and shared the data?
- Build Internet health measurements that tell us what we need to know about the network as end users experience it

Simple Data

- Using “libcurl” and http as the closest approximation of end users’ experience, each collection point could gather:
  - IPv6 DNS lookup to each target website
  - Time to connect to the target website over IPv6
  - Total time for each target website over IPv6
  - IPv4 DNS lookup to each target website
  - Time to connect to the target website over IPv4
  - Total time for each target website over IPv4
  - Traceroute to IP address of each target website

Composed to Insight

- For example: v4/v6 ratio gives you a simple test of whether v4 or v6 is performing better
- For total-time measurements
  - < 1 means v4 activity is faster
  - > 1 means v6 activity is faster
  - = 1 means v4 & v6 are the same

2016 Work

- TL;DR
  - Some operators think this is interesting
  - Getting it started is the hardest part
- Details for further reading
  - Invitational workshop in June 2016
    - http://www.techark.org/wp-content/uploads/2016/09/20160831-WorkshopReport-Final.pdf
  - Survey of existing measurements activities, and framework for comparison
    - http://www.techark.org/internet-measurements-landscape-2016-systems-approaches-and-a-comparative-framework/
  - "Template" for operator measurements activities
    - http://www.techark.org/noma-measurements-template/

How to get data?

- /me suddenly remembers the RIPE Atlas probe that has been hanging out in the basement for years, contributing data and collecting points...
- With RIPE Atlas infrastructure, it's possible to
  - do a simulation of the operators' activity, and
  - get real data (with some limitations)
- I.e., do the measurements outlined, on a per-operator basis

Data!

Locality	Execution (excl DNS) (v4)	Execution (excl DNS) (v6)	Execution (excl DNS) (v4/v6)
San Francisco	278.0186036	300.600902	0.924876146
Washington	234.1762633	175.1095543	1.337312885
Seattle	258.0885657	200.16237	1.289396032
Eugene	500.9710182	556.9739695	0.899451403
Denver	213.5601736	159.0784051	1.342483749
San Diego	219.1249776	218.2439884	1.004036717
Chicago	193.6913491	160.6142147	1.205941513
Miami	192.2035983	103.0980178	1.864280249
Dallas	204.7865375	211.3463822	0.968961642
Boston	289.2981992	113.8183904	2.541752683
Atlanta	117.000087	107.116316	1.092271387
Nashville	127.788385	92.0437175	1.388344457

That was...

- Data from RIPE Atlas probes
  - Showing Total Execution time, but also collected
    - DNS resolution (not on probe)
    - Time to First Byte
    - Time to Connection
- In two networks
  - Comcast
  - Charter
- Associated with “closest” of the 12 localities
  - The “localities” are geographic, not network topology – I have no insight into network topology
- Averaged
- HTTP measurements to one RIPE Anchor
  - “Centrally” located
  - Reston, VA, as it happens

What does it say?

- Note that this is not (yet) a rigorous study
  - One data run
- Nonetheless, there are some interesting things to note when considering the v4/v6 ratio as a metric

Good v6 in Boston; Quite a variety across the US; Eugene has issues...

Locality	Execution (excl DNS) (v4)	Execution (excl DNS) (v6)	Execution (excl DNS) (v4/v6)
San Francisco	278.0186036	300.600902	0.924876146
Washington	234.1762633	175.1095543	1.337312885
Seattle	258.0885657	200.16237	1.289396032
Eugene	500.9710182	556.9739695	0.899451403
Denver	213.5601736	159.0784051	1.342483749
San Diego	219.1249776	218.2439884	1.004036717
Chicago	193.6913491	160.6142147	1.205941513
Miami	192.2035983	103.0980178	1.864280249
Dallas	204.7865375	211.3463822	0.968961642
Boston	289.2981992	113.8183904	2.541752683
Atlanta	117.000087	107.116316	1.092271387
Nashville	127.788385	92.0437175	1.388344457

Initial Reactions

- v4/v6 ratio is nice, but it hides important information about the state of the network.
  - E.g.,

v4	v6	v4/v6
20.0 ms	25.0 ms	0.8
1200.0 ms	1500.0 ms	0.8

- An alternative is to look (also) at the v4-v6 difference

Differences

Locality	Execution (excl DNS) (v4)	Execution (excl DNS) (v6)	Execution (excl DNS) (v4/v6)	Execution (excl DNS) (v4-v6)
San Francisco	278.0186036	300.600902	0.924876146	-22.58229842
Washington	234.1762633	175.1095543	1.337312885	59.06670900
Seattle	258.0885657	200.16237	1.289396032	57.92619572
Eugene	500.9710182	556.9739695	0.899451403	-56.00295133
Denver	213.5601736	159.0784051	1.342483749	54.48176855
San Diego	219.1249776	218.2439884	1.004036717	0.880989143
Chicago	193.6913491	160.6142147	1.205941513	33.077134390
Miami	192.2035983	103.0980178	1.864280249	89.105580500
Dallas	204.7865375	211.3463822	0.968961642	-6.559844722
Boston	289.2981992	113.8183904	2.541752683	175.479808800
Atlanta	117.000087	107.116316	1.092271387	9.883771000
Nashville	127.788385	92.0437175	1.388344457	35.744667500

This is fun! What's next?

- Can keep slicing and dicing with RIPE Atlas probes to vet out the v4/v6 ratio and v4-v6 difference as metrics of IPv6 network health
  - More geographies
  - More networks
  - More targets
  - Repeating the measurements over time
  - Comparing measurements
    - between measurement runs
    - between operators
    - between operators and averages
- Also, look at DNS resolution on probe, not in Atlas network
- More data, more tables, more graphs, more fun!

BUT...

- While this is real data from real networks, it is still only a simulation of what we could do with real operator involvement
  - Atlas coverage is uneven – e.g., ~400 probes in Comcast network

Probes in Network 1

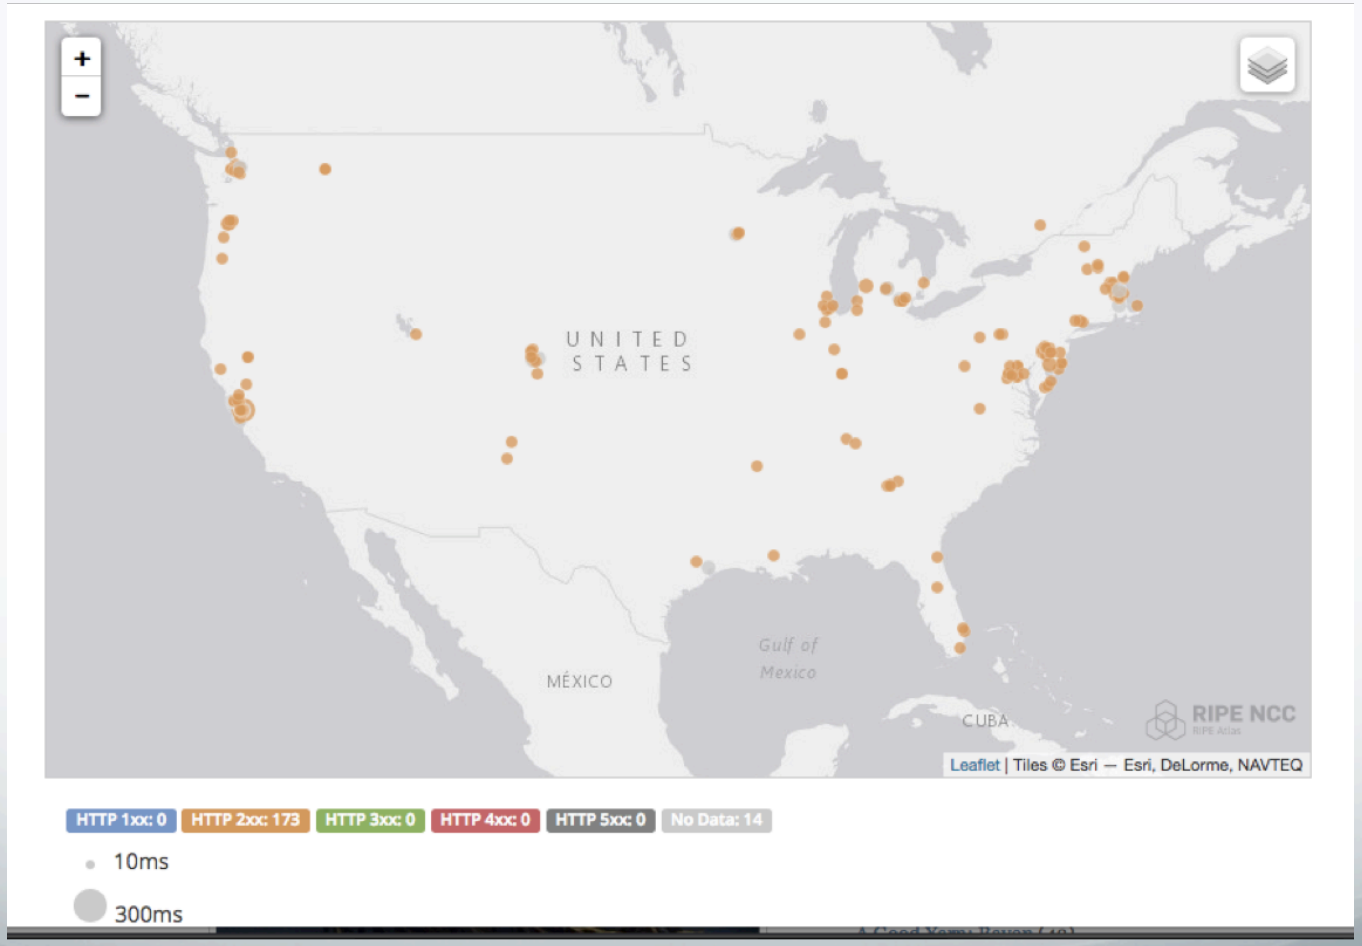

Probes in Network 2

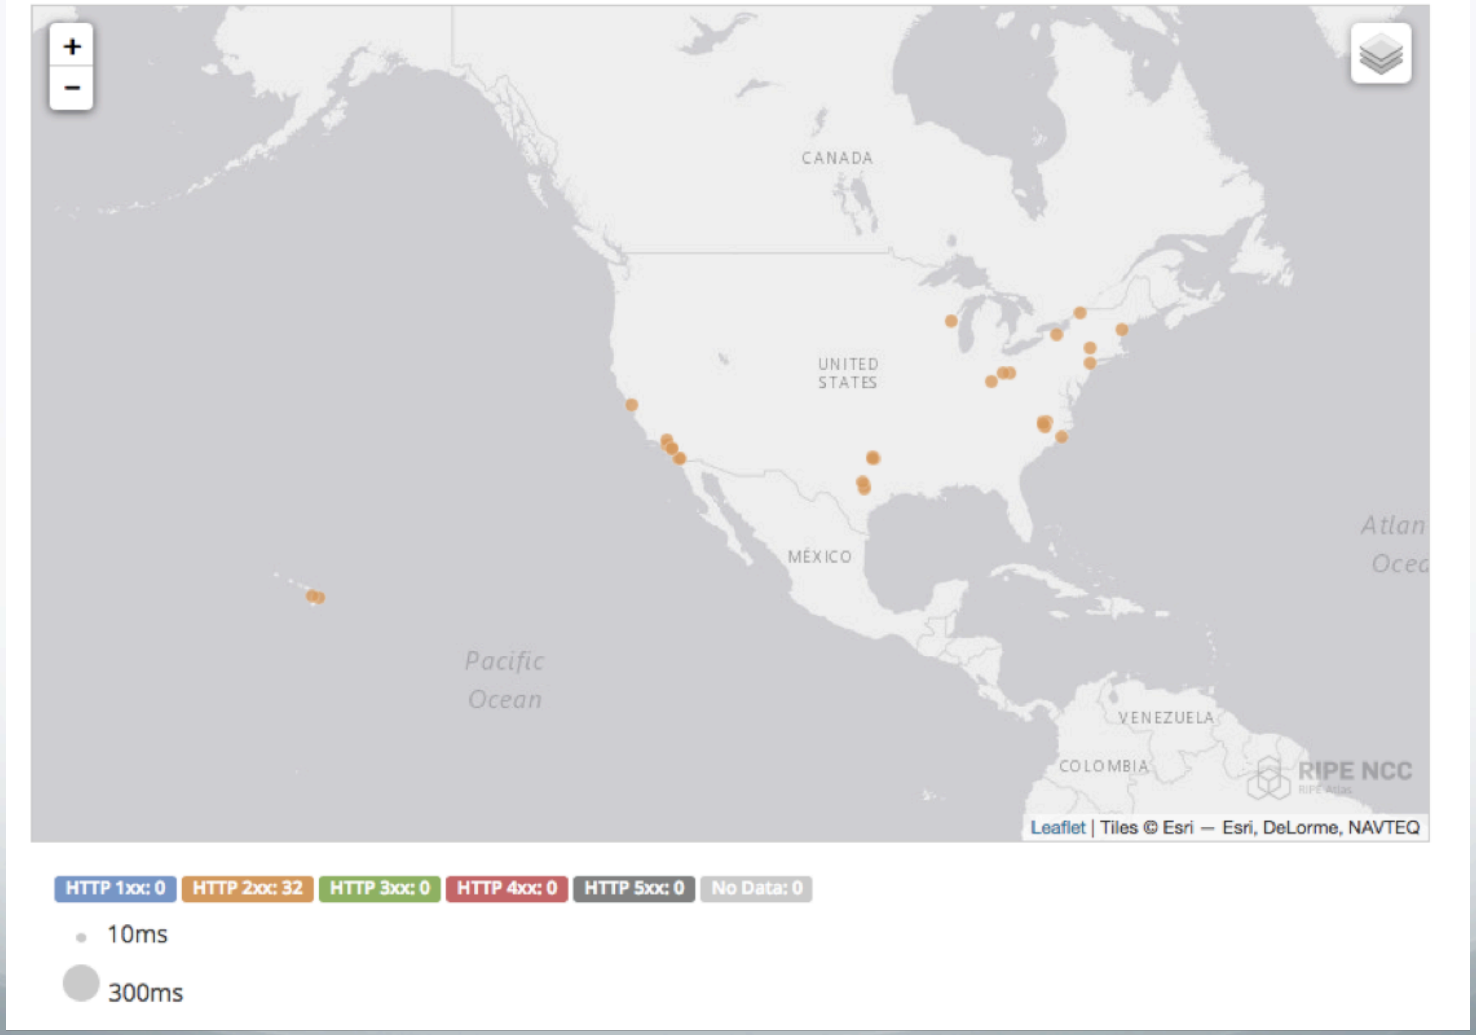

BUT… (continuted)

- Ideally, a network's access endpoints should be completely covered by this measurements framework.
  - At "full layer 3 device closest to customers" – might or might not be DSLAM,
- Atlas measurements include last mile (which is good and bad)
- Anchors are fixed points in networks – but not really exemplary of the user experience in accessing content sites
  - including things like geo-locating closest servers

This NOMA thing....

- Collaborative industry activity to share measurements of network function
  - Measured by the network operator itself
    - In return – improved data about own network
    - E.g., the test run with DNS resolution on the Atlas probe highlighted some (network) failures
  - Contributed to share a collective picture of the Internet's health
- Intended Outcome
  - An actual measure of the Internet's stability and health.
    - Starting with v6 performance (as a ratio with IPv4)
  - A target for new operators (e.g., in developing economies) to shoot for, in terms of performance and "what good looks like"
  - Promoting more networks to be objectively introspective

Takeaways

- v4/v6 ratio and v4-v6 difference are interesting metrics for considering “IPv6 Internet health”
- It’s useful to have that information available publicly
  - Have you already been working out your own network’s numbers while I was talking? ☺
- Network operator self-instrumentation is feasible and would produce useful results
  - Talk to me!

Extra material

TechArk Activities

- Elicit collaborative development of the measurements framework and culture of collecting and sharing
- Provide a platform to collect and share basic benchmarking metrics across participating networks
- Publish a baseline of public metrics for network operators
- Publicize the Activity and its outcome
- Encourage broader participation in the Activity