

Internet Research Task Force (IRTF)

Measurement and Analysis for Protocols (MAPRG)

-- Measurement Challenges in the Gigabit Era --

5 Minute Lightning Talk

20 March 2018

Jason Livingood

Measurement Challenges in the Gigabit Era

- Background:
 - See critical paper from Bauer, Lehr, Hung, “Gigabit Broadband, Interconnection, Propositions and the Challenge of Managing Expectations”
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2586805
- Wide range of existing systems:
 - Goal of many national systems is to ensure ISPs deliver on advertised rates
 - There are also many user-facing systems (e.g. Speedtest.net)
- But technical challenges are forming on the horizon
 - Apparent issues >100 Mbps today
 - Many of our customers 100 – 500 Mbps in 2018, large numbers at 1 Gbps in 2019
- And end user expectations are changing
 - It is not just about speed as a standalone issue but QoE more broadly, including of reliability/availability/performance of application services.

Measurement Challenges in the Gigabit Era

- Existing system issues

- Bottleneck link migrating from home LAN and access network to interconnection points and server-side infrastructure
- Intended measurement point has shifted, but design has not evolved (design objective now unmet)?
- Can negatively impact results by 20% - 30% at current end user speeds

- Key questions in the gigabit era

- Does it make sense to continue to use standalone destination servers, especially since no common destination sites use 1 Gbps? Why have a test that bears no resemblance to actual user behavior?
- Alternatively, should some sort of public cloud-based or CDN-based or other design be used? Such as 100 smaller tests to 100 distinct destinations in a range of clouds/CDNs.
- How can the server infrastructure be better and more independently & transparently monitored?
- Is it better to test multiple parallel tests of a smaller size to several end sites? And how would that impact network design in response? Are there unintended consequences?
- Should we shift away from speed to test latency, reliability/uptime, latency under load, or other service attributes?
- End user QoE is impacted as much by application service performance as ISP performance. Should performance testing shift to focus on this?
- What other practices do we want to encourage, such as BCP-38? (if you measure it you will incentivize it)

THANK YOU

Feel free to reach out via email:
jason_livingood@comcast.com