On “Speed” and “Latency”

BUFFERBLOAT.NET’s response to RFC - NTIA–2021–0002

This filing focuses on two aspects of “Question 13,” not well understood on today’s internet.

“13. NTIA is committed to ensuring that networks built using taxpayer funds are capable of meeting Americans’ evolving digital needs, including broadband speeds and other essential network features. What guidance or requirements, if any, should NTIA consider with respect to network reliability and availability, cybersecurity, resiliency, *latency, or other service quality features and metrics*?

What criteria should NTIA establish to assess grant recipients’ plans to ensure that service providers maintain and/or exceed thresholds for reliability, quality of service, sustainability, upgradability and other required service characteristics?”

The long-standing single minded focus on “speed” (“Mbps”) is a thing of the past,[1] especially for voice, videoconferencing, and gaming. Consistently low latency while an Internet connection is being used (a.k.a. working latency) is as important as throughput . What matters more, once a user has enough bandwidth (a.k.a. throughput), is consistently low latency across simultaneous uses and users of the internet in a household, and the quality of experience (QoE) for the end users. Unfortunately, many ISPs’ headends and modem CPE today still have really enormous, unmanaged buffers that make efficient multiplexing of these traffic types impossible.

It is indeed possible to achieve consistent low latency across uses and users, as demonstrated by the deployment of fair queuing (FQ), Active Queue Management (AQM), and similar queue management technologies. But these solutions have been slow to deploy universally. We recommend that NTIA require the implementation of methods to reduce working latency. Otherwise, there is a high risk of building networks with large peak throughput numbers but poor day-to-day performance for end users.

With improved and modern queuing technologies in place, the amount of bandwidth actually needed is remarkably lower than what people believe. For example, it’s widely misunderstood that more bandwidth than 20 Mbits improves web page load time (PLT). It doesn’t. Web page loads are actually primarily bound by the physical path latencies of the link, and from the load

[1] https://www.bitag.org/documents/BITAG_latency_explained.pdf

from other simultaneous traffic on that link.

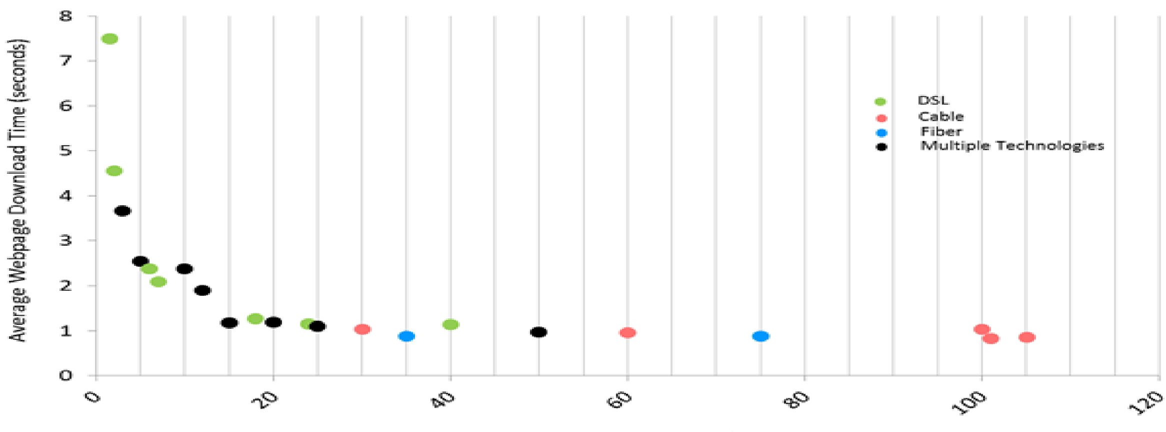

Video conferencing presently has a dynamic range of about 500 kbits to 4 Mbits per user, and while we list the bandwidth requirements of each form of traffic below, it's the interpacket latencies that matter more when two or more forms of this traffic exist on a link. With short enough delays and sufficient multiplexing, all these forms of traffic can adjust rapidly to changes in demand on the link.

Application	Examples	Upload Bitrate	Download Bitrate
Voice	VoIP	32Kbps (unless HD)	32Kbps (unless HD)
Video Conferencing	Webex, Zoom, etc	400Kbps-4Mbit	1Mbit – 20Mbit
Interactive Gaming	UDP, C&C traffic	64Kbps	64Kbps
Audio Streaming	MP3 and FLAC files	16Kbps	256Kbps
Video Streaming	MPEG-2	128Kbps-1Mbit	1.5Mbit - 35Mbit
Twitch Streaming	Game sharing	6-8Mbit	300Kbit
Web Access	HTML, images	~1/20th the download	Infinite ~2-3 sec/pg
File Transfers	Dropbox, firmware updates, SCP, etc	infinite	infinite

To manage the queuing delay problem (also known as "bufferbloat"), there have been great innovations in the last decade[2] widely available in open source, third party firmware and newer routers — and standardized in the IETF.[3] In many newer routers you will find better queue

[2] https://blog.apnic.net/2020/01/22/bufferbloat-may-be-solved-but-its-not-over-yet/
[3] See the various RFCs produced as part of the IETF AQM Working group: https://datatracker.ietf.org/group/aqm/documents/

management algorithms if you look for options like “Smart Queue Management,” or “Optimize for videoconferencing and gaming,” but these are largely left for users to discover and enable for themselves, rather than configured by the ISP[4] or mandated by a regulator.

Currently, much of the existing deployment of Internet services across the United States has yet to coherently deploy these on the network headends (including central offices and radio access networks) or CPE.

Worse, most of our measurement methodologies measure upload, download, and latency separately, rather than all at the same time. Latency tested in isolation is “idle” latency and means very little to the end user QoE, whereas latency tested in the face of upload/download traffic is “working” latency and is the key QoE indicator for most applications.

An analogy to this sad situation would be a car capable of 200 mph, but lacking adequate brakes, shock absorbers, a steering wheel, and an airbag in an age before *Unsafe at Any Speed* was published. The common “speed test” derived metrics where upload, download, and latency are tested *separately* have zero relationship to how the modern internet is actually used, or a car, actually driven.

Recommendations

“On Speeds”

1) Speed *is* important. Latency under load, or “working latency,” needs to be measured. Apple’s recent efforts on measuring “responsiveness” and providing a tool to all users to measure that is a good step forward. See Appendix A.
2) A good start would be something as simple as asking a question on a grant application form such as, “What methods and techniques do you intend to deploy to manage queuing delay across your bottleneck links?” This should also focus on both the ISP network and the CPE installed in user homes, not just within the ISP’s core network.
3) If you must focus on Mbits, some sort of ratio on a per member of household basis above a baseline makes more sense, and improving upload speed is more important than downloads. Bear in mind that many applications use “adaptive bitrate” protocols and can dynamically adjust to a wide range of bandwidth, providing a good experience at even very low bitrates (but may be highly sensitive to working latency).
4) As a baseline, the present day 25Mbit download standard for broadband (with good queueing activated) is actually a good number for a typical family of four. However, 3Mbit uploads are lower than desirable for quality videoconferencing, especially with those WFH and others attempting to telelearn. Moving forward, the standard utility baseline (lifeline) service, should start at 25Mbits down, and 10Mbits up in order to provide robust web and educational access, including videoconferencing, to all US citizens. Video

[4] With one notable exception: Comcast has rolled out the PIE AQM. Results: https://blog.apnic.net/2021/12/02/working-latency-the-next-qoe-frontier/

conference calls can utilize up to 3Mbits of upload bandwidth per end-user, with multiple video calls competing for bandwidth when uplink bandwidth is limited.[5] Households with multiple students or remote workers can struggle with simultaneous video conferencing when upload bandwidth is limited to 3Mbits, the current standard for broadband. The Consortium for School Networking recommends 12Mbits upload bandwidth per student based on their findings.[6] While that may not be achievable for all ISP networks, 10Mbits upload bandwidth may offer a more viable intermediate step to ameliorate remote learning and WFH challenges for households. Stable video conferencing is essential to ensure that remote workers and remote learning students are not left behind.

5) After establishing the baseline for upload and download bandwidth in 2022 the figure is 'indexed to inflation' in the future and will adjust annually at the nation's usage compound annual growth rate (CAGR) - such as 40% growth in download and 30% growth in upload. This eliminates qualitative decision-making and political debate and simply drives future baseline numbers based on actual growth in user demand, reflective of new applications and uses, new technologies and devices, and so on.
6) A family of two or more MUST have a network that provides reasonable multiplexing between different forms of simultaneous traffic. With the rise of remote work, this is even more important from an economic perspective. While one family member is trying to have a video conference, another may be downloading video assets for editing, reviewing large files, or even just trying to have a second video conference, are all situations where FQ+AQM mechanisms prevent both activities from crawling to a halt. Great benefits could be had by the NTIA establishing guidance and programs to encourage the upgrade and the default configuration of better queue management strategies of the routers along the edges of both existing and new networks. Upgrading the routers in areas that have existing (but poor) service would be an inexpensive, effective, and near immediate network improvement for many households.

Further Suggestions

- Equipment deployed henceforth should apply IETF standards for AQM (Active Queue Management) and/or fair queuing (FQ) along the edge, and should apply subsequent standards in the future.
- Old networks can be improved with better routers also. Short term fixes will help!
- Funds could be provided to subsidize replacement of legacy CPE to spur a shift to AQM and/or FQ capable routers.
- Funds to raise awareness of the working latency problem would be helpful.
- Funds for better training and understanding of networking basics, such as how to read a packet capture, would be good.
- In locations that cannot be quickly or easily upgraded, a rapid deployment of upgraded routers and headend equipment with these modern queuing algorithms should be encouraged.

[5] https://arxiv.org/pdf/2105.13478.pdf
[6] https://thejournal.com/articles/2021/05/04/landmark-study-calls-for-increased-bandwidth-for-at-home-learning.aspx

- Upgrades for new and existing CPE can range from a “bump in the wire”[7] to QoE middleboxes (Preseem, Paraqum, LibreQoS).[8] But the best upgrade would be using newer upgradable routers[9] with smarter queue management systems built in and enabled by default. Finding a way to retrofit older routers at a larger scale, “an Upgrade in Place,” is a proposal still under development that we are seeking more comments on.[10]

In summary, a great deal of pressure towards more bandwidth can be better put towards establishing policies and regulations that encourage the use of new technologies that will enable *better* bandwidth, and this will do wonders for end users’ quality of experience.

A busy network doesn’t have to be a bad network.

Sincerely,

Dave Taht
Director, Bufferbloat project
dave.taht@gmail.com <Dave Taht>
408-613-0871

---

[7] https://apenwarr.ca/log/?m=201808
[8] https://github.com/rchac/LibreQoS
[9] https://www.bufferbloat.net/projects/bloat/wiki/What_can_I_do_about_Bufferbloat/
[10] https://docs.google.com/document/d/1T21on7g1MqQZoK91epUdxLYFGdtyLRgBat0VXoC9e3I/edit

Appendix A - Examples of Over-Buffering

Apple Responsiveness Test Results

Excessive or ineffectively managed buffering is a key factor in why DSL, in particular, is perceived as “slow.” One upload or download, a sufficiently large web page, or a burst from a videoconferencing app can render a link nearly unusable for other members of a household.

This is a test run of Apple’s new “responsiveness” test, with an “Evenroute” router, one of many brands that leverages the open source “sch_cake”[11] algorithm to manage queue delay much better under working conditions on an otherwise very poor 7Mbit down DSL connection.

From the command line, under OSX, this test is available via:

networkQuality -s -v

And it’s also available under developer settings on the iPhones, with an open source version under development.

==== SUMMARY ====
Upload capacity: 497.994 Kbps
Download capacity: 7.440 Mbps
Upload flows: 12
Download flows: 12
Upload Responsiveness: Medium (259 RPM)
Download Responsiveness: Medium (303 RPM)
Base RTT: 36
Start: 2/1/22, 3:55:36 PM
End: 2/1/22, 3:56:14 PM

The ISP supplied DSL router, not running a modern FQ+AQM algorithm at all, won't even finish the test. It times out and cancels! The DSLReports score was an F, with greater than 5 seconds of bloat. A link with more than 100ms of delay is nearly unusable in the first place!

This sort of bad behavior is not limited to DSL (although nowadays it is frequently the worst), but is nearly omnipresent on other access technologies, and is especially bad on wireless networks such as WiFi, LTE, and 5G.

[11] https://arxiv.org/abs/1804.07617 fq_codel, cake, and pie are open source software.

SpaceX Starlink's Over Buffering

For the most recent access technology to get queuing delay wrong, we measure the uplink bufferbloat and jitter on SpaceX Starlink's new services below. The test in this case is one of the simplest tests in the commonly available flent test suite, a single tcp upload with "ping" simultaneously.

flent -H somewhere –socket-stats –step-size=.05 -te=upload_streams=1 -t test_param tcp_nup

There are many other tests in this suite. Of note are the rtt_fair, rrul, and tcp_ndown tests.

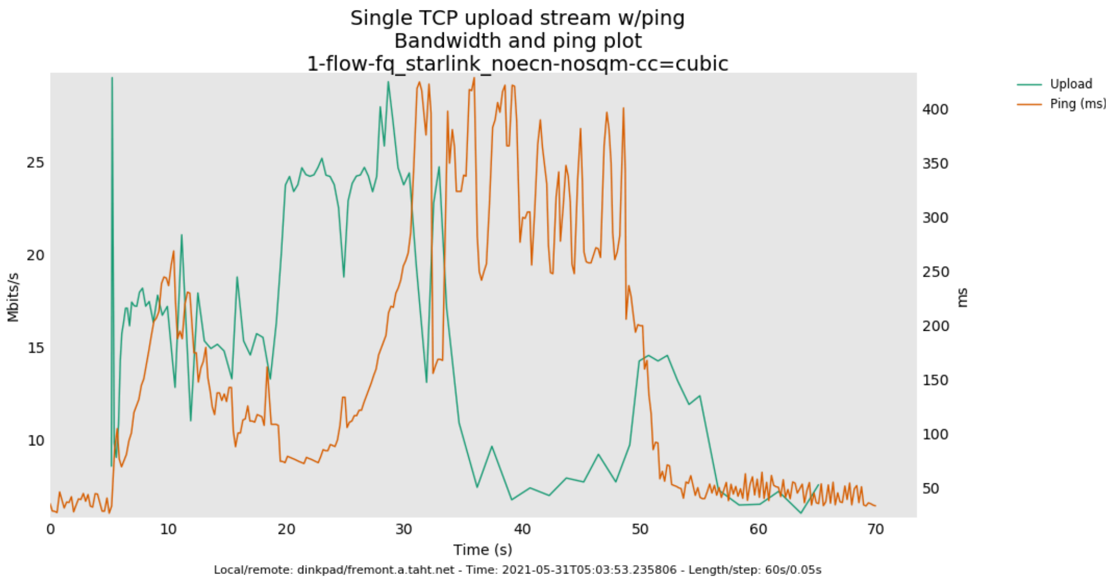

This is the typical behavior of the most common TCP congestion control algorithm, "Cubic" on a beta Starlink terminal in May 2021. This **single flow** (one upload) interacts badly with the bandwidth allocation schemes supplied by Starlink and with what appears to be about a 400ms long FIFO queue at a 2Mbit rate. The cost of a fixed length queue of this sort is obvious. The FIFO induced inverse relationship between latency and throughput is clearly shown at T-20 through T-30, where a rate of approximately 24Mbits is achieved at ~80ms latency, and a ~2Mbit rate (T-30-T50), at 400ms latency.

Taking an "average" bandwidth here is futile, and misrepresentative of actual QoE. A VoIP call that goes from 20ms (shouting across a room) to 400ms (shouting across a football field) is very undesirable.

Bufferbloat on Cellular

It would be unfair to not show a 5G result. This is the same test, run via a tether, through an Android 5G phone.

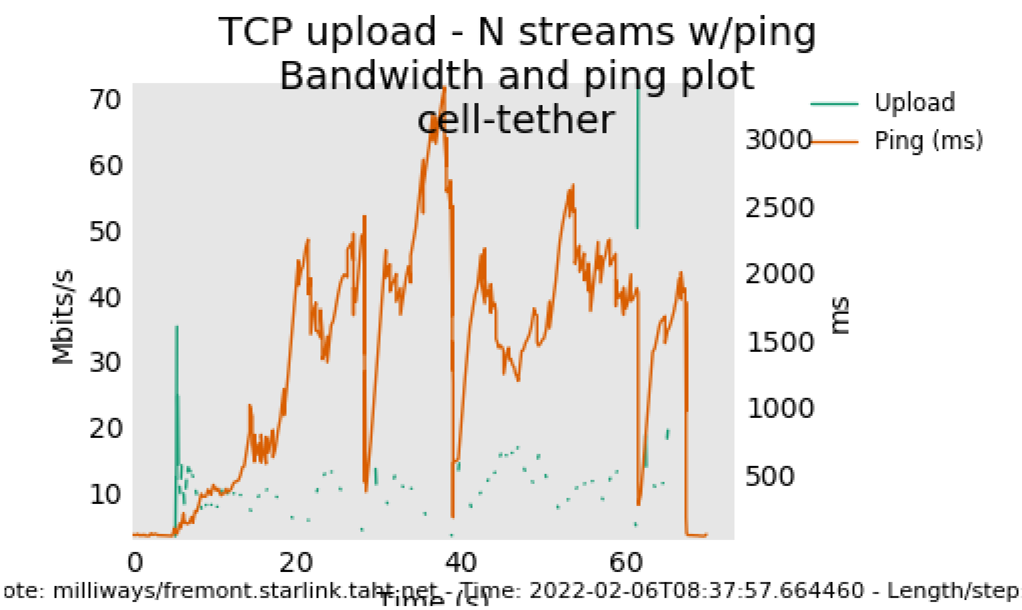

Starlink is actually vastly superior to a modern day cell phone in these regards. Downloads on cellular also induce a great deal of extra latency compared to other technologies.

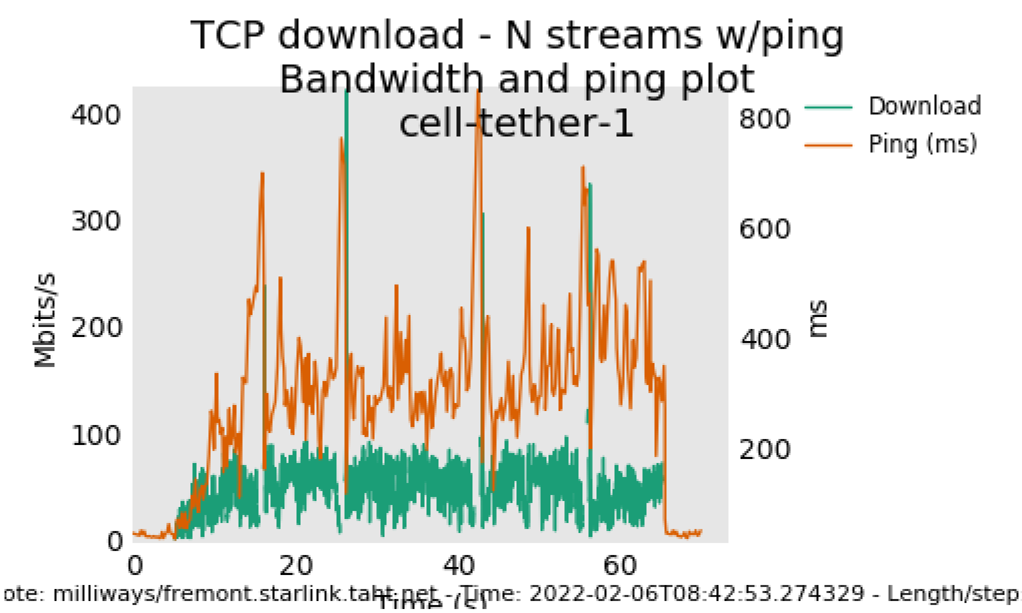

100/100Mbit Fiber Tests

Lastly, we can see fiber alone is no answer. Recent measurements of 100/100 fiber show a disturbing trend towards over 250ms of working latency on that technology also, where well under 16ms is feasible with modern queuing techniques. This data is produced by users of the DSLReports speed test, which also measures bufferbloat. Routers that have quality queue management typically exhibit working latencies well below 60ms.

From: https://www.dslreports.com/speedtest/results/bufferbloat?up=1

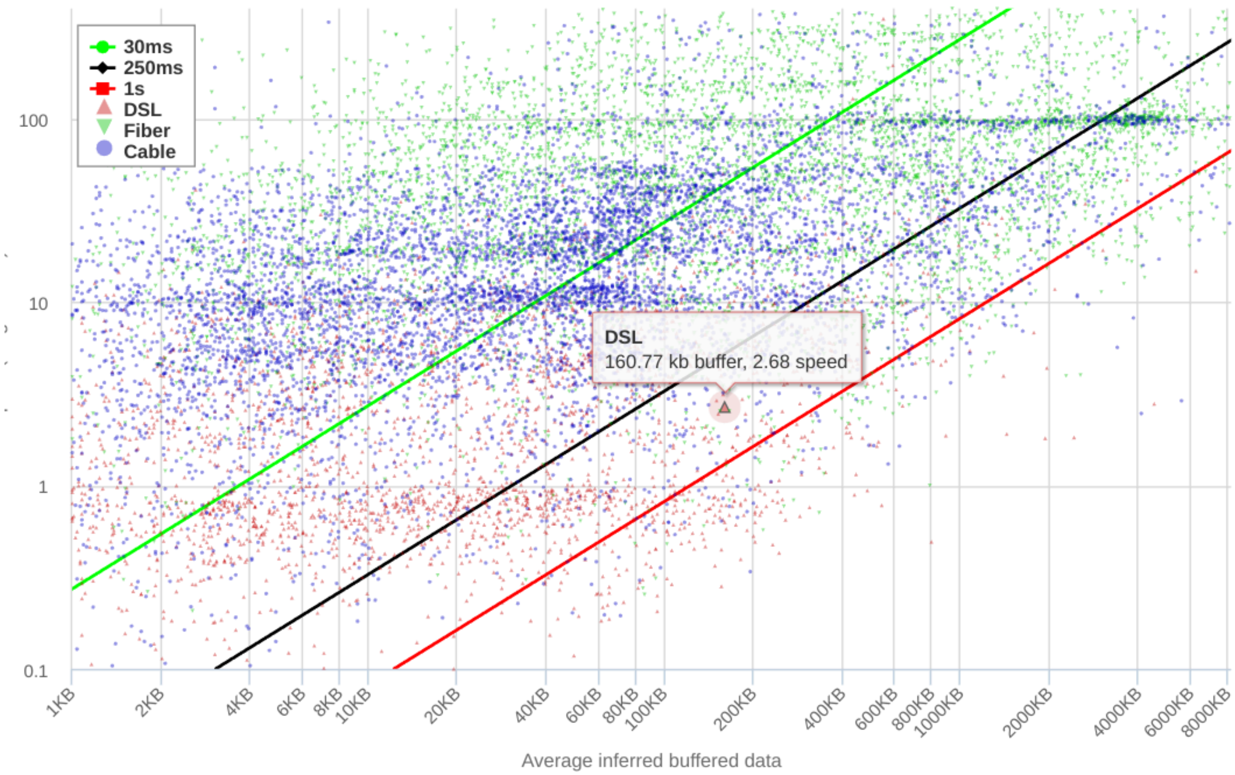

Note: In the field, working latencies of well over 1 second (the red line) have been observed. This test doesn't run long enough to find those.

There needs to be a focus on moving America forward, not just on bandwidth, but working latency!

APPENDIX B - Non-Bufferbloated behaviors

APPENDIX C - WEB AND OTHER TRAFFIC BANDWIDTH REQUIREMENTS

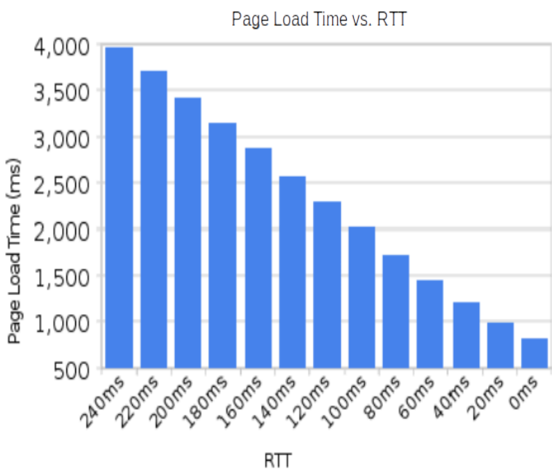

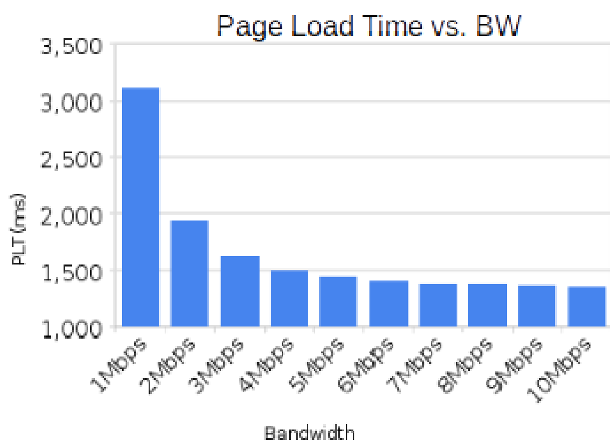

Source: SPDYEssentials, Roberto Peon & William Chan, Google Tech Talk, 12/8/11