Before the Federal Communications Commission
Washington, D.C. 20554

November 4, 2010

| | | |
|---|---|---|
| In the Matter of | ) | |
| | ) | |
| Preserving the Open Internet | ) | GN Docket No. 09-191 |
| | ) | |
| Broadband Industry Practices | ) | WC Docket No. 07-52 |

Reply Comments of Bret T. Swanson[1], responding to the *Further Inquiry*

In this perhaps final round of comments on the Commission's inquiries into the Open Internet and Broadband Industry Practices, we offer some new data and analysis relevant to the state of American broadband and then address several specific questions posed by the *Further Inquiry*.

**State of American Broadband**

Does America really rank 15th in the world in broadband? Or even 26th, as asserted by one recent report?[2] These metrics have become conventional wisdom and have driven much of the policy debate these last few years. But is it true that America lags the world? Or more to the point, is the particular metric correctly defined, is it meaningful, and are there metrics that better capture America's broadband standing?

The familiar international broadband rankings, based on broadband connections per 100 inhabitants, are not very meaningful. As numerous critiques have found, the "connections per capita" figure may tell us more about household size than broadband. The Phoenix Center showed that if every household in every OECD nation enjoyed a broadband connection – thus achieving 100% broadband "Nirvana," in Phoenix's words – the U.S. would fall from 15th to 20th in the world broadband rankings.[3] All this means is that the U.S. has larger households than many other nations. Clearly, better measures are needed.

We think a more effective way to gauge relative broadband success is to measure how much people actually use the Internet. Using figures from Cisco's Visual Networking Index, which measures actual IP traffic, and Internet user figures from the ITU, we developed metrics of IP traffic per Internet user and per capita. As shown in the chart below, the U.S. generates substantially more IP traffic per user than any other region of the world.

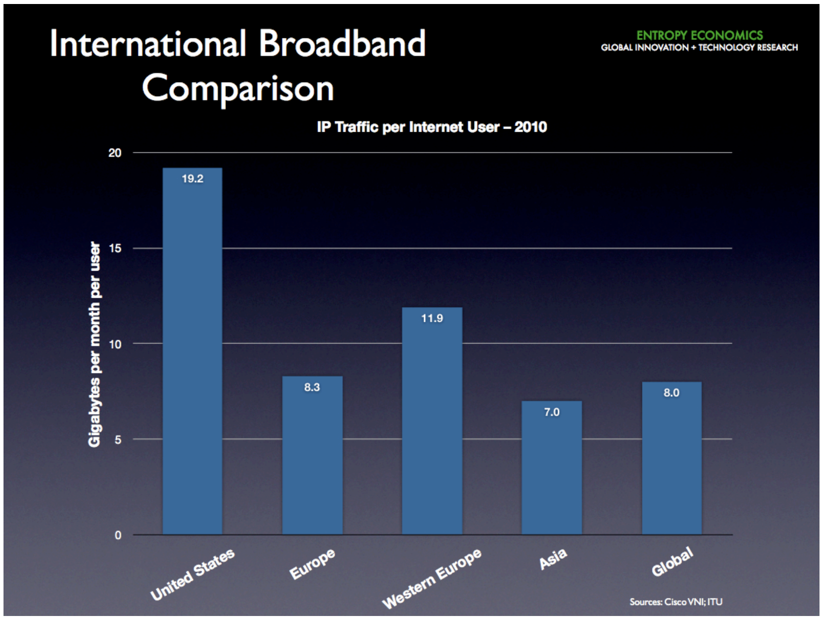


In a more granular measure among nations, only South Korea generates significantly more IP traffic per user than the U.S. (Canada and the U.S. generate essentially the same amount). South Korea's performance in this metric is not surprising. It was the first nation to widely deploy fiber-to-the-x broadband and both 2G and 3G mobile networks, and it enjoys a vibrant and youthful Internet culture.

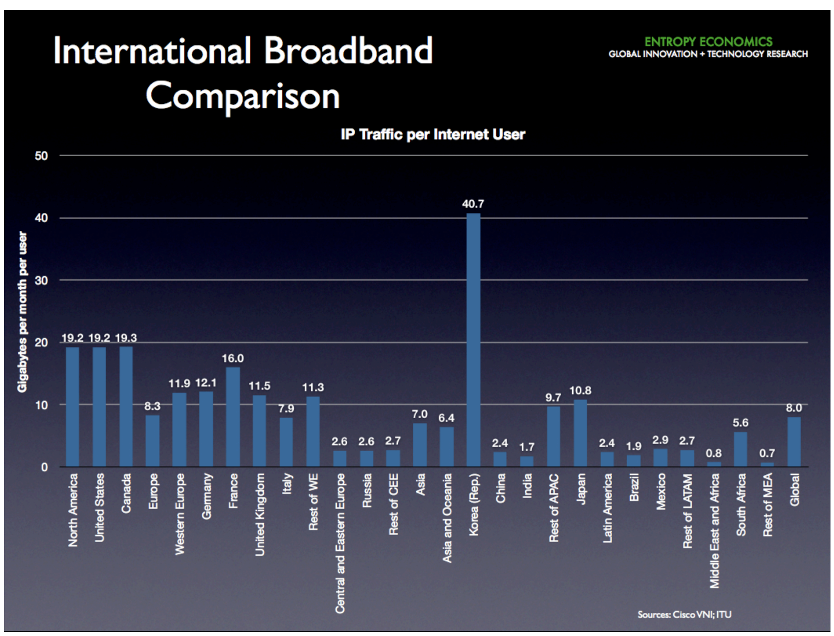


These figures are not simply an artifact of some small number of U.S. users generating a high per-user figure. The U.S. enjoys a similarly high ranking when dividing total traffic by entire national populations, yielding high per capita Internet traffic as well. The U.S. simply could not generate around twice the Internet traffic per user (and per capita) as other advanced nations if it did not compare favorably in both the number and quality (robustness, capacity, speed, prices) of its broadband connections and of its overall network. These metrics confirm other qualitative measures of American broadband innovation – such as the native development of Web video apps like YouTube, infrastructure innovations like content delivery networks, social networks like Facebook and Twitter, and the iPhone and its "app store." The metrics suggest the U.S. boasts a very large number of Internet users who enjoy broad access to, and make good use of, fast and robust wired and wireless broadband networks.

There are, of course, other metrics that might shed further light on relative broadband success. But this particular measure suggests the U.S. broadband Internet ecosystem is healthy.

**Resources and Constraints**

Before we address several issues raised in the *Further Inquiry*, it may help to outline the broad set of resources and constraints that govern the building and operation of the Internet.

Our digital world is built using three basic resources: computing power, digital storage, and bandwidth. These three key technologies – processing, storage, and transmission – are the building blocks of the larger Internet – all the PCs, mobile devices, wired and wireless networks, servers, data centers, switches, routers, and endless network nodes. We deploy these building blocks in varied quantities, mixes, and physical locations. Depending on the situation, we can substitute processing power for bandwidth, and vice versa. Sometimes it makes sense to store bits rather than transmit them. Often the reverse is true. A silicon chip itself is an extraordinarily complex, highly designed device that processes, stores, and transmits electrons. We can think of a microprocessor as a network on a chip or, as George Gilder wrote, the Internet as a chip on a planet.

We blend these resources to create functioning machines and networks and to deliver useful services to consumers and enterprises. Depending on the state of each of these technologies – their inherent technical virtues and (as we will discuss below) their cost-performance attributes – we change the amounts, types, and locations of bandwidth, storage, and processing in all our evolving digital recipes. There is no single best recipe, no one solution. All these technologies are advancing at differing rates, as are the applications, services, software, and business models built on top of our digital infrastructure.

We also operate under three fundamental constraints:

$$\boldsymbol{c}\text{, W, and \$}$$

The speed of light, ***c***, determines not just how fast a digital bit can traverse the country or the globe but also the processing speed of silicon chips. So far, we've not exceeded Einstein's limit. Power, represented in watts, W, requires energy, either from the grid or a battery. Power generates heat. Power, in a data center, needs air conditioning. Energy can only be stored and generate power for so long. And of course, power costs money. Which happens to be our final constraint, $. All of these technologies cost money – chips, phones, routers; cell towers, data centers, digging ditches to lay optical fiber. The cost-performance of the various technologies is always shifting. Moreover, the financial outlays to build and maintain our digital infrastructure must conform to or produce products and services that are valued and paid for by consumers and businesses. These are not easy calculations in the rapidly shifting marketplace of digital media.

**The Geodesic Network**

In September 2010, a new network company that had operated in stealth mode digging ditches and boring tunnels for the previous 24 months, emerged on the scene. As *Forbes* magazine described it, this tiny new company, Spread Networks

> spent the last two years secretly digging a gopher hole from Chicago to New York, usurping the erstwhile fastest paths.
>
> Spread's one-inch cable is the latest weapon in the technology arms race among Wall Street houses that use algorithms to make lightning-fast trades. Every day these outfits control bigger stakes of the markets – up to 70% now. "Anybody pinging both markets has to be on this line, or they're dead," says Jon A. Najarian, cofounder of OptionMonster, which tracks high-frequency trading.
>
> Spread's advantage lies in its route, which makes nearly a straight line from a data center in Chicago's South Loop to a building across the street from Nasdaq's servers in Carteret, N.J. Older routes largely follow railroad rights-of-way through Indiana, Ohio and Pennsylvania.
>
> At 825 miles and 13.3 milliseconds, Spread's circuit shaves 100 miles and 3 milliseconds off of the previous route of lowest latency, engineer-talk for length of delay.[4]

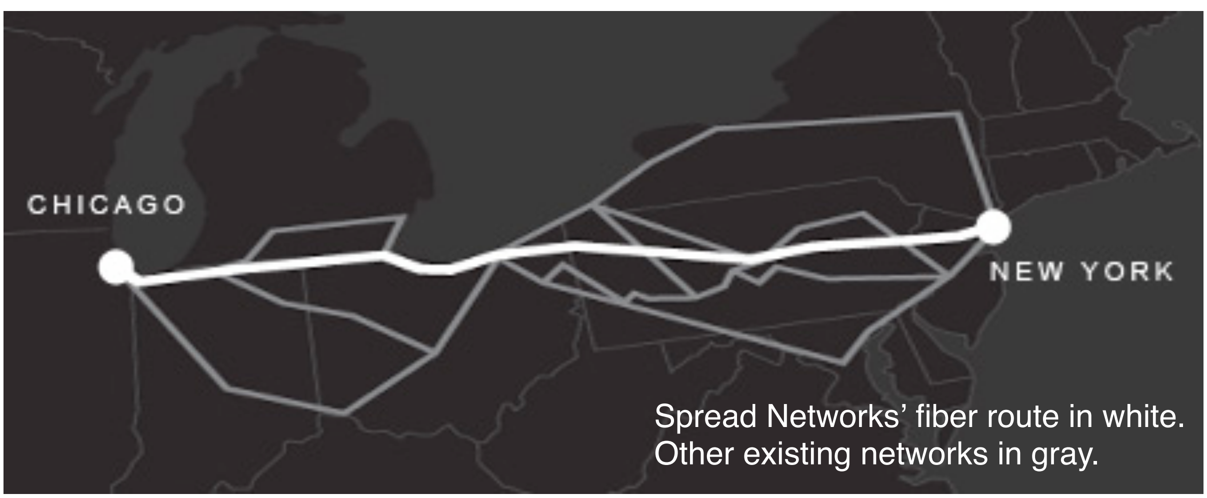

Spread Networks' fiber route in white.
Other existing networks in gray.

Why spend an estimated $300 million on an apparently duplicative route when numerous seemingly similar networks already exist? Because, Spread says, three milliseconds matters.

Spread offers guaranteed latency on its dark fiber product of no more than 13.33 milliseconds. Its managed wave product is guaranteed at no more than 15.75 milliseconds. It says competitors' routes between Chicago and New York range from 16 to 20 milliseconds. We don't know if Spread will succeed financially. But Spread is yet another demonstration that latency is of enormous and increasing importance. From entertainment to finance to medicine, the old saw is truer than ever: time is money. It can even mean life or death.

A policy implication arises. The Spread service is, of course, a form a "paid prioritization." Companies are paying "eight to 10 times the going rate" to get their bits where they want them, when they want them.[5] It is not only a demonstration of the heroic technical feats required to increase the power and diversity of our networks. It is also a prime example that numerous network users want to and will pay money to achieve better service.

One way to achieve better service is to deploy more capacity on certain links. But capacity is not always the problem. As Spread shows, another way to achieve better service is to build an entirely new 750-mile fiber route through mountains to minimize laser light delay. Or we might deploy a network of server caches that store non-realtime data closer to the end points of networks, as many Content Delivery Networks (CDNs) have done. But when we can't build a new fiber route or store data – say, when we need to get real-time packets from point to point over the existing network – yet another option might be to route packets more efficiently with sophisticated QoS technologies.

Each of these solutions fits a particular situation. They take advantage of, or submit to, the technological and economic trade-offs of the moment or the era.[6] They are all legitimate options. Policy simply must allow for the diversity and flexibility of technical and economic options – including paid prioritization – needed to manage networks and deliver value to end-users.

**Wired or Wireless? Most Likely, Both**

The Commission has inquired about the wisdom of applying its proposed Net Neutrality rules to wireless networks. As we[7] and many other commenters have argued, and as the Commission has acknowledged, wireless networks are a special case. They are more delicate, tricky, and capacity constrained than wired networks. Wireless networks, therefore, would be especially hard hit by any regulations that inhibit network management and/or business practices. Recently, the case of T-Mobile's network succumbing to one aggressive app confirmed the severe problems that would arise if a restriction on wireless network management were imposed.[8] We should thus refrain from imposing Net Neutrality regulation on wireless.

*But an exemption of wireless from Net Neutrality regulation is not sufficient to protect wireless from the possible harmful effects of new wireline regulation*.

Wireless access is expanding at an astounding rate. Even in our homes and offices, we now increasingly access the Internet via wireless devices, whether via a Wi-Fi connected notebook, desktop, smartphone, or tablet. Almost all wireless devices, however, apart from "walkie-talkie" radios and wireless appliances, attach to wired networks sooner or later. They certainly do if we want to reach the Internet.

The number of Wi-Fi and femtocell nodes will only continue to grow. It is important that they do, so that we might offload a substantial portion of traffic from our mobile cell sites and thus improve service for users in mobile environments. We will expect our wireless devices to achieve nearly the robustness and capacity of our wired devices. But for this to happen, our wireless and wired networks will often have to be integrated and optimized. Wireline backhaul – whether from the cell site or via a residential or office broadband connection – may require special prioritization to offset the inherent deficiencies of wireless. Already, wireline broadband companies are prioritizing femtocell traffic[9], and such practices will only grow. If such wireline prioritization is restricted, crucial new wireless connectivity and services could falter or slow.

Wired and wireless need each other. They are deeply interconnected, technically and economically. Applying an onerous regulatory regime to wireline, in the hopes of exempting wireless, would probably fail to achieve the Commission's goal.

**Specialized Services**

There has also been discussion of an exemption for "specialized services." Like wireless, it is important that such specialized services avoid the possible innovation-sapping effects of a Net Neutrality regulatory regime. But the Commission should consider several unintended consequences of moving down the path of explicitly defining, and then exempting, particular "specialized" services while choosing to regulate the so-called "basic," "best-effort," or "entry level" "open Internet."

Regulating the "basic" Internet but not "specialized" services will surely push most of the network and application innovation and investment into the unregulated sphere. A "specialized" exemption, although far preferable to a Net Neutrality world without such an exemption, would tend to incentivize both CAS providers and ISPs service providers to target the "specialized" category and thus shrink the scope of the "open Internet."

In fact, although specialized services should and will exist, they often will interact with or be based on the "basic" Internet. Finding demarcation lines will be difficult if not impossible. In a world of vast overlap, convergence, integration, and modularity, attempting to decide what is and is not "the Internet" is probably futile and counterproductive. The very genius of the Internet is its ability to connect to, absorb, accommodate, and spawn new networks, applications and services. In a great compliment to its virtues, the definition of the Internet is constantly changing.

Moreover, a regime of rigid quarantine would not be good for consumers. If a CAS provider or ISP has to build a new physical or logical network, segregate services and software, or develop new products and marketing for a specifically defined "specialized" service, there would be a very large disincentive to develop and offer simple innovations and new services to customers over the regulated "basic" Internet. Perhaps a consumer does not want to spend the extra money to jump to the next tier of specialized service. Perhaps she only wants the service for a specific event or a brief period of time. Perhaps the CAS provider or ISP can far more economically offer a compelling service over the "basic" Internet with just a small technical tweak, where a leap to a full-blown specialized service would require more time and money, and push the service beyond the reach of the consumer. The transactions costs of imposing a "specialized" quarantine would reduce technical and economic flexibility on both CAS providers and ISPs and, most crucially, on consumers.

Or, as we wrote in our previous Reply Comments about a related circumstance, "A prohibition of the voluntary partnerships that are likely to add so much value to all sides of the market – service provider, content creator, and consumer – would incentivize the service provider to close greater portions of its networks to outside content, acquire more content for internal distribution, create more closely held 'managed services' that meet the standards of the government's 'exclusions,' and build a new generation of larger, more exclusive 'walled gardens' than would otherwise be the case. The result would be to frustrate the objective of the proceeding. The result would be a less open Internet."

It is thus possible that a policy seeking to maintain some pure notion of a basic "open Internet" could severely devalue the open Internet the Commission is seeking to preserve.

---

[1] Bret Swanson is president of the technology research firm Entropy Economics LLC. http://entropyeconomics.com

[2] Lawson, Stephen. "U.S. Ranks 26th In New Broadband Index." Computerworld. May 25, 2010. http://www.computerworld.com/s/article/9177351/US_ranks_26th_in_new_broadband_index

[3] Ford, George S. "Fabricating a Broadband Crisis? More Evidence on the Misleading Inferences from OECD Rankings." Phoenix Center for Advanced Legal and Economic Public Policy Studies. July 7, 2010. http://www.phoenix-center.org/perspectives/Perspective10-05Final.pdf

[4] Steiner, Christopher. "Wall Street's Speed War." Forbes. September 9, 2010. http://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-street-speed-war.html

[5] Ibid.

[6] For a highly instructive guide to a range of technical (and, by inference, policy) debates over the structure of the Internet, read the excellent series of papers by Alderson, Doyle, and Willinger, et al, which outline the concept of "HOT" networks – Highly optimized / Organized tolerance / Trade-offs. See, for example, Willinger, Walter, David Alderson, and John C. Doyle. "Mathematics and the Internet: A Source of Enormous Confusion and Great Potential." Notices of the AMS. Volume 56, Number 5. 2009. http://www.ams.org/notices/200905/rtx090500586p.pdf; and Doyle, John C., et al. "The 'Robust Yet Fragile' Nature of the Internet." Proceedings of the National Academy of Sciences. October 11, 2005. http://www.pnas.org/content/102/41/14497.short

[7] Reply Comments of Bret T. Swanson. April 26, 2010. http://entropyeconomics.com/wp-content/uploads/2010/05/bret-swanson-nn-reply-comments-042610e.pdf

[8] Dano, Mike. "The Android IM App that Brought T-Mobile's Network to Its Knees." FierceWireless. October 14, 2010. http://www.fiercewireless.com/story/android-im-app-brought-t-mobiles-network-its-knees/2010-10-14

[9] Presentation of Bill Smith, President, Local Network Operations, AT&T. FCC Workshop. December 8, 2009. See, p. 9. http://www.openinternet.gov/workshops/docs/ws_tech_advisory_process/Bill%20Smith%20FCC%20Panel%20Discussion%20SLIDES%20120709.pdf