American Enterprise Institute for Public Policy Research

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

In the Matter of:	
Protecting and Promoting the Open Internet	GN Docket No. 14-28

Comments of Richard Bennett[1]

July 13, 2014

Ten years after Chairman Powell spoke about Internet Freedom at the Silicon Flatirons Center in Colorado, the FCC continues to struggle with casting his aspirations into legally binding regulations. The agency offers a simple set of proposed rules expanding the transparency principle, recasting the no-blocking rule on firmer legal ground, and similarly recasting the anti-discrimination rule within the limits of the commercial reasonableness standard affirmed by the courts in the Commission's Data Roaming Order.

If we accept the premise that the Internet is in imminent danger – as we've been told for ten years – the new rules proposed by the Commission are rational and sensible. They are certainly brief, comprising only two pages of text. Yet this compact set of rules inspires to Commission to write nearly 200 pages of explanations, questions, and invitations to comment. It appears that the Commission has been engaged in this process for so long that it is in danger of losing sight of the goal. It's certainly the case that the Internet is doing well despite the forecasts of doom; the non-factual claims that U. S. broadband is falling behind (it isn't); and the blank record of lawful net neutrality regulations in the United States.

[1] I am an independent network engineering consultant and policy analyst, presently working at the American Enterprise Institute as a Visiting Scholar. These remarks are offered in my personal capacity and do not necessarily represent the opinions of AEI or any client or sponsor. I have previously offered comments in the "Preserving the Open Internet" and "Broadband Industry Practices" dockets, GN 09-191 and WC 07-52 respectively, and offered testimony at the FCC En Banc Public Hearing on Broadband Network Management Practices in Cambridge on February 25, 2008 as an invited technical expert. My CV is available at http://www.bennett.com/resume.pdf.

At the urging of a handful of influential partisans and non-expert public figures, the Commission seriously considers applying common carrier telephone network regulations to a dynamic industry that develops and builds on technologies that are as far removed from telephony as any communication technology can possibly be. Moreover, the history of open Internet regulations shows they impair actual telephone services over the Internet such as Vonage and Skype. Like Sisyphus, the Commission pushes on.

In the name of preserving openness, the FCC has given us network fragmentation and created barriers to convergence. This is neither consistent with common engineering sense nor with the Section 706 mandate, as it slows the deployment of advanced networks.

I write these comments to urge the Commission to reframe the issue in terms that are more consistent with its legal mandate and more likely to stimulate further improvement in both broadband networks and the applications and services that depend upon them, especially real-time applications disadvantaged by Content Delivery Networks.

1. Frame the Issue in the Correct Context

There is no doubt that "Internet freedom", "Internet openness", and even "net neutrality" are important aspirations, but in the overall scheme of things they are secondary to progress writ large in the development of the essential technologies that make the Internet possible: dynamic, high capacity broadband networks, both wired and wireless, stationary and mobile; fast, resilient routing; high-speed packet switching; and a diverse and expanding menu of appealing applications.

Whether an IP network is open or closed, public or private, its power and utility depends on the ability of these fundamental technologies to power the applications that enable users to extract value from the network by doing things they want to do. End users also depend, quite keenly, on the ongoing deployment, improvement, expansion, and redeployment of infrastructure elements that increase speed, capacity, and resiliency and reduce packet loss and delay. It's not enough for equipment producers to build upgraded

switches; they need to be deployed in the field and continually upgraded, and people have to use them to realize their benefits.

Regardless of the outcome of this docket, the FCC has a binding statutory obligation under Section 706 to do all that it can to ensure that the cycle of network advancement proceeds in a reasonable and timely manner. If the agency's (self-assigned) Open Internet program conflicts with its obligation to push for network advancement, network advancement has to win; this is the law.

"Openness" is not so much a feature of networking technology as it is of the markets in which networking services are bought and sold. It is therefore an error to impose conditions on the technology itself that are more appropriate to marketing practices. Using technology restrictions to bring about marketplace goals is one of the most pernicious errors in the network neutrality debate. To paraphrase the poet Rilke, banishing the devils will make the angels take flight as well.

While it may be appealing to naïve regulators to ban broad classes of technology on flimsy pretexts (such as their alleged incompatibility with vague "end-to-end principles" or grossly simplified historical Internet operations norms), there are few technologies that lack legitimate application.

Consequently, the FCC's primary statutory mission with respect to broadband is to promote its continual improvement. Promoting an Open Internet can, if done correctly, advance this primary goal. If done incorrectly, it will inhibit broadband progress. As the Commission struggles with its questions about Open Internet means and definitions, it must bear in mind the fact that the ultimate goal is broadband advancement across the entire range of applications, not simply for video streaming despite its "flavor of the week" status.

2. Avoid Technology Demonization

It's fashionable in net neutrality circles to castigate something called "paid prioritization" by its critics as contrary to various historical norms of Internet practice, yet there is no

doubt in the communications and network engineering community that prioritization systems and their close relatives, bandwidth reservation systems, are often extremely helpful.

The quality of traditional telephone calls, for example, is assured by a bandwidth reservation protocol, and the application of the Wi-Fi Multimedia system based on the IEEE 802.11e standard's prioritization system allows Wi-Fi networks to carry four times as many Voice over Wi-Fi calls as it can with nothing more than the default best-efforts service.

Prioritization is essential to voice applications on mobile wireless networks as well. The NPRM mentions Voice over LTE (VoLTE) in footnotes 75 and 115, observing in the latter instance that VoLTE relies on prioritization. When voice calls are sold over LTE networks, people are paying for prioritization. When VoLTE provides High Definition Voice, they may even do so happily because they're enjoying a better communication service than they're had before.

It should be clear that a ban on prioritization and other means of achieving low-latency Quality of Service is effectively a ban on VoIP over the open public Internet. We already see this in the way that VoIP has become segregated onto parallel, separate networks that do not intersect with the Internet for the most part. If the Commission favors video streaming, the effect of the Title II net neutrality regulations it is urged to adopt by misguided advocates, it disfavors VoIP.

The call to demonize prioritization is fueled by misunderstanding. For example, a consortium of libraries and universities organized by the American Library Association urges the FCC to prohibit paid prioritization because it believes:

> *Prioritizing certain Internet traffic inherently disadvantages other content, applications, and service providers—including those from higher education and libraries that serve vital public interests.*[2]

While perverse inter-application effects often take place in networks, the legitimate goal of prioritization and other management systems is to minimize them, not to exaggerate them. It is not the case, as a technical matter, that boosting the priority of certain phone call packets over certain web page packets "disadvantages" the overall load time of web pages in the general case. The reasons why this is so should be obvious to the FCC's engineers.

On shared network facilities, such as radio frequencies or wired broadband connections between a home, campus, or office and an Internet Service Provider, Transit Provider, or backbone provider, all applications are modestly affected by each other according to the volume of load they place on the shared facility. This is the case with or without prioritization.

Thus, the load time of a web page is influenced by a concurrent VoIP call to a degree determined by the volume of traffic the VoIP call generates while the page is loading. This influence is very slight: the average web page is 1.8 megabytes in size, and it loads in approximately one second.[3] The average VoIP call consumes no more than 8 kilobytes per second, regardless of how long it lasts. Consequently, the VoIP call will degrade the web page's loading time by no more than one half of one percent in any case.[4]

---

[2] Jazzy Wright, "Higher Education, Library Groups Release Net Neutrality Principles," Press Release, *American Library Association*, (July 10, 2014), http://www.ala.org/news/press-releases/2014/07/higher-education-library-groups-release-net-neutrality-principles.

[3] HTTP Archive, "Average Bytes per Page by Content Type," web statistics, *HTTP Archive*, (May 29, 2014), http://httparchive.org/index.php.

[4] In fact, the data on web page load times in the FCC's report, *Measuring Broadband America – 2014* indicates that two thirds of web page load time is caused by non-network factors such as server capacity; all web pages loaded much slower than the effective speed provided by the network connection. See: FCC Office of Engineering and Technology and Consumer and Governmental Affairs Bureau, *Measuring Broadband America - 2014*, Measuring Broadband America (Washington, DC: Federal Communications Commission, 2014), http://www.fcc.gov/reports/measuring-broadband-america-2014.

If some of the VoIP call's packets are moved ahead of some of the web page's packets in the order of transmission, the degradation of the web page's overall load time will not be affected: it remains at less than one half of one percent, but the quality of the VoIP call will be improved.

I find the consortium's remarks especially puzzling because I have personally developed prioritization mechanisms for the enterprise Wi-Fi systems employed by some universities that are members of the consortium. Therefore, I am aware that the call for a ban on paid prioritization by this group is a double standard.

It is not reasonable for entities that practice prioritization within their own networks to insist that others should not be able to practice it within their networks as well. It is certainly reasonable to ask for limits on the practice to ensure it's not used in perverse ways, but it's unreasonable to ask for an absolute ban.

Similarly, a number of blog posts, satirical comedy skits, and news articles denouncing so-called "Internet fast lanes" have been presented to the public over various "fast lane" services such as Content Delivery Networks, bandwidth reservation systems, and IP prioritization.

Articles in the *New York Times*, *Washington Post*, and *Los Angeles Times* are delivered by CDNs that promise to boost the speed of their users' content. The *Washington Post* uses a service provided by Instart Logic touted by the provider as outperforming the traditional CDN: "Delivers radically faster web application performance, especially on mobile devices and across wireless access networks (3G/4G/WiFi)."[5] This system supplements common CDN services with "software-defined application delivery".

If Instart Logic pushes more bytes through any shared connection than a conventional web site in a given unit of time, it has to degrade other applications if we're to believe the

[5] Instart Logic, Inc., "What Is Software-Defined Application Delivery?," *CDN Alternative*, accessed July 11, 2014, http://instartlogic.com/technology/software-defined-application-delivery/.

rhetoric of the *Post's* reporters, who insist that the FCC's proposed rules simply "allow Internet service providers to charge content companies such as Netflix and Google for faster access into U.S. homes".[6]

Network neutrality godfather Tim Wu's claim that "with broadband, there is no such thing as accelerating some traffic without degrading other traffic" in his *New Yorker* column, was delivered by Akamai's CDN without the slightest trace of irony.[7]

The *Boing Boing* blog displayed a comic distortion of the FCC's proposed rules, claiming they're a plan to "sabotage existing infrastructure".[8] Amazon's CloudFront CDN delivers this falsehood for the *Boing Boing* blog.

Similarly, the John Oliver program that declared Chairman Wheeler a dingo was delivered to American homes served by cable companies over MPEG Transport, a protocol that runs over reserved bandwidth, and to homes served by telcos over IPTV utilizing IETF DiffServ, a simple prioritizing protocol.[9]

Fast lanes are more common and more beneficial than the Commission's critics realize. Given the less-than-factual nature of the criticism these fast lanes deliver to readers, they seem to be serving First Amendment rights quite effectively.

If in fact the FCC were to enact a wholesale ban on "fast lanes" and "paid prioritization", the proponents of such a rule would be among the first parties affected; if history is a guide, at that point they might well accuse the FCC of attempting to silence them.

---

[6] Cecilia Kang, "FCC Chair Tries to Salvage Net Neutrality Plan, Promises to Be Strong Cop in Revised Rules," *Washington Post*, June 11, 2014, http://www.washingtonpost.com/blogs/the-switch/wp/2014/05/11/fcc-chair-tries-to-salvage-net-neutrality-plan-promises-to-be-strong-cop-in-revised-rules/.

[7] Tim Wu, "Goodbye, Net Neutrality; Hello, Net Discrimination," *New Yorker*, April 24, 2014, http://www.newyorker.com/online/blogs/elements/2014/04/the-end-of-net-neutrality.html.

[8] Michael Goodwin, "Net Neutrality: What It Is, and Why You Should Care," blog, *Boing Boing*, (July 7, 2014), http://boingboing.net/2014/07/07/net-neutrality-what-it-is-an.html.

[9] S. Blake et al., "RFC 2475 - An Architecture for Differentiated Services," Internet RFC, (December 1998), http://tools.ietf.org/rfc/rfc2475.txt.

Leaving the hilarity and hypocrisy aside, however, there is certainly room for oversight around the edges of such services. There are practical limits for the allocation of best-efforts, better-than-best-efforts, and cheaper-than-best-efforts services in any network, depending on its technology and it's not unreasonable to establish guidelines to prevent the abuse of such technologies if a market were to be created for them.

These limits are best discovered through case-by-case review and by economic and engineering analysis. It may very well be the case that a simple rule limiting fast lane services to some fraction of overall capacity is enough; if it's not, it might be supplemented by a secondary rule regarding the performance of best-efforts service over a reasonable sampling period. This topic certainly warrants dispassionate analysis, but that's unlikely to take place in the current climate of hysteria.

3. Utilize the Power of Two-Sided Markets

Critics of fast lanes claim that Content Delivery Networks are different from hypothetical prioritization services potentially provided by ISPs consistent with the DC Circuit's suggestion. Indeed, there are subtle differences: video streaming over Content Delivery Networks such as YouTube and Netflix deliver packets in clumps, but ISP prioritization probably would not.[10] Clumps of traffic degrade other applications more than systems that limit packet jitter do.

CDN clumps are often rate-limited because video streaming does not need to download streams as quickly as possible (the way file transfer does), it simply needs to deliver them slightly faster than they can be rendered and viewed. Thus, even when Netflix runs over a gigabit network such as Google's Kansas City experiment, it consumes less than 4 Mbps.[11] Consequently, video streaming over a CDN leaves last mile bandwidth available for other applications between clumps. However, this unused bandwidth is most readily used by other CDNs.

---

[10] Shane Alcock and Richard Nelson, "Application Flow Control in YouTube Video Streams," *ACM SIGCOMM Computer Communication Review* 41, no. 2 (April 15, 2011): 24, doi:10.1145/1971162.1971166.

[11] Netflix, "USA ISP Speed Index Results," Netflix ISP Speed Index, *Netflix*, accessed May 30, 2014, http://ispspeedindex.netflix.com/usa.

The system effects of heavy CDN use are less clear for file transfer applications, and clumping has adverse effects on applications that require access to network links mid-clump, such as VoIP. In fact, clumping strategies are chosen by CDNs for their efficiency at transporting data from file system to network, the path of highest constraint in these systems.

Bear in mind that *Measuring Broadband America 2014* finds that web page load times are determined more by non-network factors than by network capacity: on 50 Mbps networks, only 30% of the load time of the typical web page – one second – can be explained by network capacity; the remaining 70% comes down to server capacity.

But the CDN only has control of a portion of network capacity, while the ISP controls all of it and thus has a greater ability to abuse its position. This is certainly true, provided the ISP is free to manage CDN traffic as it sees fit. But it is also true that only the ISP has the power to make real-time applications work better by overcoming their degradation by CDN traffic. This would be a good thing.

From the standpoint of application performance, the entire Internet – last mile, middle mile, and core – is a two-sided technical system in which packet throughput, delay and loss are determined by the ratio of offered load to network capacity in relevant units of time on each and every link.

Because the Internet is a statistical system, the sale of transport services with defined throughput, loss, and delay characteristics to a user or service provider implies the provisioning of dedicated capacity in all the links that comprise the path from producer to consumer or from peer to peer. When the purchaser does not use this capacity, it becomes part of the general pool available to all network users who happen to be situated on all or part of the path.

Consequently, there is a happy symmetry between the Internet's two-sided technical structure and the notion of two-sided markets: each path through an ISP network carries traffic for both an end user and a service provider, and each should be entitled to contribute toward its cost, for its own benefit. At the same time, the ISP has control over the path from public exchange point to consumer; that's their job, after all.

The ISP does not have monopoly control, as some assert, because consumers often use broadband services provided by more than one ISP: she probably has a wired broadband connection at home, another at work, one or more for mobile devices, and various Wi-Fi enabled connections at coffee shops, restaurants, and similar places. While none of these providers has an actual termination monopoly as this term was understood during the heyday of the telephone monopoly, it is nevertheless significant.

The Commission proposes that the sale of enhanced transmission should be conditioned on a verifiable generic level of service for data flows that don't come from enhanced service purchasers. Generic service is a capital notion, but it must be precisely defined to take account of the Internet's statistical nature and the fact that the service levels are only significant in the context of application requirements.

Consequently, the Commission could develop guidelines for web service that would ensure that typical pages load at a rate consistent with the capacity of the typical web site to provide data to an end user, but this guideline wouldn't necessarily ensure that Netflix works correctly. Similarly, a guideline for Netflix wouldn't ensure that Skype video calls work correctly, and none of these guidelines would ensure that Dropcam uploads would work as desired. Moreover, one can easily imagine future applications that have entirely different delivery requirements than currently popular ones.

When faced with the complexity of defining acceptable minimum service levels, net neutrality advocates tend to fall back on history and demand that all flows proceed on an unexamined "best efforts" basis for no particular reason other than ease of regulation. Some have attempted to prove that "best efforts" provides the greatest possible ongoing

stimulus to innovation, but such attempts are shallow, inconsistent, and unconvincing because they compare a fictitious Internet to a non-existent alternative.[12]

The most sensible way to progress on this question is simply to require that once web pages are delivered to the ISP at the nearest Internet Exchange to the consumer they shall be delivered to the consumer at a level of service consistent with required ISP disclosures of service characteristics specified in terms of data volume, throughput, and delay according to a relevant sampling period.

In cases where the web site does not deliver at the location most convenient to the ISP but to a location most convenient to itself, different standards apply. And in cases where the web site delivers its data more slowly than published performance standards, the ISP cannot be held responsible. The standard SamKnows test measures the path from consumer to IXP-located service.

4. Curb Vertical Integration by Expanding Effective Competition

One of the motives for net neutrality regulations is the fear that vertical integration is a barrier to effective competition for the services offered over broadband networks. Indeed, the triple- and quad-play services provided by broadband services have significant technical advantages over similar Internet-based services: linear TV programming is more bandwidth-efficient than video streaming services, and it provides a faster channel-surfing experience. Similarly, telephone service provided over broadband networks by MVPDs provides higher call quality than Internet-based services such as Vonage. Comcast is now the nation's largest telephone company in part because its service is higher quality than the Vonage service.[13]

The ability of broadband service providers to provide higher quality services than best-efforts alternative was recognized by Tim Wu in his seminal paper on net neutrality, but

[12] Barbara Van Schewick, *Internet Architecture and Innovation* (Cambridge MA: The MIT Press, 2010).
[13] This isn't definitive, but I use both, this is my experience, and I've developed voice services.

Wu let his fears get the better of his reason.[14] If Quality of Service enables better quality for real-time and specialized services such as linear TV, we increase effective competition by offering it for sale, not by creating a de facto monopoly for QoS through a "best-efforts" mandate that effectively limits Internet innovation to web sites and services such as Netflix streaming video that use protocols designed for web sites. The net neutrality demand for a single service level doesn't even accomplish the goals that its proponents claim they want.

The proposed "Commercial Reasonableness" standard achieves the goals of a competitively neutral and dynamic Internet better than the rules demanded by high volume web sites: best efforts with no data caps at the sole expense of the consumer.

Different applications require different treatment. The FCC's exception for "reasonable network management" can't be completely effective without a financial component, however. If the ISP is required to carry voice traffic more expeditiously than other traffic without charge, service providers and users will game the system by asserting that all traffic is voice.

5. Remove Barriers to Convergence

Network architecture in the past was very simple: We started with an application, such as telegraphy, standard telephony, television, computer-to-computer data exchange, or mobile telephony, and then we designed a network tuned for that one application. High-capacity packet switching – the "advanced network" described in Section 706 of the Telecom Act – enables one network to be selectively tuned for multiple applications at the same time.

David Isenberg's paper "The Rise of the Stupid Network" described selective tuning in terms of *idiot savant behaviors*: "because the data is the boss, it can tell the network, in real time, what kind of service it needs".[15] This is a practical, achievable, and worthy

[14] Tim Wu, "Network Neutrality, Broadband Discrimination," *SSRN Electronic Journal*, 2003, doi:10.2139/ssrn.388863.
[15] David Isenberg, "The Rise of the Stupid Network," *Computer Telephony*, August 1997, 16–24.

goal. A single converged network benefits from investment more efficiently than a fragmented set of different ones does.

But convergence remains an elusive goal. Many carriers transport telephone calls over IP networks today, but these IP networks are distinct from the open public Internet. This is not a matter of choice, but one of necessity: the interconnection norms widely used by the open public Internet do not yet preserve Quality of Service, and the FCC dictated that they should not in its Comcast order and Open Internet order.

In the name of preserving openness, the FCC has given us network fragmentation and created barriers to convergence. This is neither consistent with common engineering sense nor with the Section 706 mandate, as it slows the deployment of advanced networks.

A commercial reasonableness standard that permits individualized negotiation for interconnection and for end-to-end service corrects this error, if applied correctly. It also permits more effective competition in the real-time services space that is clearly distinct from the web space in terms of technical requirements.

Most advocates of strict net neutrality regulations reject the notion of individualized services and "idiot savant behaviors". Isenberg himself recanted the notion in a revision to his seminal paper titled "The Dawn of the Stupid Network":[16]

> *But suppose technology improves so much that the worst QOS is perfectly fine for all kinds of traffic, without a repertoire of different data handling techniques. Suppose, for example, that everyday normal latency becomes low enough to support voice telephony, while at the same time allowing enough capacity for video, plus data integrity strong enough for financial transactions. This would be a true Stupid Network – one treatment for all kinds of traffic.*

Isenberg asserted, in 1998, that improvements in routing and switching made the "true" stupid network imminent:[17]

[16] David Isenberg, "The Dawn of the Stupid Network," *ACM Networker* 2, no. 1 (March 1998): 24–31.

> *Routing switches from Madge and Foundry recently showed performance impressive enough to conclude that routing latency and jitter (variation in packet arrival time) may soon be a negligible issue. But these were lab tests, not field usage, and packet losses were as high as 1% under some conditions. So we are not there yet - but perhaps we will be soon."*

It's now 2014, and we are still "not there yet." What has happened is this: networks have doubled in capacity every three years, but best efforts regulation and management have limited the benefits of bandwidth abundance to high-volume, low-value data streams such as long-tail TV reruns, mediocre movies, and advertising networks. The delivery methods embedded in video streaming services by end-to-end control have impaired the Internet (even for the video streaming services themselves) under broad conditions, as Alcock and Nelson prove:[18]

> *Firstly, YouTube implements a previously undocumented form of flow control at the application level, which we call block sending, that operates in addition to traditional TCP flow control mechanisms...*
>
> *The second conclusion is that block sending can have a detrimental effect on YouTube flow performance, particularly if the client is streaming the video over a congested link. Blocks are typically transmitted as a large burst of packets, creating additional congestion and often leading to packet loss and significantly reduced throughput. Our analysis showed that over 40% of the packet loss events observed by YouTube clients using residential DSL could be attributed to congestion caused by block sending. These loss events resulted in a data retransmission rate of 1.5% of all bytes sent once block sending began. Given the popularity of YouTube, this is a significant quantity of data.*

End-to-end traffic shaping isn't enough, and it's often done poorly; but this is what the Open Internet order required. Effectively, network engineering has done its job, but regulation has transformed network advances into subsidies for *businesses that don't even know how to use them effectively.*

---

[17] Ibid.
[18] Alcock and Nelson, "Application Flow Control in YouTube Video Streams."

Throwing bandwidth at the problems caused by packet loss and delay under clumping conditions has not made Vonage and Skype video calling work at the level of reliability and quality that consumers want. The FCC's *Open Internet Report and Order* did not create effective competition in the calling space, and it did not increase the rate of the TDM-to-IP transition; it simply created a roadblock to real-time applications.

It's time to rethink the means and to clarify the goal that this exercise is supposed to achieve.

6. Conclusions

When Chairman Powell spoke on Internet Freedom in Boulder in 2004, Americans who used the Internet were as likely to use dial-up as broadband; broadband speeds were less than 3 Mbps; and the predominant Internet application by volume was web surfing.

Today, more than 70% of American homes have broadband connections, at an average peak capacity of 40 Mbps.[19] Web surfing now accounts for less than 10% of Internet traffic, and the predominant application by volume is video streaming from Netflix, YouTube, and similar services.[20] Voice is still an uncertain application, despite massive improvements in broadband capacity, sometimes it works and sometimes it doesn't.

Regulators have never correctly understood the reasons for the stagnation of VoIP; the FCC itself ignored the impact that P2P usage has on Vonage when it issued its Comcast order in 2008.[21] The transition from TDM telephony to IP is taking place largely over facilitates separate from the open public Internet because the Internet's added bandwidth

[19] Pew Research Center's Internet & American Life Project, "Broadband Technology Fact Sheet," accessed July 14, 2014, http://www.pewinternet.org/fact-sheets/broadband-technology-fact-sheet/; Akamai, "State of the Internet," archive, *State of the Internet*, accessed February 1, 2013, http://www.akamai.com/stateoftheinternet/?WT.ac=soti_banner.
[20] Sandvine, "Global Internet Phenomena Report," 1H 2014, https://www.sandvine.com/downloads/general/global-internet-phenomena/2014/1h-2014-global-internet-phenomena-report.pdf.
[21] Federal Communications Commission, "Memorandum Report and Order in the Matter of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications et Al.," August 1, 2008, https://apps.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf.

is consumed by video streaming services. These services don't use bandwidth efficiently, nor do they have incentive to do so as long as ISPs bear the full financial responsibility" and neutral network rules.

One ultra-high volume service – Netflix – demands totally free bandwidth from ISPs, which would only make matters worse for voice users and for entrepreneurs in the real-time and/or Internet of Things space.[22] The Commission is urged to reclassify broadband as a Title II service, without the least trace of irony, by activists, advocates, commercial interests, and average citizens who lack appreciation for the consequences.

Shall the FCC impose rules on broadband services that come from telephone network regulation in order to continue a status quo that disadvantages telephony-like service on the open public Internet? This question practically invites a comedic response.

The FCC is supposed to be an expert agency that enacts the will of Congress as expressed in the Communications Act and related measures. It is not an alternative Congress meant to substitute its judgment in policy matters for that of elected representatives.

Congress has directed the FCC to "encourage the deployment on a reasonable and timely basis of advanced telecommunications capability to all Americans" in Section 706 of the Telecom Act. Systems of broadband use that permit the commercially reasonable interconnections that permit real-time applications to function well fulfill this mandate; a net neutrality regime that sacrifices real-time services on the altar of streaming video does not. Continuing the same actions and hoping for a different result will not get us where we need to be.

The FCC must explicitly permit more expansive bargaining for transmission services between application entities and networks. This is the only path that can successfully

[22] Reed Hastings, "Internet Tolls And The Case For Strong Net Neutrality," corporate blog, *Netflix US & Canada Blog*, (March 20, 2014), http://blog.netflix.com/2014/03/internet-tolls-and-case-for-strong-net.html.

unlock the capabilities latent in widely deployed, widely used, and constantly improving broadband networks.