



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

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

**COMMENTS OF
INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION**

Robert D. Atkinson
Richard Bennett

Information Technology and Innovation Foundation¹
1101 K St N.W.
Suite 610
Washington, DC 20005

¹ ITIF is a nonprofit, non-partisan public policy think tank committed to articulating and advancing a pro-productivity, pro-innovation and pro-technology public policy agenda internationally, in Washington and in the states. Through its research, policy proposals, and commentary, ITIF is working to advance and support public policies that boost innovation, e-transformation and productivity.

The Information Technology and Innovation Foundation (ITIF) is pleased to offer the following comments on the FCC’s Notice of Proposed Rulemaking, “In the Matter of Preserving the Open Internet”, GN Docket No. 09-191, and “Broadband Industry Practices”, WC Docket No. 07-52. ITIF has long advocated protecting and preserving the aspects of the Internet that enable it to serve as an engine of innovation and vehicle for advancing the public interest. We have previously filed comments in the Commission’s related proceedings on the Internet and its constituent networks, participated in the Commission’s workshops on the National Broadband Plan, and written reports on the Internet architecture² and network management practices,³ as well as on challenges and opportunities facing the Internet and its constituent networks. In the interest of furthering the dialog, we offer the following comments, observations, and recommendations on the questions asked by the NPRM.

Table of Contents

OVERVIEW.....3

QUESTIONS.....4

PARAGRAPH 19, INTERNET ARCHITECTURE.....4

 Virtual Network.....4

 End-to-End Design Principle.....6

 Network Optimization.....8

PARAGRAPH 21, OPEN PROTOCOLS.....9

PARAGRAPHS 24-47, PSTN.....10

PARAGRAPH 50, JUSTIFICATION.....10

PARAGRAPH 56, BEST-EFFORTS.....12

PARAGRAPH 63, INNOVATION.....13

PARAGRAPHS 67-74, COMPETITION.....15

PARAGRAPHS 79 AND 80, CONGESTION.....19

PARAGRAPHS 81-82, VARIOUS ISSUES.....19

 Competition.....20

 Architecture.....21

 Digital Divide.....21

PARAGRAPH 92, INTERNET POLICY STATEMENT.....22

PARAGRAPH 100, WIRELESS.....22

PARAGRAPHS 103-117, DISCRIMINATION.....23

PARAGRAPHS 118-132, TRANSPARENCY.....25

PARAGRAPHS 135-141, REASONABLE NETWORK MANAGEMENT.....25

PARAGRAPH 173, MOBILE BROADBAND.....27

CONCLUSION.....27

² Richard Bennett, *Designed for Change: End-to-End Arguments, Internet Innovation, and the Net Neutrality Debate* (Washington, DC: Information Technology and Innovation Foundation, September 2009), <http://www.itif.org/index.php?id=294>.

³ George Ou, *Managing Broadband Networks: A Policymaker's Guide* (Washington, DC: Information Technology and Innovation Foundation, December 2008), http://www.itif.org/files/Network_Management.pdf.

Overview

The Internet and its constituent networks form vital parts of the information and communication infrastructure of modern international society, and as such are items of great concern to citizens and national regulators alike. It's not inaccurate to say that the entire history of network regulation has concerned physical networks built for specific purposes. As the Internet is a virtual network that depends entirely on physical networks for its operation, it represents a conceptual challenge to regulators who must strive to understand not only how it operates with today's applications over today's networks, but how it might operate with future applications and with networks not yet developed. Consequently, the result of recent reviews of the Internet regulatory framework conducted in Canada and the European Union have concluded that the proper emphasis, for the time being, is to stress consumer disclosure of Internet access services over sweeping restrictions on specific network management practices or business models. We recommend this approach to the FCC.

By all accounts, the U. S. Internet ecosystem is thriving: Internet Service Providers are upgrading the speed and reach of their first- and last-mile networks, the deployment of new undersea and transcontinental cables continues, investment in advanced network routing and switching equipment continues apace, and new devices, applications, and services emerge on a regular basis. 4G wireless networks are scheduled to come on-line in the United States this year, along with a new category of consumer device, the tablet computer, as well as a rich new application category, Mobile Augmented Reality, integrating location and information in the real and virtual worlds.

Alongside this explosion of investment, innovation, and consumer benefit, it's difficult to say that there's a need for a significant departure from the light regulatory stance the FCC has taken toward the Internet since reclassifying DSL a Title I service. In the entire history of broadband Internet services in the United States, we've only seen one case of clearly anti-competitive discrimination, the blocking of VoIP services by the Madison River Telephone Company, which was resolved by a Consent Decree. It's therefore likely that stringent new Internet regulations would be a cure more harmful than the disease.

At a bare minimum, Internet rule making should be predicated on clear and unambiguous evidence of harm to the growth of the Internet ecosystem and/or harm to consumers or competition. The Comcast case doesn't meet this threshold, for a number of reasons. In the first place, the ruling against the network management practices once used by Comcast wasn't built on a sound factual foundation: the Commission didn't conduct an independent fact-finding, but relied instead on anecdotal accounts of the system offered by interested parties. In the second place, the ruling relied on a Policy Statement that was explicitly identified as something less than a rule when adopted by the Commission, and in the third place, the ruling was inconsistent. Commissioners voting in favor of the Comcast decision declared their support for systems of management that allowed latency-sensitive VoIP applications to meet customer requirements while voting to sanction Comcast for implementing just such a system.

The Comcast case illustrated the gap between ISP management practices and the public's understanding of the Internet, which was certainly aggravated by Comcast's lack of full disclosure during the initial discussion period. We propose that the principal lesson to be gained from the Comcast network management case is the value of full, fair, and timely disclosure.

A secondary lesson concerns the public understanding of the Internet. Consequently, a portion of our remarks will address the Internet architecture and suggest a series of questions that are not asked by the NPRM but which must be answered to move the proceeding forward in a constructive direction. These remarks are addressed more fully in our report on the Internet architecture, [*Designed for Change: End-to-End Arguments, Internet Innovation, and the Net Neutrality Debate*](#), and in a forthcoming ITIF report on the relationship of Internet architecture to wireless networks.

We applaud the Commission's commitment to address these questions in a fair and open set of proceedings, but urge caution against departing significantly from the current *status quo* into the uncharted waters of a strict new system of Internet regulation that's likely to do more harm than good.

Questions

The NPRM raises a number of questions and makes a number of assertions about the Internet, a few of which are somewhat off the mark. As the conceptualization of the Internet forms an important part of this proceeding, we will attempt to identify and clarify the gaps we perceive in the proper understanding of the subject matter. These are identified in the particular paragraphs in the NPRM noted below.

Paragraph 19, Internet Architecture

Paragraph 19 states: Another more technical aspect of Internet openness has had the effect of empowering entrepreneurs to innovate without needing to seek permission. TCP/IP reflects a so-called "end-to-end" system design, in which the routers in the middle of the network are not optimized toward the handling of any particular application, while network endpoints (the user's computer or other communicating device) are expected to perform the functions necessary to support specific networked applications. The practical implication of this design philosophy has been that a software developer can develop new networked applications by writing programs only for end-user computers, without needing to modify the far more specialized programs running on network routing equipment. As the diagram below illustrates, this design differs from how the Public Switched Telephone Network has historically operated. By allowing innovation to be easily implemented at the edge of the network, the end-to-end design of the Internet has lowered technical, financial, and administrative barriers to entry for entrepreneurs with technical skill and bright ideas.

Virtual Network

This description of the Internet is not entirely correct. In the first place, it treats the Internet as a physical network, when it is actually a virtual network. A virtual network is

an abstraction that depends on physical networks to operate. From the point of view of any physical network (such as the PSTN, Ethernet, or Wi-Fi,) the Internet is nothing more than content formatted in a particular way. Therefore we can't directly compare the Internet with a physical network. The PSTN, for example, completes calls between one subscriber and another and ensures their reliability, while the Internet simply formats Internet Protocol (IP) datagrams and hands them off to a physical network for transmission.

Because it is simply a virtual network, the Internet doesn't need to concern itself with the means by which packets are transported from one location to another; transporting IP datagrams from one point to another is a physical network problem, not an Internetwork problem. It is therefore fallacious to assert that the Internet mandates a particular method of resolving contention or facilitating packet delivery: contention and delivery rules are properties of physical networks, and the Internet has no choice but to defer to them. The Internet is best understood as a property of certain types of data flowing over one or more physical networks; it adds constraints to these networks, it does not and cannot relax constraints.

Many Internet users attach through the PSTN, and are therefore just as constrained with respect to innovation and permission as those who rely solely on the PSTN for end-to-end communication. Overlaying the PSTN with IP datagrams does not alter its dynamics, in other words. The elements of the Internet that are unique concern the formatting and treatment of IP datagrams, not such broad terms as "permission." The elements of the Internet that are not unique are those that concern the logical and physical transmission and reception of the Link Layer frames for which IP datagrams are simply a form of content.

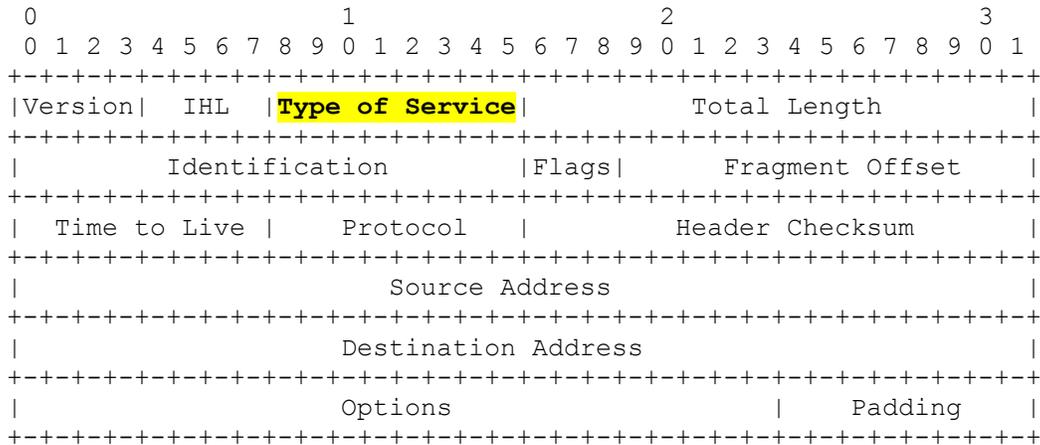
The Internet doesn't erase the economic and physical realities of network operation; it simply attempts to add a layer of uniformity over them. This uniformity – the ability to create applications that operate over a variety of physical networks with a high degree of predictability – is the Internet's innovative essence, but Internet uniformity is not absolute. Just as physical networks differ from one another with respect to cost, bandwidth, latency, jitter, and reliability when they support applications that are not Internet-enabled, they also differ from each other in terms of their effectiveness as vehicles for IP datagram transport. A 40 Gigabit/sec. fiber optic Ethernet network is a better vehicle for HD video conferencing than a 1 Megabit/sec. DSL network, whether IP-enabled or not; a GSM or CDMA network is a better network for roaming communication than a wireline network, whether IP-enabled or not.

Similarly, applications differ from each other with respect to a variety of requirements in terms of cost, bandwidth, and latency; many of these application requirements are determined by the laws of the physical world, the biological nature of the human ear, eye, and brain, and by the nature of digital content itself. Against these application constraints, the Internet is not so much a rule-maker as a follower; users and regulators are more properly concerned with adapting the Internet to application needs than with a fruitless quest to make networks and applications conform to Internet traditions, real or imaginary.

End-to-End Design Principle

The End-to-End system design principle is generally misstated in policy literature, and the NPRM does not depart from this tradition. The Internet does not differ from its constituent networks by virtue of the fact that *a software developer can develop new networked applications by writing programs only for end-user computers* as the NPRM asserts. The ability of programmers to implement applications by writing programs for end-user computers (and for the server computers not controlled by typical end-users) is obviously not unique to the Internet. In fact, it’s a side effect of the intelligent devices attached to Internet access networks; these devices themselves permit the running of programs, whether networked or not. The PSTN is actually every bit as “end-to-end” as this description entails, since it not only allows but *forces* application developers to stay out of its internal elements and confine their innovation to devices attached to the network. Something else entirely is going on with the Internet.

The PSTN prevents users from altering network internals, but the Internet’s architecture explicitly *permits and encourages* experimentation with protocols and other elements of its internal structure. The IP datagram is a minimalist network design element that permits the development of rich network services such as delivery options. The original specification for Internet Protocol assigned bits eight through fifteen to identify “Type of Service,” as we can see in the following description of the Internet Protocol header taken directly from the initial IP specification, RFC 760:⁴



Example Internet Datagram Header

IP Type of Service is not self-executing, since the Internet is a virtual network; but IP can request any physical network capable of honoring some or all of its delivery options to

⁴ J. Postel, Ed., “DoD Standard Internet Protocol,” Internet RFC, January 1980, <http://tools.ietf.org/rfc/rfc760.txt>.

differentiate packet delivery according to expressed user and application desires. This is not a new capability enabled by Deep Packet Inspection; it's a feature of IP as old as IP itself.

Since the original specification of IP 35 years ago, the engineering community has produced a long series of additions and enhancements enabling applications to exercise control over network functions, such as:

- Multiprotocol Label Switching (MPLS)⁵
- Differentiated Services (DiffServ)⁶
- Integrated Services (IntServ)⁷
- Real Time Protocol (RTP)⁸
- Real-Time Streaming Protocol (RTSP)⁹
- Resource Reservation Protocol (RSVP)¹⁰
- Explicit Congestion Notification (ECN)¹¹

Most of these enhancements have modified the code resident in network routers as well as network endpoints, and in so doing have enabled new categories of applications. The feature that sets the Internet apart from conventional networks is actually the relentless enrichment of its basic capabilities in order to enable new classes of applications to flourish at the network edge. The Internet has no barrier forbidding the application programmer working within the edge system from collaborating with network operators and equipment vendors to improve network services. Among the Internet's prime virtues is the fact that "Anyone can improve it."¹²

The Internet is about dynamism, not dogma, so any regulatory formula that aims to preserve its free and open character has to recognize that its dynamic character must continue to exist in order for the Internet to serve the needs of people and devices at the outer edge. The Internet has been able to support new types of applications because it has a structure that is friendly to change.

⁵ E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," Internet RFC, January 2001, <http://tools.ietf.org/rfc/rfc3031.txt>.

⁶ S. Blake et al., "An Architecture for Differentiated Services," Internet RFC, December 1998, <http://tools.ietf.org/rfc/rfc2475.txt>.

⁷ R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," Internet RFC, June 1994, <http://tools.ietf.org/rfc/rfc1633.txt>.

⁸ H. Schulzrinne et al., "RTP: A Transport Protocol for Real-Time Applications," Internet RFC, January 1996, <http://www.ietf.org/rfc/rfc1889.txt>.

⁹ H. Schulzrinne, A. Rao, and R. Lanphier, "Real Time Streaming Protocol (RTSP)," Internet RFC, April 1998, <http://www.ietf.org/rfc/rfc2326.txt>.

¹⁰ R. Braden et al., "Resource ReSerVation Protocol (RSVP)," Internet RFC, September 1997, <http://tools.ietf.org/rfcs/rfc2205.txt>.

¹¹ K. Ramakrishnan, S. Floyd, and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," Internet RFC, September 2001, <http://tools.ietf.org/rfc/rfc3168.txt>.

¹² Doc Searls and David Weinberger, "World of Ends: What the Internet Is and How to Stop Mistaking It for Something Else.," *The World of Ends*, <http://www.worldofends.com/>.

Network Optimization

One of the most important benefits the Internet offers is the ability to run both old and new applications over both old and new networks. It is therefore incorrect to assert, as the NPRM does, “the routers in the middle of the network are not optimized toward the handling of any particular application.” A positive statement is more appropriate: “the routers in the middle of the virtual network are capable of providing near-optimum service to a wide variety of applications.” This comes from the ability of Internet protocols and network management systems to signal application requirements to physical networks. This point may appear pedantic, but it’s particularly important given the common desire to converge diverse physical networks within a common IP infrastructure.

The historically separate networks built for television, radio, telephony, and digital content are each capable of providing near-optimum service to a broad category of applications. The converged IP network – the Internet of the Future – must also meet diverse application needs by providing near-optimum transport and economics. It will do this by exploiting elements of Internet architecture that have been present from the beginning, even though they haven’t been universally implemented. This point requires clarification.

There is a large gap between Internet architecture and Internet practice. Advocates of the far-reaching Internet Service Provider (ISP) regulatory system known as net neutrality stress, correctly, that differentiated Type of Service has not been a major element, in practice, of inter-domain carriage agreements among network operators. For the most part, organizations employ Type of Service (and its successor, Differentiated Services) inside their private networks, but pass packets to other networks according to the default Type of Service class known as “Best Efforts.” There has been a historic difference between Internet architecture and Internet practice, with practice lagging behind theory.

Advocates of strong net neutrality regulations insist that there is no good reason to permit network interconnection agreements to honor Type of Service in the future because it hasn’t been necessary in the past, or at least no good reason to permit charging in inter-domain packet transfers on such a basis. In many instances, “preserving the Internet’s openness” means little more than “banning Type of Service differentiation for a fee.” Language to this effect is to be found in many of the Congressional net neutrality bills that have appeared since 2005¹³ and is found in Paragraph 106 of the NPRM.

One problem with such a “Type of Service Pricing Ban” is that it would shackle the Internet of the Future to a set of requirements established by the applications and routing technology of the past. For most of its history, the Internet has not been as much a

¹³ One typical example is: Sen. Olympia Snowe and Byron Dorgan, *Internet Freedom Preservation Act, To amend the Communications Act of 1934 to ensure net neutrality*, 2006, <http://thomas.loc.gov/cgi-bin/query/z?c109:S.2917>: (*[Each broadband service provider may] only prioritize content, applications, or services accessed by a user that is made available via the Internet within the network of such broadband service provider based on the type of content, applications, or services and the level of service purchased by the user, without charge for such prioritization.*)

“general-purpose network” as a “content-oriented network system” that only did one thing well. This bears explaining. While it appears to the technically naïve Internet user that familiar applications such as e-mail, Web surfing, traditional file transfer and Peer-to-Peer file transfer are very different, they are highly uniform at the level of IP usage. Under the hood, each of these applications entails one or more file transfer operations, and each uses TCP. TCP is a protocol that is entirely implemented in end-user computers and is therefore not a part of the Internet infrastructure in the same way that IP is; there is no TCP in an Internet router, for example.¹⁴ Each of these applications is termed “elastic” by network engineers, which simply means that it’s highly tolerant of delay, both persistent and variable. A web page that loads in two seconds is generally considered successful; for some other types of applications (such as real-time securities trading systems and certain types of device control, for example) an end-to-end delay of two milliseconds per packet is considered fatal.

Rather than providing general-purpose network services, the Internet has provided, in practice, a *single* service that happens to work well for a very broad *single* category of applications. While this single service is very broad, it’s not broad enough to encompass the entire range of human communication needs currently enabled by broadband networks or likely to emerge in the future. The Internet is a content network today, but we need it *become* a truly general-purpose network in the future, one that is as efficient when transferring files as it is when providing extra-low-delay packet transport to video callers, television channel surfers, telemedicine participants, gamers, and security service providers.

It’s reasonable to believe that network operators may need to employ a larger subset of the Internet architecture in the future than they have in the past; the freedom to implement the full Internet architecture as well as its future extensions is among the features of the Internet that need to be preserved and protected.

Paragraph 21, Open Protocols

Paragraph 21 states: *The Internet’s openness has also created transformational commercial opportunities for network operators. Because TCP/IP is just as accessible to telephone, cable, and wireless network operators as it is to innovators and entrepreneurs at the edge of the network, these infrastructure companies have been able to adopt Internet technologies and invest to develop the \$130 billion annual business that broadband Internet access has become today.*

TCP, of course, is a technology that exists solely in end systems, and is therefore not part of the systems deployed by network operators, who have nevertheless invested enormous sums around the world to provide a variety of IP services to 1.6 billion users at both the retail and wholesale level. IP services are provided over legacy wiring systems by network equipment sold by vendors such as Cisco, Huawei, Alcatel-Lucent, Hewlett Packard, and many others to network operators under the expectation that IP is a

¹⁴ Most routers actually do have a TCP capability, but it’s used only for the so-called “control plane,” not the “data plane” that forwards IP datagrams.

financially viable service. Physical network technologies such as LTE, Wi-Fi, and MPLS are now designed with the express purpose of supporting IP and the wide range of applications that IP can support in principle.

The value of these investments and by extension the incentive to invest more in networks, is amplified by proper regulation, and similarly is retarded by inappropriate regulation. For this reason, it's imperative that the Commission gets the regulatory balance right.

Paragraphs 24-47, PSTN

These portions of the NPRM deal with the Commission's historical rulings regarding the PSTN and the early development of computer networks. While these historical developments are academically interesting, they're not directly relevant to the argument for a regulatory framework aligned with the Internet and are likely to point in the wrong direction. The Internet Policy Statement, for example, is an attempt to impose *Carterfone* principles devised for a fundamentally closed network to an Internet that is fundamentally open. This approach creates an incomplete regulatory picture since it neglects the side effects and interactions that Internet users have on each other, a phenomenon that essentially does not exist on the PSTN because of its method of allocating bandwidth. The Internet Policy Statement does not advance the debate on Internet regulation beyond *Carterfone*, which is unfortunate given the unique nature of the Internet and its value for modern society.

Paragraph 50, Justification

Paragraph 50 justifies Commission action on Internet regulation in the following terms: *Despite our efforts to date, some conduct is occurring in the marketplace that warrants closer attention and could call for additional action by the Commission, including instances in which some Internet access service providers have been blocking or degrading Internet traffic, and doing so without disclosing those practices to users.*

While we're sympathetic to the rights and needs of both consumers and innovators, we're unaware of any *current* behavior in the Internet marketplace that would demand immediate Commission action. Moreover, the two instances of operator behavior that the NPRM cites – Madison River's VoIP ban and the Comcast system that briefly limited upstream bandwidth consumption by unattended P2P seeders to half of available bandwidth – have long since been corrected. The Commission's ruling against Comcast is currently under review by the courts, and the consensus in recent commentary indicates that it's unlikely to survive.¹⁵ Moreover, in our view the actions Comcast took were a legitimate attempt to address a real challenge in their network that absent action would reduce the quality of broadband service for the majority of their customers. It is

¹⁵ Paul J. Feldman, "Court Challenges FCC in Early Network Neutrality Test," *CommLawBlog*, January 8, 2010, <http://www.commlawblog.com/2010/01/articles/internet/court-challenges-fcc-in-early-network-neutrality-test/>. (The judges seemed to feed "softballs" to the attorney for Comcast, while giving the FCC lawyer a much harder time... The Chief Judge asking the FCC lawyer, "How would you prefer to lose – [on due process or on jurisdiction]?")

instructive in this regard that the Canadian Radio-television and Telecommunications Commission has allowed Canadian cable ISP providers to engage in essentially the same kind of network management under the supervision of the CRTC.¹⁶

The problem here is not so much with what *Carterfone* requires – the ability to attach compatible devices to a network – as with the things that it ignores. In the PSTN, the network operator controls the allocation of network bandwidth, giving a small and equal amount to each caller regardless of the condition of the call or the application in use. A telephone caller consumes as much PSTN bandwidth while silent as a file transfer program does while active. The broadband networks that enable high-speed IP services allocate bandwidth dynamically, under the ultimate control of the totality of end user applications active at any given time, millisecond by millisecond. Thus, the most significant of the operative elements of *Carterfone* – the ban on “degrading” user experience – is largely outside operator control on broadband, IP-oriented networks. Every packet that traverses a packet-switched network consumes 100% of the bandwidth of each link for the duration of the packet, which causes any other potential access to the link during that period to wait for access to the link. The greater the number of packets offered to the network, the greater the delay.

Some net neutrality advocates insist that network operators must take a purely passive stance with respect to Internet traffic, at the risk of altering fundamental Internet dynamics. In such a scenario, many users would experience “degradation” for the simple reason that all Internet users share facilities with others, some of whom are long-term high bandwidth-using consumers. The general rule in unmanaged networks and in edge-managed networks is that each user will obtain bandwidth in proportion to the quantity that the user’s application seeks to consume: the more you want, the more you get. The mismatch between applications that consume bandwidth aggressively and those that have stringent delay requirements spreads dissatisfaction as applications become more diverse.

The Internet currently hosts a more diverse mix of applications than it has in the past, and it’s clearly on a trajectory toward an even more diverse application mix in the future. A bright line rule against differentiation for a fee has the potential to reduce application diversity.

With regard to the final sentence of Paragraph 50 (*In undertaking this examination, we seek to preserve the open, safe, and secure Internet and to promote and protect the legitimate business needs of broadband Internet access service providers and broader public interests such as innovation, investment, research and development, competition, consumer protection, speech, and democratic engagement.*) we would simply point out that the Internet is not yet as “safe and secure” as it should be, and in the course of making it so it would be beneficial if there were not too many restrictions on departing from any particular notions to the effect that the Internet’s architecture is inviolable.

¹⁶ *Telecom Decision CRTC 2009-677* (Canadian Radio-television and Telecommunications Commission 2009) <http://www.crtc.gc.ca/eng/archive/2009/2009-677.htm>.

Paragraph 56, Best-Efforts

This paragraph unfortunately confuses the end-to-end architecture with a mode of packet transport: *The Internet has traditionally relied on an end-to-end, open architecture, in which network operators use their “best effort” to deliver packets to their intended destinations without quality-of-service guarantees.* As previously noted, the end-to-end architecture does not constrain transport services to a single class, and the Internet architecture itself does not perform packet delivery. The Internet is a virtual network that defers to physical networks to deliver packets.

The term “best-efforts” is a technical term of art that actually denotes something different from a single service level in any case. It describes the method of packet transmission used by the early Ethernet system (before the advent of the Ethernet hubs and switches used today) described by the IEEE 802.3 10BASE5 standard and its precursor the so-called Ethernet Blue Book. Early Ethernet employed a method called Carrier-Sense Multiple Access with Collision Detection (CSMA/CD) that was able to deliver packets with a high degree of reliability by sensing the most common instance of non-delivery and automatically correcting it.

The term “best-efforts” originally described a network that delivered packets with a very high degree of reliability. Because “best-efforts” was part of CSMA/CD Ethernet, the term was expanded into a blanket term for all aspects of CSMA/CD Ethernet, such as the supposition that all packets on the Ethernet had the same priority. While the official standard does mandate use of a single, common priority, many vendors of Ethernet equipment, from chip manufacturers to system builders, have allowed users to control packet priority. The means by which this is accomplished is simply to increase or reduce the carrier sense interval (the period of time during which the network must appear idle before any system can initiate transmission of a packet) above or below the 9.6 microseconds specified as the standard.

Variable carrier sensing is a part of the IEEE 802.11e standard, where it is used to obtain higher priority for voice and video packets than for generic packets, and lower priority for background packets. The generic service level is now called “best-efforts”. Consequently, “best-efforts” delivery now means one service level among many, which implies the ability to seek and to use alternate service levels.

It is also misleading to describe Internet history as mandating a single service level. As previously noted, Internet Protocol is typically used within management domains at multiple service levels. It is also the case that peering and transit agreements – the terms under which the networks that form the Internet agree to connect with each other – frequently include service level requirements. One of the reasons that Content Delivery Networks (CDN) such as Akamai are so popular is that they are capable of providing a Quality of Service (QoS) boost over generic transport, and they obtain considerable compensation for doing so. Another traditional means of obtaining better than “best-efforts” QoS is the overlay network, such as the WebEx conferencing system. Overlay networks interconnect with the public Internet at many points, as CDNs do, but bypass

the Internet core in order to provide low latency routing. Finally, some peering agreements require networks to honor each other's QoS guarantees, which they communicate to each other by separating QoS classes onto separate Ethernet Virtual Local Area Networks (VLAN) at their cross-connect points. ICANN has allocated identifiers to designate Border Gateway Protocol (BGP) Community Attributes for identifying QoS levels in publicly advertised Internet routes.

While it is true that the much of the routing of packets over the publicly visible Internet has relied on a single service class for much of the Internet's history, there have always been significant exceptions. As the final portion of the paragraph points out, these exceptions have generally taken place when applications required more strict QoS than generic service could provide, or when new applications needed to transport greater volumes of data than generic applications transported. At no time has the openness of the Internet with respect to applications been predicated on prohibition of openness in terms of service levels. The Internet has always achieved openness by allowing users and operators to negotiate service levels among themselves.

It is also the case that the preservation of openness with respect to sources of content and patterns of consumption is quite separate from the question of packet delivery requirements. It's one thing to block access to a web site or to charge every application to gain access to an ISP's customers (practices that, it is important to note, have not happened – with the exception of Madison River – and in our view will not happen in the future.) It's quite another thing to enable a real-time service by allowing network operators to meet diverse application needs and to allow applications and content providers the opportunity to purchase more advanced, higher quality service to meet their unique needs as many do now when purchasing CDN services.

Paragraph 63, Innovation

This paragraph misconstrues a comment we offered in written testimony in one of the Commission's National Broadband Plan workshops. It begins: *Supporters argue that differentiation by Internet access service providers can be especially harmful to innovation by outsiders—individuals and entities unaffiliated with network owners—who have been responsible for some of the most important innovations in the history of the Internet.*

The NPRM then refers to our comments in footnote 145, which says: *Robert T. Atkinson & Richard Bennett, Information Technology and Innovation Foundation, Prepared Remarks for the Big Ideas Workshop, National Broadband Plan, Federal Communications Commission, at 4 (Sept. 3, 2009); see also, e.g., id. (“With regard to the particular area of network architecture research, it doesn't take large teams with enormous budgets to make fundamental advances. Paul Baran worked with a very small team, as did Louis Pouzin, the inventor of the framework for end-to-end networks that informs the Internet of today (as well as the four other major packet networks created during the same period as the Internet).”).*

The NPRM then goes on to say: *These outsiders, many of whom may have limited resources but can innovate on today's Internet with very low marginal costs, could choose not to innovate if faced with fees from Internet access service providers for equal access to end users.*

The remarks that we offered to the Workshop (on the need to support small groups with research grants) were in no way meant to offer an argument against differentiated pricing models for Internet access service, or to imply that differentiated pricing would reduce innovation – in fact we believe that it could boost innovation, including among new entrants. However, the NPRM's misunderstanding of our comments does illustrate a common misconception. While a differentiated pricing model is likely to raise the price for expedited delivery *by the ISP*, it wouldn't necessarily raise the *end-to-end* price (taking alternatives such as CDNs and overlay networks into account) for Internet traffic generally; it's also likely to *lower* the price for many applications. The Internet isn't a single, uniform network in which all fees are collected by the same firm; it's a federation of networks, many of whom compete with each other for IP transport. There are multiple providers of expedited IP carriage today, and allowing the ISPs to enter this market to provide a service that application or content providers voluntarily choose to purchase to receive higher quality service increases consumer choice.

Single-service level pricing is a two-edged sword. The generic service level has to meet the needs of the most demanding commonly used applications for transit latency and jitter; these are currently video calling and video streaming. ISPs meet these needs by spending on network equipment, transit services, and bandwidth to provision sufficient capacity to satisfy customers during periods of peak load.

Much peak load capacity goes unused during off-peak hours, and it's certainly plausible that ISPs would encourage users of high volume file transfer applications that don't have stringent timeliness requirements to defer their operation to off-peak hours by adjusting bandwidth cap accounting. If service levels are based on volume, this would be a price reduction.

As systems for applications to signal their QoS requirements to the network improve, these systems of deferral can become sophisticated enough to operate within much smaller peak load periods, such as those measured in seconds and milliseconds. There's no good reason for a network to be under-provisioned with respect to the real requirements of its users and their applications; there is similarly no reason to require consumers to pay the price of massively over-provisioning the networks they use.

Differentiated pricing goes both ways, and fees charged by ISPs to application providers are only part of the overall pricing schemes that pertain to the Internet. Content delivery networks don't give their services away for free: Akamai's net profit margin is 17.5%,

while AT&T's is 9.9%; another CDN, Limelight Networks, has a 22.9% net profit margin, while Verizon's is 11.6%.¹⁷

One means of long standing that enables application service providers to obtain expedited delivery to the customers of a given ISP is a practice known as "paid peering" or "extended on-net access" in which the ISP extends its infrastructure all the way to the application provider's hosting facility. This arrangement often reduces the application provider's bandwidth bill and reduces packet delay in one fell swoop.¹⁸ The status of paid peering with respect to proposed net neutrality regulations is uncertain, however.¹⁹

Paragraphs 67-74, Competition

The competitive analysis section of the NPRM raises a number of what-if economic scenarios, which we address here.

Paragraph 68: *Even where there is effective competition in the Internet access market, individual broadband Internet access service providers may charge inefficiently high prices to content, application, and service providers, even though it may be in the collective interest of all providers to charge a lower price or zero price in order to maximize innovation at the edge of the network and thereby increase the overall value of broadband Internet access...If allowed to do so, broadband Internet access service providers may attempt to extract some of the profit earned by content, application, and service providers by charging them fees for providing access (or prioritized access) to the broadband Internet access service providers' subscribers.*

Paragraph 69: *...This dynamic raises a collective action problem: Although it might be in the collective interest of competing broadband Internet access service providers to refrain from charging access or prioritization fees to content, application, and service providers, it is in the interest of each individual access provider to charge a fee, and given multiple providers, it is unlikely that access providers could tacitly agree not to charge such fees. Furthermore, it is unlikely that competitive forces are sufficient to eliminate the incentive to charge a fee, particularly where the imposition of such a fee will not cause the access provider to lose many customers...*

We have to question the premise of this what-if scenario. The Internet has been in wide public use since the early 1990s, and nothing like this frightful scenario has ever been practiced. To the contrary, web properties such as ESPN 360 charge ISPs for the privilege of making their content available to ISP customers, and ISPs comply. Other web properties charge end users directly for access, and large numbers comply. Content is king.

¹⁷ "Scottrade - Quotes & Research," <https://trading.scottrade.com/quotesresearch/ScottradeResearch.aspx?symbol=T>.

¹⁸ William Norton, "Paid Peering and Net Neutrality," *Ask Dr. Peering*, November 5, 2009, http://drpeering.net/a/Ask_DrPeering/Entries/2009/11/5_Paid_Peering_and_Net_Neutrality.html.

¹⁹ Richard Bennett, "How Video Is Changing the Internet," *GigaOm*, November 22, 2009, <http://gigaom.com/2009/11/22/how-video-is-changing-the-internet/>.

In addition, the possibility of this practice emerging even if there was only one monopoly provider is highly remote. The presence of two wireline providers in most market areas (not necessarily most homes, but the relevant market for competition is not the home, it's the overall market area), plus overbuilders in some markets, combined with the two to four emerging 4G platforms, makes this scenario far fetched at best.

Moreover, the dynamics of ISP competition are such that each positions itself as a superior platform on the basis of its ability to connect users to web sites and other services with superior price and performance. Therefore, as unlikely as this scenario is, if it did come about the response from both consumers and the public interest would be swift and severe.

In addition, the economic analysis that predicts such behavior omits the cost and value of customer acquisition for ISPs and the dynamics of customer satisfaction. While advocates of strict net neutrality regulations bemoan the state of competition in the American broadband marketplace (often with little foundation) there is always at least one alternative to the local telephone or cable provider, non-participation. Surveys suggest that the largest reason that many Americans choose not to purchase broadband services from any ISP is the perception that there's nothing interesting on the Internet. Reducing the scope of available content and services certainly isn't going to motivate this potential pool of new customers suddenly to go on-line.

There is money to be made in differentiating IP services in order to meet the needs of diverse service providers other than ISPs. Differentiated pricing for differentiated service increases overall utility. Mandating a single service level for all Internet access ignores the diverse needs of applications and reduces utility.

These are among the numerous reasons why we have never seen, nor are we likely to see, an ISP charging application or content providers access to their regular best-efforts Internet service. If such an unlikely practice were to emerge, the FCC is already in a position to address it.

Paragraph 71: *In addition, broadband Internet access service providers generally, and particularly broadband Internet access service providers with market power, may have the incentive and ability to reduce or fail to increase the transmission capacity available for standard best-effort Internet access service, particularly relative to other services they offer, in order to increase the revenues obtained from content, application, and service providers or individual users who desire a higher quality of service.*

We don't see this as a plausible scenario. From the technical perspective, it would be difficult and impractical to accomplish. Provisioning an IP access network with a QoS option is generally going to entail a permeable membrane between best-efforts and better-than-best-efforts facilities, such that any bandwidth not consumed by better-than-best-ef-

ports delivery is automatically available to best-efforts and lower delivery classes. It's not typical to provision separate frequencies or channels for IP with QoS, in other words.

The alternative to differentiated QoS on a common IP channel is an increase in overall bandwidth, which has its own set of costs and benefits. As a rule of thumb, the provider would need to increase raw IP bandwidth by three to ten times in order to provide the same benefit to VoIP that simple prioritization would. This is not always possible, especially in DSL and wireless networks in which bandwidth has a hard ceiling, and can be very expensive in other cases.

Moreover, ISPs are investing to increase network speeds, whether through fiber-to-the-home, fiber-to-the-node, DOCSIS 3, or 4G wireless networks. The evidence that ISPs are limiting bandwidth expansion because they fear cannibalization of existing non-Internet services is simply not there. The major reason why the United States does not have the high-speed networks enjoyed in nations like Japan, Korea and Sweden is that these networks are quite expensive, and absent the incentives these nations have provided, market forces alone produce a slower, but still significant, advance. Moreover, consumers show little desire to pay the modest premium higher speed networks require. Indeed, American broadband customers haven't proven themselves great fans of top-tier broadband access speeds, nor have their counterparts in other parts of the world. As we found in our study of worldwide broadband adoption [Explaining International Broadband Leadership](#), customer acceptance of Next Generation Broadband network trails implementation by a very long time; "if you build it they will come" is not the operative rule.²⁰

Paragraph 72 raises the specter of anti-competitive pricing: *Broadband Internet access service providers have an incentive to use this gatekeeper role to make it more difficult or expensive for end users to access services competing with those offered by the network operator or its affiliates. For example, a broadband Internet access service provider that is also a pay television provider could charge providers or end users more to transmit or receive video programming over the Internet in order to protect the broadband Internet access service provider's own pay television service.*

In the first place, this economic analysis, as with those produced by net neutrality advocates, is completely arbitrary. There has never been a case of an ISP requiring charges from application or content providers for best-efforts access, nor is there likely ever to be such a case. If there were such a case, we believe that it would be appropriate and possible for the FCC to step in and ban the practice. That is completely different from an ISP providing a higher quality, differentiated service for a fee for application and content providers that need this service and are willing to pay in order to innovate and be successful in the marketplace.

²⁰ Rob Atkinson, Daniel K. Correa, and Julie A. Hedlund, *Explaining International Broadband Leadership* (Washington, DC: Information Technology and Innovation Foundation, May 1, 2008), <http://www.itif.org/index.php?id=226>.

Differentiated pricing isn't arbitrary or discriminatory when the underlying services have intrinsically differentiated costs of production.²¹ In order to provide cable television-like services, AT&T and Verizon had to invest in an infrastructure capable of delivering IPTV. Their investment consisted of very expensive upgrades to physical wiring and network equipment, not to mention licensing fees for content. IPTV has stringent requirements for packet latency because it's an interactive, high bandwidth application in which the customer is a critic of the delivery service. Delays when changing channels and audio or video dropouts make customers unhappy. IPTV employs a number of strategies to achieve desirable QoS, not the least of which is packet differentiation.

The capacity of an IPTV system is limited with respect to the number of simultaneous video streams it can carry; increasing the number costs money. In order for third party, over-the-top video services to compete with Verizon's and AT&T's IPTV services, third party services need access to IPTV-friendly expedited delivery services, and they need for their IPTV streams to be differentiated from generic Internet traffic. It's not enough for the broadband provider to allow third party video to enter the consumer's home from the Internet; they do that already, and no rational ISP would prevent it or charge for it.

However, to offer a service that's equal to or better than the broadband provider's in quality, the third party needs to set his Internet TV packets apart from generic Internet traffic just as the broadband provider does. Broadband providers would not have invested in network upgrades (at the expense of profit margins) if generic Internet service was good enough. Enhanced delivery service is needed for this application, and enhanced delivery costs more to provide than undifferentiated, generic Internet access; it's a more demanding service.

Therefore, the range of competitive options available to the television consumer is actually reduced by a ban on the sale of expedited delivery services to competitive television service providers, not increased. Not all applications can be satisfied by single service level pricing, and banning QoS-based delivery pricing confines Internet TV to the slow lane when it needs a genuine fast lane.

The second oversight concerns the interaction between the expected sale of service and investment in infrastructure. The typical Internet user doesn't often perceive the need to buy the ISP's higher tiers of service, and the relatively low speed access that most Americans choose is a barrier to innovation (which is why the National Broadband Plan is focused, among other things, on how to get higher broadband speeds deployed and used).²² Allowing application providers to pay for a temporary speed boost into a customer's home allows the consumer to experience the benefit of a higher service tier

²¹ Inimai M. Chettiar and J. Scott Holladay, *Free to Invest: The Economic Benefits of Preserving Net Neutrality*, Institute for Policy Integrity (New York: New York University School of Law, January 2010), http://www.policyintegrity.org/documents/Free_to_Invest.pdf.

²² See our report: Rob Atkinson et al., *The Need for Speed: The Importance of Next-Generation Broadband Networks* (Washington, DC: Information Technology and Innovation Foundation, March 5, 2009), <http://www.itif.org/index.php?id=231>.

overall, which leads to a virtuous cycle of demand by consumers for faster services and incentives for network operator investment in even faster services.

Paragraphs 79 and 80, Congestion

This section deals with congestion, a central topic in the discussion of Internet regulation:

The existence of congestion in the network is a major motivating factor in the open Internet debate, and is central to arguments that differential pricing or service quality is necessary. Moreover, because the effects of delays or dropping of packets arising from congestion are not the same for all applications, broadband Internet access service providers and content, application, and service providers may have incentives to seek agreements for the prioritization of traffic or other quality of service guarantees. Permitting these activities without appropriate oversight could lead to a number of harms, undermining the public interest goals of the Act discussed above.

Congestion is a broad topic comprised by many sub-topics, such as delay, packet loss, server overload, choke points, and flow control. These topics pertain to physical networks, not to virtual ones such as the Internet, so we take the NPRM's use of the term as pertaining to congestion in ISP edge router queues and in switch and router queues internal to ISP networks. A more meaningful discussion could be had if congestion in Network Service Provider networks were part of the discussion, because that would allow the inclusion of congestion avoidance strategies such as Content Delivery Networks and Overlay Networks. Any expedited delivery service potentially offered for sale by an ISP would compete with alternate expedited delivery services as CDNs and Overlay Networks, in other words. The question isn't so much whether expedited delivery should be allowed for sale – it already is and has been for some 15 years – but whether the market for expedited delivery services should be expanded to allow ISPs to compete with such firms as Akamai and Limelight on a level playing field; and secondly whether such competition warrants special scrutiny by the FCC or some other branch of the federal government. We believe that such practices should be allowed as long as ISPs provide the service in a way that does not discriminate against similarly situated competitors.²³

An additional question concerns the relationship of expedited delivery to competition in specialized services such as “cable TV,” VoIP, and video conferencing. These issues are addressed in detail in the section on the non-discrimination rule, paragraphs 104-117.

Paragraphs 81-82, Various Issues

These paragraphs are a catch-all for any issues omitted in the foregoing summary, so we would like to comment on the questions they raise as well as on the questions relevant to the rulemaking that have so far not been raised.

²³ Robert D. Atkinson and Philip J. Weiser, *A “Third Way” on Network Neutrality* (Washington, DC: Information Technology and Innovation Foundation, May 30, 2006), <http://www.itif.org/files/netneutrality.pdf>.

Competition

We urge the Commission to review our study of broadband competition in OECD nations, [Explaining International Broadband Leadership](#).²⁴ The report examines the role of policy and non-policy factors in the broadband investment and competition and finds that facilities-based competition is present in countries with highest penetration of next-generation broadband networks.

One method of creating competition in the absence of multiple physical networks is mandatory unbundling of wholesale and retail Internet access services. While this strategy is helpful in the absence of facilities-based competition, it tends to retard investment in Next-Generation Networks (NGN) and creates a kind of DSL cul-de-sac, as we see in many European nations.

The competition analysis must be undertaken broadly in any event; the NPRM focuses exclusively on the potential behavior of ISPs who constitute, after all, a rather small portion of the Internet ecosystem. ISPs provide a number of services that are readily available from alternate suppliers, such as cable TV, e-mail, domain name services, spam mitigation, IP transit, and telephony, in addition to basic residential IP packet delivery at a number of service levels. The market for ISP services is expanding in the United States, thanks to fixed wireless systems such as Clearwire and other 4G wireless networks (including some that will start to be deployed this year), pioneering Wireless ISPs such as LARIAT in Laramie, Wyoming and hundreds of similar services across the country, improved satellite broadband deployments, and expanded use of high-rate 3G systems based on HSPA and HSPA+. A great deal of the congestion crunch on broadband networks comes about from the migration to IPTV, but alternate television delivery systems are emerging that address this problem through Internet bypass, such as the Sezmi²⁵ system and the Qualcomm wireless video system used by Major League Baseball.²⁶ Consumers have choices in many of the key areas in which competition has been lacking in the past.

Finally, there is little or no evidence to suggest that current markets are not competitive. One standard measure of market power is the presence of supra-normal profits. However, as described above, profits among cable and telephone providers of broadband are in fact lower than many other players in the ICT ecosystem. Nor is there evidence of anti-competitive collusive behavior among ISPs. Finally, with the exception of Madison River there has been no evidence of anticompetitive behavior among ISPs to favor their own services at the expense of competitors. Finally, as we have written about before, simply relying on standard industrial organization models that state that more competitors is always better ignores the unique characteristics of this industry, in particular its high fixed costs.

²⁴ Atkinson, Correa, and Hedlund, *Explaining International Broadband Leadership*.

²⁵ "About Sezmi," <http://www.sezmi.com/about.php>.

²⁶ "Qualcomm steps to the plate for MLB deal," *Fierce Mobile Content*, June 22, 2007, <http://www.fiercemobilecontent.com/story/qualcomm-steps-to-the-plate-for-mlb-deal/2007-06-22>.

Architecture

The NPRM doesn't ask any questions about the relationship of Internet architecture – the design principles evident in Internet protocols – and Internet practice. In particular, the NPRM stresses the notion, not entirely correct, that Internet practice dictates a single service level, “best-efforts” packet delivery system across the entire Internet. While this idea is widespread, it's not entirely true as a historical or empirical matter today, but the precise extent of its inaccuracy is hard to know. This is simply because the terms of many Internet peering agreements are transacted under non-disclosure, which firms are understandably loath to violate. There is also a considerable amount of private information around the placement of CDN server farms, both public (like Akamai and Limelight) and private (like Google and Amazon.) The placement of these facilities and terms under which they operate both have significant implications.

Conceding that the bulk of Internet traffic today moves at generic delivery levels from network to transit provider to network, it doesn't follow that one network's generic delivery is the same as all others. Application, content, and service providers make the same kinds of economic choices as consumers with respect to speed, capacity, equipment, and bandwidth, and these choices have performance implications. The Internet is often said to represent a “level playing field” in which all applications, content, and services compete on equal terms, but the ideal represented by this claim is significantly at odds with the facts.

It is simply not the case that a blogger with a \$10/month virtual host can compete with the New York Times. An individual blogger has to gain reputation, attract readers, build his brand, sell advertising, and trade up to larger and larger servers before he can attract enough readers and serve enough content to compete with large enterprises. The Internet permits enterprises to grow because it allows content and services providers to purchase bandwidth and server capacity in relation to their needs; it's not as much a level playing field as it is a free market for content and scalable service delivery that permits the right-sizing of spending in proportion to distribution.

Moreover, the single service level is not dictated by Internet architecture. Best-efforts Type of Service in terms of IP has been dominant in an era in which the Internet was primarily used for the single purpose of accessing static pages on web sites; there is no guarantee that it will be sufficient to enable a diverse mix of applications and services.

Digital Divide

One likely consequence of limiting Internet access accounts and ISP-application service provider agreements to a single service level is an increase in the retail price of Internet access. Removing this tool from the ISPs' repertoire of legitimate business practices will inevitably lead to lower efficiency networks, which will have higher prices to consumers. To the extent that broadband adoption and use is hampered by high prices (which isn't

the whole story, but research suggests is certainly part of it) such a regulation will increase the digital divide.

According to a recent Pew report, minority communities are more dependent on wireless broadband than is the majority. As a single service level mandate would have more severe negative impact on wireless than wireline broadband networks, it would harm minority broadband users more than others.²⁷ This can't be a desirable outcome.

Paragraph 92, Internet Policy Statement

This paragraph offers the proposed text for the codification of the *Carterfone*-based Internet Policy Statement. The most significant of these is Rule 4, which states:

Subject to reasonable network management, a provider of broadband Internet access service may not deprive any of its users of the user's entitlement to competition among network providers, application providers, service providers, and content providers.

For reasons already mentioned, the aspirations of this rule are contradicted by the proposed Rule Five on "Discrimination." In resolving this conflict, our recommendation is to affirm the importance of competition over that of historic practice. The surest way to deprive consumers of meaningful choice and competition among applications and network-enabled services is to eliminate or restrict the network services that make diverse applications possible, or the financial base that makes ISPs willing to offer such network services for sale. It's not enough for the Commission to express aspirations if the rules that are intended to execute these aspirations are likely to have an effect contrary to the aspirations.

Paragraph 100, Wireless

This paragraph seeks comment on the application of the Policy Statement to wireless networks. While the general spirit of the Policy Statement is laudable, the specifics of its application to wireless networks are troubling.

The first and most obvious issue is bandwidth. Wireline networks have a relatively unfettered bandwidth horizon. Wireline Internet access services are offered in some parts of the U. S. at speeds as great as 100 megabits/second, over facilities capable of increases into the gigabits as demand grows. While wireless systems have regularly increased in speed because of the greater spectral efficiencies predicted by Cooper's Law (a doubling of information capacity every 30 months) the greater directionality of wireline networks and their greater noise immunity endows them with obvious advantages in terms of capacity per equipment dollar. There is also the question of the co-existence of Internet-oriented services with legacy services on cellular wireless networks. Around the world,

²⁷ Lee Rainie, *Internet, broadband, and cell phone statistics* (Washington, DC: Pew Internet and American Life Project, January 2010), <http://www.pewinternet.org/Reports/2010/Internet-broadband-and-cell-phone-statistics/Report.aspx>.

some 4.5 billion people use cell phones for voice communication today. In many instances, the cell phone is the first and only interpersonal communication device the user has ever had.

Cell phones enable some of the poorest and most downtrodden people in the world to run their own businesses, to stay in touch with employers and clients, with friends and family, and with police and medical service providers. These simple, low-bandwidth devices have changed lives as dramatically as the telegraph did in the Nineteenth Century, but they use the same radio frequencies as wireless Internet access services. Cell networks are optimized for voice, but Internet access services are optimized for content. Consequently, providing VoIP on a cellular network imposes a greater burden on the network than traditional calling does, and this burden affects the ability of legacy users to communicate.

A declaration that Internet access users must be free to run the applications of their choice on wireless devices such as iPhones, Blackberries, Palms and Android devices implies that the network operator should be able to supply users with an acceptable experience. When the application a wireless Internet user chooses happens to be something like Skype video calling, which consumes a large amount of network bandwidth, and the immediate vicinity of the user is bandwidth-constrained, the user's experience will not be a happy one. Moreover, the use of bandwidth-intensive applications on channels occupied by large numbers of legacy users has the potential to impair simple phone calls as well. As we explained in our report on wireless handsets, [*Sharing the Risks of Wireless Innovation*](#), the application of *Carterfone* principles to wireless handsets is far from straightforward.²⁸ *Carterfone* rules were devised for networks that were relatively free of side effects among users, a very different scenario than the reality of wireless channel sharing.

Finally, with an all-IP wireless technology on the horizon – LTE – we have to question the urgency of enacting an IP-oriented revision of *Carterfone* for the 3G wireless networks that will soon enter “legacy” status. By the time these rules are enacted, new technology will be in the early stages of deployment that makes them largely moot.

Paragraphs 103-117, Discrimination

This brings us to the detailed discussion of the troubling anti-discrimination rule. As the NPRM astutely concedes, network discrimination can be beneficial: *The key issue we face is distinguishing socially beneficial discrimination from socially harmful discrimination in a workable manner.* We've said as much ourselves, in Atkinson and Weiser's paper, [*A Third Way on Network Neutrality*](#), which the NPRM cites.²⁹ The model that we propose is to permit the sale of expedited transport services on a non-

²⁸ Richard Bennett, *Sharing the Risks of Wireless Innovation* (Washington, DC: Information Technology and Innovation Foundation, October 19, 2009), <http://www.itif.org/index.php?id=303>.

²⁹ Robert D. Atkinson and Philip J. Weiser, *A “Third Way” on Network Neutrality* (Washington, DC: Information Technology and Innovation Foundation, May 30, 2006), <http://www.itif.org/files/netneutrality.pdf>.

discriminatory basis: *Notably, this model of regulation (unlike the Wyden bill) would allow quality-of-service assurances to be offered for payment, but such assurances must be offered universally unless a firm has a legitimate business purpose for offering it only on an exclusive basis. Significantly, this standard calling for reasonable access to prioritized service delivery (even for a fee) would also apply to the level of prioritization offered to an incumbent's affiliated application (say, its VoIP product).*³⁰ Such a framework ensures that better-than-best-efforts access (or cheaper-than-best-efforts bulk mail-like service, if that is what the application requires) is available to any application, content, or service provider with a competitive product to sell, and prevents the network operator from privileging its own managed service products over those offered by third parties. The correct way to draft a non-discrimination rule is to focus on discriminatory market practices, not on engineering practices that are either discriminatory or not depending on ones viewpoint on nuanced aspects of network design and operation.

Paragraph 113 asks a vital question: *Are there particular content, applications, or services whose quality and utility to end users depends on a broadband Internet access service provider's assuring a certain quality of service?* The answer to this question can be found in the way that IPTV and VoIP are provided over ISP networks today: operators either provision dedicated bandwidth for these services (cable and FiOS) or prioritize voice and video packets flowing over shared bandwidth (DSL.) Both of these methods are costly to the operators, so the question can be re-phrased "Are broadband network operators wasting the money they spend to provide QoS guarantees for their own voice and video products today?" We believe that they are not, and further believe that operators have made and continue to make enormous investments in network facilities.

The NPRM also asks *why such practices are required.* The answer to this comes from the bandwidth contention system we've already described, and from the statistical nature of Internet access. While the Internet is a virtual network rather than a physical one, it assumes that its underlying physical network will offer a packet switching service. When IP is implemented over a network such as the PSTN that does not provide a native packet switching service, engineers devise an adaptation layer that provides one, such as Switched-Line IP (SLIP) or Point-to-Point Protocol (PPP.)³¹ IP also assumes that it can offer packets to the packet switching service at any time, with a high degree of confidence they will be delivered as long as the destination address is reachable. It also assumes that packets will frequently be delayed in transit, but will only be dropped when some portion of the network is unable to save them in a *buffer* for later transmission; in some IP systems, a very large pool of buffers are present, into the hundreds of thousands.

These assumptions are suitable for applications that are relatively insensitive to delay, and the system properties that implement them have the effect of causing very large swings in the delivery time of particular packets. Internet traffic does not appear at

³⁰ Ibid, p. 13.

³¹ "Differences Between SLIP and PPP," <http://sunsite.nus.sg/pub/slip/slip-vs-ppp.html>.

switching points at regular intervals; it's highly stochastic, and appears in clumps separated by unpredictable intervals.

This system is optimized for file transfer applications that are very insensitive to delay. As Tim Wu explained in the paper that introduced the term "net neutrality," the Internet has a bias in favor of content-oriented, file transfer applications:

*Proponents of open access have generally overlooked the fact that, to the extent an open access rule inhibits vertical relationships, it can help maintain the Internet's greatest deviation from network neutrality. That deviation is favoritism of data applications, as a class, over latency-sensitive applications involving voice or video.*³²

Consequently, network operators need to prioritize voice and video packets over the packets generated by "data applications" in order to overcome the Internet's structural bias. What the NPRM describes as "discrimination" is actually a means of overcoming a built-in favoritism.

Paragraphs 118-132, Transparency

These paragraphs deal with the proposed transparency rule, one that we fully endorse. The issue that arises with the construction of such a rule is the large gap that exists between the public understanding of the Internet and its constituent networks and the engineering practices employed by network operators. It is literally the case that the public is incapable of making sense of the systems and tools needed to operate a packet switching network. While this gap is too large to bridge, the Commission could do well to convene a working group as part of the Technical Advisory Process to devise common terminology and metrics to enable consumers to make simple, apples-to-apples comparisons between various service offerings. A general set of guidelines should be devised before any enforcement action is taken; this area is too complex for case-by-case development.

Paragraphs 135-141, Reasonable Network Management

Reasonable network management is called upon by the NPRM to do the heavy lifting by way of justifying differentiated services, which is an incorrect emphasis. Network management is a reactive discipline, one that network operators employ after the fact to correct unanticipated problems and issues that arise in the day-to-day operation of a network. Networks deal with denial-of-service attacks on a daily (sometimes hourly) basis, but this is always a reaction to events, not a pro-active organization of the network's operation. The issues of making differentiated services or paid peering available for a fee is not one of network management, it's a decision that implicates network design in the service of a business model. The following portions deal with business practices for which no extra fee is involved.

³² Tim Wu, "Network Neutrality, Broadband Discrimination," *Journal of Telecommunications and High Technology Law* (2) (2003): 141, <http://ssrn.com/abstract=388863>.

Paragraph 137: *Some have suggested it would be beneficial for a broadband provider to protect the quality of service for those applications for which quality of service is important by implementing a network management practice of prioritizing classes of latency-sensitive traffic over classes of latency-insensitive traffic (such as prioritizing all VoIP, gaming, and streaming media traffic). Others have suggested that such a practice would be difficult to implement in a competitively fair manner and could undermine the benefits of a nondiscrimination rule, including keeping barriers to innovation low. We seek comment on whether these and other potential approaches to addressing congestion would be reasonable. On the other hand, we believe that it would likely not be reasonable network management to block or degrade VoIP traffic but not other services that similarly affect bandwidth usage and have similar quality-of-service requirements.*

This goes to the Madison River case in which VoIP was blocked in order to protect the operator's telephone services. We agree with the NPRM's analysis that such blocking is not beneficial and should not be permitted in the United States, even though it's unfortunately widely practiced in countries such as Korea. To be considered reasonable, a network management practice needs to benefit users, not simply the provider. VoIP is a very low bandwidth application; there is no scenario in which it would be reasonable to suppose that blocking VoIP would produce customer benefit. Blocking or managing high bandwidth applications in the interest of protecting VoIP, as Comcast did, when they managed peer-to-peer traffic is an entirely different matter, of course.

Nor would we consider the singling out of any particular content (i.e., viewpoint) for blocking or deprioritization to be reasonable, in the absence of evidence that such traffic or content was harmful.

We strongly agree. This scenario raises an interesting question about the proposed rules, however. Given that the protection of free speech is a matter of public interest, and has in fact been offered as the rationale for attaching restrictions to differentiated services, why is there not an explicit rule in the proposed regulatory framework to address viewpoint-based discrimination? Much, if not all, of the legitimate public interest outcry in favor of net neutrality regulations would be satisfied by such a condition.

We recognize that in a past adjudication, the Commission proposed that for a network management practice to be considered "reasonable," it "should further a critically important interest and be narrowly or carefully tailored to serve that interest." We believe that this standard is unnecessarily restrictive in the context of a rule that generally prohibits discrimination subject to a flexible category of reasonable network management. We seek comment on our proposal not to adopt the standard articulated in the Comcast Network Management Practices Order in this rulemaking.

We agree. The Comcast ruling was apparently drafted with legal precedent in mind rather than the requirements of robust, multi-purpose networks and their users. The relevant standard for evaluating a network management practice is simply whether it increases

network utility for the majority of users, producing “the greatest good for the greatest number.” More restrictive rules are impractical to implement and monitor, if for no other reason than the fact that network administrators are not legal scholars. A better approach is the one adopted by Canada’s CRTC and the European Union, combining broad management flexibility with full disclosure.

Paragraph 173, Mobile Broadband

The NPRM asks: *In what way do these wireless characteristics affect what kinds of network management practices are or are not reasonable? Are there particular wireless network management practices that should be identified by the Commission as reasonable? For example, are there any circumstances in which it could be reasonable for a wireless network to block video applications because they consume too much capacity? What about third-party VoIP applications or peer-to-peer applications?*

Mobile broadband networks are more resource constrained, more variable with respect to demands for bandwidth, and more error-prone than wireline broadband, so their operation requires more flexibility. Wireline networks don’t have to deal with roaming, fading, weak signals, and multipath, for example. Therefore, there are scenarios in which mobile users of video streaming applications and inefficient peer-to-peer applications will notice service degradation. The proper way for mobile network operators to deal with momentary congestion, however, isn’t reliant on outright blocking of entire classes of applications – there should not be a wholesale ban on P2P, for example – but it is appropriate to limit the bandwidth available to best-efforts applications as a whole when there is high demand for more critical voice circuits. Mobile network customers pay higher fees for voice on a per-byte basis than they do for best-efforts Internet access, so it’s appropriate for voice to have first call on network resources. Nothing about this prioritization is anti-competitive as it’s ultimately controlled by the customer. There is similarly nothing anti-competitive about allowing mobile broadband customers to use their voice minutes to call Skype or EQO³³ for cheap international calling. It’s not reasonable for customers to expect that lower-cost best-efforts Internet access will have the same QoS parameters as voice, however.

If the Commission wishes to increase competition for mobile calling services, a productive approach is to make it clear that carriers may sell better-than-best-efforts QoS to alternative voice over wireless providers. This will put alternative suppliers on the same quality threshold as operators, and will enable them to pay for a fair share of network resources.

Conclusion

We’re greatly encouraged that the Commission has committed to examining the issues raised by the NPRM in an open process. The challenging economic and technical issues raised by the NPRM potentially alter the utility of Internet access networks substantially; whether the outcome is for good or for ill remains to be seen, of course. We urge caution,

³³ “EQO - Mobile VoIP, free texts, and free mobile IM with EQO,” <http://www.eqo.com/>.

however, on the part of the Commission. As we've stated, evidence that the contemporary *status quo* harms the citizenry, society, or the economy is sorely lacking, so there is little justification to move aggressively into the Internet access marketplace with six guns blazing. Moreover, there is a real risk that any action taken could have the unintended consequence of limiting innovation and investment going forward. As such, it is commendable that the Commission is asking questions first, for example.

While the Internet has been with us, in one form or another, for more than 35 years, it's still a technology in its infancy. The present population of 1.6 billion Internet users worldwide is soon to be joined to the more than four billion users of the mobile phone network, and we're not in a position to appreciate the changes that the convergence of the Internet with mobile networks will bring. The Internet has been guided by a spirit of enquiry and experimentation throughout its brief history, one in which questions of network efficiency have often taken a back seat to the freedom of experimentation. The shortcomings in the Internet architecture, of which there are many, have largely been hidden behind the rapid pace of change in semiconductor technology guided by Moore's Law and the rapidly falling prices of wireline bandwidth.

Wireless networks have evolved along completely separate lines, with increases in spectral efficiency guided by Cooper's Law bearing the responsibility for every generational change in network utility. With the advent of LTE, these two cultures of network engineering are set to merge within the next two years. This is a momentous time in the evolution of networking, and a particularly inappropriate moment for telecom regulators to impose new structures on network engineering, operations, and use. For this reason, the Commission's counterparts in Canada and Europe have prudently chosen to adopt a disclosure-oriented approach, keeping a watchful eye on network developments rather than creating new frameworks for regulation.

There is little demand for wholesale new regulations on Internet access networks, especially wireless ones, outside the small circle of professional critics in the public interest lobby. The "movement" for net neutrality regulations has been ongoing since the turn of the century, if not longer. It has found little traction with the public despite the efforts of some to make it a mass movement following the remarks that former Bell South CTO Bill Smith offered about potential sources of new revenue four years ago. The Internet has continued to thrive despite forecasts of imminent doom.

There is now a need for the Commission to clarify the status of Internet in terms of disclosure and regulation. We hope that this can be accomplished without onerous regulations that harm innovation and investment. A critical insight in this proceeding concerns the relationship of innovation in the core of the Internet to innovation at the edge: they both need each other, they will both advance or both retreat in tandem, and there is no way of protecting one by harming the other. It's a serious mistake to pit one against the other, regardless of which one is favored. As long as the Commission retains its commitment to an open process, and conducts the inquiry in a fair, curious, and probing manner, we're confident that a reasonable outcome will come to pass. As ever, we stand ready to lend a hand.

Information Technology and Innovation Foundation
1101 K Street, N.W., Suite 610
Washington, DC 20005