

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

In the Matter of)	
)	
)	GN Docket No. 09-191
Further Inquiry Into Two Under-Developed)	WC Docket No. 07-52
Issues in the Open Internet Proceeding)	

COMMENTS OF SCOTT JORDAN, Ph.D. and GWEN SHAFFER, Ph.D.

Scott Jordan,
Professor of computer science

Gwen Shaffer,
Post-doctoral fellow

Department of Computer Science
University of California, Irvine
3029 Bren Hall
Irvine, CA 92697

October 12, 2010

TABLE OF CONTENTS

INTRODUCTION AND BACKGROUND.....	4
DISCUSSION.....	5
I. It is not possible, nor is it necessary, to delineate Internet traffic by creating a “managed services” category	6
A. Regardless of how the Commission chooses to define managed services, the category is nearly certain to lose relevance as technology convergence continues to blur distinctions among telephone, cable, broadband and wireless networks.....	8
B. Even industry experts fail to reach consensus on a definition for “specialized” or “managed” services.....	9
C. Some applications require QoS. However, QoS should be available to these applications regardless of the application provider.	11
II. Among the six policy approaches highlighted by the Commission in the NOI, the sole viable option would mandate non-exclusivity in specialized services.....	11
A. The goal of regulation is to foster an environment that fosters innovation and ensures a competitive marketplace.....	12
B. While we do not dismiss out-of-hand the remaining five policy approaches floated by the Commission, we do feel strongly that these approaches fail to go far enough.....	13
III. Rather than creating a managed services category, a simpler and more efficient solution involves requiring an <i>open interface</i> to QoS mechanisms in lower layers.....	14
A. The need for a separate class of managed services can be eliminated by applying a layered approach to traffic management practices.....	14
B. The Commission must re-evaluate the layering model currently under consideration.....	16
C. Broadband consumers already choose service levels from a menu of tiered services offered by the ISP.....	17
IV. Should Commissioners exempt wireless networks from net neutrality requirements, they are simply prolonging the inevitable.....	17

A. As with wired applications and devices, technological convergence is erasing traditional distinctions between communication and information services.....	18
B. From the end-user perspective, the functionality of the mobile web increasingly resembles the wired Internet.....	18
C. Reasonable network management requires limiting only traffic, not restricting applications.....	19
D. Stronger traffic management at or below the network layer does not justify different treatment above the network layer.....	19
E. The open interface represents a contract between the wireless carrier and the user. The wireless carrier offers certain services and the user requests (or purchases) those services desired at the volume desired.....	21
CONCLUSION.....	22

INTRODUCTION AND BACKGROUND

These comments are filed by University of California, Irvine Professor Scott Jordan and post-doctoral fellow Gwen Shaffer, in response to the Federal Communication Commission's ("Commission") *Notice of Inquiry*¹ in the above captioned docket. As academic researchers, we take an objective approach to this analysis, without the biases or motives that may influence many commenters. We believe that this kind of independent analysis can help inform policy decisions in a way that is distinct from guidance provided by industry practitioners and public interest groups. As a Professor with a research focus in telecommunications policy and in Quality-of-Service (QoS) issues, Professor Jordan offers a unique perspective on how to effectively manage traffic on both wireline and wireless networks. He is also participating in the Open Internet Technical Advisory Process created as part of the Open Internet proceeding.

Internet applications are becoming increasingly bandwidth intensive and, without reasonable traffic management, the Internet's usefulness and reliability cannot be maintained. At the same time, the Internet was founded—and has thrived—on the principle that users have equitable ability to access and transmit online content. This can be considered as the virtual equivalent of Americans' right to travel to anywhere in the country, from the sleepiest town to the most bustling city, and to stay in each destination for as long as they wish. Of course, flying first-class and staying in a 5-star hotel will cost more than driving to a KOA campsite. Similarly, as we explain in detail later, allowing end users to pay for QoS is in line with the call for a neutral Internet. These comments also reflect our conviction that it is possible for ISPs to manage traffic without employing techniques that discriminate against users, content, or applications. Given that the Commission requests feedback, specifically, on so-called "specialized" or "managed" services, this filing offers recommendations for how ISPs may accommodate these transmissions in a way that allows for *differentiation*,

¹ In the Matter of Further Inquiry Into Two Under-developed Issues in the Open Internet Proceedings, GN Docket No. 09-191, *Notice of Inquiry*, (rel. September 1, 2010) ("*Notice of Inquiry*").

but not *discrimination*. For instance, when broadband subscribers contact long-distance friends over Skype, they trust that their ISP will not impose unreasonable discrimination against the Skype call in comparison with the ISP's own voice over IP (VoIP) service. When Internet users download the most recent *Harry Potter* film from Netflix, they expect the option for it to stream as quickly as content offered by their broadband providers' own video-on-demand content. For over two decades, these open Internet principles have successfully fostered innovation and investment in both broadband infrastructure and applications. The Commission must undertake policy reforms that continue these positive outcomes, while also taking into account technical challenges and economic realities.

In previous research, we proposed that traffic management practices are reasonable if they are implemented at endpoints; are chosen by the user; are based on reasonable application provider payment; or involve providing QoS to traffic chosen by the user.² We proposed that traffic management practices implemented in transit nodes without user choice are unreasonable if they block applications or violate the net neutrality policy discussed above (e.g. provide QoS based on source or on unreasonably discriminatory application provider payment). We suggest that QoS based on the application can be more effectively implemented by allowing users to determine the priority of their applications, and we suggested that any charges for QoS can be most effectively implemented by integrating them into subscriber contracts and into the Service Level Agreements (SLA) between ISPs, rather than by charging application providers that are not subscribers.³ The policy guidance we offer here builds on these same principles, which attempt to find a middle-ground on the highly divisive issue of net neutrality.

DISCUSSION

² Scott Jordan: Do wireless networks merit different net neutrality than wired networks? (2010) Available at <http://www.ics.uci.edu/~sjordan/papers/tprc2010.pdf>.

³ Scott Jordan: Do wireless networks merit different net neutrality than wired networks? (2010) Available at <http://www.ics.uci.edu/~sjordan/papers/tprc2010.pdf>.

In its recent Notice of Further Inquiry into *The Open Internet Proceeding* (NOI), the Commission seeks comment on how the concerns regarding managed services can be addressed, while “promoting private investment and encouraging the development and deployment of new services that benefit consumers.”⁴ We focus on these issues in the following section.

I. It is not possible, nor is it necessary, to delineate Internet traffic by creating a “managed services” category.

The NOI requests comment on how carving out policy exceptions for specialized services would protect open Internet principles and encourage innovation. Because any definition of specialized or managed services would, ultimately, be arbitrary, we strongly discourage the Commission against establishing policy grounded in this concept.

The dominant ISPs in the United States offer plans with varying amounts of bandwidth. Cox’s plan is typical of this tiered pricing scheme, with customers choosing from 3 Mbps, 12.5 Mbps, and 25 Mbps of maximum access transmission rate.⁵ If a subscriber opts for a more expensive plan, she is paying not only for additional bandwidth, but also for QoS. In addition, in order to limit latency or jitter, ISPs use traffic management to prioritize their own delay sensitive applications such as VoIP and video streaming.

A few proponents of net neutrality argue that the mere offering of QoS to subscribers, as ISPs typically do today, is a violation of open Internet principles. We do not agree with this interpretation of net neutrality. By analogy, the U.S. Postal Service employs a similar pricing structure without criticism. Postal service customers may choose to send a book across the country via “express mail” to guarantee next-day delivery. Of course,

⁴ *Notice of Inquiry*, at 4

⁵ Cox High-speed Internet Preferred, Cox High-speed Internet Premier, Cox High-speed Internet Essential. Menu of Cox broadband service plans for Orange County, California. Available at https://secure.cox.com/Service/Store/OrderNow.aspx?hsi=w&cc=oc_hsi-half_ordernow-btn2&campcode=splash_r_oc_hsi-half_ordernow-btn2-090110

postage will cost more than had they opted to send the book “priority,” for arrival within two days. USPS offers an even cheaper ground delivery rate for customers willing to wait a week for the package to arrive. Everyone agrees it is not problematic for the postal service to establish pricing this way.

Most net neutrality proponents find the offering of QoS to be reasonable as long as it is made available on a nondiscriminatory basis. We agree. What would not be acceptable is if the Postal Service peered into customers’ envelopes and charged them more, or slowed down delivery, based on whether the book was about a politician or sports, whether it was fiction or non-fiction. Similarly, we find it unacceptable for an ISP to use traffic management to prioritize its own delay sensitive applications such as VoIP and video streaming, while neglecting to offer the same QoS to competitive applications.

A category of managed services violates consumer protection principles, which ensure that Internet subscribers can use any application. The *NOI* describes a managed service model as one that offers services that are “substantially similar to, but do not technically meet the definition of, broadband Internet access service,” with the potential for consumer protections to “not apply to such services.”⁶ A managed services category would, intentionally or unintentionally, allow ISPs to avoid compliance with nondiscrimination rules. We believe this to be the outcome regardless of how the Commission ultimately chooses to define managed services. The distinction between the concept of tiered pricing and managed services is clear—the end-user determines the value of the bandwidth in the first scenario, while the ISP determines which traffic receives QoS in the latter.

In practice, a managed services network would function separately from the “public Internet.” Telemedicine, videoconferencing, financial transactions and distance learning are frequently cited as examples of the kinds of traffic that necessitate priority delivery and, thus, should be labeled managed services. However, as numerous commenters have previously pointed out, the fuzzy boundaries inherent in this designation invite the opportunity for broadband providers to favor their own applications and services over content delivered by

⁶ See *NOI* at 2.

competitors. Netflix makes this point well: “The fact that network operators control the delivery pipes and generate significant revenue from content that travels over those pipes provides both the means and motive for discriminating against new ventures that might threaten revenue sources of the network operators.”⁷ If cable and telephone companies are able to designate their own VoIP and video services as *managed services*, while declining to offer QoS to competitive applications, consumers will experience poorer quality. This would completely undermine the Commission’s goal of promoting competition in services that require QoS.

- A. Regardless of how the Commission chooses to define managed services or specialized services, the category is nearly certain to lose relevance as technology convergence continues to blur distinctions among telephone, cable, broadband and wireless networks.

We are convinced that the distinction between managed or specialized services and the “public” Internet that consumers rely on today is nearly impossible to clarify. The challenge of crafting such a definition is amplified by technological convergence, or “everything over IP.” Today, landline telephone, cable television and cellular phone networks are all migrating to applications commonly associated with broadband.

For about a century, telephone exchanges were a complex series of electromechanical selectors. In the 1960s, software-controlled digital exchanges began replacing the old analog system—improving both the capacity and quality of the network. Another major shift occurred in 2004, when carriers began marketing VoIP to residential and business consumers. With increasing frequency, telecommunications providers use VoIP over dedicated and public IP networks to connect switching stations, as well as to interconnect with other telephony network providers. This progression is not unique to telephony. Traditionally, television programming has been delivered to American audiences through radio frequency broadcast, satellite signal, and coaxial cable. Over the past decade, however, IP multicast has emerged to challenge these formats. AT&T launched the first commercial IPTV service in the United States in 2006. AT&T U-Verse, which delivers video, voice and data

⁷ Comments of Netflix (filed January 14, 2010), at 8.

over fiber-optic cable, is now available to about 23 million homes.⁸ Verizon also offers video-on-demand content over IP through its FiOS network, which claims more than 3 million subscribers.⁹ Bill Lake, the Commission's Media Bureau Chief, predicted the distinction between TV and Internet services are likely to remain separate for only a short time. During an agency meeting in December, Lake suggested that the telecommunications market is searching for a superior way to connect the 99% of households that have TVs with the 75% of homes that have computers.¹⁰

The trend toward convergence is extending further, as IP becomes the dominant platform for mobile device connectivity. When second generation (2G) technology for wireless telephones hit the commercial market in the early 1990s, it was widely hailed for its efficient use of spectrum, as well as its ability to transmit text messages. A few years later, this digital standard evolved into 3G, which offered even higher data rates and a slew of bandwidth-intensive applications such as streaming video. However, today wireless phone carriers are investing heavily in Long Term Evolution (LTE), a mobile network technology that uses a flat IP-based networking architecture. LTE is meant to make existing applications faster, as well as to host countless new applications—from high-definition movies to photo-sharing—previously available only through a wired broadband connection.¹¹

Although lower layer protocols differ based on the characteristics of the physical connection, these four types of networks will continue to converge at the IP layer and above, and to offer an increasingly similar set of highly diversified services and applications. The integrated communications described here slash costs, simplify network maintenance, spur innovation, and enhance user mobility. They may also make the Commission's attempt to define a separate category dedicated to high-speed access and bandwidth intensive applications a

⁸ *AT&T 2009 Annual Report*, AT&T U-verse, at 4. Available at <http://www.att.com/gen/investor-relations?pid=17388#top>.

⁹ *Broadband DSL Reports*, Verizon: 3.1 Million FiOS Customers. Available at <http://www.dslreports.com/shownews/Verizon-31-Million-FiOS-Customers-103626>.

¹⁰ John Eggerton, *Broadcasting & Cable*, Broadcasters Squeezed by Convergence Push. Available at http://www.broadcastingcable.com/article/439909-Broadcasters_Squeezed_by_Convergence_Push.php?rssid=20068&q=digital+tv.

¹¹ Motorola, Beyond Mobile Broadband. Available at <http://business.motorola.com/experience/lte-experience.html>

futile endeavor. Within a short time, as new services and applications are created, any definition is likely to become obsolete. However, as we discuss in detail below, the need for a separate class of managed services can be eliminated.

B. Even industry experts fail to reach consensus on a definition for “specialized” or “managed” services.

Our concerns are supported by discussions that occurred during the Open Internet Technical Advisory Process.¹² Professor Jordan participated in these meetings, hosted by the Commission. At the Commission’s request, telecommunication company representatives spent two days “brainstorming”¹³ network management policies from the perspective of end-users, application providers and ISPs. In order to help guide the Commission’s effort to classify managed services, participants looked at historical precedents and examined current traffic prioritization techniques. Ultimately, however, these industry and regulatory experts failed to reach consensus on how to define managed services, as described in this excerpt from the February 25, 2010 meeting ex-parte notice:¹⁴

Attempts were made to define “managed services.” One participant suggested that managed services were those services that “don’t touch the Internet.” This definition was countered by examples. Others suggested managed services reserve bandwidth or receive priority or QOS from the network. Some suggested that only open, managed services with the same capabilities available to all competitors should be supported, though participants noted that, as previously described, scaling to a general offering is difficult and requires the service provider to assume risk. Another participant advanced an argument that all services should be required to be provisioned over the Internet, while others countered that purpose-built service bundles with characteristics such as QOS and bandwidth are required by the market. Another suggested that managed services might be distinguished by whether they are subject to existing regulatory categories such as Title II and Title VI.

¹² Notice of Ex Parte Communication, GN Docket No. 09-191. Letter dated April 8, 2010 and signed by Julius Knapp, Chief of FCC’s Office of Engineering and Technology.

Notice of Ex Parte Communication, GN Docket No. 09-191. Letter dated March 18, 2010 and signed by Julius Knapp, Chief of FCC’s Office of Engineering and Technology.

¹³ *Id.*

¹⁴ *Id.*

If technical experts from leading telecommunications companies including AT&T, Cox Communications, Verizon, Cisco, Intel and Clearwire cannot agree on what offerings qualify as managed services, we question the feasibility of creating such a category at all.

C. Some applications require QoS. However, QoS should be available to these applications regardless of the application provider.

We agree that applications such as telemedicine, distance learning and videoconferencing services require some form of prioritization. In the vast majority of cases, this special treatment means the broadband provider will apply QoS. However, any application that requires QoS should be able to access it without unreasonable discrimination. We strongly believe that consumers of such applications will be best served if there is robust competition between ISPs that offer such applications and new market entrants that lack access to QoS.. This competition, and the innovation accompanying it, will occur only if QoS is available to applications other than those offered by an ISP through “managed” or “specialized” services. This competition will also only occur if QoS is offered end-to-end cooperatively by multiple ISPs, rather than solely within a single ISP’s network as a “managed” or “specialized” service.

Consequently, we conclude that any definition of managed services, no matter how narrow in scope, is unworkable. At the same time, we agree that broadband providers maintain the right to apply network management mechanisms that do not threaten a level playing field, including the use of QoS. In fact, as many emerging Internet applications—such as games and video-conferencing—are time-sensitive, we view QoS as a necessary component of network management.

II. Among the six policy approaches highlighted by the Commission in the *NOI*, the sole viable option would mandate non-exclusivity in specialized services.

In proposing an approach to managed services centered on non-exclusivity, the Commission cites comments previously submitted by Vonage, the Independent Film & Television Alliance (IFTA) and the

Computer & Communications Industry Association (CCIA)¹⁵. As these organizations point out, the Commission must ensure parity between a service providers own applications and competitive offerings transmitted on the network.

A. The goal of regulation is to foster an environment that fosters innovation and ensures a competitive marketplace.

Some high quality services require end-to-end QoS, meaning that they traverse more than one broadband provider's network. However, the proposed rules neither require nor encourage ISPs to peer with each other to offer end-to-end QoS.¹⁶ The lack of availability of end-to-end QoS in the proposed policy approaches may cripple competitors' ability to offer new services in the future. This, of course, is counter to the Commission's goal of promoting competition. In fact, we believe the majority of policy approaches laid out in the *NOI* would create a less competitive marketplace and hamper innovation. Specifically, if managed services are exempt from the rules governing the bulk of Internet traffic, then ISPs may be allowed to implement QoS solely to support their own services, such as voice and video applications. As a result, if subscribers obtain voice or video services from providers other than their broadband provider, they will experience poorer quality. This would completely undermine the Commission's stated objective.

Our concerns are compounded by the reality that nearly all broadband providers are, to varying degrees, vertically integrated. They control proprietary voice and video services, as well as "walled garden" content on their portals. We urge the Commission to mandate that if an ISP uses QoS for its own services or provides QoS

¹⁵ See, e.g., Vonage Comments at 28-29; IFTA Comments at 18; CCIA Reply Comments at 21.

¹⁶ The term "peer" is used here to include both transit arrangements and settlement-free peering. The rules should encourage broadband Internet service providers that have settlement-free peering to incorporate settlement-free honoring of each other's QoS requests. Similarly, the rules should encourage broadband Internet service providers that have a transit for payment relationship to honor each other's QoS requests as part of the transit contract and corresponding payment. With respect to charges, if a broadband Internet access provider uses QoS to support its own services, then it should make available to its residential and business subscribers QoS without unreasonable discrimination. We envision this offering would likely be in the form of a QoS add-on option to a broadband service plan for a small additional fee. The fee would entitle the subscriber to QoS treatment for a specified amount of user-marked packets, and would cover the costs of QoS within the service provider's network as well as any QoS transit charges that the provider may incur.

to a third party, the ISP must make those same services available—without unreasonable discrimination—to any other third parties requesting them. This condition is necessary to ensure equal access to content and services among small competitors and new entrants to the market. A net neutrality policy must prohibit ISPs from making exclusive deals to provide telecommunications services to select application providers, or from offering telecommunications service to competitors at inflated prices.

B. While we do not dismiss out-of-hand the remaining five policy approaches floated by the Commission, we do feel strongly that these approaches fail to go far enough.

We have yet to see a definition of “specialized” or “managed” services that would not create an opportunity for ISPs to evade—directly or indirectly—rules designed to protect against discrimination. Under several of the policy paradigms presented in the *NOI*, broadband providers would have free reign to discriminate on the basis of application, in the name of “reasonable network management.” The first policy approach in the *NOI* calls for “definitional clarity”¹⁷ of what distinguishes broadband access, as it is widely understood, from specialized services. A similar policy suggestion proposes limiting specialized service offerings to services that are not functional via public broadband networks.¹⁸ We reject both of these approaches because they are predicated on the notion of establishing a separate class of specialized services. For reasons previously mentioned, we are convinced that any classification would unacceptably weaken open Internet protections. In addition, by simply protecting best effort traffic, they fail to go far enough.

We laud the potential approach described in the *NOI* that would prohibit ISPs from degrading the performance of their broadband networks in order to accommodate specialized services.¹⁹ This same proposal would also require Internet providers to continue expanding network capacity, regardless of specialized service offerings. Future investment in infrastructure and equitable treatment of network traffic are critical goals.

¹⁷ See *NOI* at 3.

¹⁸ See *NOI* at 4.

¹⁹ See *NOI* at 4

Again, however, we cannot support a plan that prioritizes packets based on application type, or that does not allow the end user to determine which services and applications are subject to QoS.

We champion the two policy approaches floated by the Commission that focus on the need for “truth in advertising” and “disclosure.”²⁰ Currently, broadband providers are not required to disclose network management practices that monitor or interfere with customers’ Internet use. Furthermore, subscribers may assume their transmission speeds are significantly higher than they actually are because ISPs routinely disclose the maximum potential speed offered—but often not the typical broadband speed expected.²¹ Ultimately, however, we cannot accept these requirements as substitutes for ensuring ISPs use QoS on competitive applications and services. Also implicit in these approaches is creation of a managed services class of traffic, which, again, we oppose. Simply protecting best effort traffic fails to accomplish the Commission’s goal of fostering competition and innovation.

III. Rather than creating a managed services category, a simpler and more efficient solution involves requiring an *open interface* to QoS mechanisms in lower layers.

In light of challenges associated with establishing a “managed” class of traffic, we propose a far simpler method to ensure that QoS is widely available without unreasonable discrimination. The idea comes from separating networking functions into two subsets: OSI layers 1-3 (network infrastructure) and OSI layers 4-7 (network applications). Network infrastructure layers exhibit a high barrier-to-entry and hence have minimal providers, while the application layers exhibit a low barrier-to-entry and hence often have numerous providers.

A. The need for a separate class of managed services can be eliminated by applying a layered approach to traffic management practices.

²⁰ See NOI at 3.

²¹ Free Press Calls for ISP Transparency. Available at <http://www.freepress.net/node/45366>.

Net neutrality goals may be achieved by a properly regulated open interface *from information services to communication services*. First, the Commission should reclassify a component of broadband Internet access as a communications service. The Commission should be careful to delineate the communications service component as consisting of *services that reside at or below the network layer or are required to manage the network*. If the Commission does not explicitly define the communications component, it may under-reach by excluding QoS techniques at layer 3 (the network layer), or it may over-reach by including prioritization techniques used at higher layers.

Second, the Commission should require an open interface to the communications component. This interface is already defined by IETF standards. The requirement of an open interface captures the central tenet of a layered Internet architecture, yet is far less intrusive than an open access mandate. An open interface between the networking and transport layer requires only standardization of the *Application Programming Interface* (API) for services offered by the network and lower layers. (By contrast, a standardized wireless air interface encompasses all physical and data link layer protocols and their use.) In previous filings,²² we proposed that any QoS mechanisms that an ISP implements in communications services should be available to application providers and peering ISPs without unreasonable discrimination.²³ Requiring such an open interface would prohibit ISPs from refusing to provide enabling Internet infrastructure services to competing application providers in order to differentiate the ISP's own application offerings. The condition would also bar ISPs from charging competing application providers inflated fees for Internet infrastructure services. Furthermore, an open interface policy would prevent ISPs from making exclusive deals to provide communications services only to certain application providers.

Significantly, this approach would ensure ISPs have the right to apply network management mechanisms that do not threaten a level playing field. Broadband providers could also arrange communications

²² See Scott Jordan, NOI on Preserving the Open Internet, GN Docket No. 09-191 (January 13, 2010)

²³ See Scott Jordan, NOI on Framework for Broadband Internet Service, GN Docket No. 10-127 (July 13, 2010).

services with consumers, application providers, and peering ISPs in a manner that does not conflict with the above goals. Finally, forbearance should be applied to these regulations where sufficient competition exists—such as outside access networks. Access to lower layer Internet services such as QoS should then be guaranteed without unreasonable discrimination to all high layer services and applications. With this refinement, broadband Internet access providers would be allowed to implement QoS below the transport layer to support voice and video services without being required to designate them as managed services. If a broadband provider uses QoS to support its own services, then it would be required to make available QoS without unreasonable discrimination to its own subscribers and to networks with which it peers.

B. The Commission must re-evaluate the layering model currently under consideration.

We note that the Commission’s original *NOI* on a Framework for Broadband Internet Service²⁴ requested comment on whether regulations should take into account the various network layers that compose the Internet. In this context, the Commission proposes distinguishing between physical, logical, and content and application layers. (By contrast, regulations historically differentiate between Internet connectivity functions and Internet applications.) Clearly, we believe in the soundness of this regulatory approach. We stress, however, that the significant boundary lies between the network and transport layers, as opposed to the physical and logical layers or between the logical and content layers. This delineation is parallel with the intent of historical use of the terms “telecommunication services” and “information services.” The network layer and the layers below it do not change the content of packets being sent—similar to telecommunications services. Comparatively, the transport layer and layers above it create, store, or change the presentation of data, similar to information services. A boundary adhering to this description is widely recognized by a host of Internet standards.

²⁴ FCC, *Framework for Broadband Internet Service*, GN Docket No. 10-127, June 17, 2010 at para. 60.

C. Broadband consumers already choose service levels from a menu of tiered services offered by the ISP.

Applying a layered approach to traffic management techniques does not violate open Internet principles and, in fact, is in line with current practices. The Commission does not prohibit carriers from providing customers the option of purchasing a higher quality of service. Under the existing paradigm, ISPs offer customers a choice of broadband plans with varying upload speeds, ranging from about 1.5 Mbps to 50 Mbps. Subscribers willing to pay for the ability to quickly download music and movies sign on to a service level agreement (SLA) that demands a comparatively high monthly fee. Conversely, Internet users interested solely in web browsing and email applications will opt for a less expensive low-bandwidth plan. A key distinction exists between prioritizing applications based on a contract negotiated between the ISP and subscriber, and prioritizing applications based on a managed services approach. Under the current scheme, *end users* themselves determine which applications are subject to QoS. This should be preserved.

IV. Should Commissioners exempt wireless networks from net neutrality requirements, they are simply prolonging the inevitable.

We recognize that the wireless industry is characterized by a great deal of innovation and that mobile consumers are able to choose from a variety of providers—each offering an array of pricing plans and devices. The fact that these business models and handsets exist, however, should not exclude wireless platforms from being subject to the same openness principles the Commission is considering for wired Internet access. In the near future, a majority of Americans will access the Internet from a wireless device. Forty percent of adults already use the Internet, email or instant messaging on a mobile phone (up from the 32 percent of Americans who participated in these activities in 2009), according to the Pew Internet and

American Life Project²⁵. To exclude these Internet users from net neutrality requirements means neglecting to address the needs of about 145 million people who access the Internet using their wireless devices.

A. As with wired applications and devices, technological convergence is erasing traditional distinctions between communication and information services.

As a result of technology convergence, wireless handsets function not only as phones, but also as web browsers, movie screens, music players and game devices. Traditional cellular phone companies now offer video and broadband services. As wireless carriers shift their networks to IP, technological distinctions between companies regulated under Title I and Title II will disappear. At the same time that technology convergence is happening, Americans increasingly rely on mobile devices—from smartphones and e-readers, to netbooks and set-top boxes—to access voice, video and Internet services. The largest Internet browser firm, Opera Software ASA, reported that global data traffic through its mobile browser rose 10 percent in August over the previous month²⁶. According to research firm Bango, Inc., the volume of mobile web browsing from Android phones in the United States increased by 400% during the first three months of 2010. During this same period, the mobile web volume from Apple devices grew about 13%.²⁷ These statistics exemplify the trend toward mobile Internet access. Therefore, for the FCC to exclude wireless carriers from complying with a net neutrality framework is to simply kick the can down the road.

B. From the end-user perspective, the functionality of the mobile web increasingly resembles the wired Internet.

Consumers expect to use the device of their choosing and to download applications, as well as to view the same legal content they access on a PC. Within a short time, the typical consumer may not even

²⁵ Aaron Smith, Mobile Access 2010, Pew Internet and American Life Project. Available at <http://www.pewinternet.org/Reports/2010/Mobile-Access-2010.aspx>.

²⁶ Tarmo Virki, Mobile Web Usage Jumped in August: Opera, Reuters. Available at <http://ca.reuters.com/article/technologyNews/idCATRE68T18P20100930>.

²⁷ Bango, Web Browsing from Android Phones Grows 400% at the Expense of Apple. Available at <http://news.bango.com/2010/08/24/web-browsing-from-android-phones-grows-400-percent/>.

distinguish between “online,” “cell phone,” “telephone,” and “cable.” To users, they’re all forms of communication that can be accessed from a single device. Yet the antiquated structure of the Communications Act allows for differential treatment of these interactions.

C. Reasonable network management requires limiting only traffic, not restricting applications.

Wireless networks require stronger forms of traffic management at or below the network layer, and regulation should ensure a level playing field between providers of various applications (including those that require QoS). Wireless providers currently may limit the applications a wireless device runs, often ostensibly to curb traffic or congestion. However, since an open interface (as discussed above) allows an ISP to regulate traffic without unreasonably discriminating against particular applications, traffic congestion does not justify restricting applications on any device. Segmentation of network applications between managed and unmanaged services, or between Internet and non-Internet offerings, is problematic as technical convergence nullifies such distinctions. As we argued for wireline networks, requiring an *open interface* to QoS mechanisms in lower layers is more practical than banning the use of reservation and prioritization practices, and more efficient than defining *managed services*.

D. Stronger traffic management at or below the network layer does not justify different treatment above the network layer.

Wireless networks often are justified in using different traffic management techniques, but only at the lower layers (at OSI layers 1-3). Protocols that respond to wireless signal variation, that limit interference, that schedule transmissions, and that reserve or prioritize resources, are all in these lower layers. However, the lack of differences in upper layers merits a definition of *reasonable network management* that enables similar competition at the application layer (OSI layers 4-7) independent of the type of network. Wireless networks should address their greater challenges at lower layers either by exercising stronger lower layer traffic

management techniques than used by wired networks or by exercising reasonable techniques used in wired networks but to a greater extent. However, technical differences do not justify wireless networks that implement different traffic management practices above the network layer than those used in wired networks. Just as we argued previously for wired networks, wireless net neutrality can be effectively achieved by a properly regulated open interface *from information services to communications services*. Our approach could apply to either web-based applications or native applications.

This interface-based approach is consistent with current wireless device and operating system architecture. While wireless devices run on a variety of operating systems, they are all built using layered architectures. As in wired networks, common wireless network device operating system APIs offer access to lower layer functionalities at the networking layer or above; they do not offer direct access to protocols at the physical and data link layers. This architecture implies that *open access* to all lower layer protocols is not necessary for net neutrality; all that is required is an *open interface*. Similarly, the architecture implies that lower layer protocols have no need to know what application a packet belongs to; they only need know what type of traffic management to apply. The goal of an open interface should be to allow application access on a level playing field to lower layer QoS.

Comments submitted by a number of wireless carriers²⁸ suggest that the end-user devices in mobile networks are “part of” the network. They contend that vital network control functionality is embodied in the handsets used by the public. From this perspective, there is no demarcation between end user devices and the network. We believe this assertion to be technically inaccurate. Wireless devices include both a wireless radio and/or wireless modem to implement lower layer functionality and other hardware and software to implement higher layer functionality. While the carrier does control functionality of the wireless radio and/or wireless modem, it does not control the other functionality. Furthermore, as long as the wireless radio and/or wireless modem comply with industry standards, the handset poses no threat to the network. Consequently, we believe

²⁸ Comments from AT&T, Verizon, etc.

that it is no more logical to view the entire wireless devices as “part of” the network than it is to view a Wi-Fi enabled laptop as a part of the network.

- E. The open interface represents a contract between the wireless carrier and the user. The wireless carrier offers certain services and the user requests (or purchases) those services at the volume desired.

The central idea is that the wireless carrier will use traffic management techniques to ensure that services are rendered and, hence, contracts are satisfied. Since the services are offered at an interface between the network and transport layers, they cannot be based on the application. However, the carrier likely will offer a range of QoS options. Some of these options may be targeted to support particular applications efficiently, such as high throughput service for file-sharing, traditional best effort service for email and web browsing, bounded delay service for streaming, and guaranteed low delay service for VoIP and video conferencing. However, the user determines which application to transmit over which QoS; if a user wants traditional best effort service for file-sharing, she can elect this option.

Two different pricing schemes in the terms of the contract at the open interface are typical. In the quota form, a user can purchase access to a specified amount of a service at a specified price, such as 5GB/month of best effort service with a peak transmission rate of 1Mbps for \$60/month. In the volume form, a user can purchase access on a per volume basis, such as \$.01/min/(8kbps) for guaranteed low delay service (where 8kbps is chosen to accommodate one VoIP stream). These two pricing plans are often combined. For instance, the carrier may charge \$60/month for 5GB/month, with an additional charge when the user exceeds volume. The most likely deployments consistent with this approach may be user tiering and/or application provider payment. Under user tiering, a user may purchase a higher tier that includes the ability to transmit and/or receive a certain amount of QoS prioritization, such as enough to support up to 1,000 minutes of VoIP prioritization (if that is how the subscriber chooses to use it). Under application provider payment, a VoIP provider other than the user’s ISP may wish to purchase a VoIP prioritization option on behalf of the user (perhaps bundled with the

VoIP service offered to the user); any reasonable payment is consistent with our approach. In contrast, many current plans are not application-agnostic and are hence not consistent with an open interface.

Some plans for smartphones include unlimited amounts of data, but restrict use to certain devices (e.g. prohibit tethering to a laptop) and to certain applications (e.g. permit web browsing and email, but prohibit file sharing, streaming, and VoIP). The goals of traffic management can be more efficiently obtained through an application-agnostic interface that allows users to choose their own applications and to match these applications to QoS options based on price. The form of the contract at the open interface and the prices charged will likely relate to usage of resources in lower layers. In wireless networks the key transmission resources are bandwidth and transmission power. The bandwidth and transmission power allocated to a user or a flow largely determine the low level performance experienced. This low level performance is usually described as a combination of throughput, delay, and packet loss. Various lower layer traffic management practices control the bandwidth and transmission power allocated to each user or flow in order to create the desired variety of QoS service offerings. The price for each QoS service offering thus depends in part on the resources used to achieve that QoS.

Placing the contract at the interface between the network and transport layer in no way unreasonably restricts an ISP's ability to charge additional amounts for services that are not part of OSI layers 1-3. An ISP may continue to subsidize the purchase price of a wireless device and to recoup that subsidy over time. For an additional fee, an ISP may also offer high layer services,, e.g. voice mail, voice dialing, navigation, ring tones, and roadside assistance.

CONCLUSION

We are well-aware of the far-reaching implications that will accompany the Commission's eventual rulemaking on net neutrality. Future opportunities for broadband innovation and competition depend heavily on how the Commission treats managed services, as well as whether Open Internet rules apply to wireless

platforms. Opponents of regulatory reform neglect to acknowledge the challenges certain to accompany increasing Internet traffic and technology convergence. As Commissioner Michael Copps noted in his supporting remarks, “What is reasonable today might be unreasonable tomorrow—and vice versa.”²⁹ Our proposal, based on the distinct functions of the network layer and application layer, offers a regulatory middle-ground. It protects the principles of openness by ensuring that QoS is application and content agnostic, while acknowledging the need for traffic management.

²⁹ Michael Copps, *In the Matter of Preserving the Open Internet*, GN Docket No. 09-191; *Broadband Industry Practices*, WC Docket No. 07-52.