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

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

COMMENTS OF
INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION

Richard Bennett

Information Technology and Innovation Foundation[1]
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 regarding the role of specialized services in the contemporary and future Internets, and the applicability of Open Internet principles to wireless broadband networks. ITIF has conducted in-depth analysis of both issues in a series of reports on Internet evolution and regulation, and has also hosted a number of public fora at which they've been discussed, most recently on October 1st.[2] ITIF has long advocated a Third Way approach to Internet regulation in the context of the net neutrality controversy and is deeply engaged in the matter of regulatory frameworks that recognize the Internet's unique characteristics and value to society.[3]

1. <u>Summary</u>

We summarize our recommendations on each of the two questions separately.

1.1 Managed Services

As presently operated, the Internet is a best-effort, content-oriented network that lacks the Quality of Service provisions needed by wideband communication-oriented applications such as telepresence. Those who desire to use an application such as telepresence therefore have no choice but to bypass the Internet in whole or in part. This form of Internet bypass is known as "Specialized" or "Managed" service.

A mistaken idea suggests that the Quality of Service provisions that enable such applications as telepresence are unnecessary provided that network operators simply increase the capacity of the last-mile networks that service residential users and businesses. Advocates of this approach ignore the fact that the content-oriented applications and protocols consume all available bandwidth on shared facilities by design, while communication-oriented applications moderate their own bandwidth consumption.

---

[2] "Role of Managed Services on Broadband Networks" (presented at the Information Technology and Innovation Foundation, Washington, DC, October 1, 2010), http://itif.org/events/role-managed-services-broadband-networks.

[3] Robert D. Atkinson and Philip J. Weiser, *ITIF: A "Third Way" on Network Neutrality*, report (Washington, DC: Information Technology and Innovation Foundation, May 30, 2006), http://www.itif.org/index.php?id=63.

Internet standards such as Integrated Services (IntServ,) Differentiated Services (DiffServ,) and Multiprotocol Label Switching (MPLS) specify means of coordinating use of shared network facilities both above the IP layer (IntServ and DiffServ) and below it (MPLS). Until such standards as IntServ and DiffServ become operational norms across the Internet, telepresence users have no choice but to employ Internet bypass.

It's reasonable to enact regulations and restrictions on the use of Internet bypass services at the financial level, but not at the technical level. The appropriate body of law circumscribes the conditions on which such services are offered for sale rather than their deployment and use.

Telepresence is not the only application to employ Internet bypass today. Content Delivery Networks such as Akamai and Limelight are bypass services that converge with ISP networks at their Internet-facing edge. The placement of these services enables them to move larger volumes of data across ISP networks that far-away servers can offer, and helps create the requirement for communication-oriented bypass services. This form of Internet bypass is uncontroversial.

The FCC should not ban or limit the use of CDNs, but it should recognize that their widespread use contributes to last-mile ISP congestion that must be alleviated by technical means other than simple capacity increases in the last mile.

1.2 Mobile Broadband

In a perfect world, a set of regulations might be constructed that treated wireline, fixed wireless, and mobile wireless network services the same way. Such regulations would contain sufficient nuance to account for the technical and market differences among these platforms without imposing technology-specific regulations. It would recognize that the requirement for spectrum and the uncertainty of roaming are parallel to constraints on the deployment of optical cables and the variations in end-user application use and it would recognize that the more competitive market for mobile broadband permits more pricing flexibility than the relatively constrained market for wireline broadband does.

We don't live in a perfect world, however, and it's therefore likely that Open Internet regulations will be more restrictive than they should be in the wireline domain and therefore inappropriate in the wireless domain. We've seen this tendency in the proposed regulatory guidelines offered by Verizon and Google that establish a rebuttable presumption against traffic engineering on wireline networks. Well-provisioned wireline networks that move a traffic mix predominated by content function reasonably well with limited management most of the time, and only need to consider traffic engineering for leading-edge applications such as telepresence.

Mobile broadband networks are in a different situation not only due to their spectrum-related bandwidth constraints, but also because their bandwidth costs are higher and the application mix they support is more heavily tilted toward communication than content. Applying a framework such as that proposed by Verizon and Google exclusively to mobile broadband would generally require the opposite presumption on traffic engineering – that it should be permitted unless shown to be specifically harmful – and a process for hearing complaints. Mobile broadband is also more competitive that wireline broadband, especially in the United States where we have four or five national carriers and a number of regional and virtual carriers.

We're also at the dawn of the mobile broadband era and therefore have little insight into the applications that will predominate in this space and the informal and organic means of oversight that might develop if the market is left to its own devices. In his infamous "Wireless Carterfone" paper, Professor Tim Wu complained about "application stall" and asked the FCC to force device and application portability.[4] The market for wireless handsets, operating systems, and applications has developed on its own without such regulation and produced hundreds of thousands of applications despite its "stalled" condition. Forecasting the development of ground-breaking new technologies is more art

[4] Tim Wu, "Wireless Carterfone," *International Journal of Communication* 1, no. 1932-8036/20070389 (2007): 389-426.

than science, and any heavy-handed regulations such as those proposed by Wu run the risk of strangling the innovative potential of mobile broadband technology in the crib.

Therefore, the FCC would be wise to rely on the traditional Four Freedoms supplemented by a transparency requirement than to impose a hard-to-define stricture against traffic engineering in the wireless space or by attempting to enact restrictions on app stores. The development of organic institutions such as the Broadband Internet Technical Advisory Group (BITAG) should also be encouraged.

Table of Contents

2. Specialized Services

The Internet is a federation of privately owned and operated networks that loosely conform to Internet standards. There is a gap between the internal performance and capability of these networks and their external performance and capability through the Internet. By operational convention, the Internet is a best-efforts system. This means that end-to-end performance between any Internet source and any Internet destination is highly variable and statistical. Therefore, no Internet user and no Internet application can expect predictable performance, a bounded guarantee of performance, or even successful communication all the time. Best-efforts networks are ideal for content-oriented applications, and challenging for communications-oriented applications. Hence, the Internet is predominately a content-oriented system. The increased popularity of communication-oriented applications, with their demanding performance requirements, challenges the Internet's operational norms.

The tension between the content-oriented applications of the present and the communication-oriented applications of the future is a major source of the conflict that must be resolved by the FCC's Open Internet initiative. The most pragmatic way for network users and operators to resolve this tension today is to separate communications from content and route communications traffic over network facilities that are designed and operated in accordance with the requirements of communications applications. These separate facilities are often called "Advanced," "Specialized," or "Managed" services. Managed Services are effectively mandated by the Internet's current best-efforts operational norms and can only be eliminated when the entire Internet adopts Quality of Service norms in their place.

2.1 Internet Performance

Internet performance depends on two primary factors: the capacity of the 40,000 Autonomous Systems (networks) that compose the Internet and the behavior of Internet users and applications as a whole. In the typical experience of residential Internet users, round-trip transit times to destinations in North America range from 50 ms. to 150 ms.; to overseas destinations the delay increases to the 200-400 ms. range; Wi-Fi at a network edge adds an additional 25-50 ms. of delay.

The variable performance of the Internet makes it a friendly place only for applications that can thrive within its loose performance parameters. These applications tend to emphasize the distribution of static content, e. g., Web applications and Netflix streaming, or low-bandwidth interpersonal communications such as e-mail and Twitter. The likelihood of success for medium to high bandwidth personal communication such as video Skype and Telepresence is highest over short distances and extremely low over long distances.

The public Internet's performance characteristics are shaped by a number of factors such as the willingness of users to pay for the various service tiers offered by ISPs, inter-ISP competition, user and ISP investment, the proximity of the service area to an Internet Exchange Point, and by the advance of technology generally. As such, the performance of the public Internet at any moment in time is largely a matter of user consensus; it's therefore descriptive to term the public Internet the "Consensus Internet." Some network users have needs that fall outside the bell curve of the consensus, and their needs form the business case for the service tiers designated as Advanced, Managed, or Specialized services.

2.2 The Internet's Content Bias

One clear difference between the Consensus Internet and advanced networks is the relative importance of content and communication in the two realms: the Consensus Internet is content-oriented system that competes with cable TV and publishing as a largely one-way system with limited interactivity, but advanced networks are vehicles for interpersonal communication, telepresence, virtual reality, and other forms of highly interactive, high-bandwidth communication.

The difference between these two networks illustrates science fiction writer William Gibson's aphorism "The future is already here – it's just not very evenly distributed." The Consensus Internet, which dominates the market for residential Internet services, is the Internet of the past, and the advanced networks more commonly used by businesses and universities are the Internet of the future. Internet 2 is an advanced network, even though it connects to the Consensus Internet. If history is any guide, the advanced networks of today will become the Consensus Internet of tomorrow. Prior to the decommissioning of NSFNET backbone in 1995, the entire Internet was "advanced."

Research tells us that humans place more value on communication than on content; according to Andrew Odlyzko's seminal 2001 essay, "Content is not king" communication is.[5] Communication-oriented networking serves the needs of those users who are further along the technology curve than their neighbors, more sophisticated in their abilities to use technology, and more willing to pay for advanced services. Advanced network services define a target toward which the Consensus Internet is likely to develop.

The nature of advanced networks also provides a target for policy makers concerned with enabling the Consensus Internet to become a more capable system. Unlike the Consensus Internet, advanced networks are not "best-efforts" systems characterized by "all you can eat" pricing plans. The pricing of advanced networking facilities is governed by Service Level Agreements between network operators and users, and their operation is governed by the careful use of traffic engineering by operators and the explicit selection of desired service levels by users.

2.3 The Need for Traffic Engineering

Traffic engineering is an imperative when a network such as the Internet supports a mixture of content and communications applications; simply adding bandwidth is not a substitute for traffic engineering, as content applications, by design, do not moderate their bandwidth consumption and will always take the lion's share of any capacity increase. Content applications use TCP, which relies on the network to signal congestion to the endpoint, and takes no responsibility for network congestion other than responding to network signals (packet loss) by temporarily reducing its offered load rate. TCP never stabilizes its bandwidth demands; it constantly seeks higher throughput. TCP's constantly cycling between low and high consumption is the worst possible background load scenario for communication applications that seek to adapt their data formats and coding

[5] Andrew Odlyzko, "Content is not king" (AT&T Labs - Research, January 3, 2001), http://www.dtc.umn.edu/~odlyzko/doc/history.communications2.pdf.

to network conditions. Internet bypass is unfortunately a very pragmatic response to the desire to employ advanced communications-oriented applications and the relative absence of effective traffic engineering on the Consensus Internet; it often takes the form of Advanced, Specialized, or Differentiated Services.

Consequently, public policy cannot circumscribe the use of Specialized Services without harming the Consensus Internet's capacity for improvement. This is not to say that Specialized Services that converge on the public Internet at the first and last mile must be forever free of regulatory scrutiny, nor to say that providers of first and last mile Internet services must be allowed to treat bypass networks, private networks, and specialized networks in any way they see fit.

The Internet is a system of private network interconnection; for purposes of this inquiry, Specialized Services are private networks. The public policy nexus between the public Consensus Internet and the private networks enabled by Specialized Services exists on both sides of the interconnection point between public and private networks. Advocates of net neutrality fear the rise of the Specialized Services network, but they fail to realize that the Internet is now and has been a collection of Specialized Services networks of a kind, engineered to optimize content delivery, since the decommissioning of the NSFNET backbone.

2.4 Enabling Communication-Oriented Applications

The questions about Specialized Services currently on the table relate to the question of adding additional forms of specialization to the Internet ecosystem to enhance the effectiveness of communication-oriented applications. The Consensus Internet is heavily biased toward content delivery, and this bias has significant innovation-limiting effects that public policy must ameliorate.

It's doubtful that there is any desire on the part of network operators to use Specialized Services as means of providing "prioritized access to a particular website" as the Further

Inquiry suggests.[6] To do so would be unproductive. Web site owners who seek "prioritized access" obtain it from Content Delivery Network operators such as Akamai, as Netflix has done, by replicating content around the Internet at locations close to end users, as Google has done, or simply by purchasing high performance servers and high bandwidth circuits, as all owners of high volume web sites have done. ISP-provided Specialized Services are desired for a completely different purpose: Specialized Services make communications-oriented applications that *cannot be accelerated* by CDNs or other replication strategies work successfully.

At a high level, ISP-provided Specialized Services enable communications-oriented applications by controlling inter-packet gaps. To understand how this works, we need only consider the case of a single content application and a single communication application sharing a circuit. Each application is affected by the volume of packets offered by the other application, and the communication application is additionally affected by the pattern of the other application's packet delivery. The communication application (comm app) controls the rate at which it offers packets to the network, but the content application (content app) offers packets as fast as the network and the receiving system allow. The content app is constrained by Jacobson's Algorithm in TCP to offer packets inside the TCP receive window, so it tends to offer packets in clumps. Content apps move information in large packets, while comm apps tend to group information in smaller packets.

Graphically, the content app's traffic looks something like this:

[6] Federal Communications Commission, "FURTHER INQUIRY INTO TWO UNDER-DEVELOPED ISSUES IN THE OPEN INTERNET PROCEEDING," Public Notice (Washington, DC, September 1, 2010), at 2.

While the comm app's traffic looks more like this:

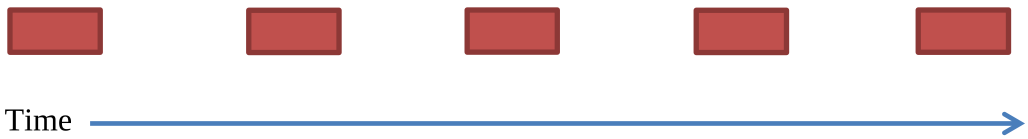

In this example, prioritizing the comm app's packets relative to the content app's packets doesn't substantially affect the performance of the content app. The content app's overall performance is determined by the delivery time of its *last* packet, which is determined by the total number of packets it transmits and the total load on the network during the transmission of its packet stream. The performance of the comm app is determined by the *gaps* in the delivery times of each of its packets, which is determined by the clumping of other traffic on the network between its packets.

A best-efforts network might combine these two packet streams in the following way:

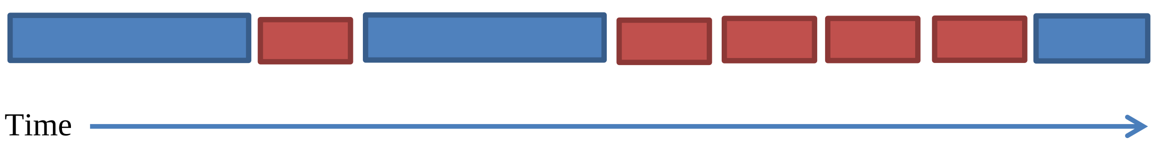

While a more optimal combination would look more like this:

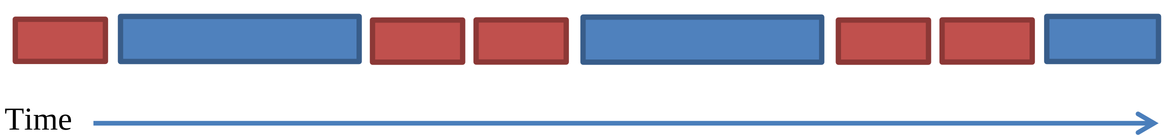

In order to co-exist with the content app, the comm app needs the network to schedule packets in such a way as to minimize gap variations, a service that the Consensus Internet does not provide due to its content bias. It therefore turns to a Specialized Services network to obtain the performance it requires, even though any network that conforms to the IETF standards for Differentiated Services or Integrated Services could meet its needs. Ironically, comm apps need to depart from the public Internet to obtain the full range of services defined by Internet standards.

Questions (E) and (F) in section I. of the Further Inquiry suggest a number of limitations on Specialized Services on the apparent assumption that Specialized Services limit the utility and performance of the Consensus Internet. Question (E) suggests limiting the types of Specialized Services to a narrow group, and Question (F) suggests the completely impractical step of guaranteeing the performance of the Consensus Internet. These suggestions are unproductive.

In fact, every user of the Internet degrades the performance of every other concurrent user of shared facilities in proportion to the load that he or she offers to the network; there is nothing unique about Specialized Services in this respect. In fact, the most significant performance degradation that the current residential Internet user experiences today is caused by neighbors streaming movies from a local CDN. CDNs bypass the Internet core and present more load to the last mile network than conventional uses do. Performance degradation is caused by the sheer volume of concurrent traffic, not by measures that limit inter-packet gaps.

The Commission's concern that investment in a parallel network of Specialized Services will starve the Consensus Internet is understandable, but probably misguided. The faulty analysis that reduces all applications to content and the corresponding blindness to the requirements of communication-oriented applications is the most substantial barrier to the successful resolution of questions about Internet services regulation. Specialized Services are necessary for applications that require better than best-efforts transport, and no restriction should be placed on their sale or use for such applications other than restrictions against economically discriminatory sale or exclusive sale. The problem of unfair discrimination in network services is economic, not technical.

3. Mobile Broadband

The broadband networks of the 20^{th} century were exclusively fixed location systems, but the broadband networks of the 21^{st} century will be increasingly (and probably predominately) mobile. This is an easy prediction to make because the uptake of mobile technology has been so much faster than the uptake of fixed location broadband: there are currently two billion Internet users worldwide, and five billion cellular users. Many cellular users never had any kind of electronic communication capability before they acquired their cell phones, and will never use a fixed location service in the future. Mobile broadband is in its infancy, and it's impossible to tell which range of applications and services will predominate once it matures. It's not unlikely that mobile broadband will develop along communication-oriented lines radically different from the content orientation that characterizes fixed broadband. Applications such as Mobile Augmented Reality, machine-to-machine communication, location-based services, and voice, text, and video communication will feature strongly in this space. Policy analysts frequently observe that mobile broadband is not a perfect substitute for fixed broadband, but fixed broadband is no substitute at all for mobile because you can't take it with you.

Wireless bandwidth is more expensive than wireline bandwidth. While both systems rely on wireline backhaul, for the most part, the mobile device-facing edge of wireless networks is more expensive to provision and to improve than is the wireline-facing edge. This is predictable from the Signal to Noise Ratio component of Shannon's Law. Cooper's Law predicts doubling of wireless capacity every 30 months, while Butter's Law predicts that fiber optic capacity doubles every eight months.[7] In the absence of increased coding efficiency, wireless operators can only increase capacity by using more spectrum or by deploying antennas more densely. Antenna deployment is labor-intensive, while wireline bandwidth increases are often accomplished simply by upgrading the electronics on existing facilities in the course of routine maintenance.

Despite the continuous output of standards bodies such as the IETF, Internet operational norms have remained largely stagnant since the de-commissioning of the NSFNET backbone. With the advent of LTE with its all-IP orientation, Internet standards for

[7] Richard Bennett, *Going Mobile: Technology and Policy Issues in the Mobile Internet* (Washington, DC: Information Technology and Innovation Foundation, March 2010), http://itif.org/publications/going-mobile-technology-and-policy-issues-mobile-internet, at 2.

Quality of Service such as RSVP, IntServ, and DiffServ are likely to be deployed on a large, inter-domain scale for the first time. The deployment of these technologies will come at a fee to end users, as envisioned by the authors of Internet RFC 2475, and not as a free networking service in which operators move packets out of pure altruism as Jonathan Zittrain has claimed they do today.[8]

Inter-domain QoS has never been employed over the Internet on a large scale. Therefore, we can't readily predict what sort of operational norms will be required to conform day-to-day network operations to IETF QoS standards, which standards will prove most valuable, and what sorts of new standards will be required in order to make such a system operate successfully. If history is any guide, some operators will be forced to employ temporary expedients from time to time, as they currently to in order to address Denial of Service attacks, and some operators will blunder, as they often have. Operators will certainly compete with each other on the basis of service plans, with some offering large quantities of best-efforts service alongside smaller quantities of better-than-best-efforts delivery, or vice versa. Some operators may choose to remain voice-oriented, and others will offer predominately data-oriented best-efforts services. The first 4G service offerings are data-only USB Aircards and we don't know what the second wave will bring.[9]

Fortunately for consumers, the healthy, competitive market for wireless services is becoming even more competitive with the advent of the FCC's White Spaces order, the deployment of auctioned 700 MHz services, and the 4G rollout. As we explained in our

---

[8] S. Blake et al., "RFC 2475 - An Architecture for Differentiated Services," Internet RFC, December 1998, http://tools.ietf.org/rfc/rfc2475.txt says: "Service differentiation is desired to accommodate heterogeneous application requirements and user expectations, and to permit differentiated pricing of Internet service." Jonathan Zittrain, "Connected Consequences" (presented at the TED Conference, Oxford, England, July 22, 2009), http://conferences.ted.com/TEDGlobal2009/program/guide.php, said that Internet packet routing as an uncompensated "neighborly duty." The RFC is right and Zittrain is wrong.

[9] Wayne Rash, "Verizon Wireless Confirms Broad U.S. LTE Rollout for 2010," *eWeek*, October 6, 2010, http://www.eweek.com/c/a/Mobile-and-Wireless/Verizon-Wireless-Confirms-Broad-US-LTE-Rollout-for-2010-652205/. "McAdam declined to be specific about which devices will be the first to appear with Verizon Wireless' LTE, except to say that the first products this year will be USB Aircards."

analysis of the FCC's most recent Mobile Competition Report, the United States has the most competitive market for wireless services in the world, and it's getting better.[10]

The mobile broadband ecosystem therefore differs from the fixed broadband ecosystem in four significant respects:

1. The mobile application mix is heavily weighted toward latency-sensitive communication applications.
2. Bandwidth is much more expensive in the mobile space.
3. Mobile broadband is in its infancy while fixed is relatively mature.
4. Mobile broadband is more competitive than fixed broadband.

3.1 Transparency Rules

The Further Inquiry asks for guidance on transparency rules. We believe transparency rules are important, and agree with Verizon and Google that they provide most of the needed protections for consumers at the dawn of the 4G era. Transparency rules are of two types:

1. Consumer-oriented rules that ensure consumers are informed regarding the fees they pay and the services these fees entitle them to receive. Where wireless services are organized into QoS tiers, as they have traditionally been, users are entitled to know the limits in each tier. This distinction has traditionally been captured as minutes of voice and megabytes of packets. It can also be conveyed as megabytes of high-priority packets and megabytes of low-priority packets, for example.
2. Developer-oriented rules that mandate fuller, more detailed, but not contractually binding information about the expected performance of the network. Our

[10] Richard Bennett, "Looking for Trouble: The FCC's Mobile Competition Report" (Information Technology and Innovation Foundation, June 2, 2010), http://www.itif.org/files/2010-looking-for-trouble.pdf.

comments to the FCC on wireless network measurement address specifics and the BITAG can be expected to provide more.[11]

3.2 Device Portability Mandate

The Further Inquiry asks if it should impose a device portability mandate, to which we must ask in return: Why? Given that wireless networks will be changing rapidly over the next five years as they upgrade to LTE, any device portability mandate will require network operators to provide, at a minimum, a certification program to identify harmful devices.

Unlike the FCC's familiar PSTN domain, the mobile wireless world recognizes degrees of harm rather than a binary, "Harms the Network – Doesn't Harm the Network" dichotomy. Networks advance by incorporating improvements in coding and modulation, which are ineffective unless supported by devices.

Older devices can still operate on advanced networks, but in a substandard manner that compromises network performance and efficiency. The presence of inefficient devices reduces the quality of experience for the neighbors, so at what point is inefficiency deemed "harmful?" The "harm" standard is typical PSTN-era formulation which is nonsensical when applied to advanced technologies.

Moreover, the device portability mandate assumes that consumers have some burning desire to retain obsolete devices when the marketing evidence suggest they're eager to adopt the latest technology to obtain the benefits it provides. Consumer advocates complain about two year contracts for subsidized devices, and with some reason. By some estimate, the typical American smartphone consumer upgrades every 18 months. Fortunately, most network operators provide a limited upgrade subsidy after 12 months to accommodate this wish.

[11] Richard Bennett, "Measuring Wireless Performance" (Information Technology and Innovation Foundation, July 29, 2010), http://www.itif.org/files/ITIF-Reply-Mobile-Broadband-Measurement.pdf.

Attempts at exploring the market for unsubsidized, network portable devices in the United States have typically failed to produce the desired result. The Google Nexus One was announced as a portable, unsubsidized, direct-to-consumer offering in January, and sales were so disappointing that Google cancelled web sales of the device by May in favor of traditional network operator sales, subsidy, and support.[12] One of the greatest sources of consumer frustration with network portable handsets is the issue of service and support. The consumer calls the network operator when he has a problem, but the operator can't support a device she doesn't understand. The device distributor is clueless about network conditions, and can't provide much help to the consumer either.

As a practical matter, the downsides of device portability outweigh the upsides. Compromise positions are possible for up-to-date devices, but they require manageability mandates in devices that don't presently exist. For a device to be portable without causing harm or degradation, in other words, the device must be diagnosable and controllable by the network operator. Old and essentially obsolete devices should certainly not be allowed to move across networks unless the operator is willing to accept the consequences (see consumer transparency requirements above.)

3.3 App Store Regulation

Finally, the Further Inquiry seeks guidance with respect to app store policies and application limits generally. Current market share data indicates that the leading smartphone platforms by operating system are Google Android, RIM Blackberry, and Apple iPhone, respectively.[13] Operating systems are significant in this context as the only significant app stores are operated by OS vendors, not by network operators or anyone else. Most complaints about app store policies have been directed against Apple, the number three platform vendor, because of their insistence on application quality and non-

[12] Official Google Blog, "Nexus One changes in availability," blog, May 14, 2010, http://googleblog.blogspot.com/2010/05/nexus-one-changes-in-availability.html; "While the global adoption of the Android platform has exceeded our expectations, the web store has not."

[13] Andrew Hachman, "Nielsen: Android Is Most Popular Smartphone OS," *PC Magazine.com*, October 5, 2010, http://www.pcmag.com/article2/0,2817,2370259,00.asp.

competition with the iPhone's core functions.[14] Apple's policies are consistent with a business model that aims to capture a high-end market segment more concerned with quality, consistent user experience, and protection from malicious applications than with "openness" and low price.

The FCC's jurisdiction over OS-sponsored app stores is certainly questionable, but if there is some sort of jurisdictional nexus between "communication by wire and radio" and app stores curated by vendors of operating systems for mobile handsets we have to wonder on what basis the Commission should be concerned about the policies of the market's number three player.

There is certainly an anomaly here, insofar as the app store that serves the smallest market and has the most restrictive policies has fostered the greatest degree of innovation, and has done so by an enormous margin. By most estimates, there are ten times as many iPhone apps as Android apps, after all. The Android app store is extremely open, and Android users are able to download apps outside the official Android store altogether if they wish. So what sense would there be in forcing the app store that has been far and away the market leader in innovation to adopt the policies of a firm that is not leading the application market and that has so far failed to attract more than a fraction of the former's creative activity? One might suggest, somewhat humorously, that the FCC should consider requiring the Android store to adopt Apple's more restrictive policies, with their track record of proven success, in order to become a more able competitor. We don't suggest that, however; the Android store is doing quite well with its own policies, of course.

It appears that the Commission's primary concerns are with overly-restrictive app store policies couched in the guise of protecting networks from overload and abuse that are

---

[14] See for example: Eliot Van Buskirk, "Google Voice Web App Circumvents Apple's Blockade," *Epicenter*, January 26, 2010, http://www.wired.com/epicenter/2010/01/google-voice-web-app-circumvents-apples-blockade/.

actually stalking horses for network operator desires to retain exclusive control over the use of their networks for high-value and high-traffic applications such as YouTube video streaming. This concern echoes the first of the two questions posed by the Further Inquiry, the Specialized Services matter, and goes to a deep public policy question: will the mobile broadband networks of the future be used primarily to serve low-value, advertising supported video content, or will they be used for two-way communication purposes such as video calling? It may be quite some time before these networks have sufficient capacity to handle both applications concurrently, so how will the choice be made in the interim? And perhaps most significantly, how can public policy accelerate progress toward the condition in which consumers can use both types of applications without restriction at all times?

We endeavored to address these questions – the ones that we believe the FCC should be asking – in our report “Going Mobile: Technology and Policy Issues in the Mobile Internet.”[15] Network technology provides efficient means of streaming one-way television over the air. Valuable and constrained cellular bandwidth is best used, in our estimation, by applications that are not amenable to broadcast and multicast distribution, such as Mobile Augmented Reality, navigation, location-based services, and interpersonal communication. We’re sufficiently confident that communication-oriented applications are more valuable to consumers than advertising-supported television that we’re happy for consumers to make the choice.

As Odlyzko points out, public policy has long sought to impose a content bias on networks, and network users have long rebelled. As long ago as 1832, when newspapers were subsidized on the postal network by letters, “newspapers generated no more than 15 percent of total postal revenues, while making up as much as 95 percent of the weight.”[16] Today, the heaviest users of networks are video consumers, many of them engaging in

---

[15] Bennett, *Going Mobile: Technology and Policy Issues in the Mobile Internet*.
[16] Richard John, *Spreading the News : The American Postal System from Franklin to Morse*, 1st ed. (Cambridge, Massachusetts: Harvard University Press, 1998).

unlawful transactions by streaming or downloading copyrighted material without permission of the copyright holder. If network operators are allowed to price network usage as they see fit in the highly competitive market for mobile broadband, it's likely that consumers will express a preference for communication and the engineering effort to create low latency, all-IP LTE networks for communication purposes will not have been in vain.

Whether we're right or wrong about consumer preference, we feel that the dawn of the LTE era is the wrong time for regulators to put a finger on the scale and attempt to predetermine the application mix that will prove most valuable once the LTE rollout is complete. Regulators should not choose between content and communication, they should only insist on high-function, high-efficiency networks capable of supporting a wide range of applications, and on terms of use that unlock their potential. In the LTE space, as in the wireline broadband space, the most urgent current need is for networks to more completely implement the backlog of QoS-oriented IETF standards that have been log-jammed since the Border Gateway Protocol (BGP) took the place of the NSFNET backbone. A more complete Internet will allow consumers to choose from a wider menu of services.

Many of the questions raised by app stores can also be handled by organic Internet institutions such as the BITAG, and that path should be encouraged.