

August 20, 2013

Dear Chairwoman Clyburn, Commissioner Rosenworcel, and Commissioner Pai:

We are pleased to attach a report of the year's activities of the FCC's Open Internet Advisory Committee, created to advise the Commission how to enforce, reflect upon, and improve its Open Internet Report and Order, approved in 2010.<sup>1</sup>

Given our diverse membership, and the correspondingly broad set of viewpoints and interests represented, we knew that achieving consensus on concrete changes to the Open Internet Report and Order would be a tall order. In order to delve into real issues, the Committee sought to clearly articulate viewpoints where judgments diverged, and to help flesh out some of the more loaded terms in the OIO, such as "specialized services," which underlies one of the exceptions to the rules for wireline service providers.

Accordingly, the documents produced by each of our Committee working groups are best understood as attempts to lay out a useful spectrum of opinions associated with particular stakeholders, rather than to come to clear conclusions about next steps. Our work also makes note of areas in which more research or information-gathering by outside parties or Commission staff would be helpful.

The Committee's work was undertaken through four working groups which met by teleconference and through e-mail lists, as well as in-person meetings over the course of the year in Washington, DC; Cambridge; Palo Alto; and Chicago. These gatherings included meetings of the full Committee, made available to the public on location and by webcast.

We thank all of the committee members and the FCC staff who devoted time to producing this report and the work it describes. We hope it will help define the landscape in which the OIO is taking place, informing judgments in this space for the months and years to come.

It is the consensus of the Committee to seek your feedback on this work, with an eye towards a constructive agenda and priorities for the next year, before we undertake further major work.

Sincerely,

Jonathan Zittrain  
Open Internet Advisory Committee Chair

David Clark  
Open Internet Advisory Committee Vice-Chair

---

<sup>1</sup> "The Committee, to be created in consultation with the General Services Administration pursuant to the Federal Advisory Committee Act, will be an inclusive and transparent body that will hold public meetings. It will be comprised of a balanced group including consumer advocates; Internet engineering experts; content, application, and service providers; network equipment and enduser-device manufacturers and suppliers; investors; broadband service providers; and other parties the Commission may deem appropriate. The Committee will aid the Commission in tracking developments with respect to the freedom and openness of the Internet, in particular with respect to issues discussed in this Order, including technical standards and issues relating to mobile broadband and specialized services."

# Open Internet Advisory Committee 2013 Annual Report

Open Internet Advisory Committee  
Federal Communications Commission

Released August 20, 2013

# Open Internet Advisory Committee - 2013 Annual Report

## **Members of the Open Internet Advisory Committee**

Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla

Brad Burnham, Founding Partner, Union Square Ventures

David Clark, Senior Research Scientist, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (Committee Vice-Chair)

Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology

Leslie Daigle, Chief Internet Technology Officer, Internet Society

Maurice Dean, Director, Open Connect Product Management, Netflix Inc.

Jessica Gonzalez, Executive Board, Media and Democracy Coalition; Vice President for Policy & Legal Affairs, National Hispanic Media Coalition (representing NHMC)

Shane Greenstein, Professor and Kellogg Chair of Information Technology, Kellogg School of Management, Northwestern University

Russell Housley, Chair, Internet Engineering Task Force; Founder of Vigil Security, LLC (representing Vigil Security, LLC)

Neil Hunt, Chief Product Officer, Netflix

Charles Kalmanek, Vice President of Research, AT&T

Matthew Larsen, CEO, Vistabeam

Kevin McElearney, Senior Vice President for Network Engineering, Comcast

Marc Morial, President & CEO, National Urban League

Matt Murphy, Senior Vice President, Digital Video Distribution, ESPN Media Networks, Walt Disney Company

Elaine Paul, Senior Vice President, Strategic Planning, Walt Disney Company

Jennifer Rexford, Professor of Computer Science, Princeton University

Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)

Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems

Charles Slocum, Assistant Executive Director, Writers Guild of America, West

Marcus Weldon, Chief Technology Officer, Alcatel-Lucent

Michelle Zatlyn, Co-Founder & Head of User Experience, CloudFlare

Jonathan Zittrain, Professor of Law and Computer Science and Co-Founder of the Berkman Center for Internet and Society, Harvard University (Committee Chair)

## **Designated Federal Officer**

Tejas Narechania, Special Counsel, Federal Communications Commission

## Table of Contents

- Working Group Membership.....Page 4
- Executive Summaries.....Page 6
- Economic Impacts of Open Internet Frameworks Working Group.....Page 14
  - Policy Issues in Data Caps and Usage-Based Pricing.....Page 14
- Mobile Ecosystem Working Group.....Page 39
  - AT&T FaceTime Case Study.....Page 39
  - Openness in the Mobile Broadband Ecosystem.....Page 47
- Specialized Services Working Group.....Page 66
  - Summary of Findings and Conclusions.....Page 66
  - Appendix 1: Case Study of IPTV.....Page 72
  - Appendix 2: Specialized Services Case Study.....Page 77
- Transparency Working Group.....Page 82
  - Open Internet Label Study.....Page 82
- Appendix: Committee Member Contributions.....Page 89

## Working Group Membership

- **Economic Impact of Open Internet Frameworks**
  - Chair: Shane Greenstein, Professor and Kellogg Chair of Information Technology, Kellogg School of Management, Northwestern University
  - Brad Burnham, Founding Partner, Union Square Ventures
  - Maurice Dean, Director, Open Connect Product Management, Netflix Inc.
  - Neil Hunt, Chief Product Officer, Netflix
  - Kevin McElearney, Senior Vice President for Network Engineering, Comcast
  - Marc Morial, President & CEO, National Urban League
  - Matt Murphy, Senior Vice President, Digital Video Distribution, ESPN Media Networks, Walt Disney Company
  - Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)
  - Charles Slocum, Assistant Executive Director, Writers Guild of America, West
  - Jonathan Zittrain, Professor of Law and Computer Science and Co-Founder of the Berkman Center for Internet and Society, Harvard University (*ex officio*)
  - David Clark, Senior Research Scientist, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (*ex officio*)
- **Mobile Broadband**
  - Chair: Jennifer Rexford, Professor of Computer Science, Princeton University
  - Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla
  - Brad Burnham, Founding Partner, Union Square Ventures
  - Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
  - Charles Kalmanek, Vice President of Research, AT&T
  - Matthew Larsen, CEO, Vistabeam
  - Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)
  - Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
  - Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
  - Jonathan Zittrain, Professor of Law and Computer Science and Co-Founder of the Berkman Center for Internet and Society, Harvard University (*ex officio*)
  - David Clark, Senior Research Scientist, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (*ex officio*)
- **Specialized Services**
  - Chair: David Clark, Senior Research Scientist, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory
  - Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
  - Leslie Daigle, Chief Internet Technology Officer, Internet Society
  - Maurice Dean, Director, Open Connect Product Management, Netflix Inc.
  - Neil Hunt, Chief Product Officer, Netflix
  - Kevin McElearney, Senior Vice President for Network Engineering, Comcast
  - Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
  - Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
  - Jonathan Zittrain, Professor of Law and Computer Science and Co-Founder of the Berkman Center for Internet and Society, Harvard University (*ex officio*)
- **Transparency**
  - Chair: Russell Housley, Chair, Internet Engineering Task Force; Founder of Vigil Security, LLC (representing Vigil Security, LLC)
  - Leslie Daigle, Chief Internet Technology Officer, Internet Society
  - Jessica Gonzalez, Executive Board, Media and Democracy Coalition; Vice President for Policy & Legal Affairs, National Hispanic Media Coalition (representing NHMC)
  - Charles Kalmanek, Vice President of Research, AT&T
  - Matthew Larsen, CEO, Vistabeam

## Open Internet Advisory Committee - 2013 Annual Report

- Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
- Michelle Zatlyn, Co-Founder & Head of User Experience, CloudFlare
- Jonathan Zittrain, Professor of Law and Computer Science and Co-Founder of the Berkman Center for Internet and Society, Harvard University (*ex officio*)
- David Clark, Senior Research Scientist, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (*ex officio*)

## Executive Summaries

### **1. Data Caps Report (Economic Impacts working group)**

The report aims to analyze data caps in the context of the Open Internet Report and Order. The Open Internet Report and Order discusses usage-based pricing (UBP), but does not expressly mention data caps except by implication in that data caps can be considered a form of UBP. The Order left open the possibility of many experiments in business models and pricing. Moreover, the Internet had evolved over time, and the Order anticipated that the Internet would continue to evolve in unexpected ways. The Order set up the advisory group to consider whether aspects of the Order remain consistent in its effects on the Internet as the Internet evolves, and it is in that spirit that this conversation was undertaken.

The report seeks to clarify relevant terminology (e.g., cap, UBP, thresholds), identify a common fact-basis for discussion, analyze different perspectives, and identify unaddressed open questions.

The Report concludes that there is considerable variance and experimentation in the market by ISPs. It is difficult to interpret even the highest thresholds in the situations in which they arise, as there is no definitive public source on household usage per month to use as a benchmark. In addition, usage varies depending on ISP and technology. All public measurements show great skew in usage, and suggest that caps do not yet impact users other than the highest users.

The committee could reach only tentative conclusions. Although caps do not seem to be affecting a large number of US users now, the situation may change in the future, as user habits, supplier experimentation, vendor policy, and applications all change.

The report also elaborates on many of the key concerns of three stake-holders prominently identified in the Order, namely, users, broadband providers, and edge-providers.

The discussion about users focuses on user understanding about perceptions of caps and thresholds. The report concludes that this topic may require future monitoring, especially given the importance of consumer education to user perceptions of caps and thresholds. It is not yet apparent whether the issues in this topic are a transitory or permanent concern. The experience of ISPs with providing customers with tools to monitor or control data usage could also be valuable to insights about the perceptions of caps by consumers.

The discussion about broadband providers focuses on many divergent perspectives: whether data caps, tiers and related forms of UBP may encourage end users nearing that cap to act efficiently; whether data caps, tiers and related forms of UBP may spur efficiency and innovation on the delivery of services; whether data caps, tiers and related forms of UBP may help manage network growth; whether data caps, tiers and related forms of UBP might encourage heavy users to change their usage, and if so, in what way; whether data

caps may shape the future and conduct of other service providers (i.e. application developers).

The discussion about edge-providers considers how data caps, tiers and UBP can shape other providers of services in broadband ecosystem, e.g., entrepreneurs who provide applications, build web pages, and operate other services in the cloud. This part of the report identifies areas where ISPs and edge providers have different perspectives on open questions. It also examines competition policy for specialized services, recognizing that this topic is also covered by other working groups. In general, competition policy is concerned about situations where one firm, such as a broadband provider, supplies a service and also controls aspects affiliated with the cost, performance, and user-experience in a competing service, provided by an edge-provider. The report identifies how the ISP's perspective and the edge provider's perspective diverge on this topic. The report concludes the situation yields no easy answers in general, and, at a minimum, merits further monitoring.

In general, the committee concluded that these debates cannot be easily summarized in a brief set of bullets or summary paragraphs. The report contains many perspectives, as well as many open questions, and it identifies many issues that the FCC could further monitor.

## **2. FaceTime Case Study (Mobile working group)**

Mobile broadband networks and traditional fixed networks are treated differently in the Open Internet Report and Order. Mobile broadband providers can, more easily than fixed providers, (1) block devices and applications which do not compete with voice or video telephony services of those providers and (2) discriminate in traffic service. Under certain circumstances, this differential treatment might obstruct a free and open Internet, which is why the Mobile broadband working group of the Open Internet Advisory Committee (OIAC) decided to investigate it through a case study. The working group looked into how AT&T restricted the cellular data usage of Apple's FaceTime application to only AT&T customers who used the "MobileShare" plans (instead of "unlimited" data plans). AT&T had disagreed with claims that it had violated the FCC's Open Internet Report and Order. In October 2012, during the working group's work, AT&T agreed on its own accord to support FaceTime on all of its tiered data plans.

The case study raises the following points:

- 1.) Pre-loaded applications, such as FaceTime, are more readily adopted than downloadable applications.
- 2.) FaceTime appears to have been designed in a way that generates a substantial amount of traffic and consumes more bandwidth than comparable applications (e.g., Skype), raising questions about whether FaceTime could feasibly adapt to congestion like other comparable applications.
- 3.) Restricting application usage to customers of a particular data subscription could actually, for the *benefit* of an open Internet, limit the number of users in an initial deployment of a new application, and limit the total amount of traffic.

4.) It is important to determine, in advance, *where* an application-management decision should be enforced and *who* should enforce such decisions (i.e. currently, a smart phone can block users from running an application).

The working group came to different opinions about AT&T's restriction of FaceTime usage on its network. Overall, the group agreed that blocking applications can discourage innovation, but that carriers should also have the freedom to manage their limited cellular network resources. More specifically, three main opinions emerged:

- 1.) Blocking an application from some users under a certain pricing plan could stifle the vibrancy of the mobile application market.
- 2.) AT&T's approach of permitting FaceTime on either Wi-Fi or within shared data plans was a logical way of managing network congestion.
- 3.) Encoding video frames at lower bit rates and adapting to changing network conditions (which Skype, unlike FaceTime, was capable of doing) is central to the use of video or voice calling applications.

### **3. Openness in the Mobile Broadband Ecosystem (Mobile Broadband working group)**

This report analyzes how different actors in the mobile broadband ecosystem have each influenced Internet openness—as well as each other. These actors, not all of whom are subject to the Open Internet Report and Order, include:

- 1.) Mobile broadband providers (e.g. Verizon, AT&T, Spring, and T-Mobile);
- 2.) Device vendors (e.g. Apple and Samsung);
- 3.) Operating system developers (e.g. Apple iOS and Google Android);
- 4.) Network equipment vendors (e.g. Ericsson, Alcatel-Lucent, and Nokia-Siemens);
- 5.) Application developers and content providers

The mobile broadband system is theorized as a “virtuous cycle,” in which fast and widely available networks encourage the creation of mobile devices to connect to these networks. In a “virtuous cycle,” connectivity spurs innovation of applications and content, while encouraging users to adopt technologies and promoting further investment in the networks.

Multiple obstructions to the “virtuous cycle” exist. Most immediately, the nature of relationships between actors (listed above) might inhibit innovation and investment. Additionally, some companies hold more advantageous roles in world communications, while other companies hold significant roles in multiple parts of the mobile broadband ecosystem, which can lead to inconsistent incentives throughout the mobile ecosystem (see Section 1.2).

Four case studies demonstrate how relationships between actors within the mobile broadband ecosystem can affect the incentives of actors to invest and innovate:

1.) **App Stores:** Application stores, while useful for consumers, can also restrict the development of mobile applications by influencing which applications are made available under varying conditions. HTML5 technologies, however, may provide an alternative model to the current application store model by granting application developers access to device functionality (Section 2.1).

2.) **Service Agreements:** Mobile broadband providers can directly influence their customers' access to networked services. Different service agreements, which shape how customers are able to use their mobile devices, demonstrate tensions between the financial risks of providers and the flexibility of the user experience (Section 2.2).

3.) **Network Unfriendly Apps:** Mobile broadband networks face several challenges to minimizing network congestion, including (1) mobile applications written by software developers who are unaware of how high-level designs affect network usage or battery resources, (2) radio access networks with limited bandwidth, permitting one application to consume the majority of available resources, (3) the "bearer" that mobile devices must establish with the cell tower, and (4) the substantial upfront investment necessary to expand the capacity of a cellular network, since it is expensive to acquire spectrum licenses, deploy cell towers, and transition to new technologies (Section 2.3).

4.) **WiFi Offloading:** Mobile wireless data traffic is increasingly shifting from mobile broadband services to Wi-Fi access, which is cheaper and more accessible. Accordingly, Wi-Fi is becoming an essential part of providing mobile broadband services to users. However, users of Wi-Fi networks may experience interferences from users of neighboring access points. There are different categories of Wi-Fi solutions, each of which vary in their benefits and limitations. Licensed and unlicensed spectrum solutions should be considered in the future (Section 2.4).

The report puts forth the following conclusions:

- 1.) The FCC should consider all of the interactions between different actors in the mobile broadband ecosystem, even actors which are not subject to the Open Internet Report and Order.
- 2.) The FCC should pay attention to new trends, such as HTML5 and Wi-Fi offloading, both of which might increase competition as they impact the mobile landscape.
- 3.) Transparency, education, and competition will all contribute to a healthy mobile broadband ecosystem.

#### **4. Specialized Services Report (Specialized Services working group)**

The specialized services subgroup within the Open Internet Advisory Committee (OIAC) had two tasks: (1) to agree upon a definition of "specialized services," and (2) to provide the FCC with advice about how they should oversee broadband Internet access service (BIAS) in light of specialized services. Two concerns about specialized services in the Open Internet Report and Order (R&O) are: (1) that broadband providers might label services as specialized services that would normally be labeled as Internet access services to evade Open Internet rules; and (2) that

broadband providers might stop expanding network capacity allocated to broadband Internet access service to allow more space for specialized services. The Open Internet Report does not specifically examine the impact of specialized services.

Defining “specialized services” proved to be difficult for the subgroup, and the agreed upon definition of the term is meaningful only within the context of the R&O. In that context, the definition of “specialized services” sets a limit on which IP-based services are subject to the Open Internet rules, as services labeled as “specialized” are not subject to further regulation under the R&O. The working group concluded that the primary criteria proposed by the FCC to classify a service as specialized are that (1) it is not used to reach large parts of the Internet, and that (2) it is not a generic platform—but rather a specific “application level” service. The committee identified one additional criterion that might classify a service as specialized: capacity isolation from BIAS.

Three high-level principles concerning specialized services that the FCC should consider are:

- Regulation should not create a perverse incentive for operators to move away from a converged IP infrastructure
- A service should not be able to escape regulatory burden or acquire a burden by moving to IP
- Proposals for regulation should be tested by applying them to varying technologies used for broadband

Two approaches may be used to address the FCC’s concern that specialized services might deter or limit investment in Internet services, though they both have risks associated with them. The first approach is that the FCC could define how much Internet service is “enough” and compare actual offerings to this minimum standard. However, this minimum standard will likely change over time as consumption habits shift. The second approach is that the FCC could examine what innovators can accomplish using specialized services compared with what they can accomplish with the public Internet, thereby revealing raw capacity as well as quality of service concerns.

In order to better understand the impact of specialized services on BIAS, and to understand when an Internet service is “good enough,” this subgroup advocates for examining the quality of the user experience rather than technical parameters.

## **A. Appendix 1: IPTV**

This paper examines the effects of video services (including IP-based video services) on broadband Internet access service (BIAS) and more generally in today’s marketplace.

### **High Level Overview Of Broadband Access Network Architectures**

The delivery of services over varied network architectures are surveyed, along with their potential repercussions for BIAS. This paper focuses on the access network, the portion of the network closest to the customer. Three commonly used access networks are (1) Hybrid Fiber Coax (HFC), typically used by modern cable systems, (2) Digital Subscriber Line (DSL), and (3) Passive Optical Networking (PON) based technology, typically used by telecommunications service providers.

### **Service Delivery Methods**

Services provided over the aforementioned architectures generally include: video (provided by Multichannel Video Programming Distributors, or MVPDs), voice, and BIAS services, which typically use separate channels from the linear video services. Through IPTV, which is another means of service delivery, all services are carried using IP on the same physical network. Different methods of delivering services are chosen partially based on how closely connected the physical access path is to the various services.

### **Capacity Isolation**

IP bandwidth in a household is dynamically allocated to different services, varying based on exact usage at a given time. However, capacity isolation is often used to ensure that IPTV bandwidth does not interfere with bandwidth used for BIAS services. The degree of isolation varies from service to service. This discussion is important because the degree of capacity isolation between a video service and BIAS service has implications for whether the video service should fall under the rules of the Report and Order.

### **Differences Between MVPDS' IP-Video and Over-the-Top Video**

One consequence of higher-speed broadband networks has been the proliferation of Over the Top (OTT) video services, which deliver content through the BIAS service of the end user. Examples of OTT video services include Netflix, Amazon Prime, and Vudu. These services differ from IPTV systems in the following ways:

- 1.) Customer Expectations: Customer support is less extensive with OTT services than with MVPD services.
- 2.) System Design: OTT services are generally provided via a third-party content delivery network, while MVPD services are generally provided over a privately owned and managed network within the service provider's infrastructure.
- 3.) Equipment: OTT services can be accessed through a number of retail consumer devices in the home, such as computers, tablets, and special OTT devices from cable operators. MVPD services, however, are usually accessed on equipment leased from the service provider.
- 4.) Regulatory Requirements: Devices and video services of OTT providers are not subject to the same regulatory obligations as MVPD services (except for the requirement of closed captioning support).
- 5.) Video Quality: Unlike OTT services, MVPD services generally do not need adaptive coding to preserve the user experience.

## **B. Appendix 2: Third-Party Purchasing of Services for Their Customers**

This case study examines how the increasing online service requirements on network performance might affect broadband Internet customers. The Internet provides "best effort" delivery of packets with no guarantees of delivery, delivery time of packets, and no guarantees one packet will have the same path/fate as the next. However, guaranteed quality of service from servers could be useful for customers. The subgroup explored four examples of third-party purchased quality of service:

- 1.) Establishing a separate specialized service to carry traffic between third-party services and its customers on the access ISP

- 2.) Prioritizing of OTT service traffic from certain third-party service providers amongst all general Internet traffic going to users over their Broadband Internet Access Service, either at the customer's request or at the third-party provider's request
- 3.) Establishing a dedicated core transit network to connect third-party service servers and access ISP networks
- 4.) Ensuring that there are open standards and best practices developed to support highly interactive traffic

In sum, the FCC should think readily about the distinction between challenges and solutions today, and opportunities tomorrow.

## **5. Open Internet Label Study (Transparency working group)**

### **Introduction**

This paper is concerned with transparency in the context of the Open Internet Report and Order, which mandates that fixed and mobile broadband providers be transparent in their management practices, performance characteristics, and terms and conditions of services. Specifically, this paper examines how ISPs present performance characteristics and pricing of their service offerings, and proposes a labeling system that would allow consumers to more easily compare services across companies.

### **Motivation**

A voluntary open Internet labeling program would help consumers select Internet services by clearly delineating points of comparison between Internet service providers. The main reason for this program is that many consumers are confused about how and why to choose a particular wireless service provider. In addition to facilitating comparisons, labels would provide access to test sites and to third-party analyses of performance parameters to customers.

### **The Proposal**

The suggested labeling program, through which data labels would correspond to each active service offering, would offer information pertaining to performance, price, and usage restrictions.

To partake in the suggested labeling program, ISPs would self-report data pertaining to upload speed and download speed (both reflecting the performance delivered by the ISP to a consumer's broadband modem), as well as the average monthly price over 36 months (which is designed to reflect both initial discounts or promotions and the long-term costs to the consumer). The label data could be published through (1) the ISP website, (2) an API provided by the ISP, or (3) periodic filings with a third party. Given the strengths and weaknesses of the respective publishing options, the working group recommends that the FCC pursue option (1).

### **Complexities**

Various complexities nonetheless remain, including those related to service offerings (i.e. bundling and promotions), customers (i.e. customer location, variability of Internet usage throughout the day, and thresholds where customers do not see a difference between two offerings), and companies (i.e. quality of service, ease of use, and setup time), all of which must be taken into account in order to understand the label program.

**Benefit**

The label program could provide a number of benefits, including (1) awareness, (2) consumer clarity, (3) competition, (4) incentivized open Internet practices, (5) marketing, (6) improved customer loyalty, and (7) global applicability. However, the label program could also introduce problems, including (1) misled consumers, (2) increased governmental costs, and (3) slow adoption of the label program.

**Summary**

In summary, the Transparency Working Group encourages the FCC to collaborate with the industry to develop a voluntary labeling program, through which ISPs would provide information to consumers about their services.

*Executive summaries were prepared by the Office of Professor Jonathan Zittrain.*

# Policy Issues in Data Caps and Usage-Based Pricing

FCC Open Internet Advisory Committee

Working Group on Economic Impacts of Open Internet Frameworks

Prepared for the meeting on July 9, 2013

The following report on Data Caps was prepared by the Economic Impacts working group in reaction to the press coverage and strong consumer sentiment regarding caps on data plans.

## Table of Contents

Topics Covered.....	Page 15
Definitions.....	Page 15
The Report and Order on UBP.....	Page 17
Competition.....	Page 18
Caps: The Facts.....	Page 19
The Perception of Users.....	Page 25
User Control.....	Page 26
The Perception of ISPs.....	Page 27
UBP and Price Discrimination.....	Page 28
Managing Network Growth.....	Page 28
Managing Instantaneous Congestion.....	Page 29
Perception of Edge Providers.....	Page 30
Specialized Services and Edge Providers.....	Page 31
Summary.....	Page 33
Further Reading.....	Page 34
Appendix.....	Page 35

## Topics Covered

The report examines data caps within the context of the Open Internet Order, primarily in wire-line, non-specialized services Internet access, and seeks to bridge the divide between the vernacular conversation surrounding caps and the perspectives from various stakeholders. Thresholds, caps, and usage-based pricing have been implemented in a variety of ways. This study focuses on providing definitions and identifying concerns/questions, with an emphasis on highlighting concerns and questions of the Open Internet Advisory Committee members.

The working group has chosen to focus on caps, thresholds and usage-based pricing because of questions raised about caps and tiers in many public forums and working papers. The Order expressly approves of usage-based pricing and experiments in pricing. Some members are concerned that this report could be construed as the working group second-guessing the FCC's decision. The Order set up the advisory group to consider whether aspects of the Order remain consistent in its effects on the Internet as the Internet evolves, and it is in that spirit that this conversation was undertaken.

The report considers only one part of a larger topic in detail, while aspiring to summarize many important aspects of this topic. However, it recognizes that it may be difficult or impossible to be comprehensive. Accordingly, the study ends with a section of further reading.

## Definitions

**Specialized Services** – The Order offers a rough definition on paragraph 112.

“...services that share capacity with broadband Internet access service over providers' last-mile facilities, and may develop and offer other such services in the future. These 'specialized services,' such as some broadband providers' existing facilities-based VoIP and Internet Protocol-video offerings, differ from broadband Internet access service and may drive additional private investment in broadband networks and provide end users valued services, supplementing the benefits of the open Internet.”

This report uses these terms merely for one pragmatic purpose, namely, to discuss the policy issues raised by data caps. Further discussion of the exact boundaries of this term are the province of the Specialized Services working group and are beyond the scope of this report.

**Usage-based pricing** - Usage-Based Pricing (UBP) takes many forms. It includes a continuum of practices from metering to discrete steps in price levels. In addition, volume-based pricing can discount or increase with volume. UBP appears in many economic settings and no single characterization will capture all these settings. For example, it describes metered pricing in electricity, as well as tiered pricing in cellular telephony. In general, UBP in the Internet context is based on amount of time online and/or volume of data transmitted. The working group uses UBP as a technical term that includes all form of charging functions that incorporate volume, whether linear or not.

**Data caps** - Data caps are often considered to be a form of UBP. The term “data cap” is characterized by several phenomena. In general, if a user is within a cap, he or she pays a set price. That is, the cap defines a limit on amount of data per month per household (today

expressed in gigabytes). Exceeding the cap could subject a household to alterations to its Internet access, possibly after one or more warnings, such as reduction of access speed, additional charges, suspension of service, or even termination of service.

The termination of service has received particular attention in public discussion, though to date, this appears to be a rare event, as noted below. A cap is rarely, if ever, a hard and fast ceiling on a customer's ability to access the network. A cap is usually better understood as a threshold after which the user is subject to a different set of conditions for access, such as movement to a higher priced tier, different product or different speeds. As discussed below, another way of thinking of this is as the boundary between different "tiers" of service.

The history of dial-up Internet access accounts for the present ambiguity in language. Historically caps referred to limitations on hours of use. It was quite common for dial-up ISPs to place capacity limitations based on hours of use of the ISP service per month, even for services sold as "unlimited." A common level for a cap was 100 to 120 hours of use per month. After exceeding that cap, certain ISPs would discontinue service altogether. Other ISPs used an early version of UBP instead and, rather than terminating service, would simply charge extra additional hour of service. One asserted basis for this practice was that UBP was needed to address capacity issues related to the fixed capacity of modem banks.<sup>2</sup>

Modern caps refer to limitations on downloading and uploading of data. Today, as the tables below show, hourly use is not restricted by any major ISP. Instead, thresholds, if they exist, pertain to monthly limits or tier thresholds on the total transmission and reception of data, and, moreover, the draconian features of historical caps, such as abrupt termination of service, are largely absent from the modern version. Within the United States, no major ISP stops providing service to consumers without notifying consumers and providing additional options in the way of tier upgrades or overage charges.

There are a variety of viewpoints about caps. Mirroring the different perspectives used throughout this document, the following perspectives may be helpful as a start to the discussion:

*From the user viewpoint:* The viewpoints vary depending on if caps or thresholds are actually impacting the user. However, the difference between a high threshold and a cap may be a semantic distinction without a meaningful difference, particularly if the threshold appears to be abrupt, and there is little perceived difference between being terminated, and the alternatives, such as overage charges or throttling. Lack of consumer understanding of how a data caps are impacted by use of various services may impose mental transaction costs that could dissuade consumers from using Internet-delivered services – even if a user does not come near to exceeding a cap. These concerns are particularly acute if the user perceives little option to contract with alternative suppliers of Internet access. Additional questions also arise: can cap information be difficult to find,

---

<sup>2</sup> Providers justified these policies by noting that modem banks were dimensioned assuming statistical multiplexing and specific usage patterns. For a history of dial-up access business in the United States, see e.g., Greenstein, Shane. 2008. "The Evolution of Market Structure for Internet Access in the United States." in William Aspray and Paul Ceruzzi, editors, *The Commercialization of the Internet and its Impact on American Business*, MIT Press. pp 47-104.

and relatively opaque to users, who may believe that they are contracting for unlimited Internet access?

*From an ISP's viewpoint:* Usage thresholds in most US broadband ISPs are set so high that they impact very few customers (around 1-2% depending on the ISP). Under most usage thresholds, a broadband user can successfully run many applications, stream video, download music, share photos, surf the web, play games online, etc. The concept of ultra-high end thresholds is to ensure that the low end (1G-10G), average (15G-50G) and even the high end user (100G-250G) is not subsidizing the most extreme bandwidth user (250G-1000G+). Network resources are not unlimited, and the ISP's viewpoint is that, as the Open Internet Order explains, "lighter end users of the network" should not be forced "to subsidize heavier end users" who require more of a dedicated commercial level of service vs. residential broadband.

*From an edge provider's viewpoint:* (An edge provider is a firm that provides online content, applications, or services to end users.<sup>3</sup>) When users and edge providers exchange traffic, the traffic goes over an ISP's facilities. A high threshold or cap may represent an additional factor that shapes the ability of an edge provider to supply its service or conduct business with a user. If an ISP imposes a data cap or other form of UBP, this could affect user demand for the edge provider's service, which, in turn, may shape the ability of the edge provider to market and deliver its service. This is especially so if the ISP offers specialized services that compete with the edge provider, and for which a cap or other UBP does not apply.

The discussion will focus on the implications of these thresholds as one form of UBP, and expand on the different points of view. The study will occasionally use the phrase "caps" or "threshold," depending on context and point of view.

Two words of caution are warranted at the outset. First, assessment of caps is not synonymous with assessment of all forms of thresholds within UBP. This discussion leaves many other topics about UBP uncovered. Second, the study initially will focus on issues in the absence of competing specialized services. In the presence of specialized services, there are additional issues raised concerning selective applications of thresholds to some types of traffic, which will be discussed below.

## **The Report and Order on UBP**

The Open Internet Report and Order discusses usage-based pricing, but does not expressly mention data caps except by implication in that data caps can be considered a form of UBP. The most direct mention of UBP is in Paragraph 72 of the Order:

"Some commenters suggest that open Internet protections would prohibit broadband providers from offering their subscribers different tiers of service or

---

<sup>3</sup> See footnote 2 of the Order. The Order uses "... 'edge provider' to refer to content, application, service, and device providers, because they generally operate at the edge rather than the core of the network."

from charging their subscribers based on bandwidth consumed. We are, of course, always concerned about anti-consumer or anticompetitive practices, and we remain so here. However, prohibiting tiered or usage-based pricing and requiring all subscribers to pay the same amount for broadband service, regardless of the performance or usage of the service, would force lighter end users of the network to subsidize heavier end users. It would also foreclose practices that may appropriately align incentives to encourage efficient use of networks. The framework we adopt today does not prevent broadband providers from asking subscribers who use the network less to pay less, and subscribers who use the network more to pay more.”

The Order left open the possibility of many experiments in business models and pricing. Moreover, the Internet had evolved over time, and the Order anticipated that the Internet would continue to evolve in unexpected ways, including in pricing for mobile broadband services (see especially paragraph 94).

## **Competition**

Data caps are a source of concern in settings where there are no or few substitutes for Internet access. That reduces the discipline affiliated with competitive markets. Limited competition gives a supplier the ability to make take-it or-leave-it offers to users, and users cannot leave for another supplier if they find the service or contracts unsatisfactory. As noted in the data section, there is no indication that ISPs are offering different policies in areas with limited competition.

Resolving any such question, however, requires defining the extent of competition, which, in turn, requires a precise definition of the size of the market. It is the typical first step in any textbook policy analysis. In practice, however, a precise definition can be elusive.<sup>4</sup>

That matters for discussions of caps, thresholds, and UBP. While there are a variety of issues with UBP, most of the issues with thresholds do not arise when the prices are low. Many interesting policy questions concern the highest thresholds and the biggest charges, especially those that (effectively) determine the difference between unlimited service and limited service.<sup>5</sup>

While that makes it seem like it might be possible to reduce many questions to a narrow issue, it turns out that even narrow questions contain challenges. For example, there is simply no general definition for “demand for high bandwidth,” which varies by supplier, by geography, and technology. No simple definition – e.g., all markets for services above 5GB, 20GB or 50GB or some other arbitrary floor – will work in all settings. In addition, as will be shown below, because demand is growing rapidly, policy is shooting at a moving target, so it is also hard to describe a general rule for the size and scope of the market in which the policy issues arise.

Consider concerns about caps and thresholds that focus on the “high end,” or users who consume a significant amount of data. There is a perception that users at the “high end” are more likely to

---

<sup>4</sup> See the National Broadband Plan, particularly chapters 3 and 4, for an extensive discussion of questions pertaining to defining the structure of the market. See <http://www.broadband.gov/plan/>.

<sup>5</sup> This section focuses on the policy issues at “the high end” for purposes of illustration. The discussion below will discuss further issues about thresholds across a range of bandwidth levels.

exceed caps and find no alternative source of Internet access that meets their high-usage demands. This usage pattern could be considered more typical of business-class users. However, even this perception is difficult to substantiate, partially because it is difficult to estimate what “high end” usage consists of now, or what it will consist of in the future. The size and definition of “high” is a moving target. It is also difficult to estimate what high or low end use consists of because estimations of usage distributions also vary widely, with no definitive standard. In addition, the lack of definitive data reflects real underlying variance in situations in which firms deploy wireline broadband in the United States – variance in access technology (cable, DSL and FTTH), vendors (different local pairings of rivals, if any), regulatory treatment, and geographic features (city/rural and flat/hilly). The National Broadband Plan discusses this variance extensively, as does the Order.

Growth in data traffic also reflects real underlying variance in the data-intensive applications that users deploy (e.g., YouTube, Hulu, Netflix, peer-to-peer, multiplayer gaming). Usage of data by these applications grows at different rates because there is variance in the rate of adoption – and intensity of use – of these and related applications. All of these variations confirm the need to refrain from sweeping generalities for all settings and times about the state of competitive alternatives.

Hence, there is no consensus on the definition for “high” either now or in the near future. This means that it’s very difficult to draw conclusions about whether high end users would switch from wireline broadband providers with a lower cap to ones with a higher cap. This lack of data about even the user population, let alone their behavior in the marketplace makes it difficult to draw conclusions about the role of data caps in competition.

This does not mean it is impossible to discuss and analyze caps and related matters. However, it does imply that it is usually challenging to come to sweeping and general conclusions. This theme will arise in several places throughout the report.

### **Caps: The Facts**

Many types of data charges exist in United States residential wireline Internet access. Table 1.1 shows data from an October 2012 article in GigaOm. Table 1.2 shows data collected by a working group member in February of 2013, based on publicly available data, which breaks out some of the thresholds by pricing tiers. The section will present these facts, and later sections will offer overlapping and competing interpretations.

Examination of the tables shows several things. First, the highest thresholds typically range between 150 and 300GB per month. Second, a number of ISPs do not have any caps at all. Third, many thresholds that resemble caps are part of a system of many-step thresholds, often within one pricing plan or tier. Fourth, some ISPs offer many tiers, and the highest thresholds vary by tiers. Fifth, when an overage charge arises (see appendix), firms tend to use similar levels, generally around \$10 for 50 additional GBs beyond the threshold (See appendix. This is not reflected in the Tables).

These observations reinforce the conclusion that there is considerable variance and experimentation in the market by ISPs. Note, however, that these are observations of firms and

contracts, not subscribers. This table does not address questions about how much data most users actually consume and what thresholds, if any, most users actually face.

**Table 1.1. Caps quoted in GigaOm**

• ISP	Cap	• ISP	Cap
Comcast	300GB per month	Charter	100GB – 500 GB per month
AT&T	250GB or 150 GB per month	Frontier	No
TWC	No	Windstream	No
Verizon	No	SuddenLink	150GB to 350 GB per month
CenturyLink	150 GB per month to 250 GB per month	MediaCom	150 GB to 999 GB per month
Cox	30GB-400GB per month	Cable One	1GB, 50 GB and 100 GB per month
Cablevision	No	FairPoint	No
		Cincinnati Bell	No

Source: See <http://gigaom.com/2012/10/01/data-caps-chart/>.<sup>6</sup>

**Table 1.2. Highest thresholds, Recent sampling of Company sites**

Provider	Use Threshold - GBs <sup>7</sup>
Comcast	min 300 GB (increasing by speed tier) <sup>8</sup>
AT&T - U-Verse HSIA	250
AT&T – DSL	150
Time Warner Cable	None
Verizon - FiOS / DSL	None
CenturyLink - 1.5 Mbps	150
CenturyLink - >1.5 Mbps	250
Cox - Ultimate (100 Mbps)	400

<sup>6</sup> The article includes additional details on exceptions, tiers, and overages. The appendix consists of more recent and accurate data, and corrects several inaccuracies in this article.

<sup>7</sup> Gigabytes per month, unless otherwise noted.

<sup>8</sup> At the time of writing Comcast does not have any caps in place but is trialing two UBP plans. See appendix for further details.

## Open Internet Advisory Committee - 2013 Annual Report

Cox - Premier (25 Mbps)	250
Cox - Preferred (15 Mbps)	200
Cox - Essential (3 Mbps)	50
Cox - Starter (1 Mbps)	30
Cablevision	None
Charter - Lite & Express ( )	100
Charter - Plus & Max (30 Mbps )	250
Charter - Ultra100 (100 Mbps)	500
Frontier	100 / 250 in selected trial mkts
Windstream	None
SuddenLink (>30 Mbps)	350
SuddenLink (10-30 Mbps)	250
SuddenLink (<10 Mbps)	150
MediaCom - Launch (3 Mbps)	150
MediaCom - Prime (15 Mbps)	250
MediaCom - Prime Plus (30 Mbps)	350
MediaCom - Ultra/Ultra plus (50/105 Mbps)	999
Cable One – Economy	Monthly: 1GB <sup>9</sup>
Cable One - Preferred (50 Mbps)	Monthly: 50 GB <sup>5</sup>
Cable One - Elite (50 Mbps)	Monthly: 100 GB <sup>5</sup>

<sup>9</sup> Daily limits also apply. See appendix.

## Open Internet Advisory Committee - 2013 Annual Report

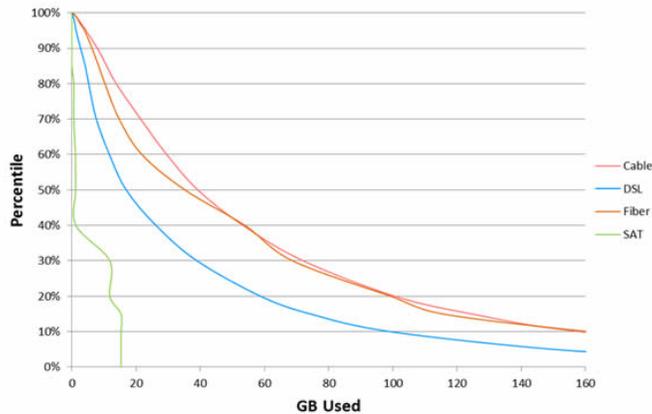
Cable One - Standard (5 Mbps)	Daily: 3 GB
Cable One - Premium (10 Mbps)	Daily: 5 GB
Cable One - Ultra (12 Mbps)	Daily: 10 GB
FairPoint	None
Cincinnati Bell	None
Google Fiber	None

Sources: See Appendix.

It is difficult to interpret even the highest thresholds in the situations in which they arise, as there is no definitive public source on household usage per month to use as a benchmark. Several different sources are available. Usage varies depending on ISP and technology. All public measurements show great skew in usage, and suggest that caps do not yet impact users other than the highest users. A first look at the usage distribution is offered by Figure 1.1., which comes from the July 2012 Broadband Report.

Figure 1.1 puts the median at approximately 15 GB for DSL, 25 GB for Fiber, and 30 GB for cable users. Other estimates vary, but are in a “similar neighborhood.” For example, another estimate puts the median at 14 GB, and an average at 47 GB. (Bauer, Clark, Lehr, 2012). A Cisco study last year put the average at 26.2 GB average in 2011, with a forecast of 84 GB by 2016.

**Figure 1.1. Distribution of monthly use of data**



Source: <http://www.fcc.gov/measuring-broadband-america/2013/February#Chart20>.

In addition to data collected by various studies, it is also possible to think of caps in terms of hypothetical use rates. One committee member offered a “cord cutter” benchmark: the Internet usage equivalent of the five hours per TV per day. Consider the following: 5 hr/day (television viewing) x 2 GB/hr (high quality video) x 30 days. That would yield 300 GB/month in use. Recognize that this is a speculative simulation, and considerable variance is possible. Many factors could change the outcome at a household – e.g., DVR use with HD antenna, type of viewing, live news and sports over the air. This also does not include other Internet usage.

This leads to a number of conclusions. For one, most thresholds in wire-line today in the US appear to affect only high end users. The lack of subscriber data makes it impossible to provide an estimate of the precise percentage of users affected by high thresholds, but at this point a high threshold, such as 150-300 GB, appears to affect a small percentage of households.

Despite that, there is some evidence that caps may be binding on users, if set low enough. For example, many Canadian ISPs have set caps in the range of 25/40/60 GB per month.<sup>10</sup> According to Netflix, streaming video at normal or high definition quality caused users to exceed their data allowances. Netflix reported that these low caps seemed to have an effect on household demand for its services and that it observed a noticeable response in its business. The same response would have been anticipated in the best of circumstances, but it was further magnified by the poor measurement of traffic at the household level and the lack of transparency to users. In reaction to these low caps, Netflix reduced the default quality of the videos it sent to Canadian users. Netflix set a lower quality bitrate limit (625kbps vs. 4800kbps) as the default for all users, to prevent users from accidentally hitting their caps. According to Netflix, streaming of high-definition content on the ISPs that cap in Canada is essentially non-existent, and the quality of the user experience has been reduced.

Will caps within the United States ever affect more than a small percentage of US households? Here we review two perspectives.

To begin, experts disagree on predictions for the likely rate of future growth in data usage due to (expected) growth in cloud-based services and video services at the level of household and in the marketplace overall (more discussion below). Even predictions for the near future vary heavily. Committee members were familiar with predictions as low as 20% and as high as 40-50% growth per year. This report draws from Sandvine Global Broadband trends, Cisco Visual Networking Index, SamKnows, and the FCC’s Measuring Broadband Report. While all such reports provide a similar outlook of the broad picture, these reports can differ significantly in the specific numbers provided.

Even this simple presentation of facts illustrates a point of disagreement between distinct perspectives. Though more will come in later sections of this report, here is a brief illustration:

---

<sup>10</sup> The outline of these events has been reported in the trade press. See, e.g., Nate Anderson, March 29, 2011, “Data caps claim a victim: Netflix cuts streaming video quality,” *Arstechnica*. <http://arstechnica.com/tech-policy/2011/03/data-caps-claim-a-victim-netflix-streaming-video/>, and Richard Lawler, March 28, 2011, *Engadget*, “Netflix Canada announces new bandwidth management settings for capped users,” <http://www.engadget.com/2011/03/28/netflix-canada-announces-new-bandwidth-management-settings-for-c/>,

<sup>10</sup>both accessed April 28, 2013.

Some non-profit advocacy groups argue that caps will become binding assuming a constant rate of growth of bandwidth usage without corresponding cap adjustment. Some point out that “yesterday’s so called “bandwidth hogs” are today’s typical users.”<sup>11</sup> A bit of simple speculation can illustrate the circumstances in which the claim is valid or not. If growth rates are at the lower end of projections, say, 20% growth rates, there would be a doubling of use in a little less than four years. With such growth rate, a 150GB cap would become relevant to the behavior of much more than 10% of cable and fiber households portrayed in figure 1.1. Additionally, advocacy groups express concern that so called “extreme” users tend to be disproportionately early adopters of new technologies, and as such, caps that affect them may prove to have a large impact on innovation in the field, independent of the sheer number of users they affected.

Suppliers counter that the highest thresholds are unlikely to ever affect more than “extreme” users. Some ISPs determine their thresholds in reference to usage—often the threshold is either explicitly set as a certain percentage of their subscriber base’s usage, or is set so as to only affect an estimated percentage of the subscriber base. Under either methodology, by definition, the threshold can only affect that top percentage of users that are using the most bandwidth, and will not affect the vast majority of subscribers.<sup>12</sup> These thresholds are often established and periodically re-assessed, specifically to focus any effect on only the uppermost percentile of users. Therefore, by definition, these will only affect “extreme” users. For example, Comcast has raised its thresholds over time.<sup>13</sup> In addition, some ISPs have stated publicly that these “extreme” users tend to be those that are utilizing 24x7 file sharing or operating content or application servers from their homes. This usage pattern ties up infrastructure in a dedicated fashion that is similar to a reserved capacity of commercial service offering.

From the facts and examples listed above, we can reach only tentative conclusions. Although caps do not seem to be affecting a large number of US users now, the situation may change in the future, as user habits, supplier experimentation, vendor policy, and applications all change. As such, the FCC should monitor the situation. The committee makes no recommendation about which, of many factors, would be the most useful to monitor. Among the candidates for potential monitoring: definitions of tiers by data download limit; whether those limits are packaged with other features of a contract, such as bandwidth and speed; contractual provisions for what happens when users bump up against a tier (see, e.g., the appendix); and whether systematic differences arise across categories of service (fiber, DSL, etc).

In addition, it may be valuable to consider what warning signs of increasing effects by caps would look like. In addition, the reports about the Canadian experience with caps generally lacked verifiable data or other surveys of user response. It would be interesting to compare usage

---

<sup>11</sup> The New America Foundation, “Capping the Nation’s Broadband Future?” [http://www.newamerica.net/publications/policy/capping\\_the\\_nation\\_s\\_broadband\\_future](http://www.newamerica.net/publications/policy/capping_the_nation_s_broadband_future), accessed May, 17, 2013.

<sup>12</sup> For example, AT&T describes: “In fact, less than 2% of AT&T High Speed Internet users utilize more than 150GB per month. We estimate that 98% of our customers will not be affected by this change because our data plans include so much bandwidth.” (<http://www.att.com/esupport/article.jsp?sid=KB409045#fbid=kiJ0SSZjH9l>).

<sup>13</sup> See e.g., Nate Andresen, May 17, 2012, “Comcast suspends 250 GB cap for now,” *Ars Technica*, <http://arstechnica.com/business/2012/05/comcast-suspends-data-caps-for-now/>, accessed April 29, 2012.

before and after caps were imposed, and to further learn what general lessons, if any, this situation can teach.

### **The Perception of Users**

From a consumer standpoint, caps and high thresholds are generally more appealing when their properties are knowable and predictable. Additionally, user behavior may be impacted substantially by incorrect understanding of contractual obligations or data use. Thus, the questions about the effects of caps can only really be answered if we understand what users know and think.

Policies for caps and thresholds should be concerned about user understanding because many household surveys find rather poor knowledge of speed/usage of own broadband and applications (see, e.g., the Pew Surveys<sup>14</sup>). In addition, there are changing norms for software usage and users may have a limited ability to understand the typical GB per hour of use of an application. Application and service owners bear some responsibility here as well, as they can make efforts to understand their own efficient and inefficient use of network resources, and its cost.

Lack of user understanding of how many GB may be used by applications could lead to two sub-optimal scenarios: (1) Users could underestimate the amount of data consumed and exceed their monthly data allotment, thereby incurring penalty fees or unanticipated upgrade charges; (2) Users could overestimate the amount of data consumed, thereby dissuading them from using Internet-delivered services even though they are well below their cap thresholds.

The history of unlimited dial-up can possibly explain some of the lack of user understanding of data use. The lack of limitation (i.e., unlimited use) is usually regarded as better for users than the presence of a limitation (i.e., a cap on use). Some commentators perceive an association between the lack of unlimited pricing and the lack of competitive alternative. In part, one of the most prominent historical examples reinforces the perception, namely, AOL's experience moving from usage-based pricing (specifically, metering of hours of use) to unlimited contracts. This change came about in response to competitive pressure.<sup>15</sup> Hence, in the minds of some commentators the increasing use of usage-based pricing with thresholds is affiliated with the decreasing use of unlimited plans, which, in turn, is presumptively affiliated with a decrease in competitive alternatives.

Unfortunately, much information about user understanding of caps and thresholds is missing. Some open questions that could be useful to answer: Do users have an ability to measure their

---

<sup>14</sup> See the Pew Internet and American Life Project. <http://www.pewinternet.org/>.

<sup>15</sup> The perception partly arises from the reminiscing many years later. The CEO for AOL at the time, Steve Case, states that AOL had studied the potential switch for quite some time, but not acted on it because management could anticipate a difficult transition. Competition eventually forced his hand. Said Case, "It came to a head over a weekend as Microsoft announced they were offering MSN on a flat rate basis, and it was clear they were planning to steal a lot of market share from AOL. So I decided within hours of their announcement that we had to match them, and the company worked throughout a weekend so we could make an announcement." See <http://www.quora.com/AOL/How-did-AOL-make-the-decision-to-go-to-an-all-you-can-eat-pricing-strategy/>. For a longer account of these events, see Swisher, Kara, 1998, *aol.com: How Steve Case Beat Bill Gates, Nailed the Netheads, and Made Millions in the War for the Web*, Random House; New York.

own data use in real time? While some tools for aiding user measurement are beginning to emerge, how widely are they used and are they effective? What is the accuracy of some typical data meters?<sup>16</sup> Can users measure own usage by application? If so, how to encourage their use? Can users manage to monitor their use in households with multiple users and multiple devices?

The move from unlimited data to capped plans in wireless suggests some users can adjust over time to caps. However, it is difficult to predict whether that experience would carry over to wireline households, with its different applications, and in particular, whether households where multiple users of different ages occupy the residence will be able to adjust to a communal limit. However these questions of user experiences and ability to control raise questions about whether caps or thresholds that are set too low could lead to a world where the average user carefully monitors her bandwidth use, rather than leaving the average user well enough alone while only forcing “extreme” users to make changes to their use.

This topic also has implications for common notions of fairness. Typical users may be paying the same price for their Internet access as heavy users. Caps also need to be updated to match current usage patterns in order to continue to only impact “high users.” From an ISP’s perspective, someone who uses a steady and moderate stream of data is very different from someone who uses heavy data at peak moments of heavy use of capacity. Yet, a threshold pricing scheme hits them the same.

Another equity concern from the user perspective has to do with some models of steady data use, such as for medical purposes, which also can have implications for peak load and non-peak load use. These questions require more information about peak load pricing, a topic we take up below. For the time being, we defer more discussion.

To conclude in a similar manner to the previous section, this topic may require future monitoring, especially given the importance of consumer education to user perceptions of caps and thresholds. It is not yet apparent whether the issues in this topic are a transitory or permanent concern. The experience of ISPs with providing customers with tools to monitor or control data usage could also be valuable to insights about the perceptions of caps by consumers.

## **User Control**

If users do not have enough control over their data usage to adequately respond, even if well informed, to caps and thresholds set by ISPs, “punishment” of users by caps or thresholds may become a problem. For example, data-intensive video commercials are increasingly being embedded in web pages by edge providers. Automated nightly/weekly updates of software are also increasingly common from software vendors. In addition, most users operate software over which the user has little control.<sup>17</sup>

---

<sup>16</sup> See, e.g., Stacey Higginbotham, Feb. 7, 2013, “More bad news about broadband caps: many meters are inaccurate,” *GigaOM*, <http://gigaom.com/2013/02/07/more-bad-news-about-broadband-caps-many-meters-are-inaccurate/>, accessed May, 17, 2013.

<sup>17</sup> See e.g., Peter Sevcik, 2012, “Empowering Internet users to manage broadband consumption,” Netforecast, [http://www.netforecast.com/wp-content/uploads/2012/06/NFR5109\\_Empowering\\_Internet\\_Users\\_to\\_Manage\\_Broadband\\_Consumption.pdf](http://www.netforecast.com/wp-content/uploads/2012/06/NFR5109_Empowering_Internet_Users_to_Manage_Broadband_Consumption.pdf), accessed April 28, 2013.

Conversely, some available tools -- today used by some sophisticated users -- allow ad-blocking and other user-traffic management. Ad-blocking and flash-blocking tools are the methods most commonly discussed in online forums.<sup>18</sup>

User control also plays a role in discussions about overage charges. Overages only arise when a threshold is exceeded, and actual charges can depend on specific details about how overages are enforced.<sup>19</sup> For many users there is only downside as that threshold becomes closer. Do households consider that monitoring burdensome, particularly multi-dweller households? Do multi-dweller households perceive the monitoring as a hassle or perceive the increased uncertainty in billing as a burden? There is not enough experience yet to suggest how to characterize most households.

The working group did not further explore this topic. This issue seems largely irrelevant for the average user, as few users are affected by caps, as a practical matter. In addition, many issues in user control are too small to matter, and if they become a problem, providers typically have conversations with users, and offer amnesty. This includes issues linked to several phenomena, such as automated syncing, spam, denial of service, and compromised machines that send out messages as part of denial of service attacks. Generally speaking, the committee did not perceive these issues to be big at this time.

This may change over time. If data use grows without a commensurate increase in caps, these concerns may become urgent for policy deliberation. If this occurs, a more accurate labeling system for software applications and monitoring system that take into account caps may be a way to educate users and increase awareness of the necessity of controlling bandwidth use.

### **The Perception of ISPs**

ISPs generally explain the use of thresholds (caps) as providing a simple pricing mechanism for matching demand for bandwidth consumption with purchasing behavior. ISPs view pricing and product choices as consumer options that are just as important to the delivery of Internet services to end users as content or technical innovations in those services.

Speed tiers also match demand for bandwidth, and most ISPs correlate speed tiers with usage thresholds. Suppliers argue that UBP with a few thresholds balances the efficiency of metered pricing without creating the stress or mental costs associated with such metering. Thus, suppliers emphasize that UBP with a few thresholds, or some forms of tiered pricing, provides a measure of bill stability, predictability, and “peace of mind” to the vast majority of consumers relative to more linear usage pricing (i.e. metered, or per KB/MB/GB, or finer-grained use tiers).

Depending on how it is structured, UBP can also enable additional lower-cost broadband plans to be offered to consumers, spurring adoption or better meeting the underserved demand from the

---

<sup>18</sup> For example, see the second comment at <http://www.dslreports.com/shownews/Why-is-ATT-Capping-DSL-Users-but-Not-UVerse-Users-123692>, accessed April 29, 2013.

<sup>19</sup> Some care is required in drawing sweeping conclusions without precise data. For example, in plans being trialed by Comcast (at the time of this writing) a user must exceed a threshold for three months in a twelve month period before overages are imposed.

low-end of the market. ISPs can afford to offer lower usage plans at a lower price point (e.g. Cable One's Economy plan): they do not add as much to the aggregate bandwidth demand for the ISP. This is one approach to manage long run bandwidth-sensitive costs.

In this sense, UBP generally serves two functions. It may affect a small number of users who use large amounts of resources. It also may shape the use of resources among the vast majority of users. As the tables showed, there are examples of ISPs pursuing policies that lend themselves to each interpretation in wireline broadband today. However, the most common so far is the use of UBP to limit a small number of users who use a large amount of resources.<sup>20</sup>

Beyond these generalities, more detailed analysis of the issues from a supplier's perspective falls into three categories: how to arrange prices so "high end users" pay for the additional investment they use (i.e., price discrimination in the economics literature), managing network growth (e.g., managing long run capacity investment) and managing instantaneous congestion (e.g., managing peak load pricing). The report summarizes each of these in turn.

### **UBP and price discrimination**

Generally, in a high fixed and high sunk cost setting (such as network provision), usage based pricing is about raising revenue over incremental costs and recouping substantial fixed costs. This is generally called the economics of non-linear pricing, or price discrimination in common economic parlance.<sup>21</sup>

The economics literature on price discrimination provides two motives for UBP: (1) associating higher prices with higher costs and higher willingness to pay, while (2) avoiding the potential losses when some users do not buy at all. Such association can come closer to common notions of fairness and also reinforces the incentives to save on costs by showing users the price of inputs.

The Order has already made clear that usage-based pricing ensures that lighter end users are not forced to subsidize heavier end users. Charging distinct prices aligns incentives to encourage efficient use of networks. The Order also has made clear that the FCC will continue to monitor the marketplace. Thus, as the marketplace continues to develop, presumably the FCC will take these issues into account in its decision-making.

### **Managing Network Growth**

If measurement and transparency issues were satisfactorily addressed, could a cap or threshold at a high end of downloading (e.g., less than 1% or 2% of households) reduce data use? There is little evidence (outside of Canada, as noted), so it is difficult to judge. The answer is necessarily speculative.

---

<sup>20</sup> We note the interesting contrast with the use of UBP in wireless contracts, where it is much more common to use UBP to shape the use of resources among the majority of users. This difference motivates open questions about why the difference arises, and what lessons can be learned from those differences.

<sup>21</sup> A side note about vernacular interpretation of economic terms: The word "discrimination" has a pejorative meaning in common language, though none is meant in the economics literature on price discrimination.

Generally speaking, it is thought that a data cap (in this context, a threshold with discrete changes in speed) can incentivize those near the cap to behave differently. If so, then a household that uses much more than a typical user can build more efficient usage into its own network or decrease its usage upon nearing the cap. Some of these changes may not interfere with normal Internet usage by merely optimizing bandwidth heavy behavior— for example, users streaming video footage could alter the use of uncompressed HD cameras streaming to the Internet 24x7, when on-demand will do. Users also can reduce use of peer-to-peer servers, e.g., BitTorrent, substituting partial uptime for full time. Users who run servers out of the house on a household contract could switch to a business offering that better matches needs and usage. Also, as discussed below, potentially users can take advantage of bandwidth efficiencies as they become available from edge providers. However, households that are already using bandwidth efficiently may be forced to make changes that do impact day-to-day usage.

Access providers also ask whether data caps and related means of linking price to use can encourage edge providers to innovate more efficient means of delivering their services. There is a perception that data caps and usage-based billing are not potential barriers to entry but, rather, potential drivers of greater efficiency in the delivery of edge services. They point to the incentives on Netflix and other edge service providers to innovate their services, for example, Netflix improving efficiency in Canada and licensing innovative technology like EyeIO.<sup>22</sup> Access providers also raise questions about the extent to which prices are misaligned and resources are misallocated because all the obligations for carriage of content is passed onto consumers (and the ISP) by edge providers. (Edge providers have a different perspective, which is discussed below.)

At most, we can draw a tentative conclusion. Over the long run a data cap or a UBP threshold can help manage network growth if users and/or edge service providers respond to the cap or threshold with less or more efficient data use; a carrier would then incur less costly operations and may be able to make less expensive infrastructure upgrades over longer periods. However, this conclusion is mostly theoretical – there is no quantitative data to suggest to what extent how much long run costs increase with growth in use or how much of a difference carrier contributions to provisioning have made to growth over time. Both the broad and specific questions cannot be answered because there is no quantitative evidence – to accept or refute – propositions about how caps and thresholds shape usage.

### **Managing Instantaneous Congestion**

Generally speaking, instantaneous congestion management is not a stated rationale behind use of tiers, metering, or caps. There are other techniques in TCP/IP to address congestion caused by unexpected demand, outages, or major traffic shifts. Caps provide no direct incentive to heavy users to reduce traffic at peak times because there is no differential pricing across time periods. For example, monthly caps generally count traffic from the middle of the night (when traffic in general is low) against a cap.<sup>23</sup>

---

<sup>22</sup> See e.g., Janko Roettgers, 2012, “EyeIO: Netflix’s secret weapon against bandwidth caps?” GigaOm, Feb 1, 2012, <http://gigaom.com/2012/02/01/eyeio-video-encoding-netflix/>, accessed April 28, 2013.

<sup>23</sup> There has been some experimentation with time-sensitive lifting of cap restrictions. See for example, this description of a satellite broadband provider’s recent policy. <http://www.dslreports.com/shownews/Exede-Caps-Lifted-For-Overnight-Use-120776>.

However, if there is a rough correlation between total use and peak use – i.e., the largest total users over the month are also the biggest users at peak moments – then a data threshold might have some of the properties of a peak load pricing scheme by inducing a large data user to reduce their data usage. This is an open question, as there is little public analysis of the correspondence between data consumption and bandwidth usage.<sup>24</sup> In addition, there is little experience with alternative arrangements, as many ISPs do not perceive users calling for the option to manage data use over time.

There is no evidence, one way or another, that caps leads heavy users to reduce activity at peak moments any more than at any other moment. It would be illustrative to see if there are systematic differences between usage in the United States and Canada because of the imposition of caps and thresholds. Again, no particular data speaks to this specific question one way or another, or to the broad questions motivating it. There has not been much experience with peak load capacity management thresholds for users. Historical experience with peak load management suggests the timing for data usage and peaks would shift, but there is no evidence to suggest which applications will shift their usage patterns, or by how much they would shift them.<sup>25</sup>

### **Perception of Edge Providers<sup>26</sup>**

A data cap or high threshold from broadband provider can shape other providers of services in broadband ecosystem, e.g., entrepreneurs who provide applications, build web pages, and operate other services in the cloud. Edge providers are concerned that a widely used cap reduced – rationally or irrationally – demand for data-intensive services and reduced entry of new data-intensive software firms, decreasing the commercialization of innovation. This concern is partially motivated by Netflix’s example in Canada, which illustrates the phenomenon when a cap does bind.

Some of the power of data caps to affect edge providers that serve video or other high bandwidth media content might be offset by improvements in codecs. A codec encodes a data stream of signal for transmission, storage or encryption, and decodes it for playback and editing. (The word is a portmanteau of COder and DECoder.) There are many codecs in use today. Would improvement in codecs – i.e., to higher resolution using fewer resources with more efficiency – occur regardless of the presence/absence of caps? ISPs argue that edge providers have incentives to improve codecs when faced with caps and high thresholds. The alternative view argues that improvements arise for largely exogenous reasons, and have little relationship with the policies of ISPs.<sup>27</sup>

---

<sup>24</sup> One of the earliest studies of the correspondence between data consumption and bandwidth usage examined one ISP’s traffic in 2011. It found a small correlation, not consistent with using caps to manage bandwidth. It is at <http://www.fiberevolution.com/2011/11/do-data-caps-punish-the-wrong-users.html>.

<sup>25</sup> For example, AOL experimented changes in pricing for different times of the day in order to save on phone line costs, and experienced changes in the time of day in which the “peak” usage occurred.

<sup>26</sup> As elsewhere in this study, we focus on the perception of “Edge Providers,” as in the Order, rather than focusing on other groups of providers, such as “over the top providers,” or “application service providers.”

<sup>27</sup> The working group noted that parallel arguments take place in wireless applications.

Once again, these questions are necessarily speculative, as caps do not yet bind most households in the US, and, at present there is no decline in the demand for data-intensive services. In addition, as noted above, the experience with data limits in Canada has not been measured, so there is no data to assess the impact the caps had in that setting.

It is unclear how much entrepreneurs target already-data-intensive users. For such open questions, it is also important to recognize an asymmetry between the perspectives of edge providers and ISPs – namely, what is small to an ISP may be large to an edge provider. For example, thresholds or caps applied to a small number of households for an ISP, such as 10% to 20% of access users, can have substantial impact on the business of edge providers. A small fraction of customers to an ISP can be a large fraction of demand to a provider of data-intensive services. Fear and uncertainty could exacerbate any response, which appears to have occurred in Canada. Hence, the answer from an edge provider to these open questions could diverge from the answer from an ISP to the same open questions.

Edge providers also express a different perspective on the effects of data caps on their incentive to innovate more efficient means of delivering their services. They stress that caps could impact the deployment of new innovative services and competitors because caps disincentivize the use of more data-intensive applications. For example, in 2012, a Sony executive suggested that the company was holding off its release of an Internet video service because of ISPs data cap implementation practices.<sup>28</sup> Edge providers also stress that the services provided by Internet applications and websites create the value from the broadband access product offered by ISPs. Edge providers do not deliver data unless it is requested by the customers of ISPs. ISPs have an obligation for carriage of content.

We have noted elsewhere that the user response to a data cap could be exacerbated by the absence of widely used measurement tools. Here too the perspective of an edge provider may differ from that of an access provider. If users knew the “data-intensity” for various applications, they could use that information to measure the incremental contribution of each application to additional capacity use and, accordingly, adapt their own use. So there may be a consumer information dimension to this topic. For example, many edge providers offer streams of content at multiple bitrates and detect connection speed to show users a higher or lower bitrate. It’s possible that edge providers could experiment with charging different prices for streams with different bitrates. What can be learned from experiments with such programs in mobile and low-bit-rate DSL?

These questions may become salient at some point for entrants who might anticipate growth in data use among US households. At what point do these concerns become urgent? If so, whose responsibility are they?

### **Specialized services and edge providers**

In some settings, an ISP is vertically integrated into the provision of services that substitute for services a user may access over the public Internet. Thus caps may provide a method for

---

<sup>28</sup> Timothy B. Lee, May 2, 2012, “Sony: Internet video service on hold due to Comcast data cap,” *Ars Technica* <http://arstechnica.com/tech-policy/2012/05/sony-warns-comcast-cap-will-hamper-video-competition/>, accessed May 17, 2013.

differential treatment of traffic or partners' traffic in order to favor certain applications provided by the ISP, like Voice Over IP (VoIP is a low-bandwidth application, in general).

Many aspects of this topic have been discussed by the Specialized Services working group, and we do not seek to replicate those findings here. That group has discussed questions related to incentives to build specialized services, different traffic metering to reflect different costs, difficulties with benchmarking performance in specialized services, and the different needs of distinct applications.

Here we focus on one key concern for competition policy. In general, competition policy is concerned about situations where one firm provides a service and also controls aspects affiliated with the cost, performance, and user-experience in a competing service. In public conversation this concern is often framed as a metaphor about the slope of the pitch: Does a cap or threshold tip the playing field by slanting consumers to an ISP or another online supplier? Said another way, what is a "level playing field" when a specialized service competes with an edge provider attempting to sell services that operate over the public Internet?

Despite the generality of the concerns, the answers are not sweeping or general. The specific details of this situation play an important role in determining appropriate policy. These concerns arise in a setting where managed service and Internet service use similar infrastructure, and the threshold or cap does not apply to a managed service but does apply to a range of arguably substitutable services. In such a setting, there is one set of prices and conditions for broadband service and another for the specialized service. Users pay a different price for each and have a different experience. Data caps may play a role in the prices users face and the experience they have between the two services.

This is another place where the ISP's perspective and the edge provider's perspective diverge. To see the divergence, it is useful to contrast these perspectives side-by-side.

From an ISP's perspective, since limitations do not apply to any but a small percentage of users, there is plenty of headroom for growth in competing services today and tomorrow. There is a rationale for separately provisioning between the specialized and non-specialized services, usually to achieve some engineering or market objective, such as improve the quality of service (e.g., reduce user perceptions of delay). In addition, one service often has a set of regulatory requirements associated with it, and one often does not. ISPs also note that the environment should promote innovation. For example, an ISP that is also an Incumbent Local Exchange Carrier (ILEC) transitioning to Voice over Internet Protocol (VoIP) may prioritize its VoIP traffic and exempt it from any usage threshold. In these instances, that ISP's exemption of its VoIP traffic is entirely consistent with how its traditional telephone service traffic has always been treated and should not be counted toward a cap. Any contrary conclusion would create a disincentive for the ILEC to migrate to IP and potentially stifle that migration.

From the perspective of an edge provider, similar services compete, using similar capacity, and the edge providers are providing innovative services. However, one has a threshold – say, from Hulu, Netflix, YouTube, Crackle, and competitors – and the other does not – from the ISP. The key concern is whether the rationale for distinct treatment of traffic in specialized services and

non-specialized services makes sense for the improvement of user performance, or is merely an excuse to put an edge provider competitor at disadvantage.

Does the concern arise when the thresholds are set comparatively high, as they tend to be for most ISPs today? The competition policy questions appear to be most salient in streaming of video services today, but may arise in services other than streaming. Similar issues may arise in home security systems and home video conferencing, for example. What is a level playing field in those cases?

It is difficult to forecast what users will want in a few years, and whether data caps will have any impact on those demands. It is also difficult to forecast what new applications edge providers will invent, what new specialized services ISPs will invent, and whether data caps will be relevant to their market experiences. There are both gains from flexible policy – to allow for new invention and the new situations created by invention – and gains from certainty – to allow edge providers and ISPs to plan for long-term investments. Therefore, the situation yields no easy answers in general, and, at a minimum, merits further monitoring.

## **Summary**

This study reviewed concerns with data caps and thresholds in the context of usage-based pricing in wire-line broadband services. The report focused on providing definitions, identifying the concerns of participants, and identifying the policy issues these raised. Many open questions emerged, and full or complete answers would require considerably more discussion.

## **Working Group on Economic Impacts of Open Internet Frameworks**

### **Chair:**

Shane Greenstein, Professor and Kellogg Chair of Information Technology, Kellogg School of Management, Northwestern University

### **Members:**

Brad Burnham, Founding Partner, Union Square Ventures

Neil Hunt, Chief Product Officer, Netflix

Kevin McElearney, Senior Vice President for Network Engineering, Comcast

Marc Morial, President & CEO, National Urban League

Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology  
(representing TMobile)

Charles Slocum, Assistant Executive Director, Writers Guild of America, West

## Further Reading

- Atkinson, Robert C., Ivy E. Schultz, Travis Korte, and Timothy Krompinger, 2011, *Broadband in America, 2<sup>nd</sup> Edition, Where It Is and Where It Is Going (According to Broadband Service Providers), An Update of the 2009 Report Originally Prepared for the Staff of the FCC's Omnibus Broadband Initiative*, CITI, Columbia University. May 2011.
- Bauer, Johannes, and Steven Wildman, 2012, *The Economics of Usage Based Pricing in Broadband Markets*, Michigan State University, December 14, 2012.  
[http://i.ncta.com/ncta\\_com/PDFs/Wildmanreport\\_web.pdf](http://i.ncta.com/ncta_com/PDFs/Wildmanreport_web.pdf).
- Bauer, Steve, David Clark, and William Lehr, 2012. *Understanding Broadband Speed Measurements*, (August 15, 2010), TPRC 2010, Available at SSRN:  
<http://ssrn.com/abstract=1988332>
- Cisco, 2013, *Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012 - 17*, White Paper, Cisco Public, San Jose, CA. February 6, 2013.  
[http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-520862.html](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html)
- Federal Communications Commission, 2010. *National Broadband Plan: Connecting America*. Washington, D.C. March, 2010. <http://www.broadband.gov/>.
- Federal Communications Commission, 2013, *2013 Measuring Broadband America February Report A Report on Consumer Wireline Broadband Performance in the U.S.*, Washington, D.C. February 2013. <http://www.fcc.gov/measuring-broadband-america/2013/February>.
- Higginbotham, Stacey, 2012, "Which ISPs are Capping Your Broadband and Why?" GigaOm, October 1, 2012. <http://gigaom.com/2012/10/01/data-caps-chart/>
- Hussain, Hibah, Kehl, Danielle, Lennett, Benjamin and, Lucey Patrick (2013), *Capped Internet: No Bargain for the American Public*, The New America Foundation.  
[http://www.newamerica.net/publications/policy/capped\\_internet\\_no\\_bargain\\_for\\_the\\_american\\_public](http://www.newamerica.net/publications/policy/capped_internet_no_bargain_for_the_american_public)
- Hussain, Hibah, Kehl, Danielle, Lennett, Benjamin and , Lucey Patrick (2012), *Capping the Nation's Future?* The New America Foundation.  
[http://www.newamerica.net/publications/policy/capping\\_the\\_nation\\_s\\_broadband\\_future](http://www.newamerica.net/publications/policy/capping_the_nation_s_broadband_future)
- Lyons, Daniel A., 2012, "*The Impact of data caps and other forms of usage based pricing for broadband access.*" Mercatus Center, George Mason University, Working Paper 12-27, October, 2012.
- Odlyzko, A., Arnaud, B. S., Stallman, E., and Weinberg, M. (2012). *Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service*. Washington, DC: Public Knowledge. [http://publicknowledge.org/files/UBP\\_paper\\_FINAL.pdf](http://publicknowledge.org/files/UBP_paper_FINAL.pdf)
- SandVine, 2012, *Global Internet Phenomena Report*, Revision 11-06-2013, Waterloo, Ontario, Canada.

## Open Internet Advisory Committee - 2013 Annual Report

Wallsten, Scott, and James Riso, 2010, *Residential and Business Broadband Prices Part 1: An Empirical Analysis of Metering and Other Price Determinants*, Technology Policy Institute, Washington, D.C., November 2010.

### Appendix 1.

	<b>Provider</b>	<b>Use Threshold – GBs</b>	<b>Excepted Traffic</b>	<b>Overage Charge</b>
1	Comcast	min 300GB (increasing by speed tier)	XFINITY Voice or Comcast Digital Voice (VoIP)	\$10 / 50GB (per tier)
2	AT&T - U-Verse HSIA	250	AT&T 3G MicroCell	\$10 / 50GB
	AT&T – DSL	150		
3	TWC	None	n/a	n/a
4	Verizon - FiOS / DSL	None	n/a	n/a
5	CenturyLink - 1.5 Mbps	150	Upload	None
	CenturyLink - >1.5 Mbps	250		
6	Cox - Ultimate (100 Mbps)	400	Cox Digital Voice (VoIP)	None
	Cox - Premier (25 Mbps)	250		
	Cox - Preferred (15 Mbps)	200		
	Cox - Essential (3 Mbps)	50		
	Cox - Starter (1 Mbps)	30		
7	Cablevision	None	n/a	n/a
8	Charter - Lite & Express ( )	100	None	None
	Charter - Plus & Max (30 Mbps )	250		

Open Internet Advisory Committee - 2013 Annual Report

	Charter - Ultra100 (100 Mbps)	500		
9	Frontier	100 / 250 in selected trial markets	None	None
10	Windstream	None	n/a	n/a
11	SuddenLink (>30 Mbps)	350	We prioritize Suddenlink voice packets in order to provide quality service to our phone customers.	\$10 / 50GB
	SuddenLink (10-30 Mbps)	250		
	SuddenLink (<10 Mbps)	150		
12	MediaCom - Launch (3 Mbps)	150	None	\$10 / 50GB
	MediaCom - Prime (15 Mbps)	250		
	MediaCom - Prime Plus (30 Mbps)	350		
	MediaCom - Ultra/Ultra Plus (50/105 Mbps)	999		
13	Cable One – Economy	Monthly: 1GB	0000-1200 Daily	?
	Cable One - Standard (5 Mbps)	Daily: 3GB	None	None
	Cable One - Preferred (50 Mbps)	Monthly: 50 GB	0000-0800 Daily	\$0.50 / 1 GB
	Cable One - Elite (50 Mbps)	Monthly: 100 GB	0000-0800 Daily	\$0.50 / 1 GB
	Cable One - Premium (10 Mbps)	Daily: 5GB	None	None
	Cable One - Ultra (12 Mbps)	Daily: 10GB	None	None

## Open Internet Advisory Committee - 2013 Annual Report

14	FairPoint	None	n/a	n/a
15	Cincinnati Bell	None	n/a	n/a

15	Google Fiber	None	n/a	n/a
----	--------------	------	-----	-----

	<b>Overage Treatment</b>	<b>Cite</b>
1	Comcast does not have a cap or usage threshold but is trialing two usage based pricing plans: one with a 300 GB threshold and another with varying thresholds (the lowest being 300 GB) based on service tier.	<a href="http://corporate.comcast.com/comcast-voices/comcast-to-replace-usage-cap-with-improved-data-usage-management-approaches">http://corporate.comcast.com/comcast-voices/comcast-to-replace-usage-cap-with-improved-data-usage-management-approaches</a>
2	Notice after 1st month; notices @ 65% & 90% in following months	<a href="http://www.att.com/esupport/article.jsp?sid=KB409045#fbid=kiJ0SSZjH9l">http://www.att.com/esupport/article.jsp?sid=KB409045#fbid=kiJ0SSZjH9l</a>
3	n/a	<a href="http://help.twcable.com/html/twc_sub_agreement.html">http://help.twcable.com/html/twc_sub_agreement.html</a>
4	n/a	<a href="http://www22.verizon.com/about/terms/networkmanagementguide/">http://www22.verizon.com/about/terms/networkmanagementguide/</a>
5	"Customers will be given options to reduce their usage, subscribe to a higher speed residential plan, or migrate to an alternative business class high-speed Internet service."	<a href="http://www.centurylink.com/Pages/AboutUs/Legal/InternetServiceManagement/">http://www.centurylink.com/Pages/AboutUs/Legal/InternetServiceManagement/</a>
6	If you do exceed your allowance, Cox will attempt to notify you by one or more methods: email, phone, or message on your computer before action is taken. We will then work proactively with you to resolve the problem. In many cases, customers are not even aware of their usage because they have an unsecured Wi-Fi network used by others or a computer virus. Cox can work with you to ensure that these issues are identified and corrected. In other cases, customers may choose to reduce their usage or switch to another plan that provides a higher usage allowance as Cox has assigned a different usage allowances to each of its Internet packages. In rare cases of extremely high usage Cox will suspend the user's service until they call Cox. In even rarer cases, Cox will terminate a customer's service if they do not decrease their usage after consultation with Cox.	<a href="http://ww2.cox.com/aboutus/northernvirginia/policies/speedsusage.cox">http://ww2.cox.com/aboutus/northernvirginia/policies/speedsusage.cox</a>
7	n/a	<a href="http://www.optimum.net/Privacy/AUP">http://www.optimum.net/Privacy/AUP</a>
8	Customers who exceed the "No Excessive Use of Bandwidth" section in the AUP may be notified by Charter that they have exceeded their monthly threshold and informed of Charter's Excessive Use policy. Charter Customer Care Representatives will help identify possible causes and offer suggested ways the customer can reduce bandwidth consumption. If the customer exceeds the "No Excessive Use of Bandwidth" policy and is notified three times in a six-month period, the customer's Internet service may be suspended after the delivery of the third notice.	<a href="http://myaccount.charter.com/customers/support.aspx?supportarticleid=2124">http://myaccount.charter.com/customers/support.aspx?supportarticleid=2124</a>
9	"In the affected markets, high bandwidth users (e.g. usage over 100Gb or 250Gb of data per month) are advised to either limit usage or convert to a high user service plan."	<a href="http://www.frontier.com/networkmanagement">http://www.frontier.com/networkmanagement</a>
10	n/a	<a href="http://www.windstream.com/Terms-and-Conditions/">http://www.windstream.com/Terms-and-Conditions/</a>

## Open Internet Advisory Committee - 2013 Annual Report

11	<p>After the first overage, the customer's Web browser will be directed to a Suddenlink notification page. The customer will be required to read that page, select how he or she wants to receive future notifications (by Web browser or email), enter the account number, and then save the information.</p> <p>From that point forward, future notifications on this subject will be sent each time an account reaches 80% of its monthly allowance and again when it exceeds 100%. Those notifications will be delivered through the means selected on the first overage, unless customers change their notification preference by visiting their Internet usage summary page at Suddenlink.net.</p> <p>Customer accounts will not be billed for exceeding their monthly allowance until the third overage. On the third and subsequent overages, the monthly allowance will be increased in installments of 50 GB at a cost of \$10 per installment.</p>	<a href="http://www.suddenlink.com/allowanceplan/">http://www.suddenlink.com/allowanceplan/</a>
12	<p>The data customers send and receive each month will contribute to monthly data usage. Speeds and usage allowances remain subject to change. Greater usage will result in additional charges of \$10, excluding taxes and fees, for every increment of up to 50 additional Gigabytes used. For example, if usage exceeds the allowance by 51 Gigabytes, an additional charge of \$20 will result.</p>	<a href="http://mediacomcable.com/site/internet.html">http://mediacomcable.com/site/internet.html</a>
13	<p>If a user that subscribes to the Economy plan exceeds the allocated monthly bandwidth of one gigabyte, Cable One automatically will allocate a second Gigabyte to the user for a fee set forth in the subscriber agreement. If the user exceeds the bandwidth allocated by this second Gigabyte, then Cable One automatically will allocate a third Gigabyte to the user for a fee set forth in the subscriber agreement, and so on. This incremental allocation of gigabytes is valid only for the billing cycle during which it was allocated and cannot be carried forward. The total number of Gigabyte allocations and related fees charged to the user in the Economy plan is capped in the subscriber agreement.</p>	<a href="http://www.cableone.net/Pages/internetaup.aspx">http://www.cableone.net/Pages/internetaup.aspx</a>
	<p>[I]f Cable One in its sole but reasonable discretion determines that a customer has exceeded the Excessive Use threshold or is using the Service in a manner significantly uncharacteristic of a typical residential user, Cable One reserves the right to (a) adjust, suspend or terminate Service accounts at any time and without notice; or (b) require the user to upgrade his service level or pay additional fees in accordance with Cable One's then-current, applicable rates and charges for such Service; or (c) use any technology to be chosen by Cable One at its sole discretion to slow the user's service for purposes of conserving bandwidth.</p>	
14	n/a	<a href="http://www.fairpoint.com/document/Residential_HSI_Terms_of_Service_tcm12-4842.pdf">http://www.fairpoint.com/document/Residential_HSI_Terms_of_Service_tcm12-4842.pdf</a>
15	n/a	<a href="http://www.cincinnati-bell.com/customer_support/consumer_information/network_management/wireline.pdf">http://www.cincinnati-bell.com/customer_support/consumer_information/network_management/wireline.pdf</a>
16	n/a	<a href="https://fiber.google.com/legal/network.html">https://fiber.google.com/legal/network.html</a>

Source: First fifteen observations accessed on February 6, 2013. Observation 16 accessed May 3, 2013.

## **AT&T/FaceTime Case Study**

Mobile Broadband Working Group  
Open Internet Advisory Committee  
Federal Communications Commission  
Released 1-17-2013

**The Mobile Broadband group created a document explaining the facts behind AT&T's limited rollout of FaceTime on its mobile network, and included a number of different opinions on whether the limitations were appropriate.**

The Mobile Broadband working group of the Open Internet Advisory Committee (OIAC) was formed to review the state of mobile broadband networks and assess how well Open Internet principles are working in practice. Although this report does not attempt to engage in any legal interpretations of the Open Internet Order, we do note that the Order [[http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/FCC-10-201A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf)] treats these mobile broadband networks differently from traditional fixed networks. While both fixed and mobile broadband providers must disclose their management practices, mobile broadband providers have greater latitude for blocking devices and applications (as long as they do not compete with the provider's own voice or video telephony services) and discriminating in how they serve traffic, in accordance with reasonable network-management practices.

The working group is investigating the tension between the goals of a free and open Internet, and the very real challenges that arise in managing mobile broadband networks. Such an investigation can easily devolve into vague discussions of high-level concepts or principles that may not be realizable in practice. To ground the discussion, the group started by considering several concrete case studies to help identify important trade-offs, principles, and other issues warranting further study, rather than trying to reach consensus on specific policy recommendations. The group explored one timely case study concerning how AT&T restricted the use of Apple's FaceTime application over its cellular data network to customers subscribed to a particular pricing plan. Video communication is widely viewed as the logical next step beyond the delivery of voice, text, and images over cellular data networks. Yet, these applications consume significant bandwidth and often have strict performance requirements, making them especially challenging for carriers to support efficiently. In the rest of this report, we discuss the specifics of the case study, analyze the high-level issues it raises, and present several possible conclusions from the unique perspectives of application developers, carriers, and equipment vendors.

### **AT&T and FaceTime**

FaceTime is a high-quality video-calling service created by Apple for use on the iPhone, iPad, and Mac. On the iPhone, rather than operating as a separate application, FaceTime is automatically integrated into the normal calling features of the user device. A user can upgrade a conventional phone call to include video simply by pressing a FaceTime button. Originally, Apple made FaceTime available only over wireless (WiFi) connections to the Internet, and the FaceTime calling features could not be used when devices were connected to a cellular network; however, that restriction was recently lifted, in part.

## Open Internet Advisory Committee - 2013 Annual Report

In June 2012, Apple announced that FaceTime would be available over cellular data networks, though Apple acknowledged that carrier restrictions may apply. In August 2012, AT&T announced that, in the wake of Apple's lifting of its restriction on FaceTime use, AT&T would limit the use of FaceTime over its cellular data network to customers of its MobileShare plans, in which multiple devices share a single limit for total data usage. Customers with "unlimited" data plans would not be able to use FaceTime on AT&T's cellular data network. The requirement for a specific plan would be enforced directly by the device, based on carrier settings [<http://support.apple.com/kb/HT1970>] (such as the current data plan or other eligibility information) learned from the carrier when the device authenticates with the cellular network.

Other providers, such as Sprint and Verizon, announced that FaceTime would operate over their cellular data networks for users of all billing plans [<http://9to5mac.com/2012/07/18/sprint-says-it-will-not-charge-for-facetime-over-cellular-verizon-calls-talk-premature/>, <http://arstechnica.com/apple/2012/09/verizon-will-enable-iphones-facetime-on-all-data-plans-unlike-att/>].

Some advocates and press denounced AT&T's decision, claiming that AT&T was violating the FCC's Open Internet Order [<http://www.savetheinternet.com/press-release/99480/att-blocking-iphones-facetime-app-would-harm-consumers-and-break-net-neutrality>, <http://publicknowledge.org/att-facetime>]. They argued that AT&T was blocking an application competing with its own voice or video telephony services, and that reasonable network management practices do not include favoring one pricing plan over another.

Responding to these claims, a blog post by AT&T [<http://attpublicpolicy.com/fcc/enabling-facetime-over-our-mobile-broadband-network/>] argued that AT&T's policy was fully transparent, and that AT&T does not have a competitive video calling application. AT&T also argued that the FCC's Open Internet Order does not regulate the handling of pre-loaded applications (i.e., applications integrated into the device's operating system, rather than installed manually by a user). AT&T also noted that all customers can continue running FaceTime over WiFi connections to the Internet.

In September 2012, several public interest groups announced their intent to file a formal complaint with the FCC [<http://arstechnica.com/tech-policy/2012/09/att-faces-formal-fcc-complaint-for-blocking-cellular-facetime-use/>], arguing that AT&T's restrictions of FaceTime usage violate the Open Internet Order. In October 2012, an AT&T customer in San Francisco filed a consumer complaint with the FCC concerning AT&T's blocking of FaceTime on his "unlimited" data plan [<http://www.businessinsider.com/consumer-fcc-complaint-att-facetime-2012-10>].

On November 8, 2012, AT&T announced [<http://attpublicpolicy.com/consumers-2/a-few-thoughts-on-facetime/>] plans to support FaceTime on all of its tiered data plans for users with an LTE device, over the next 8-10 weeks. AT&T customers with non-LTE devices or unlimited data plans would still not have access to FaceTime over the cellular network. AT&T also began rolling out new billing plans to enable deaf and hard-of-hearing customers to use FaceTime.

## Main Issues

AT&T's restrictions on the FaceTime application raise several interesting issues:

**Pre-loaded application:** Unlike many applications, FaceTime comes pre-loaded on a very popular phone. The application is immediately available to all users of the phone without requiring purchase or download, and is accessed via the core calling functions of the device. Every time a customer makes a phone call, the option of using FaceTime is immediately available. This makes it much more likely that the application would enjoy large-scale adoption very quickly. In addition, simultaneous use of the application (say, by spectators at a sporting event) could overwhelm the available radio network capacity, with its finite spectrum. In contrast, applications that require a manual download typically see lower penetration, even for popular applications that can be downloaded free of charge. For example, while around 75 million iPhones were sold in 2010, Skype was downloaded to only 7 million iPhones, resulting in less than 10% penetration [<http://www.statisticbrain.com/skype-statistics/>]. The rapid availability of FaceTime is said to be a particular challenge for AT&T, which historically has a much larger penetration of Apple iPhones among its customers, compared to other carriers [[http://news.cnet.com/8301-13579\\_3-57492508-37/iphone-owned-63-percent-of-smartphone-marketshare-at-at-t/](http://news.cnet.com/8301-13579_3-57492508-37/iphone-owned-63-percent-of-smartphone-marketshare-at-at-t/)]; today, more than half of AT&T's cellular data-network subscribers use an iPhone.

**High bandwidth requirements:** Cellular data networks have limited capacity, particularly in the "upstream" direction from user devices to the Internet; as such, carriers must carefully manage the shared "up-link" bandwidth to ensure reasonable performance for all users. While most content-delivery applications primarily impose load on the "down link," high-quality, video-telephony applications (like FaceTime) typically generate a large amount of traffic in both directions to deliver high-quality video to both participants in a video phone call. The quality of a multimedia application depends on the available bandwidth. Most popular applications adapt automatically in the presence of congestion, to decrease the quality of the audio or video stream to share bandwidth fairly with other applications. For example, data from Skype suggests that 128-300kbps is required for a standard video call [<https://support.skype.com/en/faq/FA1417/how-much-bandwidth-does-skype-need>], whereas various online reports suggest that FaceTime consumes around 100kbps - 1000kbps [<http://www.tested.com/news/254277-why-is-att-doing-you-a-favor-by-blocking-facetime/>, <http://www.padgadget.com/2012/06/20/concerns-about-facetime-over-cellular-will-you-max-out-your-data-limits>, <http://appadvice.com/appnn/2012/10/its-pretty-stupid-ridiculous-how-much-data-netflix-uses-over-lte>, [http://www.nokiasiemensnetworks.com/system/files/document/smart\\_labs\\_-\\_facetime\\_over\\_cellular\\_in\\_iphone\\_ios6\\_final\\_0.pdf](http://www.nokiasiemensnetworks.com/system/files/document/smart_labs_-_facetime_over_cellular_in_iphone_ios6_final_0.pdf)], consistent a limited set of measurements conducted at Bell Labs at the request of this working group. It therefore seems to be the case that FaceTime currently consumes on average 2-4 times more bandwidth than a similar Skype video call. It is important to note that there is no fundamental reason why FaceTime could not adapt to congestion the same way as other applications, and the way FaceTime behaves in the presence of congestion may easily change in the future.

**Staged deployment of new applications:** Rapid adoption of a new application might lead to large and unpredictable changes in the traffic load on a cellular data network. Carriers may want to

start with a limited trial deployment of a new application to better understand its effects before wide-scale deployment. This can provide measurement data and operational experience that carriers and application developers can use to make the most effective use of limited resources, or to identify appropriate policies for sharing resources with other applications. The AT&T/FaceTime case study raises an interesting question of whether or not restricting usage to customers of a particular pricing plan is a good way to limit (i) the number of users in an initial deployment (i.e., to users of a particular plan) or (ii) the total volume of traffic (i.e., by denying access to users with unlimited data plans), and what other alternatives might exist.

Application management on the device vs. the network: A carrier can block an application by discarding the packets it sends or receives; alternatively, a device such as a smart phone can prevent users from running a particular application, thereby keeping the traffic from ever reaching the network. In the AT&T/FaceTime case study, the usage of FaceTime on AT&T's network was limited directly on the device, rather than inside the network. An interesting policy question is whether it matters where an application-management decision is enforced, and which organization decides what policies to place on an application's use. In some cases, the creator of an application may want its users to enjoy unfettered access to the application, but in others the application developer may prefer to limit usage to ensure that supported users enjoy good performance; distinguishing between these two situations is surprisingly difficult. In this case, Apple and AT&T have not commented on which organization initiated the restrictions, and whether or not this was a collaborative decision.

These issues demonstrate the subtle trade-offs that arise in determining whether restricting FaceTime usage over AT&T's network constitutes blocking and/or reasonable network management.

### **Summary Opinions**

Different members of the working group came to different opinions about the restriction of FaceTime usage on AT&T's network. Generally, the working-group members agreed that blocking applications runs the risk of discouraging innovation, but that carriers also need effective ways to manage the limited resources in cellular networks. This led to three main opinions about AT&T's decision to restrict customer access to the FaceTime application over its cellular network, presented from the perspectives of different parts of the mobile broadband ecosystem -- application developers, carriers, and network equipment vendors. These opinions convey the conclusions of advocates for these perspectives among the working-group members, but do not attempt to fully represent each community.

#### **- From the perspective of application developers:**

AT&T did not choose the optimal approach by blocking access to the FaceTime application for customers on certain data plans. By singling out one popular application, the door is opened for carriers to block lawful use of applications, require customers to upgrade to potentially costlier, limited plans, and justify their actions by claiming to be engaged in reasonable network-management practices. Unfortunately, blocking a specific application for a large number of users on certain pricing plans, instead of managing the congestion that application and others might cause, sets a precedent that could have very negative consequences for the vibrant market for mobile applications. Allowing application blocking means that no developer could be sure that

his or her mobile application will be able to reach customers. If a carrier can block an application entirely at its discretion, investors will have to consider a new risk in addition to the normal risks faced by any start up. Unlike technical risk, financial risk, or organizational risk, the risk of being blocked cannot be mitigated. The existence of that risk will limit the investment available to applications developers, limiting the number of applications created, slowing innovation, and limiting consumer choice.

AT&T may have chosen to block FaceTime because it was a simple way to manage the potential congestion that could have occurred if the application were widely used. The carrier may have chosen to block FaceTime because it was concerned that broad use of a high-bandwidth data application by users of unlimited pricing plans would impact its profitability. Managing congestion and profitability are legitimate objectives for AT&T, but furthering those objectives by blocking specific applications is not the way to do it. There are many ways AT&T could have managed the roll out of FaceTime over cellular without taking the kind of application-specific action that harms applications developers and ultimately consumers. For example, AT&T could have instituted rate-limiting of individual customers, applied in a neutral manner, to limit congestion. Rate limits could be imposed at peak times or in response to congestion. In the medium- or long-term, AT&T could more aggressively scale up network capacity or apply other bandwidth-management techniques (such as WiFi offload) in localized hot spots where FaceTime and other high-bandwidth applications create congestion problems. AT&T can also charge users for the amount of data they consume, independent of the application. We recognize that these approaches require AT&T to deploy the technology in the network to actually manage the network, or to make the investment to market a new pricing plan to consumers. We understand that blocking FaceTime may be simpler and cheaper than deploying new network-management technology, increasing capacity, or changing pricing, but blocking a specific application chills investment, harms application developers, and reduces consumer choice. That is too high a price to pay when other alternatives are readily available.

In short, network management should focus on the underlying conditions that cause degraded performance of the network and address those conditions with solutions that optimize performance in a neutral manner for all users and applications. Such approaches -- indeed, all aspects of traffic management and engineering -- may require advanced planning to ensure that they are available when network conditions require them, but that fact makes them no less appropriate from a technical perspective. Application-agnostic network-management approaches should be considered and exhausted before application-specific approaches are even considered on a temporary basis, and customers should be able to have their choice of applications without having to change their data plans. Giving customers choice includes the option for user-controlled quality of service, where users decide to favor traffic from one application over another, in allocating whatever share of network bandwidth they receive from the carrier.

**- From the perspective of carriers:**

Given the bandwidth-intensive nature of the FaceTime application and AT&T's significant base of iPhone subscribers, AT&T has good reasons to be concerned about the potential for FaceTime to cause a focused, or localized, overload condition in its network. AT&T's approach of enabling FaceTime on Wi-Fi and on cellular for shared data plan subscribers is a reasonable way of managing the risk of network congestion. As data about FaceTime usage becomes available and

as its network evolves, AT&T has indicated that it may further expand the availability of the application. In fact, AT&T has already expanded the availability of the application to users with LTE devices on tiered service plans and on new custom plans for the hearing impaired.

AT&T's approach reduces the probability of a focused overload of its network due to FaceTime usage. By requiring a usage-based plan to access FaceTime over the cellular network, AT&T's approach both encourages use of the FaceTime service in a manner that is less likely to adversely impact the experience of other users on the network, and manages the number of subscribers that are likely to use such a bandwidth-intensive application. Usage-based data plans provide an incentive for users to manage their consumption of network bandwidth, and ensure that heavier bandwidth users pay a proportionate amount for their usage when compared to lighter bandwidth users. Unlimited data plans provide no incentive to users to manage the data consumed by bandwidth-intensive applications. Unlike some of its competitors, AT&T continues to offer unlimited data plans to existing subscribers to those plans, even when those subscribers upgrade to a new subsidized device. Since some carriers mandate that subscribers switch to a shared data plan when upgrading to a new device, AT&T's approach gives customers more flexibility than some of its competitors in choosing pricing plans and services that meet their needs. AT&T's announced expansion of FaceTime availability to LTE devices on individual tiered plans recognizes the increased capacity of its LTE network which, unlike its UMTS network, is not currently carrying voice calls, thus balancing the overall service quality for all of its customers.

While critics of AT&T's approach have described possible alternative approaches to the situation, none of the alternatives would effectively address AT&T's concerns. AT&T is aggressively expanding its cellular network capacity, and its devices are configured to support offload of data traffic to Wi-Fi networks where possible. AT&T currently operates over 30,000 Wi-Fi hotspots freely available to its data plan subscribers. While some have proposed rate limiting subscribers during periods of congestion, this approach is problematic for two reasons. One reason is that dynamic rate limiting is a complex mechanism that is not currently supported by wireless standards and vendor equipment. While dynamic rate limiting might be an option in the future, it is not an option that is available to AT&T today. The second reason is that dynamic rate limiting has the potential to degrade performance for both FaceTime and other applications. As a result, rate limiting may lead to more user dissatisfaction than AT&T's approach. This does not rule out dynamic rate limiting as a potential solution. However, it illustrates the complexity of providing good quality mobile broadband services.

While some have argued that AT&T's approach may adversely affect innovation, this risk can be mitigated by application developers by working cooperatively with carriers to build applications that do not risk harm to the network. In the case of FaceTime, the company developing the application built a mechanism into its operating system that enables operators to require certain plans. Other non-US carriers have used the same mechanism. Apple's page at <http://support.apple.com/kb/ht1937> shows the carrier-by-carrier breakdown of features supported by carriers world-wide. This specific example does not support the "chill to investment" argument, as the dominant player allowed its offering to be managed, which is rather different from a new entrant struggling to break in to a market.

In making these types of decisions, carriers are weighing multiple factors and taking competitive risks that may or may not succeed in the marketplace, but the marketplace can and should determine the success of these approaches. These decisions and the set of available techniques are not static and cannot be proscribed or regulated with any reasonable degree of applicability or validity over time.

**- From the perspective of network equipment vendors:**

Applications supporting real-time, two-way video calling such as Skype have become increasingly popular (more than 100 million logins/month and 30 million simultaneously active calls [<http://www.statisticbrain.com/skype-statistics/>]) and this popularity has increased with the availability of mobile clients for these applications. Given the significant additional bandwidth requirements of video sessions over voice calling, encoding the video frames at lower bit rates and the ability to adapt to changing network conditions such as the available bandwidth is key to the successful deployment or use of such applications. This is particularly true for mobile networks which represent a highly constrained and shared resource in both the uplink and downlink directions. For these reasons Skype utilizes adaptive session control techniques to constantly adjust the bit rate of the video stream transmitted between the two endpoints.

Apple's Facetime application is targeted to the same video calling market segment, but as noted above does not seem to adapt as readily/aggressively to changing network conditions. To illustrate the additional potential consumption compared to Skype usage, consider the following: if, as stated above, 10% of iPhone users were Skype users. When one compares this to the 100% of iPhone users who have access to the Facetime client and the at least 2x additional bandwidth consumption by the iPhone Facetime client compared to the Skype client, it is reasonable to conclude that the total network bandwidth usage (across all users and sessions) of Facetime could be as much as 20x higher than that of Skype, for operators who have a significant proportion of iPhones in their network.

In this context, it is reasonable to conclude that AT&T, with the largest number of iPhone users and largest fraction of their subscribers as iPhone users would have particular concerns about the load that the Facetime application would put on their network, with the potential to significantly degrade the available bandwidth for all other applications. Moreover, the concern would be most prevalent with respect to the most scarce resource -- the cellular network (which typically has ~20Mhz of spectrum compared to the more than 100Mhz of WiFi spectrum [[http://en.wikipedia.org/wiki/List\\_of\\_WLAN\\_channels](http://en.wikipedia.org/wiki/List_of_WLAN_channels)]) -- which is also the resource for which users have the highest service expectation. Given this, there would be a clear imperative to manage the usage of FaceTime application on AT&T's cellular network, with the option for unmanaged usage of FaceTime over their network. This is precisely the behavior that AT&T exhibited by limiting the usage of FaceTime to only a subset of their pricing plans, whilst making FaceTime available to all users over the WiFi interface. As such, it is reasonable to conclude that AT&T was trying to employ reasonable network management to the use of FaceTime over their network, albeit it in a relatively crude form.

It is interesting to contemplate whether there are alternative means by which the usage of FaceTime could have been managed in a way that would have made it available to all cellular users but in a scalable way. Clearly, if FaceTime was similar to Skype in terms of its bandwidth

utilization per session, or having the same device penetration (available on 10% of iPhones), no specific network management would have been required relative to that applied for Skype. Therefore an alternative approach would have been for AT&T to work with Apple to improve the bandwidth adaptation capabilities of the FaceTime application. Another alternative approach could have been to rate limit the usage of FaceTime in the network both on an individual session basis (per user), and an aggregate basis (total bandwidth allowed for all FaceTime users) using rate limiting techniques employed by some operators when usage caps have been reached, or for service plans that explicitly exclude usage of certain applications. Last, a non-application-specific rate-limiting approach could have been employed whereby the peak bandwidth usage by each user was limited when the network was congested. These approaches would have been reasonable and preferable in terms of the universal applicability and equanimity of the solution. It is important to note, however, that these alternative approaches may actually have resulted in a less satisfactory experience for all FaceTime users, or across all applications being used (for the non-application-specific approach), in contrast to the approach that AT&T took which likely resulted in a more satisfactory FaceTime experience, but for a subset of users. In other words, non-application specific approaches can appear 'fair' as they apply a 'one size fits all' philosophy whereby all users receive the same treatment for all applications. But, in some cases, and at some points in time, users may have a preference for a certain application (e.g. a FaceTime session for an important call) and would prefer it to be prioritized over other internet-based services when the network is congested.

### **Conclusion**

The three summary opinions capture different perspectives, with some overlapping points and differences in emphasis. Most members of the working group agreed with aspects of all three opinions, with some aligning more strongly with one view over the others. The case study also highlights the need for future cellular networking equipment and management systems to offer greater flexibility in managing the fine-grain sharing of limited network resources. This would make it easier for carriers to limit the impact new applications have on the performance experienced by other users using application-neutral techniques.

### **FCC Open Internet Advisory Committee Mobile Broadband Working Group**

- Chair: Jennifer Rexford, Professor of Computer Science, Princeton University
- Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla
- Brad Burnham, Founding Partner, Union Square Ventures
- Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
- Charles Kalmanek, Vice President of Research, AT&T
- Matthew Larsen, CEO, Vistabeam
- Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)
- Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
- Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
- Jonathan Zittrain, Professor of Law and Professor of Computer Science, Harvard; OIAC Chair (ex officio)
- David Clark, Senior Research Scientist, MIT Computer Science and Artificial Intelligence Research Laboratory; OIAC Vice Chair (ex officio)

# Openness in the Mobile Broadband Ecosystem

Mobile Broadband Working Group  
Open Internet Advisory Committee  
Federal Communications Commission

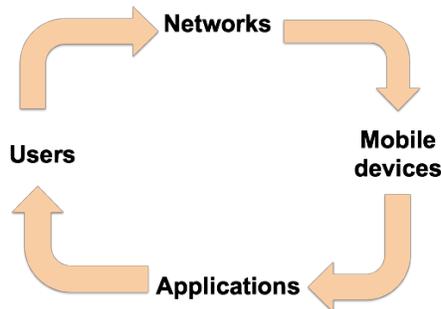
**The Mobile Broadband group also created an analysis of the mobile broadband ecosystem, identifying key players and articulating their relationships.**

The FCC’s Open Internet Order<sup>29</sup> characterizes “openness” as “the absence of any gatekeeper blocking lawful uses of the network or picking winners and losers online” and indicates that the openness of the Internet promotes a self-reinforcing “cycle of investment and innovation” (p. 3). In the mobile broadband ecosystem, a variety of players have significant roles in shaping the opportunities that the Internet provides, including mobile broadband providers (e.g., Verizon, AT&T, Sprint, and T-Mobile), device vendors (e.g., Apple, Samsung, and LG), operating system developers (e.g., Apple iOS and Google Android), network equipment vendors (e.g., Ericsson, Alcatel-Lucent, and Nokia-Siemens), and application developers and content providers.

This report examines the relationships between these parties and highlights the different kinds of influence they can have over openness, broadly defined. While many of these parties are not subject to the Open Internet Order, understanding the impact they can have on openness provides a more complete picture of the mobile broadband ecosystem. Because of our specific focus on mobile broadband, our analysis inherently reflects business and technical dynamics that may differ from those for fixed broadband networks. Also, while mobile broadband networks carry a variety of traffic (e.g., downloading e-books to Kindle devices, machine-to-machine communication, connected cars, etc.), this report focuses on the general, universal service that connects end-user mobile devices to the Internet.

## 1. Mobile Broadband Ecosystem

The mobile broadband ecosystem is built on a seemingly “virtuous cycle,” where *networks* that are fast, reliable, and widely available encourage the creation of *mobile devices* that connect to these networks, which spurs innovation in compelling *applications and content*, which in turn motivate more *users* to adopt the technology, spurring further investment in the underlying networks.



<sup>29</sup> FCC Open Internet Report and Order, December 2012.  
[http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/FCC-10-201A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf)

## Open Internet Advisory Committee - 2013 Annual Report

Yet, the players in the mobile broadband ecosystem have complex relationships that can cause tensions that can dampen the incentives for innovation and investment. The main parties include the network (i.e., *mobile broadband providers* and network *equipment vendors*), the devices (i.e., *device manufacturers* and *operating-system developers*), the applications (i.e., *application developers*), and the *component manufacturers* who make the components used in mobile devices and network equipment.

### 1.1 Major Mobile Broadband Companies in the U.S. Market

In most sectors of the mobile broadband ecosystem, a small number of companies drive the market, as shown in the following table:

	<b>Ecosystem Players in the U.S. (1Q 2013)</b>
Smartphone vendor shipments <sup>30</sup>	Apple (38.3%), Samsung (28.8%), LG (9.9%), and many smaller players (< 5% each)
Smartphone OS market share (through 1Q13) <sup>31</sup>	Google Android (56.0%), Apple iOS (38.3%), and other smaller players (< 4%)
Mobile broadband provider market share <sup>32</sup>	Verizon Wireless (34%), AT&T Mobility (30.9%), Sprint (16%), T-Mobile USA (12.2%), and other smaller players (< 3%)
Radio access network equipment vendors <sup>33</sup>	Ericsson (50%), Alcatel-Lucent (36%), Nokia-Siemens (10%), Huawei (3%)
Application developers <sup>34</sup>	Many, diverse, most make < \$500/month

A few main vendors lead the sectors for creating smart phones (e.g., Apple, Samsung, and LG) and the operating systems that run on them (e.g., Google Android and Apple iOS), along with some smaller players. The U.S. has four main mobile broadband providers (Verizon, AT&T, Sprint, and T-Mobile). Mobile broadband providers can acquire equipment for cellular access networks from three main vendors (Ericsson, Alcatel-Lucent, and Nokia-Siemens), with Samsung a new entrant into the U.S. LTE equipment market, and Huawei a smaller player in some U.S. regional markets. In addition, a small number of companies create most of the components used in handsets (e.g., Qualcomm and Samsung) and the components used in network equipment (e.g., Texas Instruments, Broadcom, and Freescale). In contrast, the applications sector is extremely large and diverse, with many thousands of developers

<sup>30</sup> Strategy Analytics, “North America Smartphone Vendor & OS Market Share by Country: Q1 2013,” May 2013

<sup>31</sup> Ibid.

<sup>32</sup> Strategy Analytics, “Wireless Operator Performance Benchmarking Q4 2012,” April 2013

<sup>33</sup> Alcatel-Lucent internal analysis of Dell’Oro data, average over the last four quarters.

<sup>34</sup> Source: Vision Mobile

creating applications that compete for users' attention. The app market generated more than 13.4 billion downloads and \$2.2 billion of revenue<sup>35</sup> in the first quarter of 2013 alone. While most application developers operate at a very small scale (e.g., making less than \$500 per month), half of all app revenue comes from just 25 developers<sup>36</sup> --- mostly major game developers such as Zynga, Electronic Arts, Rovio, and Disney.

While mobile broadband providers are typically regional or national companies, the rest of the mobile broadband ecosystem is an international marketplace. While most of the leading companies in the U.S. have significant market share internationally, some companies play a much larger role in the rest of the world; for instance, Huawei has a much larger market share in the network equipment market internationally. Historically, the U.S. was the leader in cellular deployments, but lost the lead to Europe in 2G (GSM) and to Asia in 3G (WCDMA), before regaining the lead again with 4G (LTE). The U.S. also leads the recent innovations in smart phones, mobile operating systems, and applications. Still, the manufacturing of components and handsets mainly takes place outside the U.S., and the mobile broadband ecosystem relies heavily on international agreement for technology standards. In addition, many new mobile-broadband business trends, such as the decreasing role of carrier subsidies for mobile handsets, started outside the U.S., providing a unique opportunity to analyze the effects of emerging trends.

Some companies play a significant role in multiple parts of the mobile broadband ecosystem, giving them extra influence. While industry forces often work against having a primary “vertical player” (e.g., Motorola, in earlier days), several companies increasingly play multiple roles in the mobile broadband sector. For example, the top handset manufacturer (Samsung) also sells LTE equipment, as well as the low-level components used in other handsets (such as the Apple iPhone)<sup>37</sup>. Huawei also sells both mobile devices and network equipment. As such, Samsung and Huawei can have a unique relationship with carriers, by having bundled offerings of handsets and network equipment. Apple and Google also have significant influence in multiple parts of the ecosystem. Apple creates devices (e.g., iPhones and iPads) that are tied to its own operating system (iOS), and also develops mobile applications that come bundled with the device. Google has the lead mobile operating system (Google Android), and also creates popular applications and, recently, mobile handsets. In the next subsection, we discuss the interaction between these and other companies in the mobile broadband ecosystem.

### 1.2 Complex Inter-Relationships in the Mobile Broadband Ecosystem

Each of the players in the mobile broadband ecosystem is affected by the policies and practices of the others, including:

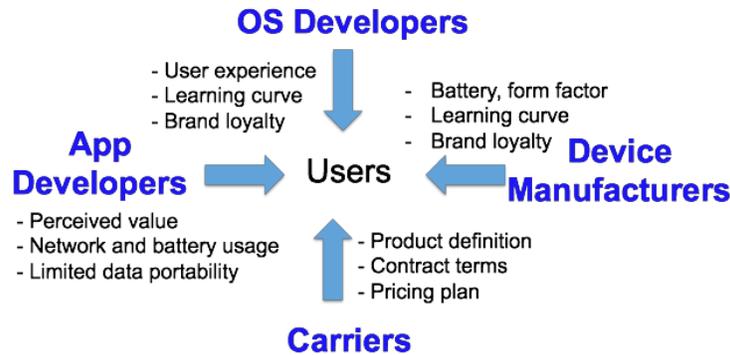
**Users:** End-users identify strongly with their mobile devices, from the early Razr flip phone to the Apple iPhone. With the emergence of smart phones, users increasingly associate their entire mobile broadband experience with their device, and particularly with the operating system (e.g., Apple iOS and Google

<sup>35</sup>[http://news.cnet.com/8301-1035\\_3-57578563-94/app-market-soars-with-13.4-billion-downloads-in-q1-2013/](http://news.cnet.com/8301-1035_3-57578563-94/app-market-soars-with-13.4-billion-downloads-in-q1-2013/)

<sup>36</sup><http://www.canalys.com/newsroom/top-25-us-developers-account-half-app-revenue>

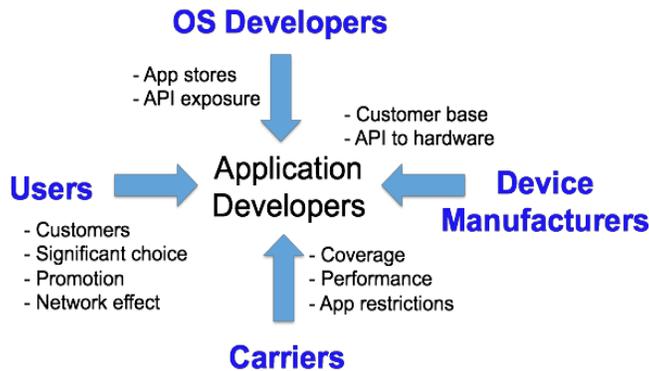
<sup>37</sup><http://www.economist.com/blogs/dailychart/2011/08/apple-and-samsungs-symbiotic-relationship>

Android) and its associated “app store”. Using the same platform as friends and family members also eases communication through instant messaging, video conferencing, and photo sharing applications bundled with the operating system.



Many users stay with the same platform over time, due to brand loyalty, adoption of built-in features like automatic syncing of data with cloud services (e.g., Apple iCloud), and the learning curve for adapting to a new operating system. The users increase the value of their mobile devices through mobile applications, some of which come pre-installed on the device; these applications may also have a significant impact on battery lifetime and bandwidth consumption, though most users have difficulty determining which of their applications are the “resource hogs.” Despite the emphasis on the device and the applications, the relationship with the mobile broadband provider is important, too. Most users receive a handset as part of the service contract from their carrier, though the emergence of tablet computers, and changes in the device pricing model being introduced by some carriers (e.g., T-Mobile), are increasing the fraction of mobile devices purchased directly. The mobile broadband provider also has significant influence over the users in terms of pricing plan (e.g., unlimited bandwidth, bandwidth caps, or usage-based billing) and contract restrictions (e.g., early-termination fees, limitations on tethering, etc.).

**Application developers:** The ecosystem includes a large and diverse group of developers creating applications for a variety of platforms (e.g., Apple iOS and Google Android).

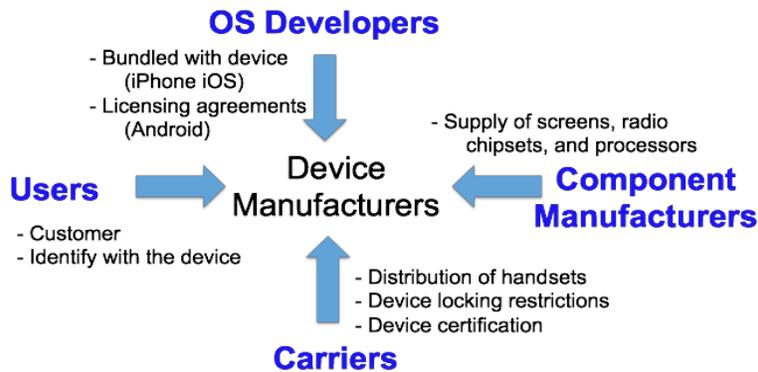


Applications range from network and device utilities, to mobile access to online content, to mobile games, and location-centric applications. Creating a successful application is challenging, and typically requires creating a separate version for each operating-system platform, and relying on whatever Application

## Open Internet Advisory Committee - 2013 Annual Report

Programming Interfaces (APIs) the operating system developers and device manufacturers make available. A range of business models have emerged, as application developers and consumers experiment with different monetization paths, including initial purchase price, “freemium” or free download with limited functionality and pay-to-upgrade charges, ad-supported, and free (or paid) download with in-app purchasing of extras or subscription services. Application developers are somewhat dependent on “App Stores” (the largest app stores are operated by Apple and Google) to distribute their applications, in exchange for a fraction (e.g., 20-30%) of their revenue. In addition, the large number of available “apps” mean that users have tremendous choice, forcing developers to keep prices low to compete with free or low-cost apps; many apps rely on advertising for revenue, and “word of mouth” from users to promote their applications. In addition, application developers rely on mobile broadband providers for good coverage and performance, and are subject to the terms and conditions of the end-user’s service contract which may restrict the use of certain apps.

**Device manufacturers:** Devices such as smart phones, tablets, and smart meters connect to mobile broadband networks. Many end-users identify more strongly with their mobile devices than with their mobile broadband provider.

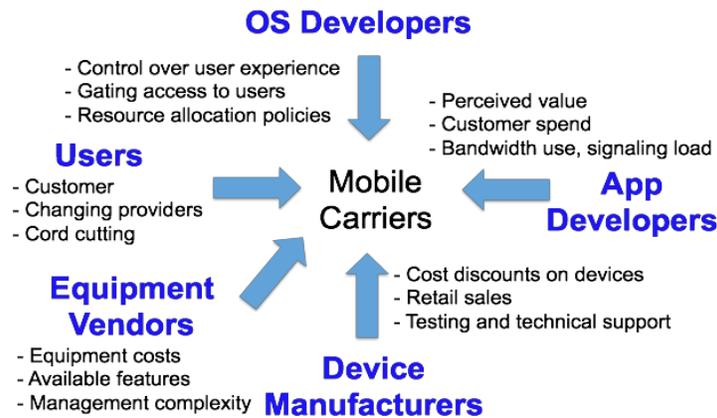


While many handset manufacturers rely on mobile providers to offer sizeable discounts on price of devices sold to consumers (colloquially known as “device subsidies”), the market increasingly includes tablet computers that are sold directly to consumers. Most mobile providers “lock” handsets on their networks, restricting their customers from using the devices with other carriers. The device manufacturers also rely heavily on the component manufacturers for a regular supply of parts. Companies like Qualcomm, Samsung, Intel, and Infineon make radio chipsets and processors that govern radio network operations and compatibility, features, and performance. Even if existing components are limited in functionality, device manufacturers typically find that building their own components is prohibitively expensive. The relationship with component manufacturers is particularly complicated if the company also sells its own mobile devices; for example, Samsung is a leading manufacturer of mobile handsets but is also the primary supplier of screens for its chief device competitor, Apple.

**Operating-system developers:** The operating system (OS) runs on the devices and provides a development platform for applications. In some cases, the operating system is provided by the device manufacturer (e.g., Apple iOS and Blackberry OS). In other cases, the operating system is provided separately (e.g., Google Android and Microsoft Windows Mobile). Some operating system developers

seek to limit the “fragmentation” of the OS software to avoid problems with interoperability, where applications work on one device but not another. Yet, device manufacturers may want to customize the software or experiment with new features. Though Google’s Android operating system is open source, recent changes in the terms of service<sup>38</sup> for the Android software development kit prevent developers from creating their own “fork” of the code, to reduce code fragmentation. Similarly, Microsoft’s Windows Mobile 7.5/8 is specifically licensed to select hardware partners under terms that greatly limit the variability of the OS implementation across devices. While Android and Apple iOS are by far the largest players in the mobile OS market, the landscape sometimes changes rapidly, as evidenced by the rapid penetration of Google Android OS in the past few years. There are also efforts to launch new, competitive operating systems, such as Mozilla’s Firefox OS and Samsung’s Tizen. Each OS platform also has very different philosophies towards “openness,” with regard to both the OS itself and the application environment it enables.

**Mobile broadband providers:** Users typically pay mobile carriers to access mobile network services, either through a “post-paid” monthly subscription or a “pre-paid” monthly purchase.

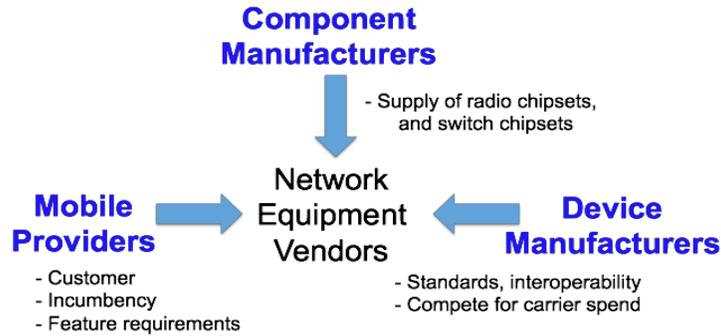


Historically, mobile carriers tightly controlled both the devices and services available to users, but the ecosystem has evolved such that operating system developers, device manufacturers, and application developers have greater control over the user experience and the consumption of network resources. Users who identify primarily with their mobile device may be more willing to change providers at the end of their service contracts, leading to competition over service plans across carriers. The design decisions by application developers influence the consumption of network bandwidth and signaling resource and can degrade performance for all users in congested cells. For example a “chatty” application that sends regular updates every 60 seconds can easily overwhelm signaling resources on the radio access network. The rapid emergence of new applications written by a large community of developers with widely varying expertise makes managing a carrier network challenging. Carriers have little ability to influence a user’s choice of applications or an application developer’s efficiency in using network resources other than through various forms of usage-based pricing. If data usage continues to grow, carriers will face significant costs to expand network capacity. Carriers’ technical options for managing network resources are also limited by the capabilities in the underlying network equipment and mobile devices. Carriers may

<sup>38</sup> [http://www.theregister.co.uk/2012/11/15/android\\_sdk\\_fragmentation\\_license\\_change/](http://www.theregister.co.uk/2012/11/15/android_sdk_fragmentation_license_change/)

also limit their experimentation with alternative network-management practices to avoid drawing attention from regulators like the FCC.

**Network equipment vendors:** Mobile broadband providers rely on equipment like cellular base stations, serving and packet gateways, and mobility control software to build and manage their mobile broadband infrastructure.



Buying this equipment is a significant capital expense for the carriers as they expand their network footprint, and the capabilities of the equipment influence how the operators can manage their customers' traffic. This, together with the entrance of low-cost players, has driven the rapid commoditization of the network equipment market, and an attendant limit in the level of research and development that can be supported. While the network equipment vendors do not interact directly with end users, or the application and operating system developers, the interplay with device manufacturers is more significant. The network equipment and mobile devices must implement the same standard protocols for the radio access network, leading to cooperation (and competition) in standards bodies leading to more complex standards and the need for extensive interoperability testing. In addition, network equipment vendors must compete with device manufacturers for the limited capital the carriers have to spend on equipment and device subsidies. The network equipment vendors are also dependent on the component manufacturers (e.g., Texas Instruments, Broadcom, and Freescale) that make the chipsets used in their equipment for the radio access and cellular core networks.

In conclusion, the mobile broadband ecosystem has complex power dynamics that affect the incentives each party has to invest in innovation. These dynamics shift rapidly over time in response to business trends (e.g., the prominence of the Blackberry giving way to the iPhone, the emergence of the open Android operating system as a replacement for Apple iOS, and the transition from circuit voice to VoIP with the attendant ecosystem changes). In the next section, we present several case studies of technology and business trends that are affecting openness in the mobile broadband ecosystem.

## 2. Case Studies

In this section, we present several case studies that illustrate how the inter-relationships in the mobile broadband ecosystem can affect the incentives of different parties to invest and innovate.

### 2.1 App Stores: App Developers and Operating System Developers

## Open Internet Advisory Committee - 2013 Annual Report

App stores have become an omnipresent feature of mobile broadband. Consumers and app developers both benefit from the convenience that they provide, but app store operators can also restrict the development of mobile applications by leveraging their control over which applications are made available and under what conditions. This section explains some of the motivations for the creation of app stores, explores how app stores may impede openness, and discusses how the trend towards web-based app development might change these dynamics in the future.

The development of mobile app stores – and the app-centric nature of the mobile environment in general – is in some ways a reaction to issues that have arisen with other common software distribution models: traditional desktop software and web-based applications.

In the desktop environment, installed programs have access to a computer's operating system under permission systems that vary as to their robustness and security properties. During the early to mid-2000s, prevalence of malware on personal computers was especially high<sup>39</sup>. The rise in malware was correlated with the emerging prevalence of downloadable, executable content and a runtime model that allowed users to easily and inadvertently introduce malicious code into their machines. Thus, the pure desktop model, with associated malware risks, was seen by some early smartphone innovators as inappropriate for smartphones<sup>40</sup>.

Web-based applications, on the other hand, are becoming increasingly robust and are generally safer to run by virtue of the fact that they are confined to the browser<sup>41</sup>. Unfortunately, web applications still lack direct access to many mobile devices' underlying functionality and hardware and thus cannot perform the same functions or provide the same performance as local apps. Although the continued development of HTML5 (discussed below), sophisticated JavaScript APIs, and other web technologies are rapidly pushing web apps forward, in-browser applications still lag behind in some cases in terms of functionality and convenience.

The app-centric model for mobile broadband has therefore been viewed as a way to combine trust and functionality. Apps often undergo review by platform providers and run in a semi-sandboxed environment on the phone's software platform, increasing trust. Because they run locally on the device, they can be hardware-accelerated and have access to a more rich suite of device features than web-based apps.

Apple, Google, Microsoft, and other app store providers shape these dynamics and the overall openness of the mobile app landscape through the policies that they set. These policies concern a variety of technical, operational, and business aspects, including:

- **Installation sources:** On some devices and operating systems (notably Apple's), going through the app store is the *only* way to install an app on non-jailbroken devices. Apple allows web-based applications to be saved as bookmarks, but the user interface and interactions with web

---

<sup>39</sup> [http://download.microsoft.com/download/1/A/7/1A76A73B-6C5B-41CF-9E8C-33F7709B870F/Microsoft\\_Security\\_Intelligence\\_Report\\_Special\\_Edition\\_10\\_Year\\_Review.pdf](http://download.microsoft.com/download/1/A/7/1A76A73B-6C5B-41CF-9E8C-33F7709B870F/Microsoft_Security_Intelligence_Report_Special_Edition_10_Year_Review.pdf)

<sup>40</sup> [http://www.nytimes.com/2007/01/11/technology/11cnd-apple.html?\\_r=0](http://www.nytimes.com/2007/01/11/technology/11cnd-apple.html?_r=0)

<sup>41</sup> <http://blog.chromium.org/2008/10/new-approach-to-browser-security-google.html>

## Open Internet Advisory Committee - 2013 Annual Report

bookmarks and installed apps are not always equivalent. In contrast, Google Play is one of many avenues for app developers to get their apps onto Android devices; Android users can download apps directly from web sites or from other app stores and the OS includes a setting that allows users to “accept apps from unknown sources.” Established providers such as Amazon have created their own app stores and developer resources to get apps onto Android-based devices, such as the Kindle Fire.

- **Screening policies:** App store providers have a variety of policies and procedures for screening apps before and after they have been placed in the store. Apps may be reviewed for performance, functionality, access to user data, security, user interface design, and content. Apple reviews all apps before they can appear in the App Store, rejects those that do not meet its App Review Guidelines<sup>42</sup>, and may remove apps even after they have been approved. Microsoft uses a similar process and policy<sup>43</sup>. Google generally does not do up-front app screening but removes apps from Play that are found to have security vulnerabilities or that violate Google’s terms<sup>44</sup>. Google has also removed specific tethering apps from its app store, reportedly at the request of carriers, because carriers forbid the use of tethering in some of their service plans<sup>45</sup>. Incidentally, the mobile OS vendors also have the capability to remotely uninstall malicious apps<sup>46</sup> directly from users’ devices.
- **Revenue-sharing requirements:** App store providers can establish terms that allow them to retain a portion of apps’ purchase prices, in-app subscription fees, or ad revenue. Apple, Google, and Microsoft generally retain a 20-30% share of app purchase prices (as does Amazon for its Android-based store)<sup>47</sup>. They may also set the terms about how subscriptions and content can be sold within apps<sup>48</sup>.
- **App store navigation:** App store providers choose which apps to feature prominently in their stores and how to categorize apps, at times making decisions that run counter to app developers’ desires<sup>49</sup>.

All of these policies have the potential to limit the openness of mobile app development. Developers that want to be able to reach users of non-jailbroken Apple devices have no choice but to comply with the terms that Apple sets for the App Store, including the revenue-sharing policies, standards concerning what Apple considers to be “objectionable” content, and technical limitations that include the inability to

---

<sup>42</sup> <https://developer.apple.com/appstore/guidelines.html>

<sup>43</sup> <http://msdn.microsoft.com/en-us/library/windows/apps/hh694083.aspx>

<sup>44</sup> <http://play.google.com/about/developer-distribution-agreement.html>

<sup>45</sup> [http://news.cnet.com/8301-30686\\_3-20059461-266.html](http://news.cnet.com/8301-30686_3-20059461-266.html)

<sup>46</sup> <http://latimesblogs.latimes.com/technology/2011/03/google-removing-virus-infected-android-apps-from-phones-tablets-promises-better-security.html>

<sup>47</sup> <https://developer.apple.com/programs/ios/distribute.html>

<sup>48</sup> <https://developer.apple.com/in-app-purchase/>

<sup>49</sup> <http://www.businesswire.com/news/home/20130314005784/en/Adblock-Reports-Removal-Google-Play-Store-Android>

## Open Internet Advisory Committee - 2013 Annual Report

obtain administrative privileges, tether, or alter the “look and feel” of the app<sup>50</sup>. The Android ecosystem is free of many of these limitations, but Google still retains the final say over which apps may appear in Google Play and how easy they are to find and use. On some devices, Google Play is a central source for Android apps despite there being other ways for users to obtain them.

In principle, the convenience and security of the app store model need not be tied to store provider policies limiting the operation or availability of certain apps. Cydia, for example, provides an app store and directory for jailbroken Apple devices, allowing users to more easily discover apps without subjecting app developers to restrictive installation policies or revenue-sharing agreements. While app stores play a pivotal role in the user experience of mobile broadband, it is important to distinguish between the barriers erected by app stores’ policies, technical limitations on app development that may be platform-specific but unrelated to app store policies, and the security properties that motivated the development of app stores in the first place. For example, operating system vendors could make the full suite of hardware APIs available to all browsers and apps while still retaining an app store model. This would ease the development of independent apps, but would still subject app developers to the terms set by the app store providers. By the same token, sandboxing and other techniques for making code execution safer could be supported by operating system vendors regardless of whether they enforce an app store model on their platforms or not.

One trend that may shift developers and users away from existing app store models is the continued maturation of the suite of HTML5 technologies<sup>51 52</sup>. HTML5 comprises the latest versions of the building blocks of the web plus a wide variety of newly developed APIs that give mobile developers access to critical device functionality, including sensors (camera, microphone, etc.), the file system, network interfaces, graphics support, and much more. Because it is based on open, interoperable web standards, the HTML5 technology suite allows developers to build applications from a single code base that work on any device with an up-to-date browser -- which means most any smartphone or tablet already in use. Thus, as HTML5 takes hold as an app development platform, developers will be able to distribute their apps across platforms, independent of whether they are also offered in app stores. HTML5 also includes a variety of security features designed to prevent the kinds of attacks that are often associated with downloadable software and that motivated the development of app stores.

Many HTML5 components are already fully functional and supported by the major browsers, but certain parts of the technology suite are still in the process of being developed and standardized, and questions remain about whether web-based apps can match the performance and user experience of platform-specific ones. As the tools that developers need to create HTML5-based apps that are equivalent or superior to platform-specific apps become increasingly available, the role of app stores in influencing which apps are available and under which conditions may be diminished.

### **2.2 Service Agreements: Users and Mobile Broadband Providers**

Mobile broadband providers have a direct influence on how their customers can access networked services. Service agreements constrain how customers can use their mobile devices. These agreements

---

<sup>50</sup> <https://www.eff.org/deeplinks/2012/05/apples-crystal-prison-and-future-open-platforms#gatekeeper>

<sup>51</sup> <http://www.w3.org/TR/html5/>

<sup>52</sup> <http://www.whatwg.org/specs/web-apps/current-work/multipage/>

## Open Internet Advisory Committee - 2013 Annual Report

illustrate the tensions between the providers' need to limit financial risks (e.g., in discounting or “subsidizing” handsets for customers willing to sign a long-term contract, expanding network capacity, and setting prices for multi-year contracts) and the benefits of giving users flexibility in how they use their mobile devices in a rapidly changing environment.

**Billing model:** Most mobile broadband providers offer service contracts with a variety of pricing plans. Over the years, unlimited, “all you can eat” data plans have largely given way to plans with bandwidth caps (where subscribers lose network speed after exceeding the cap) or additional charges for additional increments of bandwidth consumption. Still, some providers have many subscribers on “grandfathered” unlimited data plans, increasing the likelihood of high bandwidth consumption when certain applications (e.g., streaming video) or user practices (e.g., tethering) become popular. To manage traffic from these subscribers, some carriers “throttle” top users (i.e., limiting their bandwidth consumption during periods of peak load). Usage caps and usage-based billing encourage users to limit their use of network bandwidth (or defer usage until wired or WiFi connectivity is available), while only indirectly constraining usage during periods of peak load. Alternatives like time-dependent pricing, where providers offer lower prices during off-peak hours (and higher prices when the network is congested), have received significant academic attention, but to our knowledge have not been offered in the market.

**Device locking:** Many carriers provide customers with a “locked” phone that cannot work with other carriers. Software on the phone ties the subscriber ID (on the SIM card in GSM phones) to the serial number of that particular phone, preventing the customer from using the SIM card in a different phone, or using the phone with a different SIM card. While unlocked phones are common in Europe, most U.S. providers offer locked phones that prevent customers from switching service providers (without buying a new phone), temporarily using a different SIM card during international travel to avoid large roaming fees, or selling an old phone to another user. Providers vary in whether they offer unlocked cell phones (possibly at a higher price) or are willing to unlock a phone after the contract ends (i.e., after recouping the cost of the device subsidy). Recently, the Library of Congress moved to ban mobile users from unlocking their phones without the carriers' permission<sup>53</sup>, treating attempts to circumvent device locking as violating the anti-circumvention provisions of the Digital Millennium Copyright Act (DMCA). In response, some regional and rural providers have supported efforts to allow users to legally unlock their phones<sup>54</sup> without their providers' permission.

**Tethering:** Many providers restrict customers from “tethering” to share a mobile broadband connection with other devices, such as a laptop. Some providers do not allow tethering on certain data plans (e.g., unlimited plans), or require customers to pay extra (above the normal cost of their data plan) for tethering. The rationale is that tethering often leads to a substantial increase in bandwidth usage, beyond what the provider may have anticipated when designing its network and pricing structures. In 2012, Verizon was accused of requesting that Google remove tethering applications from the Android app store, so customers could not use these applications as a way to avoid paying a \$20/month tethering fee. The FCC ultimately reached a consent decree<sup>55</sup> and settlement with Verizon, under the terms of which Verizon

---

<sup>53</sup> <https://dl.dropboxusercontent.com/u/3155588/SJUD%20cell%20phone%20bill.pdf>

<sup>54</sup> <http://www.mobilenapps.com/articles/7901/20130314/phone-unlocking-small-carriers-backing-bill-for-apples-iphone-access.htm>

<sup>55</sup> <http://www.fcc.gov/document/order-and-consent-decree-verizon-wireless-pay-125-million>

## Open Internet Advisory Committee - 2013 Annual Report

could not block access to tethering applications<sup>56</sup>, making it possible for users with unlimited data plans to tether without paying extra charges; customers subject to usage caps or usage-based billing would have their tethering traffic metered just like any other data traffic. This decision by the FCC was specific to Verizon (under the conditions attached to spectrum licenses that Verizon purchased at auction), and the FCC has not taken any action as to other providers.

**Application restrictions:** Some providers impose restrictions on what mobile applications a subscriber can run under specific pricing plans. A good example is the evolution of AT&T's policies concerning Apple's FaceTime application for high-quality video calls, as discussed in an earlier report<sup>57</sup> from our OIAC working group. FaceTime is automatically integrated into the calling features of the mobile device, and makes heavy use of radio network bandwidth in both directions between the device and the cellular base station. When FaceTime first became available over cellular data networks, AT&T limited the use of FaceTime to customers of its MobileShare data plan, where multiple devices share a single limit for total data usage. Later, AT&T broadened the range of plans that support FaceTime, but still did not support the application for subscribers on its legacy unlimited data plan; recently, AT&T announced that all customers<sup>58</sup> (even those on unlimited data plans) will be able to run FaceTime over the cellular LTE network by the end of 2013. Another example of carriers imposing application restrictions occurs when they prohibit the use of tethering applications in their terms of service. These restrictions sometimes arise *after* a customer has chosen a specific service contract, when the emergence of a new application leads to heightened concerns about sudden increases in bandwidth usage.

**Two-sided pricing:** Usage caps and usage-based billing naturally make users conservative about running bandwidth-intensive applications (e.g., video streaming and online gaming). Some content providers and mobile providers may be willing to offer "toll free" or "sender pays" services, where the bandwidth consumed is sponsored or paid by the content provider, rather than counted towards the customers' usage cap. Broad use of two-sided pricing is not (yet) common in the U.S. mobile broadband market<sup>59</sup>, though several European and Asian providers have partnered with content providers to offer plans that do not count applications like Facebook and Spotify against a usage cap<sup>60</sup>. These trends raise interesting questions about openness. On the one hand, "toll-free" data may facilitate end-users' ability to access mobile content at a reasonable cost from those providers willing to subsidize the cost of delivering the data. Enabling content providers to pay for data delivery offers users an incentive to access the sponsored content. In the short run, this is beneficial for consumers of that content, particularly for budget conscious users on smaller data plans. On the other hand, sponsored delivery potentially works *against*<sup>61</sup> the goals of

---

<sup>56</sup> <http://bits.blogs.nytimes.com/2012/07/31/fcc-verizon-tethering/>

<sup>57</sup> <http://transition.fcc.gov/cgb/events/ATT-FaceTimeReport.pdf>

<sup>58</sup> <http://www.macobserver.com/tmo/article/att-opening-facetime-over-cellular-to-all>

<sup>59</sup> Discussions of two-sided pricing sometimes reference the Amazon Kindle e-reader device, which in some cases is sold to users without requiring them to purchase a separate service contract with a carrier despite the fact that the device uses a cellular network. However, e-book downloads consume relatively little bandwidth and do not constitute general, universal Internet service. As the Kindle started supporting basic Web browsing, and some users started [tethering the device](#) to use as a mobile hotspot, Amazon started [capping the free cellular bandwidth usage to 50 megabytes per month](#).

<sup>60</sup> [http://www.npt.no/marked/markedsregulering-smp/marked/marked-7/\\_attachment/2362?\\_ts=139b9fde471](http://www.npt.no/marked/markedsregulering-smp/marked/marked-7/_attachment/2362?_ts=139b9fde471)

<sup>61</sup> <http://media.law.stanford.edu/publications/archive/pdf/schewick-statement-20100428.pdf>

openness because (i) increasing the costs for content providers may reduce innovation and (ii) smaller, upstart content providers cannot easily amortize the “chargeback” costs through advertising revenue or subscription fees. Entrenching the largest content providers that have the means to strike deals for sponsored data with carriers puts new entrants at a disadvantage. This is clearly an area of ongoing debate.

The evolution of service contracts and pricing plans show that there is a great deal of experimentation in mobile business models, which is enabling innovation and value to customers and others in the ecosystem. Some business models raise concerns about carriers restricting the way consumers use their mobile devices and about long-term impacts on application and content innovation.

### **2.3 Network-Unfriendly Apps: Mobile Broadband Providers and App Developers**

The applications running on mobile devices have a profound influence on the network resource demands for mobile providers. While supporting the resource demands of applications is also important in wireline networks, mobile broadband networks raise several unique challenges. First, mobile apps are written by millions of software developers, including an unprecedented number of novice programmers who have little understanding of how high-level design decisions affect the usage of network and battery resources. Second, radio access networks have very limited bandwidth, particularly on the “uplink” from the mobile devices to the cellular base station, making it relatively easy for one rogue application to consume most of the available resources. Third, communication in cellular networks requires mobile devices to first establish a “bearer” with the base station, leading to signaling overhead. Fourth, expanding the capacity of a cellular network requires a substantial upfront investment for acquiring spectrum licenses, deploying cell towers, and transitioning to new technologies (e.g., LTE).

For mobile providers, applications that (unwittingly) consume excessive bandwidth and signaling resources cause congestion for other users in the short term, and require a larger investment in network capacity in the long term. In addition, applications that waste network bandwidth or battery lifetime limit the value of a mobile broadband service to end users, particularly if users are subject to usage caps or usage-based billing. As a result, without greater transparency to increase user awareness of an application’s efficiency -- and usage-based pricing models to incent them to choose the most efficient applications -- providers could see a limited return on the substantial investment required to expand network capacity, and still face the risk of a new mobile application swamping the available resources. Mobile applications can consume excessive network resources in several ways:

**Chatty applications consuming excessive signaling resources:** In contrast with wireline networks, mobile devices cannot communicate over a cellular network without first establishing a “bearer” to the cell tower. Establishing a bearer requires the mobile device to exchange several control messages over the cellular network. To avoid the overhead of establishing a new bearer, the mobile device continues to occupy transmission channels and codes until a period of inactivity expires. As such, transmitting a small amount of data can consume significant resources in the radio access network, as well as significant battery resources on the mobile device. The problem is exacerbated by “chatty” applications that periodically send short messages to monitor user behavior, maintain a connection for “pushing” data to the mobile device, or update the display of advertisements. Depending on the frequency of the messages, each transmission may require establishing a new bearer, at the expense of additional signaling resources.

## Open Internet Advisory Committee - 2013 Annual Report

A recent study<sup>62</sup> showed that some applications consume as little as 1.7% of network bandwidth, but up to 30% of signaling capacity. Signaling load is a low-level issue that even a seasoned application developer might not consider, and it may cause an application that worked perfectly well on a wireline network to overwhelm a cellular network.

**Aggressive applications consuming excessive bandwidth:** The Internet relies on end-host computers to adapt their sending rates in response to network congestion, to ensure fair sharing of the available bandwidth. Applications using the Transmission Control Protocol (TCP) automatically send data more quickly when the network is lightly loaded, and more slowly when the network is congested enough to drop packets. In addition to decreasing the sending rate, multimedia applications may adjust the audio or video encoding to continue streaming data quickly enough for continuous playback despite the reduced available bandwidth. However, some applications do not use TCP or perform “TCP-friendly” congestion control, open multiple parallel TCP connections to receive a larger share of the limited bandwidth, or do not use adaptive content encodings. The encoding issue was apparently at play with Apple’s FaceTime application, as discussed in an earlier report<sup>63</sup> by this OIAC working group. In addition, some operating systems are intentionally more aggressive than the protocol standards prescribe in sending data at the start of a TCP connection<sup>64</sup>, to reduce latency particularly for small transfers. Given the Internet protocols place important resource-management functionality at the end hosts, the sharing of the limited bandwidth in a cellular network is not completely within the provider’s control.

**Inefficient applications transferring redundant data:** A mobile application often needs to display the same data to the end user more than once, such as previously-downloaded images or articles. Caching content on the mobile device is an effective way to avoid duplicate transmission of the same data, reducing the consumption of battery, bandwidth, and signaling resources. Despite some support for caching on mobile devices, a recent study<sup>65</sup> found that redundant data transfers still consume 18-20% of bandwidth and 6% of signaling load. Rather than performing data transfers themselves, many mobile applications use HTTP libraries. Unfortunately, many of these libraries do not perform caching at all, or do not fully support the HTTP protocol standards for caching. Similarly, some mobile Web browsers do not make effective use of caching. In addition, cached data does not always survive an application crashing or a mobile device rebooting, leading to further wasted transfers and battery resources. In some cases, software bugs can cause excessive downloading of redundant content, as was in the case with an earlier bug in Apple iOS 6.0<sup>66</sup> that caused duplicate downloads of certain podcasts<sup>67</sup>. Enforcing usage caps and usage-based billing can help carriers recoup the cost of duplicate data transmissions, but also gives users the perception of a lower quality of experience for a given price for their mobile broadband service, and does not provide a direct incentive to app developers to reduce redundant transmissions.

---

<sup>62</sup> Feng Qian et al, “Periodic transfers in mobile applications: Network-wide origin, impact, and optimization,” in *Proceedings of the World Wide Web Conference*, May 2012.

[http://web.eecs.umich.edu/~zmao/Papers/periodic\\_www2012.pdf](http://web.eecs.umich.edu/~zmao/Papers/periodic_www2012.pdf)

<sup>63</sup> <http://transition.fcc.gov/cgb/events/ATT-FaceTimeReport.pdf>

<sup>64</sup> <http://blog.benstrong.com/2010/11/google-and-microsoft-cheat-on-slow.html>

<sup>65</sup> Feng Qian et al, “Web caching on smart phones: Ideal vs. reality,” in *Proceedings of MobiSys*, June 2012. [http://web.eecs.umich.edu/~zmao/Papers/caching\\_mobisys2012.pdf](http://web.eecs.umich.edu/~zmao/Papers/caching_mobisys2012.pdf)

<sup>66</sup> <http://venturebeat.com/2012/11/14/ios-6-0-bug-causing-massive-data-consumption-on-podcasts/#bmb=1%20%E2%80%A6>

<sup>67</sup> <http://labs.prx.org/2012/11/14/ios-6-0-devours-data-plans-causes-cdn-overages/>

Although applications may consume excessive resources, the incentives of all of the parties---application developers, mobile broadband providers, and end users---are generally aligned. More efficient applications lead to better performance (and better battery lifetime) for users, and lower loads on the network. As such, the main challenges are *education* (of application developers, so they can write network-friendly apps) and *visibility* (for users, so they know which applications are hogging resources). A good example of education of developers is AT&T's Application Resource Optimization (ARO) tool<sup>68</sup> and associated training, which helps application developers understand how their apps would behave on mobile broadband networks. ARO helped the developers of the popular Pandora application substantially reduce their consumption of energy and signaling resources by transmitting audience measurement data less frequently. A good example of visibility is the reviews of applications in app stores, which increasingly comment on an application's use of battery and bandwidth (though not signaling load). Further investment in tools, training, and rating of applications would help application developers and users alike make more informed decisions about resource consumption.

### **2.4 Wi-Fi Offloading: Competition for Mobile Providers**

One technology trend that is changing the dynamics of the mobile broadband market is the growth of non-commercial, wireless Internet access, typically provided using unlicensed spectrum approaches such as Wi-Fi, in many cases, backhauled over a pre-existing (wired) broadband connection.

Over the past 10 years, there has been an exponential growth in cellular data traffic, driven primarily by the dramatic increase in use of smart phones and tablets. As a consequence of the growth in demand, mobile broadband providers are aggressively expanding their network capacity. In addition, due to the prevalence of Wi-Fi on smart phones and tablets, and the increasing availability of Wi-Fi-enabled Internet service in public places (e.g. coffee shops, airports, campuses, hotels) and Wi-Fi-enabled routers at home and in the enterprise, an increasing fraction of mobile wireless data traffic is carried over Wi-Fi access, rather than cellular networks, with different studies suggesting that anywhere from 20-80% of wireless data traffic is carried over Wi-Fi, and ~30-50% of the 'mobile' data traffic may be *cost-effectively* offloaded from cellular networks, depending on the specific deployment scenario<sup>69</sup>.

One of the key differences between Wi-Fi networks and cellular networks is that Wi-Fi users may be subject to interference from users of neighboring access points. The quality of a Wi-Fi connection as compared to a cellular data connection may therefore suffer in the presence of interference due to a lower signal-to-noise ratio, resulting in a significantly diminished throughput relative to cellular networks in public settings; a recent paper<sup>70</sup> suggests that less than a third of mobile data traffic may be carried over Wi-Fi networks even in campus environments with dense Wi-Fi deployments. Likewise, similar Quality of Service (QoS) mechanisms that offer hierarchical or differential scheduling and queuing of data flows

---

<sup>68</sup> <http://www.att.com/gen/press-room?pid=22388>

<sup>69</sup> Randall Schwartz and Magnus Johansson, "Carrier WiFi Offload: Building a business case for carrier WiFi offload," Wireless 20/20, March 2012.

<http://www.wireless2020.com/docs/CarrierWiFiOffloadWhitePaper03202012.pdf>

<sup>70</sup> Shu Liu and Aaron Striegel, "Casting doubts on the viability of WiFi offloading," in *Proceedings of ACM SIGCOMM Workshop on Cellular Networks*, August 2012.

<http://conferences.sigcomm.org/sigcomm/2012/paper/cellnet/p25.pdf>

## Open Internet Advisory Committee - 2013 Annual Report

with different priorities may not be available on Wi-Fi and cellular connections, depending on their configuration. But the availability of cheap or free capacity (and considerable spectrum, e.g., ~400 Mhz in the 5Ghz band<sup>71</sup>) makes Wi-Fi-based solutions attractive for simple web services delivery. Furthermore, with the emergence of usage-based pricing for cellular data services, which encourages users to manage their cellular data usage, and provides unlimited access when the user is connected to certain Wi-Fi Access Points (their own at home, or in a public place), it is legitimate to ask “will Wi-Fi eventually carry a large enough share of mobile user traffic to cause a significant change in the mobile broadband market, and change the essential economics?”. This section explores some aspects of this question.

To address this question, we must first identify the types of Wi-Fi solutions. For the purposes of this discussion, we characterize three types of Wi-Fi: (i) non-public indoor (owned/operated by an individual or business), (ii) public indoor including both free or fee-based (likely owned and operated by a business, and provided to its customers) and commercial (owned and operated by a Wi-Fi network operator), and (iii) public outdoor (likely owned and operated by a network provider or campus-based business, or municipality).

These different types of Wi-Fi access points have different characteristics in terms of accessibility, security, and performance, as well as different degrees of utility to the user. They also have different economics. The benefits and limitations of each are summarized in the following table:

	<b>Cost to Operate</b>	<b>Accessibility</b>	<b>Service Continuity</b>	<b>Radio Performance</b>	<b>Commercial Service</b>	<b>Cellular Offload Potential</b>
<b>Type 1</b> (non-public indoor)	<b>Low</b> (unmanaged & connected to existing BB)	<b>Limited</b> (only to individual users or employees)	<b>Limited</b>	<b>Not managed</b>	<b>No</b>	<b>&gt; 50%</b>
<b>Type 2</b> (public indoor)	<b>Medium</b> (managed by connected to existing BB)	<b>Good</b> (subject to business rules)	<b>Some</b> (indoor continuity)	<b>Some management</b>	<b>Yes</b> (direct or indirect payment or subscription)	<b>&lt; 50%</b>
<b>Type 3</b> (public outdoor)	<b>High</b> (managed & uses new network connection)	<b>Good</b> (subject to subscription or business rules)	<b>More</b> (outdoor continuity and cellular networking)	<b>More management</b>	<b>Yes</b> (subscription service)	<b>&lt; 50%</b>

The preceding table summarizes the essential properties of the different Wifi deployment types, with the table categories and entries defined as follows:

- **Cost to Build and Operate:** This refers to existence of a backhaul network, power, and Wi-Fi access point management

<sup>71</sup> [http://en.wikipedia.org/wiki/List\\_of\\_WLAN\\_channels](http://en.wikipedia.org/wiki/List_of_WLAN_channels)

## Open Internet Advisory Committee - 2013 Annual Report

- Low cost: Pre-existing, economical backhaul and power with no AP management
- High cost: New backhaul and power network and sophisticated management
- Accessibility: This refers to the ability to connect to Wi-Fi APs
  - Limited: Restricted only to certain users (e.g. employees)
  - Good: Can be accessed by anyone willing to subscribe or agree to terms and conditions
- Service Continuity: This refers to ability to maintain a session or connectivity when moving from one location to another
  - Limited: Little or no ability to seamlessly connect to neighboring AP
  - Some: Able to maintain session between APs in similar location, from same provider
  - More: Session and service continuity by interworking with other APs and/or the cellular network
- Radio Performance: This refers to management of the Wi-Fi air interface
  - Not managed: Air interface configuration independent of all other APs
  - Some management: Some coordination of APs via common controller
  - More management: Coordination of APs via common controller, with interference management
- Commercial Service: This refers to whether a Service Provider owns and manages APs
  - No: APs owned by private individual or entity
  - Yes: APs owned by commercial entity (business, building provider) or Service Provider
- Cellular Offload Potential: This concerns the potential of a Wi-Fi AP to offload cellular network traffic
  - There are many different estimates of the how much data offload can be achieved by a Wi-Fi network (see the preceding references for examples), but it is broadly agreed that somewhere between 50-75% of time the average user is in home or in an enterprise environment where Wi-Fi experiences relatively little interference and so is highly effective at offloading data traffic, and consequently only 25-50% of the time is the user outdoors or in a public indoor location, where a combination of Wi-Fi and cellular networks would provide the solution.

What does this simple analysis suggest about the impact of Wi-Fi solutions on the mobile broadband market? The growth of these free or lower cost alternatives in any market clearly benefits consumers in terms of providing access to more wireless capacity. However, it is also the case that the user experience amongst Wi-Fi services varies widely, with registration procedures not being seamless, the network performance sometimes poor due to interference, and inconsistent deployment of recent Wi-Fi security enhancements. Some of these issues are being addressed by the Hotspot 2.0 initiative<sup>72</sup> of the Wi-Fi Alliance, which seeks to increase the degree of ‘management’ of Wi-Fi access points, and to provide seamless authentication and session continuity (between Wi-Fi access points within the same area). Based on these trends, mobile operators are increasingly integrating Wi-Fi solutions with their cellular offers and encouraging use of Wi-Fi for unlimited data offload for ‘best effort’ services. Indeed, 3GPP is working in standards to allow seamless session continuity between cellular and Wi-Fi solutions, per serving area or per cell, or even per application in future, based on the local availability of capacity, and the needs of the application, as well as user preference and services agreements.

---

<sup>72</sup>[http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns673/white\\_paper\\_c11-649337.html](http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns673/white_paper_c11-649337.html)

## Open Internet Advisory Committee - 2013 Annual Report

These emerging trends effectively mean that Wi-Fi will not just be a wireless broadband solution, but will also become an essential part of providing mobile broadband services to users. Furthermore, given the lower barrier for entry into the Wi-Fi solution space (due to the absence of the need to acquire spectrum or to support wide-area coverage, or mobility), the number of providers that can and will likely enter this space is significant and will likely therefore stimulate additional innovation in wireless data services.

So the future of mobile broadband should consider the combined roles of licensed and unlicensed spectrum solutions, as they are complementary parts of the space, with licensed spectrum approaches providing coverage and capacity with full mobility, security, and quality of service, and unlicensed approaches providing additional capacity with some (e.g., indoor) mobility and nomadicity, but with more limited QoS capabilities and inconsistent security implementation, at least in the near future.

Looking forward, there will be further evolutions of this licensed/unlicensed paradigm to include ‘shared spectrum’ approaches, based on white-space spectrum (spectrum in and around the TV frequencies that is either unused or infrequently used) or in higher frequency bands such as the 3.5GHz band currently licensed for military use, but for which the FCC has indicated the desire to make available for commercial use by multiple parties in a shared way (use it when you need it, then release it) in a Notice of Public Rule Making (NPRM)<sup>73</sup>.

Consequently, we conclude that the user mobile broadband experience will be provided by a combination of complementary approaches, and potentially a variety of different providers, indoor, outdoor, at home, and at work. This dynamism to the mobile broadband market suggests that the future of user choice and experience delivery will continue to grow and expand, with increasing value delivered by the expanded ecosystem.

### 3. Conclusions

The mobile broadband ecosystem is complex and dynamic, with a variety of players affecting the user experience and the incentives for further innovation and investment. This report encourages the FCC to take a broad view of interactions between the different players in the mobile broadband ecosystem, even though most of the parties involved are not subject to the Open Internet Order. Also, we recommend being watchful of recent trends, such as HTML5 and Wi-Fi offloading, that may lead to greater competition, as well as the emergence of several “vertical players” with growing influence spanning multiple parts of the ecosystem.

We believe that transparency, education, and competition are important complements to existing FCC oversight in helping achieve the goal of a healthy mobile broadband ecosystem. Transparency can take many forms, such as the disclosures required by the Open Internet Order, and improved communication to users (about applications’ battery and network resource consumption) and application developers (about the policies by which app stores and carriers might restrict access to their applications). Education includes teaching application developers how to create “network friendly” applications. Finally, competition includes both a healthy balance *between* the various parts of the ecosystem as well as having

---

<sup>73</sup> <http://www.fcc.gov/document/enabling-innovative-small-cell-use-35-ghz-band-nprm-order>

## Open Internet Advisory Committee - 2013 Annual Report

multiple viable choices *within* each part of the ecosystem. The combination of all these factors will help ensure all players – not just those subject to the Open Internet Order – contribute to the openness and health of the mobile Internet.

### **FCC Open Internet Advisory Committee Mobile Broadband Working Group**

- Chair: Jennifer Rexford, Professor of Computer Science, Princeton University
- Harvey Anderson, Vice President of Business Affairs & General Counsel, Mozilla
- Brad Burnham, Founding Partner, Union Square Ventures
- Alissa Cooper, Chief Computer Scientist, Center for Democracy & Technology
- Charles Kalmanek, Vice President of Research, AT&T
- Matthew Larsen, CEO, Vistabeam
- Dennis Roberson, Vice Provost & Research Professor, Illinois Institute of Technology (representing T-Mobile)
- Chip Sharp, Director, Technology Policy and Internet Governance, Cisco Systems
- Marcus Weldon, Chief Technology Officer, Alcatel-Lucent
- Jonathan Zittrain, Professor of Law and Professor of Computer Science, Harvard; OIAC Chair (ex officio)
- David Clark, Senior Research Scientist, MIT Computer Science and Artificial Intelligence Research Laboratory; OIAC Vice Chair (ex officio)

# Specialized Services: Summary of Findings and Conclusions

FCC Open Internet Advisory Committee  
Summary of findings and conclusions, July 2013

**The Specialized Services working group prepared a series of case studies to explore issues in the specialized services landscape, and created a series of conclusions based on those case studies.**

The Open Internet Report and Order (R&O) assigned to the Open Internet Advisory Committee (“OIAC”) the task of aiding the FCC in the task of monitoring specialized services for their impact on Internet access.<sup>74</sup> As part of the proceedings of the Open Internet Advisory Committee, the Specialized Services working group has met for the 12 months prior to the July 2013 meeting of the committee. This report summarizes the findings and conclusions of the working group.

We organized our work around two tasks:

- Attempting to articulate a careful definition of the term “specialized services”, and considering whether the working group has advice to the FCC on the criteria that will prove useful in practice to define and characterize a specialized service.
- Developing advice to the FCC with respect to how they should monitor the impact of specialized services on the character of broadband Internet access service (BIAS).

## Background

The ability to offer multiple services was an initial driver for many of the significant network investments made by service providers in higher capacity broadband access network architectures. For legacy telephone operators, the emergence of VDSL and ADSL2+ and MPEG-4 enabled them to leverage their existing copper infrastructure to more rapidly deliver a "triple play" of services: voice, data, and video. Similarly, the cable operators have used their platform to deliver a range of services. The current trend is that all these services will migrate to a provider platform based on the Internet protocol (IP). The R&O uses the term “specialized services” to identify those IP-based services that are not subject to the FCC’s Open Internet rules.

The R&O states that the specialized services category in the report could raise two concerns that it would monitor going forward. First, the FCC should guard against the possibility that a broadband provider might label a service as a specialized service that would otherwise be correctly identified as an Internet access service in order to evade Open Internet rules. Second, broadband providers might constrict or fail to continue expanding network capacity allocated to broadband Internet access service in order to provide relatively more capacity for specialized services.

---

<sup>74</sup> Preserving the Open Internet Broadband Industry Practices, GN Docket No. 09-191, WC Docket No. 07-52, FCC 10-201, 114 (Dec. 23, 2010) [hereinafter *R&O*].

The FCC notes that their goal is to achieve a balance of innovation in infrastructure and applications, but the report does not state any conclusions as to the impact of specialized services on that objective. On the one hand, the R&O notes that: “specialized services may raise concerns regarding bypassing open Internet protections, supplanting the open Internet, and enabling anticompetitive conduct.”<sup>75</sup> The advantages to a facilities owner of deploying a service as a specialized service, as opposed to an OTT service, is that the facility owner can offer the service with attributes such as a guaranteed quality of service not permitted today with BIAS, and thus not accessible to competitive OTT services<sup>76</sup>. On the other hand, the benefits to the consumer of specialized services are considerable. The business case to justify the investment in the expansion of fiber optics and improved DSL and cable technology which led to higher broadband speeds was fundamentally predicated upon the assumption that the operator would offer multiple services: while all offerings present uncertainty and risk, the projected value that consumers placed on multiple offerings promised an acceptable return on the investment in the expansion of the overall broadband infrastructure, while the value consumers placed on increased BIAS speeds alone did not yield acceptable projected returns.<sup>77</sup> This appears to remain true today, as even new entrants such as Google Fiber offer video services in addition to BIAS<sup>78</sup>. Accordingly, high speed internet access service has benefited from the deployment of specialized video services like IPTV, because the investment in the higher bandwidth infrastructure needed for video services brought higher capacity to more households.

### **Defining specialized services**

Our starting point in this discussion was to see if we could agree on a meaning of the term “specialized services”, as given to us by the FCC. This proved difficult. The Open Internet Report and Order defines a specialized service as a service “that broadband providers may offer... over the same last-mile connections used to provide broadband service.”<sup>79</sup> Examples of specialized services mentioned in the R&O include facilities-based VoIP, IP video,<sup>80</sup> e-reading services, heart rate monitoring, and energy sensing.<sup>81</sup>

The use of the term in the R&O is in the context of the scope of the rule-making, which is set forth as following<sup>82</sup>:

“We find that open Internet rules should apply to “broadband Internet access service,” which we define as:

---

<sup>75</sup> *Id.* at 112.

<sup>76</sup> Independent of whether it is in the business interest of a BIAS provider to offer QoS, the R&O may not permit this option.

<sup>77</sup> The FCC has concurred with this assessment in its Report and Order relating to local cable franchising: see In the Matter of Implementation of Section 621(a)(1) of the Cable Communications Policy Act of 1984 as amended by the Cable Television Consumer Protection and Competition Act of 1992, MB Docket No. 05-311, FCC 06-180, para 51.

<sup>78</sup> For a discussion of the role of video in the Google fiber offering, see [http://news.cnet.com/8301-1023\\_3-57586894-93/google-exec-sees-google-fiber-as-a-moneymaker/](http://news.cnet.com/8301-1023_3-57586894-93/google-exec-sees-google-fiber-as-a-moneymaker/)

<sup>79</sup> *Id.* at 7.

<sup>80</sup> *Id.* at 61.

<sup>81</sup> *Id.* at 33.

<sup>82</sup> *Id.* at 44.

## Open Internet Advisory Committee - 2013 Annual Report

*A mass-market retail service by wire or radio that provides the capability to transmit data to and receive data from all or substantially all Internet endpoints, including any capabilities that are incidental to and enable the operation of the communications service, but excluding dial-up Internet access service. This term also encompasses any service that the Commission finds to be providing a functional equivalent of the service described in the previous sentence, or that is used to evade the protections set forth in this Part.”*

With some informal guidance from the FCC, the working group took as a starting point that the term “specialized services” describes anything not covered by this rule. In other words, the group took the term to describe services that are “anything else”. This inclusive definition would imply that for purposes of the R&O, the category of specialized services would include services regulated in other ways by the FCC, including voice and video.

However, this inclusive definition proved very difficult for the working group to accept in our discussions, because the term has also been used by the FCC elsewhere in less inclusive ways. The R&O itself refers to specific text in the Open Internet NPRM, which defines specialized services as follows:

“As rapid innovation in Internet-related services continues, we recognize that there are and will continue to be Internet-Protocol-based offerings (including voice and subscription video services, and certain business services provided to enterprise customers), often provided over the same networks used for broadband Internet access service, *that have not been classified by the Commission*. We use the term “managed” or “specialized” services to describe these types of offerings. The existence of these services may provide consumer benefits, including greater competition among voice and subscription video providers, and may lead to increased deployment of broadband networks.<sup>83</sup>”

The italicized text might be read to suggest that if the FCC has classified some service in some other way, then it may not be considered a specialized service. This narrower use of the terms is made explicit in the merger agreement between Comcast and NBCU, which defines specialized service as follows:

“Specialized Service” means any service provided over the same last-mile facilities used to deliver Broadband Internet Access Service other than (i) Broadband Internet Access Services, (ii) services regulated either as telecommunications services under Title II of the Communications Act or as MVPD services under Title VI of the Communications Act, or (iii) Comcast’s existing VoIP telephony service<sup>84</sup>.

---

<sup>83</sup> Federal Communications Commission, Notice of Proposed Rulemaking, In the Matter of Preserving the Open Internet Broadband Industry Practices, GN Docket No. 09-191, WC Docket No. 07-52, FCC 09-93, (October 2009) 148 [italics added, footnote omitted]

<sup>84</sup> Federal Communications Commission, In the Matter of Applications of Comcast Corporation, General Electric Company and NBC Universal, Inc. For Consent to Assign Licenses and Transfer Control of Licensees, MB Docket No. 10-56, FCC 11-4, Appendix A, I (Definitions), pg. 121

This text makes explicit that in the context of the Comcast-NBCU Order, specialized service does not include Title VI MVPD service. Yet the R&O states that that IP video is explicitly included in the definition (but also, as noted above, may exclude services that are otherwise classified). These varied definitions have slowed the working group's progress, and may require future clarification by the FCC.

For the purpose of this working group, which functions in the context of the R&O, we have attempted to work with the inclusive definition of specialized service. The term as we use it is thus only meaningful within the context of the R&O. Used in this way, "Specialized services" are not a new category of items for regulation. Rather, they set a limit on which IP-based services are subject to the Open Internet rules. In this usage, some specialized services, such as VoIP and video, may already be subject to regulation under other laws and orders – the Open Internet R&O does not affect these other regulations. Rather, the labeling of a service as "specialized" would mean that that service is not subject to further regulation under the R&O.

We proceed with this definition, mindful of the fact that all such use of the term should properly be prefaced with OI, as in "OI specialized services".

### **Criteria for distinction**

Based on the reading of the R&O, and subsequent discussions with FCC staff counseling the OIAC, there are two criteria in the R&O that would move a managed service far enough away from the open Internet that the R&O would not apply.

- 1) The service is not used to reach large parts of the Internet.
- 2) The service is not a generic platform but a specific "application level" service.

Using a number of case studies, we tried to tease out other aspects of a service that would set it apart from the services covered by the rules of the R&O. We identified one other criterion that we bring to the attention of the FCC.

- 1) Capacity isolation. The criterion of "capacity isolation" came up in a number of working group case studies, including the IPTV case study, the third-party platform case, and VoIP<sup>85</sup>. The argument is that a specialized service should not take away a customer's capacity to access the Internet. Since statistical multiplexing among services is standard practice among network operators, the isolation will not be absolute in most cases. However, if a specialized service substantially degrades the BIAS service, or inhibits the growth in BIAS capacity over time, by drawing capacity away from the capacity used by the BIAS, this would warrant consideration by the FCC to further understand the implications for the consumer and the possible competitive services running on the BIAS service.

### **Distinctions between BIAS and specialized services**

---

<sup>85</sup> Voice over IP, or VoIP, is not a case study elaborated in this report, but was discussed by the working group, and shares the isolation attributes of IPTV.

The discussions concerning the differences between specialized services and a BIAS service tend to focus on the fact that specialized services, since they are not bound the requirements of the R&O, can offer different sorts of services, in particular enhanced service qualities. However, there will be other dimensions along which the services may differ; providers of BIAS who have usage tiers or usage caps need not impose those caps on specialized services, and specialized services may be priced and packaged in different ways.

### **High-level principles**

We identified three high-level principles that the FCC should consider if and as it further deliberates about specialized services:

- Open Internet regulation should not create a perverse incentive for operators to move away from a converged IP infrastructure. Using IP should not imply a regulatory burden related to any regulation of the Internet.
- A service should not be able to escape regulatory burden, or acquire a burden, by moving to IP. A service may change or evolve as it migrates to IP, and the regulatory implications of such a change should be evaluated based on its characteristics.
- Proposals for regulation should be tested by applying them to the range of technologies being used for broadband. To the extent possible, regulation should be technology-neutral. (There are painful edge-conditions to this principle, which we acknowledge.)

These seem like simple statements, but in fact they may have very powerful consequences. They are an attempt to bound the scope of regulation without the need to debate the definition of any terms such as specialized services.

### **Monitoring the Internet**

In recognizing specialized services as a category that is not subject to the Open Internet rules, the FCC also expressed the importance of ensuring that specialized services do not deter or limit investment in Internet services. The FCC expressed concern that “broadband providers may constrict or fail to continue expanding network capacity allocated to broadband Internet access service to provide more capacity for specialized services.”<sup>86</sup> The FCC has declared their intention to monitor this situation. This committee is asked to advise them as to how to undertake this task.

Two approaches may address these concerns, although neither approach is wholly satisfactory and both approaches carry the risk of unintended consequences. On the one hand, the FCC may choose to define how much Internet service is “enough”, and compare actual offerings to this standard. By setting a minimum standard for how much capacity for Internet service is available, the FCC could potentially make sure that sufficient capacity exists for providers of high-level service to innovate. It is important to note, however, that this minimum standard would likely have to change over time as consumers’ usage habits and expectations shift. Alternatively, the FCC could compare what innovators can do using a specialized service as compared with the public Internet. Such a comparison would help the FCC to determine whether ISPs are exploiting

---

<sup>86</sup> R&O at 61.

a significant set of innovative opportunities via specialized services that are not available to others who would like to innovate over the open Internet. This second approach would reveal not only raw capacity concerns, but also quality of service concerns. As is illustrated in the third-party platform case study, the issue of comparing what can be done over the Internet and as a specialized service is not a simple matter of capacity, but depends on several parameters of the service.

The FCC currently performs a range of measurements on the Internet, tracking metrics such as achieved throughput, latency, and so on. In our discussions of specialized services, we did not identify any additional technical metrics that might be usefully measured, in order to better understand the impact of specialized services on the BIAS service. Instead, we focused on the higher-level question of what to make of these measurements—what sort of results would lead to the conclusion that the Internet was “good enough”.

Exploration of this question is our tentative task for the next study period, but we have identified a possible approach to the issue. We believe that a promising approach is to start by looking at the quality of the user experience, not the technical parameters. The National Academies, in a 2002 report titled “Broadband: Bringing home the bits”<sup>87</sup>, chose not to define broadband in numerical terms, because the committee knew that the target number would change over time. Instead, they defined it in terms of the needs of the applications of the time. They offered two definitions: a baseline definition and a forward-looking definition.

- Broadband Definition 1. Local access link performance should not be the limiting factor in a user's capability for running today's applications.
- Broadband Definition 2. Broadband services should provide sufficient performance and wide enough penetration of services reaching that performance level to encourage the development of new applications.

Neither definition is quantified, and neither, as stated, could directly be used as the basis of regulatory specification. However, the view of the committee was that these definitions could be translated into numbers that would be applicable at a given time. Based on our initial discussions, we believe that there have been a number of studies that relate the various technical parameters describing broadband performance to the operation of specific applications. We plan to explore this (and potentially other) approach to answering the question of when an Internet service is “good enough”.

---

<sup>87</sup> Computer Science and Telecommunications Board, *Broadband: Bringing Home the Bits*, National Academy Press, 2002.

## **Appendix 1: Case study of IPTV**

The Specialized Services working group is examining a range of issues surrounding “specialized services” in the context of the Open Internet Order, and how they relate to broader Internet access service and innovation. This appendix looks at the role of video (including IP based video) services, in today’s marketplace and the potential effects on broadband Internet access service (BIAS). The paper provides a high-level overview of certain access network architectures, describes how services can be delivered over those architectures, and then discusses possible implications for BIAS.

### **High level overview of broadband access network architectures**

Broadband Internet networks typically have a common general structure: the network operator’s backbone connects to the networks of other operators and to its regional metro network, which in turn connects to local access facilities all of which contain fiber, optical components, routers, servers, switches and the like. The focus of this paper is on the access network, which is the portion of the network closest to the customer, and most relevant to the provision of specialized services over a shared facility that is used to deliver BIAS. Access networks typically comprise a mix of fiber and either coaxial cable (cable systems) or copper facilities (telco) to the home, and more recently, some network providers are using fiber facilities all the way to the home. Modern cable systems typically use a Hybrid Fiber Coax (HFC) access network, while telecommunications service providers typically use either a Digital Subscriber Line (DSL) or Passive Optical Networking (PON) based technology.

In a typical implementation of an HFC system, a cable operator will extend fiber from a Cable Modem Termination System (CMTS) to an Optical Node in a local neighborhood, which can serve anywhere from a few to several hundred homes. From each Optical Node, coaxial cable is then used to deliver service to the home. Services are delivered over Radio Frequency (RF) over coax typically using frequency bands from 52 MHz to 1000 MHz for downstream signals and 5 MHz to 42 MHz for upstream signals. The IP bandwidth is delivered by bonding together multiple 6 MHz RF channels, the same channels that traditionally were used to deliver a single analog video channel (explained later). With the recent DOCSIS 3.0 specification, cable operators typically bond 8 channels downstream to support a downstream channel of approximately 300 Mbps (although some cable operators are starting to bond 12 and 16 channels downstream), which is shared among a number of subscribers attached to a given node. Depending upon the details of the HFC infrastructure, the total number of subscribers connected to an Optical Node, and the number of subscribers online at a given point in time, this architecture can deliver a wide range of BIAS speeds along with specialized services.

Telecommunications service providers have typically used DSL and more recently PON systems to deliver service to the home. Similar to cable operators, over time DSL providers have extended fiber optics closer to homes, using some combination of Fiber to the Node (FTTN) and Fiber to the Home (FTTH). The emergence of next generation DSL technologies, such as Very

High Rate DSL (VDSL), ADSL2+ and techniques such as pair bonding and vectoring have enabled service providers delivery speeds much higher compared with legacy DSL technologies. In the case of a DSL implementation, the broadband connection in the access network is dedicated to an end user from the node to a user's home, rather than being shared as in typical cable HFC systems. For FTTH implementations, most service providers are using a technology referred to as Passive Optical Networks (PON). PON systems generally take one strand of fiber to a fiber splitter location, and then replicate the optical signal onto multiple separate fiber strands connected to subscriber homes. A PON system consists of an Optical Line Terminal (OLT) placed in a serving central office and an Optical Network Terminal (ONT), or electronics, at the subscriber premises. As with VDSL services, this technology can deliver speeds far in excess of traditional DSL.

### **Service delivery methods**

Services delivered over these architectures typically include video, voice, and BIAS services. Broadband providers offering video services are classified as Multichannel Video Programming Distributors (MVPDs). Different MVPDs deliver video service in a variety of ways. Most cable systems today, and in some instances PON based video services, provide live linear programming ("traditional TV") using specific frequency bands dedicated to specific channels. All channels are simultaneously delivered or "broadcast" to the subscriber's premises, and tuners in the set top box act as filters to permit display of the desired programming network. For Video on Demand (VOD) services, MVPDs typically dedicate certain channels for delivery of requested content. In some cases, cable operators are offering linear programming networks and VOD delivered using IP or another packet-based transmission system, however, the vast majority of live linear video programming continues to be delivered using specific frequency bands dedicated to specific programming networks.

Modern cable systems use a digital representation of video, either compressed Motion Picture Expert Group (MPEG)-2, or more recently MPEG-4, video modulated onto Quadrature Amplitude Modulated (QAM) RF signals. In a typical implementation, a cable operator will organize the bandwidth used for digital video into the same 6 MHz channels of frequency as it would in a traditional analog cable system and, using 256 QAM, deliver approximately 38 Mbps per 6 MHz channel. In a typical MPEG-2 configuration, a Standard Definition (SD) channel can be encoded in a range from 2-6 Mbps and High Definition Content ranging from 15-19 Mbps. MPEG-4 halves these ratios to around 2-3 Mbps for SD and 6-7 Mbps for an HD channel. Thus a single 6 MHz channel slot with 256 QAM at approximately 38 Mbps could deliver up to 2 HD channels or 10 SD channels with MPEG-2, or perhaps twice that capacity with MPEG-4. The High Efficiency Video Encoding (HVEC) currently under development by the ISO/IEC Moving Picture Expert Group (MPEG) and the ITU-T Video Encoding Expert Group is intended to be the successor standard to MPEG-4 and is projected to reduce the bandwidth requirement by 50% for the same quality picture. It can also support resolutions up to 8192x4320.

The BIAS services offered over these cable systems will typically use separate and distinct channels and frequencies from the linear video services, creating a separation between the services sharing the infrastructure and dedicating fixed amounts of bandwidth to each service. As

noted, in some markets cable operators have begun offering traditional cable video services, both linear channels as well as VOD, in IP format. These IP cable services are delivered over the IP bandwidth a cable operator creates by bonding multiple 6 MHz channels, but these IP cable services typically use a separate service flow to customers' homes – with bandwidth above and beyond the bandwidth allocated for the customer's BIAS service – that is allocated specifically for the IP cable service

Another means of service delivery is a pure IP based infrastructure where all services are carried using IP on the same physical network. In this case, all video will be carried as IPTV. Any broadband IP network, regardless of the access network infrastructure, can be used for IPTV. The continuous improvements in data transfer speeds, brought about by advancements in both Digital Subscriber Line (DSL) and cable DOCSIS technology, combined with the improvements in compression ratios (e.g. the greater bandwidth efficiency offered by MPEG-4 over MPEG-2), and the emergence of switched digital video have enabled more video streams at higher quality to be delivered over broadband than previously possible.

The broadest use of IPTV has been by telecommunications operators to enable video delivery over their existing copper loop infrastructures. In contrast to broadcast video distribution typically used by cable companies, IPTV services utilize a switched, two-way, client server based architecture. Thus when a user “tunes in” to a “channel” delivered by an IPTV system, they are actually sending a request to initiate a stream of IP packets containing the requested video, and the servers stream only the requested content.

### **Capacity isolation**

As the previous discussion suggests, one factor that distinguishes different methods of delivering services is how the overall capacity of the physical access path is allocated to the different services. On cable systems, the capacity used for traditional video (encoded over QAMs) is separate from the capacity for BIAS. When the video service migrates to IPTV, the capacity that is allocated to the IPTV service may be isolated from the BIAS capacity to different degrees. In general, IP bandwidth to the home is dynamically allocated, meaning that varying amounts of bandwidth will be allocated to different services, depending upon the exact network usage of the household at a given moment in time.

Different technologies may accomplish capacity isolation among services in different ways. Cable systems using DOCSIS may open a separate service flow for the MVPD IPTV and allocate capacity to that flow sufficient for the video. In this way, the possibility that the IPTV and the BIAS may affect each other is minimized. On some other systems the allocation of capacity between MVPD IPTV and BIAS may not be as rigid. Based on information from the members of the subgroup familiar with current practices, most schemes for delivery of MVPD IP video attempt to isolate the capacity used for MVPD and BIAS to a high degree. However, public documentation is usually not specific as to practices.

The previous discussion has focused on the access path into the residence, but issues of traffic isolation can also arise in other parts of the network. Depending where the content servers are, the IP traffic between the servers and the access network might be totally segregated from the

public Internet infrastructure, or might share some of that infrastructure (potentially generating conflict with open Internet traffic, unless sufficient capacity is provisioned).

The committee discussed whether the degree of capacity isolation between a video service and BIAS service has implications as to whether the video service should fall under the rules of the R&O. For example in the extreme case where there is no capacity isolation between the video service and the BIAS service, it might seem that this is an OTT service, even if the service met the “reach” criterion from the R&O. But as the degree of isolation increases, there is an ambiguity as to what the boundary is.

### **Differences Between MVPDs’ IP-VIDEO and Over the Top Video**

The emergence of higher speed broadband networks of all access network technology types has contributed to the emergence of Over the Top (OTT) video services that deliver content via the end users’ BIAS service. Examples of OTT video services include Netflix, YouTube, Hulu, Amazon Prime, and Vudu. While OTT services may function in a manner somewhat similar to the IPTV systems described above—i.e. they have a client-server architecture, and stream only the requested content to the user—there are a number of distinctions between MVPD IP-video and OTT services.

- 1) **Customer Expectations:** MVPD services are usually offered as an integrated service package by the MVPD, often including “truck rolls” to install in-home wiring and equipment, network monitoring, customer care and helpdesk services, etc. OTT services typically offer only online and/or phone support and in-home service is available only through 3<sup>rd</sup> party integrators, if at all.
- 2) **System Design:** MVPD services are typically engineered to provide features for the linear TV service such as Instant Channel Change that consumers have grown accustomed to. MVPD services are delivered over a privately owned and managed network within the service provider’s infrastructure, rather than over the public Internet. Of particular note, an MVPD’s IP-video services are delivered via the MVPD’s own network and generally are not available via the Internet outside of a customer’s home. This aspect of the service may relate to the “reach” criterion of the FCC. OTT services typically are delivered via a third-party (i.e., not the MVPD/ISP) content delivery network and use the subscribers’ BIAS service for access to the home.
- 3) **Equipment:** MVPD services typically are accessed on leased equipment, although increasingly operators are making it possible to access MVPD IP video services on retail equipment. OTT services can be accessed via retail consumer devices in the home such as Apple TV, Roku, and Boxee, or “smart” TVs, Blu-Ray players, AV receivers, as well as via Internet browsers on general purpose devices such as computers and tablets. Some cable operators and telco IPTV providers offer their own OTT video services that are wholly distinct from their managed in-home MVPD services. Some of these services are simply standalone third party devices that provide a hardware and software “front end” for a variety of OTT services (e.g., Roku). Others are offered by the OTT content provider as a more convenient means of accessing their own content (e.g., Apple TV), as well as other partnered providers’ content. Boxee is an example of yet a different category, a sort of hybrid device that combines non-IP broadcast and cable services (either local OTA broadcasts or basic cable video delivered by QAM) with OTT Internet

video content. Satellite TV providers are also now delivering OTT video, both on-demand type streaming and/or downloadable (to a DVR), and selected live linear TV that is concurrently being broadcast on their satellite signals (e.g., DirecTV's "DirecTV Everywhere" service).

- 4) Regulatory Requirements: MVPD services typically face local franchise requirements, EEO and other back-office requirements, PEG (public, educational, government access) programming requirements, emergency alert requirements, CALM requirements, etc. The STBs that provide the video services, and/or the services themselves, must be capable of complying with these obligations. If the OTT IPTV uses a separate STB, these devices and the video services they deliver, generally speaking, are not subject to the same set of regulatory obligations. One notable exception is that OTT services and devices are required to support closed captioning.
- 5) Video quality: OTT services typically offer a range of streaming rates or video resolutions for different content (e.g., differentiating between SD and HD content), and use adaptive bitrates that can vary, adjusting to the bandwidth available on the user's connection. Most of these services stream at bitrates ranging from less than 1 Mbps up to 5-6 Mbps. Further, many OTT service providers rely on content compression, buffering and error correction on the consumer device, as well as adaptive bit rate streaming to optimize the user experience. MVPD services are typically provisioned such that adaptive coding and similar techniques are not needed to preserve the user experience.

The committee considered this list, and concluded that these differences are typical characteristics, but were not definitional, with the exception of the relationship to the issues of reach and capacity isolation.

## Conclusions

In the end, each of the methods described above for delivering video content and other services to the user can potentially deliver the same or closely similar functionality and experience while watching video in the home. However, the underlying technical methods and requirements are significantly different, with differing benefits and limitations. In the context of the R&O, the multi-channel video service in an IPTV configuration can be considered a specialized service: they use capacity on the provider's last mile facilities, they are application level services, they are logically separate from the BIAS service, and the IP service over which they run is restricted to the facilities of the MVPD operator; it does not provide access to all of the public Internet. In contrast, the OTT video services run on top of the BIAS service, and partake of the same service as all the other Internet-based applications. The resulting differentiations are important in signaling the implications of specialized services. Providers of MVPD IPTV can make higher assurances of delivery quality, can offer different pricing packages, and assure that IPTV and OTT Internet services do not disrupt each other.

It would seem that at the present, many versions of BIAS are good enough to support innovation in TV services, and the combination of MVPD and OTT alternatives are providing competition and consumer choice in the market. Concerns about the implications of specialized services on BIAS must be forward looking and thus speculative.

## **Appendix 2: Specialized services case study**

### ***3rd party purchasing of services for their customers (e.g. games)***

This case study looks at the challenge of supporting applications that have a requirement for enhanced service qualities that cannot today be met over the Internet.

This is a forward-looking case study.

### **Background**

The Internet provides “best effort” delivery of packets – no guarantees of delivery or delivery time of packets, no guarantees one packet will have the same path/fate as the next.<sup>88</sup> This approach has meant that the Internet is resilient overall, no participating network imposes performance requirements on another, and interconnection between networks is simplified with minimal agreements and commitments required between providers.

This approach to internetworking has successfully allowed significant third-party online services to be developed for use by customers globally and independent of any individual customer’s access ISP. When these services operate over the Internet they are sometimes referred to as “over the top” – (OTT) services.<sup>89</sup> Increasingly, these services support high-performance hardware on the client end as well as the server end, with attendant expectations of network connections that support their activity. They include applications with particular performance expectations – subject to reduced quality in the face of jitter or high latency, or even any form of timing disruption. A case in point is massive multiplayer action video games, where network-induced delays not only cause deterioration in the video quality experience, but can also get a player killed in the game. A person using a network that is persistently lagging is not going to keep up their (paid) subscription to the service. Consequently, having assured quality of network service from their servers to (and from) the end user may be of considerable interest to such services.

We describe three different ways that a provider of access service<sup>90</sup> can arrange with a third-party service developer to provide enhanced quality of service. All three seem to offer a similar enhancement for the third party service, but one seems to be a specialized service, one seems to be forbidden under the rules of the R&O, and one seems to be permitted within the rules that govern BIAS service. We use these illustrations to make the point that the R&O as written may not provide the right distinction between what is permitted and what is forbidden.

### **Third-party services over the access ISP’s network.**

#### **Example 1: A separate specialized service for third-party service**

---

<sup>88</sup> Some networks might provide Service Level Agreements (SLAs) that provide bounds on service quality parameters.

<sup>89</sup> The Report and Order refers to providers of these types of services as “edge providers.”

<sup>90</sup> In this Appendix, this type of provider is called an “access ISP”

An access ISP might set aside capacity separate from the BIAS service to carry the traffic for the third-party services that are using it. For the purpose of this discussion, we will refer to this separate capacity as an Enhanced Access Channel (EAC). There are a number of ways that one might argue that an EAC is a specialized service, and thus not covered by the requirements of the R&O.

Reach: The EAC service, as described, is not intended to reach large parts of the Internet. It is designed only to reach to specific customers who subscribe to the third party service. Using the sub-group's interpretation of the R&O, this example is thus a specialized service. The third-party service is no longer considered OTT, because it is now delivered over the access ISP's EAC. (On the other hand, the packets from the third-party service provider must reach the access ISP by some means—it is a question for consideration whether the means of delivering these (across other parts of the Internet or separated in some way) is part of determining how we characterize the EAC. See Example 3 for an elaboration of this point.)

Capacity isolation: If the EAC is implemented without impacting the BIAS customers' agreed capacity to access the Internet, it can be considered "isolated" from the BIAS service. This argument is similar to the one posed in the IPTV case study.

Generic service: The third-party service is not a generic platform – it is a specific "application level" service. The EAC, as described, would be a general IP platform, but one that is specially provisioned to support such third-party services.

Business model: An access ISP might offer the EAC service independent of BIAS, with separate models for revenue generation. Customers might not need to subscribe to the BIAS service to get access to the third-party services delivered over the EAC.

In addition to the reach criterion, one or another of these reasons might be used to make the case that the EAC can be considered a "specialized service," as defined by the working group, under the Open Internet Report & Order (R&O), even though it is providing access to a third-party service that in other circumstances might be delivered over the Internet (OTT).

**Example 2a: Buying quality of service guarantee (access provider choice) – differentiated service level on BIAS**

If, in contrast, the access ISP implements enhanced access to the third-party service over BIAS by prioritizing the service's OTT traffic amongst all the general Internet traffic going to users over the BIAS, the situation is different. In this example case, there would be no capacity isolation. There is a separate business relationship and possible additional revenue stream. The OTT service is using the Internet, with its global reach. The sub-group concludes that this behavior might fall under the Open Internet rules for BIAS in the R&O and might not be allowed. The lack of capacity isolation (of the preferentially-treated OTT service and general Internet traffic) might additionally warrant consideration by the FCC to further understand the implications for the consumer and the effect on competitive services running over the BIAS.

**Example 2b: Buying quality of service guarantee (user choice) – differentiated service level on BIAS**

Like Example 2a, this scenario assumes the access ISP agrees to implement prioritization of the OTT service's traffic amongst all the BIAS traffic, but only if a given customer elects to have that prioritization of their traffic.

In this case, although there is no capacity isolation, the impact on the customer's other Internet traffic is at their election.<sup>91</sup> The sub-group believes that such a scenario would be subject to the R&O, but would be deemed an acceptable behavior under that order. It might still warrant consideration by the FCC to further understand the implications for the consumer and the possible competitive services running on the BIAS service.

### **Regulatory analysis**

The distinctions between these various approaches are subtle. In each case, the goal is to provide a differentiated experience for a specific third-party-provided application or service. Possible objections to this outcome may include:

- The new service sets a high barrier to entry for new OTT competitors, essentially requiring that they establish such delivery relationships in order to be viable in the market; and/or
- The new service reduces the access ISP's need or likelihood to improve the BIAS service with techniques and tools that might generally improve the performance of similar OTT services. (The so-called "dirt road" BIAS).

Using our proposed definition of a "specialized service," the working group believes (using Example One for illustration) that an ISP that wants to offer enhanced access service qualities to third party services can do so as a specialized service under the R&O. Since there are potential benefits as well as potential harms that might arise from these various services, as the R&O notes, these services must be monitored for their effects on the growth of Broadband Internet Access Services. The working group is of different opinions as to whether consideration of hypothetical outcomes should warrant any reconsideration of definitions at this time, or whether monitoring is the correct action.

These are potential policy considerations that might arise as the FCC considers the method for monitoring the effect of specialized service on BIAS.

### **Third-party services beyond the broadband access network**

The focus of the R&O is on broadband access—the network that provides the actual path to the end user. But the issues that distinguish specialized services from BIAS can be found in the other parts of the network.

#### **Example 3: Specialized core network support**

Assuming there are common performance characteristics and requirements for more than one third-party service, it's not unreasonable to think of a dedicated core transit network being set up to serve as "glue" between third-party service servers and access ISPs – e.g., the early model for

---

<sup>91</sup> Such a service could affect other consumers' service in the case of congestion.

Internap as “Super Performance IP”, or what content delivery networks do for accelerating static content.

In this example, then, a customer will have good performance from the third party service if their ISP interconnects with this dedicated core network. While the third-party service experiences will be different for customers of such ISPs than for their neighbors who do not use an ISP connected to the dedicated core, this is not due to a new or distinguishing feature of the access ISP (e.g., no preferential treatment is given to the 3<sup>rd</sup> party service on the ISPs network).

This is not particularly new – performance between consumers and any network endpoint is dependent on core network connections and conditions.

The working group believes that a reasonable reading of the R&O would suggest that the core of the Internet (the global interconnection of ASes) is not subject to the order. However, much discrimination might occur in that part of the Internet. The working group also asks whether different treatment of traffic in the core of the Internet might influence whether the delivery path across the access ISP’s network is a specialized service, as we question in Example One.

#### **Example 4: Open-standards based approach to signaling requests and requirements throughout the network**

Establishing prioritization of traffic at the access ISP is only going to solve part of the performance problem. Non-interactive services can couple access priority with heavy (and heavily distributed) caching, but that is not applicable in the case of massively multiplayer games. Such OTT services need to have solid network performance between all nodes involved in the interaction, including any transit links.

A future approach might be to ensure that there are open standards and best practices that are developed to support highly interactive traffic in general, and perhaps some level of mutually-cooperative signaling of performance preferences that works across network domain boundaries in the Internet.

(This is not completely theoretical – RITE (“Reducing Internet Transport Latency”) is funded by the European commission under the fp7-ICT programme, with the following focus:

RITE proposes to remove the root causes of unnecessary latency over the Internet. Whilst time-of-flight delay is inevitable, greater delays can result from interactions between transport protocols and buffers. It is these that RITE will tackle.

<http://riteproject.eu/about-2/> )

As part of ensuring that the BIAS service offerings evolve appropriately and are not unduly pushed aside by specialized services, the FCC could consider monitoring such developing technologies and whether they are being appropriately implemented in improving access ISP networks for broadband Internet access services.

### **Consider the future**

In all of this, perhaps the most important thing for the FCC to consider is the distinction between challenges and solutions for today, versus opportunities tomorrow. While the problem outlined here (high performance requirements in globally distributed services) is real, as the examples highlight there are many approaches to addressing the issue in both near and long term ways. Making a ruling to require, enable or prevent a particular behavior today may curtail some of those options.

In the case of high performance requirements of globally distributed services, there is every possibility that technologies will evolve to address the problem in general, and a general trend away from optimizing packet traffic and towards more application/service optimization is possible. This is the thrust of proposals for “Software Defined Networking”, “Information Centric Networking”, and general cloud infrastructure.

## **Open Internet Label Study**

Transparency Working Group  
Open Internet Advisory Committee  
Federal Communications Commission

**The Transparency working group has proposed a system to label Internet service with information that consumers may find useful when selecting a provider, including speed, price, and other metrics.**

The Transparency Working Group of the Open Internet Advisory Committee (OIAC) was formed to provide advice to the FCC on the transparency of offerings from Internet Service Providers (ISPs). In particular, the Open Internet Order [1] says:

“Fixed and mobile broadband providers must disclose the network management practices, performance characteristics, and terms and conditions of their broadband services”

The Transparency Working Group has studied the way that ISPs present performance characteristics and pricing of their service offerings to consumers, coming to the conclusion that presentation consistency would benefit consumers. The Transparency Working Group recommends the adoption of a voluntary open Internet labeling program as a means of helping consumers more easily compare and select Internet service offerings.

### **Motivation for an Internet Service Labeling Program**

Some consumers are not able to easily compare Internet service offerings. Organizations such as the National Hispanic Media Coalition have conducted focus groups that show that some consumers are confused when choosing an ISP. Many articles have been written to highlight that some consumers are confused when choosing a wireless service provider [2][3][4][5].

A simple and consistent label will enable consumers to make apples-to-apples comparisons when considering an Internet service selection or when considering a change.

Once the consumer has made a selection, and at any time afterwards, the label provides the information that could be used by the consumer when accessing a test site to confirm that the service is performing roughly as expected. In addition, third parties can provide consumers with performance parameters that help the consumer in determining whether their existing service fully meets their needs.

While mobile data networks are rapidly evolving, fixed and mobile connections are both a significant part of today's network experience. For this reason, service providers that do not provide access to the entire public Internet should not make use of the label at all.

### **The Proposal – A Label Similar to the Nutrition Label**

The FCC could promote a labeling program for both mobile and fixed services. Such a label program would provide the following information:

- Performance: upload speed and download speed
- Price (monthly fee averaged over three years)
- Usage Restrictions (any points at which the terms of service that apply change)

These numbers are very far from a complete picture of an Internet service offering, yet they seem to be the right level of detail for most consumers. These numbers do not capture important technical factors such as jitter, latency, and impacts of over provisioning. For this reason, the ISP might also provide a much more complete disclosure like the one recommended by BEREC [6]. These details are vital for expert analysis and service offering comparison.

## **Methodology**

To participate in the label program, ISPs self-report three pieces of data: upload speed, download speed, and price. In addition, if there are any usage restrictions, including data caps, ISPs need to report them as well.

The label data is made available for each active service offering. If a service offering is a legacy service and no longer available to new customers, the ISP can determine whether they want to report current data for the legacy service; however, ISPs are encouraged to report data for both active and legacy services.

### ***Upload and Download Speed***

The upload and download speed numbers are meant to reflect the performance delivered by the ISP to a consumer's broadband modem. Yet, it is recognized that upload and download speeds vary greatly from consumer to consumer since they depend on several factors such as geographic location, home network configuration, and time of day. These complexities are well known, and they have been discussed in the context of the FCC's Measuring Broadband America (MBA) program, which compares an ISP's advertised speed with a measured speed. It is important that the terminology and methodology used for the label program be consistent with the MBA program, allowing the two programs to reinforce and supplement each other.

It is envisioned that the label data would include the upload and download speed as determined by lab testing. ISPs measure the maximum ("up to") speeds achievable, within statistical bounds, over a segment of the access network closest to the user (e.g., DSL-capable copper loop segment, or shared DOCSIS channel).

In the near term it is not feasible to base the reported data on large-scale customer measurements. Currently, this type of data reporting is not usually available at scale due to a lack of measurement standards in deployed equipment. In order to establish the labeling program, the FCC will need to work with industry to define a measurement process for the data to be reported by ISPs. Since the upload and download speed numbers are meant to reflect the speeds that consumers can expect to receive, ISPs should take into account any short-term traffic management loads that impact consumer experience as well as long-term capacity management processes when reporting the data for the label.

Please note that outside of the label, the Open Internet Order obligates ISPs to provide relevant information about their service (e.g., upload speed, download speed, usage thresholds, latency, and price). ISPs provide this information today in a variety of ways, including their web sites. Currently, the data used for the upload and download speed inputs for the label is often the same data that the ISPs disclose on their corporate websites. Publication of label data is discussed further below.

### ***Price***

Price is an important aspect of a consumer decision. Initial price for Internet service often reflects a discount or promotion as a purchase incentive. As a result, to reflect the long-term cost to the consumer, an average monthly price reflected for 36 months is proposed. In addition, the prices should reflect all taxes and fees. Since the label shows the monthly average, this will take into account any sign-up discounts, promotions or incentives for new customers, and it reflects any rate adjustments following the expiration of any such incentives.

The price is based on a geographic location, such as the zip code or census block for each service offering. Since pricing often varies by location, it is not usually possible to provide one price for the entire country.

Bundling is a popular practice for ISPs. Bundling refers to giving a price discount to Service A if a consumer purchases both Service A and Service B from the ISP. While regional discounts are reflected in the price, the label only reflects the price for the Internet service offering. Consumers may receive a lower price for the Internet service if they choose additional services from the same ISP. The ISP can make this obvious by providing two labels, one for Service A by itself and another one for Service A and Service B together. When the consumer purchases the Internet service on an ISP's website, the label could reflect the actual price, including any bundle discounts of all of the items in the consumer's shopping basket.

If an ISP has many different service offerings, with and without bundling, in many different geographic locations, then the publication of all of these labels might become unwieldy. However, presentation on a website can be straightforward if the consumer provide their location and then the applicable labels are displayed.

At least one ISP has reservations about the inclusion of price data in the label. This ISP is concerned about the potential to increase customer confusion rather than reducing it.

### ***Publishing the Label Data***

Three alternatives were considered for ISPs to make the label data available:

- 1) The ISP posts the label data on its own web site
- 2) The ISP provides an API to obtain them
- 3) The ISP periodically files them with a third party

Choices (1) and (2) offer the opportunity to be dynamic. That is, when the ISP adds a new offering or makes a change to a current offering, the information is available to the consumer almost instantly. Further, these choices can be driven by a back-end provider database, which

allows the potential customer to provide a location (e.g., a street address) and learn the label data associated with each of the service offerings that are available.

Choice (2) is the easiest for third parties to facilitate comparative shopping using very current information.

Choice (1) is easiest for small ISPs. Choice (3) may also be acceptable for small ISPs, but a periodic filing process could be more cumbersome for consumers and analysts to obtain timely information.

The Transparency Working Group recommends that the FCC pursue choice (1).

### ***Other***

In addition to self-reporting upload speed, download speed, price, and if applicable, usage restrictions for each service offering, ISPs can provide links to the appropriate page on their company website for each offering so that customers can find additional information.

### **Complexities**

There are a number of complexities that must be taken into account when evaluating the label program. Complexities encompass service offerings, customers, and companies. Consideration of these complexities is necessary for a successful label program.

### ***Service Offerings***

Bundling: It is common for ISPs to bundle services. Often bundles provide a price benefit for customers, where the cost of the bundle is less than each service individually. The price discount in a bundle may not be broken out by service. As a result, this adds a layer of complexity when participating in the label program since the price benefit of the bundle is not easily reflected in the price data.

Promotions: Throughout the year, ISPs may choose to run promotions for new and existing customers. These promotions are often limited to a certain time period and may include restrictions such as customers committing to a certain length of service contract. The promotion is reflected in the average, but the initial lower price followed by a subsequent higher price is not reflected on the label itself.

### ***Customers***

Location: Actual download speed and upload speed will vary based on consumer location. The ISP needs reasonably accurate data for each location where the service is offered. Of course, there will be variability within the region. Measuring each zip code, for example, is not practical. Yet, the ISP needs to provide label data that will be close to the actual performance delivered to the consumer's broadband modem in that geographic area. Reasonable estimates can come from laboratory testing.

Variability: Internet usage is not constant throughout the day or week. Similar to highways or air travel, there are peak usage periods during specific times of the day or on specific days of the week. For example, Internet usage is often high during special events like the Super Bowl.

Also, Internet usage is higher between 3pm and 9pm EST than at 3am EST. As such, it is difficult to capture one download speed and upload speed to display to consumers given the variability throughout the week.

Thresholds: The label reports download speed, upload speed, price, and if appropriate usage restrictions. There is a risk that customers will look for service offerings with the highest speed numbers, perhaps greatly exceeding their needs. There is a threshold where the customer will not see a speed difference between two offerings. So, even though an ISP may offer the fastest speeds, the difference between that fast speed and a lower speed may be undetectable for the average consumer. The lack of education in the market on how much speed is sufficient may confuse some consumers.

Other Contributing Factors: Many factors contribute to end-to-end broadband performance that are beyond the control of the ISP, including the specific user application, server capacity, aged equipment, and home network configuration. If a consumer does not get the advertised performance due to these factors, this may lead to confusion and increased customer care costs for the ISP.

### ***Companies***

Beyond Speed, Price, and Usage Restrictions: The label takes into account upload speed, download speed, price, and if appropriate usage restrictions. While each of these elements of a service offering is important for consumers, these elements are not a complete picture. Key factors that also impact consumers but are omitted from the label include, but are not limited to, quality of customer service, ease of use, setup time, jitter, and latency. By not including all the factors in the label, there is a risk that ISPs will start to de-emphasize these essential factors. Creating a market where ISPs are evaluated only by the numbers included in the label may not be a market improvement.

### **Potential Benefits**

The proposed label has the potential to:

Raise Awareness: A well-branded label would raise an average consumers' awareness about the performance and cost of the Internet services that they purchase. The basic information provided in the label would help consumers perform cost-benefits analyses and make good choices based on their needs and budgets.

Reduce Consumer Confusion: The standardization provided by the label would make it easier for consumers to compare services. The simplicity of the label would help reach even the least tech-savvy consumers. In addition, a label with numbers is much easier for non-English speakers to understand than a lengthy explanation of services in point of sale contracts, bills, or advertising materials.

Promote Competition: Internet service providers, in vying to put forward the most favorable label, would be compelled to provide the fastest and most affordable service to an open Internet. Attaching speed, price, and if needed, usage restrictions in a simple and consistent label format that is easily comparable across ISPs will enhance competition.

## Open Internet Advisory Committee - 2013 Annual Report

Incentivize Open Internet Practices: The label will likely become a symbol that the provider, regardless of whether they provide fixed or mobile services, offers access to the entire open Internet. In fact, the lack of a label could be an indication that the provider is not providing access to the entire open Internet.

Marketing Tool: The label may make it clear how the selection of a service bundle impacts the price of the open Internet service.

Improve Consumer Loyalty: A label may improve consumer experience by managing expectations and building trust.

Global Applicability: If the FCC encourages the adoption a label, it could lead to an international standard for rating open Internet services. A label with numbers that are easy for non-English speakers to understand will be more palatable for global adoption.

### **Potential Concerns**

The proposed label could:

Mislead Consumers: A label does not cover all aspects of a service that a consumer might consider in selecting a service. The label does not capture the whole picture, and it might omit an attribute that is important to a particular consumer.

Government Cost: The FCC program will require a design team for the label and the development of guidance on its use. A team will be needed to manage the program over time.

Slow Adoption: The benefits will only be achieved once all ISPs embrace the label program. In addition, promotion is needed for all consumers to be aware of the label and its use.

### **Long-term Future**

The Internet Engineering Task Force (IETF) has developed a set of standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services. Network operators, end users, or independent testing groups can use these unbiased quantitative of performance measurements.

The Broadband Forum has an initiative underway to bring advertised “up to” speeds to be more in line with real-life speed data.

Specific metrics and procedures for accurately measuring and documenting these metrics are under development. Once these metrics are in widespread use, the FCC should consider migrating from service provider estimates of their offerings to actual measurements.

### **Conclusion**

The Transparency Working Group recommends that the FCC work with the industry to develop a voluntary labeling program, in which ISPs would disclose in a simple and consistent manner, relevant information about their broadband Internet access services.

The next steps in establishing the labeling program:

- Establish technical definitions for upload and download speed metrics that are consistent with the definitions used by the FCC's Measuring Broadband America (MBA) program. It is important that the terminology and methodology used by the labeling program be consistent with the MBA program so that the two programs reinforce and supplement each other. If necessary, the FCC should convene a group of subject matter experts to define the upload speed and download speed performance metrics.
- Select a measurement program that will be used in the near term while comprehensive measurement standards are developed and deployed.
- Confirm that publication of the labels on ISP websites is viable.
- Confirm that price should be a part of the label program.
- Get input from the ISP industry.
- Get input from the public and interested organizations, such as the Electronic Freedom Foundation, the Center for Democracy and Technology, and the National Hispanic Media Coalition.
- Design a proposed label as well as HTML assets for use on the ISP websites and marketing documents.
- Implement a pilot with a small number of ISPs to refine the label design, the label presentation, and the methodology. During the pilot, get feedback from consumers as well.

The Transparency Working Group is confident that the Label program will make it easier and less confusing for American consumers when choosing an Internet Service Provider.

## References

- [1] [http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/FCC-10-201A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf)
- [2] <http://www.billshrink.com/blog/press-releases/americans-overpay-336-a-year-on-wireless/>
- [3] <http://readwrite.com/2013/03/27/smartphone-data-plan-comparison-chart>
- [4] [http://news.cnet.com/8301-30686\\_3-20088415-266/ask-maggie-help-these-data-plans-are-too-confusing/](http://news.cnet.com/8301-30686_3-20088415-266/ask-maggie-help-these-data-plans-are-too-confusing/)
- [5] [http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/DOC-298516A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-298516A1.pdf)
- [6] [http://bereg.europa.eu/doc/bereg/bor/bor11\\_67\\_transparencyguide.pdf](http://bereg.europa.eu/doc/bereg/bor/bor11_67_transparencyguide.pdf)

# Committee Member Contributions

In order to provide more comprehensive insight into the individual perspectives that constitute the Open Internet Advisory Committee, members were invited to submit statements reflecting on the Committee's work.

## **Alissa Cooper**

*Chief Computer Scientist, Center for Democracy & Technology*

I was pleased to serve on the FCC OIAC this year. The committee's efforts to examine a variety of complex, contentious Internet openness issues has resulted in an annual report containing valuable insights that should be read carefully by the Commission.

I participated primarily in the mobile working group, and therefore I offer thoughts below about work in other areas. In all of the committee's work areas, it would be beneficial to obtain input from a broader array of both established and start-up companies, including fixed and mobile platform and app developers, content delivery networks, and transit providers.

## **Specialized services**

The *Open Internet Order* recognizes the possibility that specialized services have the potential to impinge on the growth of Internet services, but as the working group alludes to, it is difficult to judge whether this is taking place in the absence of rigorous metrics for assessing existing services of both kinds. For example, the working group rightly concludes that there is ongoing innovation in the delivery of Internet video, but making that observation does not answer the question of whether innovation in Internet video would be even better served if the relationship between current specialized services and Internet services were different with respect to capacity allocation, congestion management, or counting against data caps. The status quo should not necessarily be assumed to be free of openness concerns in the absence of criteria for evaluating the relative quality of the two kinds of services and the progress of both over time. It may be possible for the working group's future work to make a helpful contribution in this area.

## **Transparency**

The Transparency Working Group proposed a voluntary labeling program that would have Internet service providers (ISPs) display labels on their web sites indicating maximum upload and download speeds, prices, and usage restrictions. While the idea of a label is useful, there are several aspects of the program that deserve further consideration. First, the focus on speed may put too much weight on a metric that is not always the primary determinant of performance, particularly as more users opt for broadband products with higher maximum speeds. Second, the label may need to better account for the variability of broadband performance, particularly for mobile users, so as to avoid being more misleading than informative. Finally, recent research has indicated how difficult it can be for consumers to select the most appropriate broadband package for their needs and to understand usage restrictions.<sup>92</sup> The Commission should work jointly with

---

<sup>92</sup> See, e.g., London Economics, "Consumer information on Broadband Speed and Net Neutrality Experiment" (May 2011), [http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/statement/Consumer\\_information1.pdf](http://stakeholders.ofcom.org.uk/binaries/consultations/net-neutrality/statement/Consumer_information1.pdf) and Consumer Focus, "Lost on the broadband super highway" (Nov 2011), <http://www.consumerfocus.org.uk/files/2012/11/Lost-on-the-broadband-super-highway.pdf>.

## Open Internet Advisory Committee - 2013 Annual Report

the research community, consumer advocates, and industry to ensure that labels are effective for consumers.

### **Conclusion**

The committee has made useful contributions to ongoing discussions about Internet openness and I look forward to participating in its the future work.

## **Maurice Dean**

*Director, Open Connect Product Management, Netflix Inc.*

### **Netflix Statement on The Open Internet Advisory Committee July 2013 Report**

Netflix commends the Committee chairs and members for dedicating time and resources to produce this report. The report identifies areas of debate among stakeholders on this committee. It does not resolve these debates. Instead, it identifies areas that should be addressed by the FCC to ensure the Internet remains an open and dynamic platform for free speech and economic growth. Moving forward, we recommend the OIAC and FCC focus on the following areas:

#### **Data Caps**

Consumers pay ISPs for Internet access to connect to the amazing content and services the Internet offers. Data delivery and consumer data consumption are not cost centers for ISPs – they are revenue generators<sup>93</sup> and significant sources of support for broadband deployment. Consumer groups<sup>94</sup> and technology experts point out that data caps are hard for people to understand and are not effective for managing network congestion or costs. Concerns remain regarding the purpose and incentive for applying data caps. These are heightened when ISPs apply caps in ways that favor their own services and penalize consumers who want to use alternative online services. The Open Internet Order cautions against anti-consumer and anti-competitive billing practices. The FCC should monitor to ensure that data caps do not suppress overall Internet usage and impede the Nation’s goal of encouraging broadband adoption, usage and investment.

#### **Interconnection**

The Open Internet Order seeks to keep the Internet “open and interconnected,” yet the Interconnection policies of market-dominant ISPs may negatively impact reliable delivery of popular applications. Just four access providers control nearly 70% of the 80 million broadband subscriptions.<sup>95</sup> The OIAC should examine interconnect practices that impede the free flow of content delivery to consumers.

#### **Application Layer Performance Testing**

The OIAC should augment its current transparency focus by promoting greater consumer insight into application performance. Evaluating broadband performance based on the applications and services that people use most empowers consumers with practical information to evaluate plans and service providers. Content services collate a wealth of anonymized data that can form a wider picture of application performance. Expert input from Edge Providers & transit providers would provide clarity and direction for more open disclosure.

#### **Specialized Services**

Specialized services may promote development of innovative services, but they should not be permitted to cannibalize Open Internet capacity growth. Nor should specialized services be arbitrarily invoked to evade Open Internet protections or to unfairly disadvantage rival Internet applications. Netflix supports identifying criteria to prevent such gamesmanship, however the

---

<sup>93</sup> <http://gigaom.com/2013/02/20/say-it-with-me-now-data-caps-are-about-profits-not-recovering-fixed-costs/>

<sup>94</sup> <http://www.publicknowledge.org/blog/question-core-data-caps-debate>

<sup>95</sup> <http://gigaom.com/2012/08/14/us-added-260000-broadband-subscribers/>

working group has been challenged by definitional inconsistencies within the Open Internet Order. The development of these services should be monitored to ensure that they do not deter investment in Internet services, or limit consumer access and enjoyment of Internet services.

**Broader Stakeholder Input**

The OIAC should address current and emerging obstacles to the Open Internet by seeking input from a broader range of stakeholders in the Internet ecosystem, including transit providers, content delivery networks, and greater representation by Internet start-ups and new entrants who would be significantly impacted by barriers to an Open Internet.

**Jessica Gonzalez**

*Vice President, Policy & Legal Affairs National Hispanic Media Coalition*

This report is the result of countless hours of collaboration between open internet committee members. I congratulate the Commission, the committee and the working group chairs for this notable accomplishment and I look forward to continuing the dialogue which, thanks to these efforts, will now be more informed than ever before.

With that said, it is important to note that adoption of the recommendations in this report, alone, cannot guarantee the continued openness of the internet.

It should come as no surprise that creating this document through a consensus-driven, multi-stakeholder process was not easy. Although I appreciate the collegiality and dedication of my co-members that represent internet service providers, we cannot well expect them to put the public interest over their business interests. As a result, some recommendations are too watered down to actually serve the public interest and preserve a multiplicity of voices over the internet.

The Commission should solicit public comment from the many diverse stakeholders that could not be invited to this table. Many other interested individuals and organizations have expertise in this area and can build upon the outstanding work of this committee.

**Shane Greenstein**

*Professor and Kellogg Chair of Information Technology, Kellogg School of Management, Northwestern University*

**Thoughts on the Open Internet Advisory Committee at its first year**

It is the one year anniversary of the OIAC, and time to briefly reflect on its work. I have been happy to be part of this committee, and contribute to discussions through participation.

My own views are partially shaped by my experience in my working group. It analyzed data caps, and usage thresholds in broadband networks.

Unless I am mistaken, our working group was the first organization to map the landscape by bringing all the arguments into one place. We did not settle anything, nor was it within our purview to settle anything. Rather, we attempted to move the conversation to a more productive place.

Let's be clear about what we accomplished. The group wrestled with competing definitions, and identified bridges between different general arguments and specific observable behavior and facts. The group analyzed where the FCC or another consumer-protection policy-making body, such as FTC, might want to monitor events and where issues remained unsettled.

I have read the other reports too, and they cover novel territory, and they are thorough and careful. They too attempt to move the policy conversation to a more productive place.

There is a lot here to like. I would be happy to continue to serve.

## **Charles Kalmanek**

*Former Vice President of Research, AT&T*

### **Thoughts on the Open Internet Advisory Committee at Its First Year**

The one year anniversary of the Open Internet Advisory Committee is a good time to consider the ground that's been covered. The committee brought together a broad group of experts with different perspectives, all of whom share a commitment to an Open Internet. I was actively involved in the mobile broadband, specialized services, and transparency working groups, and participated in discussions with the economic impacts working group. The collective conclusions of the working groups indicate that the broadband Internet marketplace continues to be dynamic and beneficial for users, identified no clear issues that require regulatory intervention or rule changes, and recommend continued observation by the Commission.

The mobile broadband working group highlighted the dynamic and global nature of the mobile ecosystem. The U.S. leads innovation in this ecosystem in many ways. While mobile network providers play an important role in the ecosystem, the openness and innovation of the mobile Internet depends on the interplay between many players, including platform and device vendors. As a measure of the rate of change in the mobile ecosystem, some of these players have acquired significant influence in the market only recently. Therefore, one of the concerns about the Commission's approach to openness arises from the particular limits of its purview under the Open Internet Order, which is focused exclusively on broadband providers.

The specialized services working group recognized the trend towards network providers offering both Internet services, and specialized services that use the Internet Protocol, over a shared network infrastructure. This trend benefits the Internet by encouraging network investment in a shared IP infrastructure. Because there are functional similarities between specialized and "over the top" services, the group suggested that the Commission should continue to monitor the impact of specialized services on Internet and "over the top" services.

The transparency working group started from the premise that transparency about network management practices is already a requirement of the Open Internet Order, and is the best way to protect consumers. The group's proposal for a labeling program, similar to a nutrition label, has the potential to improve customer understanding of Internet service offers by standardizing the information disclosed by providers. However, the proposal to include an "averaged" price in such a label would likely increase customer confusion.

The economic impact working group recognized the validity of usage-based pricing, and the fact that data caps do not currently affect most households in the U.S. Usage-based pricing approaches are based on the principle that heavier data users pay somewhat more than lighter data users. Some members expressed concern that data caps could affect the growth of data intensive services in the future. The group acknowledged the need for better consumer education and consumer measurement tools, and suggested the importance of continued monitoring by the Commission.

It is gratifying that the conclusions of OIAC's first year of work affirm the Internet's continued robustness, openness, and innovation. While I am leaving the OIAC after this year, I appreciate

the opportunity that I've had to contribute.

**Kevin McElearney**

*Senior Vice President for Network Engineering, Comcast*

This Report of the Open Internet Advisory Committee (OIAC) summarizes our first year of work. The Committee considered many challenging issues from the Open Internet Order that have been the focus of speculation and much public debate. But because the Committee committed to adhering to the facts, and to considering diverse views and information from all its members, the Report generally achieves a consensus that we all can embrace and that acknowledges that the Internet is a complex, dynamic, and multi-party ecosystem.

As the Report reflects, the OIAC collaboratively considered issues regarding specialized services, data caps, consumer disclosures, and openness in the mobile ecosystem. From my perspective, there are three general findings that stand out:

- 1) Broadband Internet Access Services in the US continue to grow, evolve, and remain open to innovation.
- 2) The OIAC is an effective way to help bring the Internet community together to help understand different points of view and educate with facts and data (versus speculation) on challenges and opportunities for the future.
- 3) The FCC should continue to play a constructive role in this dialogue.

Additionally, there are some specific recommendations to guide the FCC that reflect the fact that the Internet's greatest strength is its continuing ability to evolve and that we should be cautious of anything that may unexpectedly constrain innovation and investment. For example, the Report concludes that:

- Regulations should not create perverse incentives for operators to move away from converged IP infrastructure or innovative technologies that benefit consumers.
- The existence of specialized services is not something new and provides consumer benefits, including new investment in networks that support broadband.
- A service should not escape regulatory burdens, or acquire a burden, merely by moving to IP.
- To the extent possible, regulation should be technology-neutral.

The Report also reached some conclusions about data caps and usage based pricing that are worth highlighting. As an initial matter, the concept of a take-it-or-leave-it "cap" has been replaced with products designed to offer flexibility. The reality is, with wireline broadband, most of today's *usage tiers* impact only extreme users and meet 98-99 percent of the customer demand today allowing customers to enjoy a full Internet experience. Usage tiers are designed to have the following benefits:

- Ensure that the majority of end users are not forced to subsidize the highest extreme end users;

## Open Internet Advisory Committee - 2013 Annual Report

- Enable ISPs to create lower-cost broadband plans that spur adoption while also offering the highest end services for early adopters and innovators; and
- Ensure applications and Internet services have incentives to use network resources efficiently.

It has been an honor to participate in this important work with an exceptionally talented and diverse group of Internet industry experts and advocates who share a common interest – the Open Internet. At Comcast, we will continue to work with the FCC to protect the openness of the Internet as well as the continued investment and innovation that has made the Internet the vibrant and dynamic platform that it is today.

**Charles Slocum**

*Assistant Executive Director, Writers Guild of America, West*

I joined the Open Internet Advisory Committee to participate in the important work of evaluating the effects of the Commission's Open Internet Rules. The Writers Guild of America, West, a labor union representing more than 8,000 television, film and online video writers has been an ardent supporter of the open Internet and the rules that protect it. An open Internet promises to increase the options for content distribution, a critical development for independent producers who have been all but eliminated from the television landscape following the demise of the financial interest and syndication rules and subsequent media consolidation. Already, the open Internet is delivering on this promise, with Netflix series created by Writers Guild members garnering 14 Primetime Emmy nominations. It is important to maintain this openness to encourage further development.

My participation on the Committee has been focused on the Economic Impact Working Group's review of data caps. While the Open Internet rules allows ISPs to offer usage based pricing models, I remain concerned about the potential impact of such offerings, particularly on the development of online video. As discussed in the report, capping Internet usage or imposing additional costs for higher levels of consumption could deter consumers from adopting online video viewing. This could harm the positive progress that has been made by the introduction of online video services such as Netflix and Amazon Prime and could deter new entrants, to the detriment of competition and innovation. The report produced by our Working Group is the product of varying viewpoints and interests. It examines the issue of data caps from different perspectives, but makes no recommendations for Commission action. As suggested in the report, I urge the FCC to continue to monitor ISP use of data caps and other forms of usage based pricing.