Specialized Services: Summary of Findings and Conclusions

Specialized Services Working Group

Open Internet Advisory Committee
Federal Communications Commission

Released August 20, 2013

Full Annual Report of the Open Internet Advisory Committee available here

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.[74] 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.

[74] 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."[75] 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[76]. 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.[77] This appears to remain true today, as even new entrants such as Google Fiber offer video services in addition to BIAS[78]. 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."[79] Examples of specialized services mentioned in the R&O include facilities-based VoIP, IP video,[80] e-reading services, heart rate monitoring, and energy sensing.[81]

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[82]:

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

---

[75] *Id.* at 112.
[76] Independent of whether it is in the business interest of a BIAS provider to offer QoS, the R&O may not permit this option.
[77] 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.
[78] 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/
[79] *Id.* at 7.
[80] *Id.* at 61.
[81] *Id.* at 33.
[82] *Id*. at 44.

> *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.[83]"

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[84].'

[83] 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]

[84] 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[85]. 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

[85] 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."[86] The FCC has declared their intention to monitor this situation. This committee is asked to advice 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

[86] 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"[87], 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".

[87] 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 3rd 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.[88] 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.[89] 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[90] 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

[88] Some networks might provide Service Level Agreements (SLAs) that provide bounds on service quality parameters.

[89] The Report and Order refers to providers of these types of services as "edge providers."

[90] 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.[91] 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

[91] 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 3rd 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.