

Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, DC

GN Docket No. 14-28

In the Matter of

Protecting and Promoting the Open Internet

Appropriate Rules for Managed or Specialized Services

Jon M. Peha

Professor, Carnegie Mellon University

Address: Carnegie Mellon University
Department of EPP
Pittsburgh, PA 15213-3890
peha@cmu.edu

January 5, 2015

1 The Specialized Services Loophole

Poorly written rules for *specialized services* or *managed services*, a category of services recognized in the FCC's 2010 Open Internet Order, could turn the pending Open Internet proceeding into a farce. (This comment will use the terms *specialized services* and *managed services* interchangeably.) There are people who seem to believe that any service from an Internet service provider (ISP) that has a quality of service (QoS) that is different from standard Internet service must be a specialized service, and that specialized services are exempt from all Open Internet rules. If the FCC accepts both premises, then it is possible that no form of discrimination, no matter how unreasonable or anticompetitive, could ever be prohibited under Open Internet rules. Imagine that an ISP offers a new service that complements its standard Internet service. If the new service has a different QoS from standard Internet, then the new service is labeled a specialized service and there are no applicable Open Internet Rules to break. On the other hand, if the new service has the same QoS as standard Internet, then there is probably no discrimination going on, which means there is also no chance that the ISP has violated antidiscrimination rules. There is a similar problem if specialized services are exempt from all rules, and if any service that carries traffic from just one application type is considered a specialized service. In this case, the mere act of blocking traffic from the other application types makes the service exempt from rules that prohibit unreasonable blocking based on application.

If specialized services are to exist at all, the FCC clearly cannot accept both of the premises above but it could accept one, which yields (at least) two reasonable policy options. One is to establish a much narrower definition of specialized services than was considered above, which could make it reasonable to exempt specialized services from most Open Internet rules. (Some transparency rules may still be important, so consumers know what they are paying for, but the nature of transparency also depends on the definition of specialized service.) Alternatively, specialized services could come with rules that prohibit some forms of blocking and discrimination, although these rules could differ from those that apply to standard Internet access. In this case, it might be reasonable to allow a much broader definition of specialized service. Although in previous filings [1] I have suggested the former approach, either approach is viable. This comment will explore both of these approaches. It is important to note from the above that anyone who wants to devise an effective policy cannot address the issues of what a specialized service is and what rules should apply to specialized services in isolation. These two issues are too intertwined.

It is similarly difficult to separate rules regarding specialized services from rules regarding unreasonable discrimination and unreasonable blocking, both of which are linked to the definition of reasonable network management. Any attempt to address one of these issues without consideration of the other could lead to bad policy. Applications may emerge that require better QoS than is widely available on the Internet today, such as telemedicine. If we allow some forms of discrimination that are deemed to be reasonable within an Internet access service, as I have proposed in the past [1,2], then these applications can operate under general Open Internet rules, and we can greatly limit what constitutes a specialized service without accidentally prohibiting these applications. In this case, there

may be little reason to regulate specialized services. On the other hand, if we greatly limit the discrimination allowed under Open Internet rules with the belief that applications like telemedicine will simply shift to specialized services, then we need a broad definition for specialized service. In that case, we need more rules regulating specialized services.

Section 2 explains why Open Internet rules should allow reasonable discrimination to support applications that require better QoS, and why rules are needed to prohibit unreasonable discrimination. This balance could be achieved in rules for basic Internet access, or rules for specialized services, or both. Section 3 describes a policy where many applications requiring better QoS choose specialized services, so the definition of specialized service is broad, and there are more rules regulating specialized services. Section 4 describes a policy where applications requiring better QoS can get the reasonable discrimination they need in basic Internet access, so the definition of specialized service can be much narrower, and the rules governing specialized services can be more relaxed. The options are summarized and compared in Section 5.

2 Reasonable Discrimination Can Help Consumers. Unreasonable Discrimination Can Harm Consumers.

As I've discussed in more depth elsewhere [1,2], consumers can greatly benefit from some uses of discrimination, and can be harmed by other uses of discrimination. Ideally, we should seek Open Internet rules that encourage the former and prohibit the latter.

Some applications benefit from superior QoS, and this is best provided through some form of discrimination. ("Priority" is one of the forms of discrimination that could be used, although not the only form. As discussed in [1], the FCC NPRM discusses limiting priority when it really should be limiting discrimination.) For example, someday a telemedicine application may allow patients to recover from surgery in their own homes, but with constant monitoring from medical professionals at a hospital. An ISP could use discrimination to ensure that QoS is adequate for medical monitoring. Thus, some applications that benefit from superior QoS actually involve life and death, but many do not. For example, those recovering patients and their healthy neighbors may want to listen to music streamed by "Internet radio" stations, but with guarantees that transient congestion won't ruin the sound. FCC rules should not prohibit the kind of discrimination that makes such applications possible, nor should FCC rules prevent an ISP from charging for these services, since the services may not emerge unless someone pays for them. These examples should be considered *reasonable discrimination*. In the best cases, Internet access providers would simply be allocating limited resources to improve the QoS of those applications that benefit most from those resources, and charging for the resources allocated.

Arrangements like these become problematic when an ISP with market power seeks to extract oligopoly rents through fine-tuned discrimination [2]. Perhaps the telemedicine application requires the

exact same QoS from the network as a new multiplayer game, but if the network is allowed to set prices based on the application users choose rather than based on the QoS that the network offers, the ISP will charge very differently for telemedicine and gaming. In each case, with limited competition, the ISP will charge the maximum that the market will bear for that particular application, and consumers could see higher prices and less availability for content and applications. Or perhaps someone wishes to use the same underlying service that supports the streaming of music to support the streaming of political speeches against mergers in the telecom industry. Would an ISP agree to carry the music but not the speeches? These are examples of discrimination based on application and content, and discrimination by those criteria should be considered *unreasonable discrimination*. If an ISP with market power has unrestricted ability to discriminate based on content and application, then it has the ability and the incentive to give content providers affiliated with the ISP an advantage over competitors, or to bring separate oligopoly rents onto each distinct application or content market even when that application or content market is highly competitive [2]. In accordance with the Open Internet principles, it should be considered *unreasonable* to block or discriminate with respect to QoS or price solely based on content, application, user, or type of attached device, when none of these factors affects the scarce resources that the network must allocate. It should be considered *reasonable* to discriminate between class A service and class B services with respect to QoS and price if both services are available to all, but *unreasonable* if some users can access a service and other similarly situated users cannot.

3 Option 1: Broad Definition of Specialized Service, and Meaningful Regulation of Specialized Services

First consider the case where many applications that require superior QoS would have to operate over specialized services, because the Open Internet rules for Internet access do not allow this form of discrimination, or do not allow Internet access providers to be paid for offering superior QoS. In this case, we can expect different specialized services to emerge, each with a QoS that is appropriate for a different application type, and there should be Open Internet rules governing these specialized services to prevent unreasonable discrimination.

As demonstrated in Section 2, Open Internet rules should allow ISPs to offer services of different QoS, which requires some form of discrimination, but some uses of discrimination should not be allowed. Open Internet rules should not allow Internet access providers to offer a given specialized service to some users, perhaps those affiliated with the ISP, and not to others who are similarly situated. To favor some users over others in this way should be deemed unreasonable. In addition, while Internet access providers should be able to take technically relevant factors that affect cost such as data rate or guaranteed latency into account when setting prices for specialized services, Open Internet rules should not allow these providers to base prices on content or application, nor should they be allowed to limit access to a specialized service based on content or application and block non-conforming traffic. These strategies would constitute unreasonable discrimination and unreasonable blocking, respectively.

All of the unreasonable discrimination described above would be prohibited while still allowing Internet access providers to offer superior QoS through discrimination through either of the two following policies. In a more traditional approach, ISPs could be required to post the types of specialized services they offer, including the form of discrimination/prioritization or a technical description of the QoS that is expected as a result of that discrimination, and the associated prices. The Internet access provider must make the same options available to similarly situated users, and those users would be free to choose the specialized service that meets their needs. Alternatively, in what I have called a “Most Favored Nation” approach [1], ISPs could negotiate individually to create new arrangements for specialized services, but the details of those arrangements would subsequently become public, and other similarly situated users would have the right to enter into an arrangement with the ISP that has the same technical and financial terms and conditions. Again, those arrangements should be based on the form of discrimination or a technical description of the QoS that is expected as a result, and not on content or application. The latter approach gives Internet access providers greater flexibility, but still offers protection against unreasonable discrimination and the extraction of oligopoly rents.

4 Option 2: Narrow Definition of Specialized Service, and Little Regulation of Specialized Services

Some people believe specialized services should have no Open Internet rules. Some form of transparency might be needed, but there may be little need for rules on discrimination and blocking if there are adequate limits to the definition of specialized services, and if Open Internet rules do not prohibit reasonable discrimination within basic Internet access so that applications requiring better quality of service can be supported.

In the absence of clear limits on this definition, current technology would make it easy for Internet access providers to engage in unreasonable discrimination among their offerings while maintaining the illusion that their specialized services are somehow distinct from their Internet access service. For example, as previously suggested, traffic from a supposed specialized service “may be sent over a separate virtual local area network (VLAN), or a separate service flow in a cable system operating under the Data Over Cable Service Interface Specification (DOCSIS) standard” [2]. While the VLAN or service flow for a specialized service may appear to be separate from general Internet traffic, traffic from both services travels through the same bottleneck links, and the VLAN or service flow identifier is used to give the supposed specialized service priority over other traffic during periods of congestion. In reality, this is simply a priority Internet service. There is nothing inherently wrong with priority Internet services, but they should not receive blanket exemption from Open Internet rules regarding discrimination and blocking.

If we consider the case where few rules if any are imposed on specialized services, then we must ask whether there are cases when FCC rules should allow ISPs to engage in even the most blatantly anticompetitive forms of discrimination or blocking, and then limit the definition of specialized services to just those cases. For this situation to arise, a company must offer two communications services. One is Internet access and subject to Open Internet rules, and the other is not Internet access, and imposing Open Internet rules on that service would be inappropriate.

For a service to be exempt from Open Internet discrimination and blocking rules without undue risk of the kind of oligopoly rents discussed in Section 2, that service should of course not be Internet access. Also, it should not simply be just another way to communicate with things that users would otherwise access over the Internet, as that would make it a trivial substitute for Internet.¹ Finally, it should not operate over limited resources that would otherwise be used for Internet, as that would make it a simple Internet service with preferred access to shared resources.

There are services with the properties above for which the application of Open Internet rules would be inappropriate. This can occur when Open Internet rules apply to separate networks as if they were one network. Consider a company that offers two IP-based services. One gives a user access to the entire Internet, and the other gives a user access to a closed and highly secure network with endpoints that trust each other. For example, a company might use the closed network to connect its various offices and the residences of some trusted employees. The network for Internet access and the closed network are separate, in that traffic from one cannot reach the other, cannot carry malware to the other, and cannot even cause congestion to the other, which is precisely what makes the closed network valuable to its users. To further enhance the security and productivity of the closed network, traffic from some applications is blocked. Open Internet rules should not be imposed on the closed network, as this might allow traffic from unknown sources and unknown applications into the network, thereby reducing the network's value to users. Note that this is a true private network, and not a mere *virtual* private network (VPN). Unlike private network traffic, VPN traffic is intermixed with Internet traffic, and is therefore similar to any other application running on the Internet. If a network acts to protect QoS for a VPN even when the volume of Internet traffic is high, it is discriminating between two types of Internet traffic: VPN and non-VPN. Such discrimination should be allowed, but to avoid the problems discussed in Section 2, not with a blanket exemption from Open Internet discrimination and blocking rules.

Another example occurs when a company offers both an Internet and a telephone service that are entirely separate. Consider the case where the telephone network is upgraded from circuits to voice over IP (VOIP), although it remains separate from the network that provides Internet access. This alone should not make the telephone service subject to Open Internet rules, which would prohibit blocking non-VOIP packets.

¹ As [2] warns, discrimination can be accomplished if "one can simply provide separate channels for different classes of traffic. For example, favored traffic may be sent over a lightly used wavelength in a fiberoptic cable, while other traffic goes over a heavily used wavelength."

Both of the services above are not Internet and not a trivial substitute for Internet because customers use them to access endpoints that are not on the open Internet. In the first case, those endpoints are other trusted computers within the company's private network. In the second case, the endpoint accessed is a VOIP-PSTN gateway that is operated by the carrier, and is accessible only through the carrier's closed network. In addition, central to the examples above is that the service that is exempt from Open Internet rules is truly separate from Internet access. Implicitly, some of the debate about specialized services is over what it means for two services to be separate. I propose here a simple litmus test that should be used to determine if services are separate, along with a condition to exclude services that are trivial substitutes for Internet.

A communications service can be considered a *specialized service* under Open Internet rules if (i) the primary use of the service is not to access content, services, or systems that are accessible through an Internet access service, and (ii) the service does not share capacity with Internet access.

Two services are said to share capacity if it is ever possible for utilization of one service to affect the performance of the other service.

For example, when a telephone network offers a traditional DSL service, circuit-switched telephone and Internet services do operate over the same copper wire but they do not share capacity. Internet and telephone traffic travel within separate frequency bands in the last-mile connection, arrive at separate switches at the central office, and are forwarded into separate nationwide networks (i.e. the PSTN and the Internet). If there is a high volume of Internet traffic on the last-mile link, in the IP router closest to the user, or in any link or router anywhere on the global Internet, this will not affect telephone service in any way. Similarly, a high volume of telephone traffic will not reduce the performance of the Internet service. This would not change if the telephone service was converted to VOIP, but the capacity allocated to telephone service is fixed, as is the capacity allocated to Internet service. In contrast, if VOIP packets are sent over the same last-mile link as Internet traffic, and VOIP packets are simply given transmission priority over Internet traffic based on VLAN label, then telephone utilization would somewhat degrade the quality of Internet service. This latter case is an example of shared capacity.

The impact of a specialized services rule on telephone and cable TV services deserves particular attention, in part because these two services have played a particularly important role in the business case for Internet access (e.g. as part of "triple play"), and in part because these two services have their own regulations. Some IP-based telephone and cable TV services would qualify as specialized services under the proposed rule above, but some would not. For example, with AT&T's U-verse service, a customer who begins watching a pay-per-view video may notice a sudden decrease in Internet performance because capacity is shared. AT&T adopted this architecture years ago, and if the FCC were to decide now that this technical approach makes the IP-based service used for video distribution subject to Open Internet rules, it would be disruptive. Consequently, I propose the following rule to specifically address IP-based telephone and cable TV services.

A communications service can be considered a *specialized service* under Open Internet rules if the service is only used to provide a service that is subject to telephone regulations or to cable TV regulations.

If specialized services are limited to the extent proposed in this Section, then it is important for the FCC to allow reasonable discrimination within Internet access services. That way, Open Internet rules will not prevent the emergence of those services with high QoS that benefit consumers. This is easily accomplished by applying either of the two policies described in Section 3 to basic Internet access rather than to specialized services as proposed in [1]. In addition to its greater simplicity, there are legal advantages to adopting these rules for Internet access rather than just for specialized services. If Open Internet rules required the Most Favored Nation approach for Internet access (and not just specialized services), the courts would be more likely to accept the argument that the FCC has authority under Section 706 of the 1996 Telecom Act to impose Open Internet rules, as explained in [1]. However, for good or for ill, if the FCC simply applies the definition of telecommunications services as mandated by Congress in the 1996 Telecom Act, the FCC must conclude that commercial Internet access service as it is offered today is a telecommunications service and is subject to Title II, as explained in [3]. There is a long tradition under Title II of allowing reasonable discrimination while prohibiting unreasonable discrimination (and the FCC has authority to forbear from applying regulations that would be counterproductive [4]).

5 Summary

It is important that the FCC adopt appropriate rules for specialized services. Failure to do so could provide a loophole that would allow even the most harmful forms of discrimination or “paid priority” to gain widespread use, or it could have the effect of denying Internet users access to valuable applications that need good quality of service. This comment has described two approaches that could work. Both are summarized in the following table.

	<i>Discrimination rules for basic Internet access</i>	<i>Definition of Specialized Services</i>	<i>Rules for Specialized Services</i>
Option 1	May or may not allow reasonable discrimination and the ability to charge for it.	Definition is broad to support all traffic streams that could benefit from superior QoS.	Unreasonable discrimination and blocking are prohibited.
Option 2	Allow reasonable discrimination and the ability to charge for better QoS. Prohibit unreasonable discrimination. ²	Definition is narrow, including only services that meet one of these conditions 1. The primary use of the service is not to access content, services, or systems that are accessible through an Internet access service, and (ii) the service does not share capacity with Internet access. <i>Or</i> 2. The service is only used to provide a service that is subject to telephone regulations or to cable TV regulations.	Few rules if any (other than transparency).

References

- [1] J. M. Peha, "A 'Most Favored Nation' Approach to an Open Internet," Comments in the Matter of Protecting and Promoting the Open Internet, *Federal Communications Commission GN Docket No. 14-28*, July 15, 2014. http://users.ece.cmu.edu/~peha/Peha_Open_Internet_July_2014.pdf
- [2] J. M. Peha, "The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy," *Telecom Policy Research Conference (TPRC)*, 2006, and *International Journal of Communication*, 2007. http://users.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf
- [3] B. A. Cherry and J. M. Peha, "The Telecom Act of 1996 Requires the FCC to Classify Commercial Internet Access as a Telecom Service," Comments in the Matter of Protecting and Promoting the Open Internet, *Federal Communications Commission GN Docket No. 14-28*, Dec. 22, 2014. http://users.ece.cmu.edu/~peha/IP_Transfer_is_Telecom.pdf
- [4] J. L. Mindel and M. A. Sirbu, "Regulatory Treatment of IP Transport and Services," 2000. <http://dspace.mit.edu/handle/1721.1/1514>

² See Sections 2 and 3 for discussion of the difference between reasonable and unreasonable.

Author Qualifications and Disclaimer

Jon Peha is a Full Professor at Carnegie Mellon University, with experience in industry, government, and academia. In government, he served at the FCC as Chief Technologist, in the White House as Assistant Director of OSTP, in the House Energy & Commerce Committee, and at USAID for the Telecommunications Leadership Program. In industry, he has been Chief Technical Officer for three high-tech companies, and member of technical staff at SRI International, AT&T Bell Labs, and Microsoft. At Carnegie Mellon, he is a Professor in the Dept. of Electrical & Computer Engineering and the Dept. of Engineering & Public Policy, and former Associate Director of the Center for Wireless & Broadband Networking. Dr. Peha holds a PhD in electrical engineering from Stanford. He is an *IEEE Fellow* and an *AAAS Fellow*, and was selected by AAAS as one of 40 Featured Science and Technology Policy Fellows of the last 40 years ("40@40"). Dr. Peha has received the FCC's "Excellence in Engineering Award," the IEEE Communications Society TCCN Publication Award for career contributions, and the Brown Engineering Medal. He consults on a wide range of technical and policy issues related to information and communications technology.

In writing this comment, Dr. Peha represents no one but himself.