The Exceptions Make the Rule: Reasonable Network Management in the FCC's Vision of Net Neutrality

Stephen Goldmeier, May 2010

I. Introduction

Everyone perceives of the Internet differently. Some people see it as the last best hope for humanity's salvation. Others see it as a lawless frontier where anything goes. The reality is somewhere in between. More and more users are heading to the Internet for more and more purposes, both illicit and noble. The Internet is getting bigger and busier, and companies providing connection to the Internet need to keep apace of that growth.

With that in mind, various lawmaking bodies have undertaken to try to preserve the best functions of the Internet. And one such tool for preserving the Internet's best function is network neutrality regulation. Network neutrality, or net neutrality, is the idea that content on the Internet should be delivered by providers of Internet service (ISPs) without any discrimination based on the nature of that content.[1] This principle has most recently been outlined in a Notice of Proposed Rulemaking from the FCC.[2] The bulk of the debate over the form of the FCC's rules, however, is over what would constitute good reasons to allow some forms of discrimination, or exceptions to define the limits of the rule.

In response, the FCC has specified some exceptions in the Open Internet Notice to its general rule of nondiscrimination. In this paper, I will first explain the background of the Internet, why net neutrality is so important and hotly debated, how Internet regulation has

[1] I will discuss more fully why there is so much concern over net neutrality and why this principle is so important in section II, *supra*.

[2] *Preserving the Open Internet; Broadband Industry Practices*, 24 F.C.C.R. 13064 (2009) (hereinafter "Open Internet Notice").

worked so far, and some common forms of network management. Then I will describe how the FCC attempts to define "good reasons" for network management with the exceptions in the Open Internet Notice. Next I will discuss each of the proposed exceptions in turn, looking at its form and the shortcomings of that form in meeting the FCC's goals. Finally, I will propose some alternatives to the proposed exceptions.[3]

II. Background: The Internet and Internet Regulation

Regulation of the Internet has been a rocky process. That rockiness is due to the unique character of the Internet and the structure of the computer network that makes it up. I will first discuss that uniqueness, then I will outline the history of the FCC's regulation of the Internet, and then I will briefly discuss the most common categories of network management.

A. The Unique Character of the Internet

To understand why net neutrality has become such an important issue, we must first understand what makes the Internet so different from traditional broadcast media.[4] The Internet is, most simply, an interconnected array of computers, with multiple connection pathways from one computer to another. These connected computers can be servers (like the places where website content is mostly stored) or users (the people viewing that website content). The important thing is that the network itself makes no distinction between these different types of computers; any computer with the right protocols can connect to this complex network without barrier. When these computers are connected, the data they send and receive is divided into

[3] I am aware of the recent decision in *Comcast v. FCC*, No. 08-1291, slip op. (DC Cir. Apr. 6, 2010), and its effects on the FCC's jurisdiction. I feel, however, that a full discussion of the FCC's approach to "reasonable network management" can still add to an understanding of how this term is defined in the future, either as an effective FCC rulemaking or in some other legislative form.

[4] *See generally,* Lawrence Lessig, *The Future of Ideas* (2001) at 26-99 (fully discussing the nature of the Internet, the end-to-end structure, and how this encouraged innovation on the early Internet. I use Lessig's structure of this discussion as a model for my own).

discreet packets, sent through the complex labyrinth of Internet connections, and then reassembled from these packets on the other side.

For all of their complexity and branches, however, the connections between these computers are actually very unsophisticated. They are essentially "dumb pipes" that relay data from one computer to another. Any actual processing or manipulation of this data is done at the ends, where the information is either sent or received. The pipes themselves are like any pipes, passively doing its job.

And just like real pipes, these pipes need to actually exist somewhere, eventually connecting to the user. That's where ISPs come in. ISPs are the people that connect users' computers to that larger "dumb" array. Users pay a fee, and the ISP keeps them connected to the Internet. As a result, of course, ISPs have to maintain these connections for users, making sure they don't burst from too much traffic or slow to a trickle from too many users. One of the ways they do this is by "network management."

ISPs can monitor Internet traffic in various ways, including deep packet inspection, which allows ISPs to tell not only what kind of file is being transferred, but also some content of that file, by looking at individual packets that are being transferred.[5] This inspection can even reveal key words within text as it is transmitted.[6] ISPs can, and do, use information like this to influence the way traffic flows over its networks, often slowing down traffic during peak usage times or stopping it all together if it's harming its network.

That's where the fear comes in; many commenters, scholars, experts, and even average citizens are concerned that if ISPs have uncontrolled latitude to manage their networks, they could discriminate between these different types of traffic for the wrong reasons. Some of these

[5] Jon M. Peha, *The Benefits and Risks of Mandating Network Neutrality and the Quest for a Balanced Policy*, 1 Nt'l J.of Comm. 644, 650 (2007).
[6] *Id.*

wrong reasons include viewpoint discrimination, blocking Internet services that compete with its offerings, or even impeding new technologies from development by charging them more for access to the network.[7] These fears prompted the FCC to consider its options in regulating this kind of discrimination.

B. The FCC's Internet regulations

Since those earliest days, however, regulation of the Internet has posed a problem for the FCC. Some argued that the FCC should be treated like telephone companies and regulated under common carriage law. Others wanted specific regulations that would require network neutrality. In the end, the latter camp won out, and the FCC From the beginning, the FCC elected not to regulate broadband Internet providers through common carriage or open access regulations.[8] Instead, the FCC drafted a very brief Internet policy statement, which included provisions guaranteeing competition among Internet providers and protecting access to all lawful Internet content and applications for any users.[9] The policy statement also made clear that any regulations promoting the open internet would be subject to service providers' needs to engage in reasonable network management.[10] Finally, the FCC made the adoption of network neutrality principles a condition of merger approval.[11]

Aside from the Internet policy statement, the FCC has also used its adjudication power to create a record of open Internet regulation. The FCC entered into a consent decree with Madison River Communications, requiring the company to halt its practice of blocking voice over Internet protocol (VoIP) traffic.[12] Then, the commission received a complaint about discriminatory

[7] Some types of network management are explained in section II.C., *supra.*

[8]*Appropriate Framework for Broadband Access to the Internet over Wireline Facilities*, 20 F.C.C.R. 14853 (2005).

[9] *Internet Policy Statement*, 20 F.C.C.R. 14986 (2005).

[10] *Id.* at 14988 n.15.

[11] *See, e.g.*, *SBC Communications Inc. and AT&T Corp. Applications for Approval of Transfer of Control*, 20 F.C.C.R. 18290, 18392, Para. 211 (2005).

[12] *Madison River Communications*, 20 F.C.C.R. 4295 (2005).

practices by Comcast, specifically that the company was halting any Bit Torrent traffic initiated over its network.[13] This was the FCC's opportunity to further define its stance on net neutrality and non-discrimination.

In this case, Comcast was monitoring its users for BitTorrent traffic, a kind of traffic that involves sharing files (often copyright-infringing files) and transmitting pieces of these files to other users.[14] Comcast felt that this traffic was harmful, so they injected disrupting packets into the stream of packets making up this traffic.[15] The injected packets caused the connection between these sharing users to cease, thus essentially blocking BitTorrent users from maintaining a connection over Comcast's networks.[16] In effect, Comcast was blocking anybody using BitTorrent from using its services. The Commission ruled against Comcast, using the opportunity to further define and explain how it would treat practices like Comcast's in the future. Specifically, it said that if a practice appears unreasonable, then to prove that the practice is reasonable, the ISP must demonstrate "a tight fit between its chosen practices and significant goal."[17]

The Comcast decision represented the first time the FCC had concretely exercised its power to use Internet regulation to preserve the open Internet and to foster network neutrality. Soon after the Comcast decision, the FCC used its substantial record of broadband industry information and its history of regulation through decisions and the policy statement to propose draft rules on preserving the open Internet, published in the Open Internet Notice.

[13] *Formal Complaint of Free Press and Public Knowledge; Broadband Industry Practices*, 23 F.C.C.R. 13028 (2008) (hereinafter "Comcast Order").

[14] *Id.* at Para. 7-9.

[15] *Id.* at Para. 9.

[16] *Id.* at 3 (defining "RST" packet and explaining how it effectively ceases a particular Internet connection).

[17] *Id.* at Para. 47.

C. Types of Network Management

Before embarking on a discussion of the FCC's actual proposed rules, I will briefly discuss the two major types of network management.

The first is protocol agnostic management.[18] Such management practices are those that put limits on bandwidth consumption of individual users, determined by that user's usage patterns and the normal usage of the network.[19] An example of protocol agnostic management is a system for determining if a user's bandwidth consumption is excessive and subsequently slowing down their usage to preserve the connections of other users.[20] Such protocol agnostic measures are targeted at excessive users, not specific services.

The second is called protocol specific network management. Protocol specific management would be any tools that treat traffic differently depending on what applications or services the traffic represents.[21] An example is the prioritization of latency-sensitive[22] traffic over more low-priority traffic.[23] This type of management makes decisions on priority based on the service that is being used on the network.

Clearly, both systems leave room for problematic discrimination, but protocol specific management, as it is tied to the services a user chooses to use, allows for more anticompetitive behavior. The exceptions that allow "reasonable network management" do not, on their face,

[18] Benjamin Lennett, *Dis-Empowering Users vs. Maintaining Internet Freedom: Network Management and Quality of Service (QOS)*, 18 CommLaw Conspectus 97, 119 (2009) (hereinafter "Lennett").
[19] *Id.*
[20] *Id.* (Comcast uses an approach similar to this one called Fair Share, which slows down traffic to heavy users during times of congestion).
[21] *Id.* at 123.
[22] *See* VI.A., *supra*, for a more detailed discussion of latency.
[23] Cox Communications ran a trial of network management practices using a protocol specific management mechanism. This trial is discussed in VI.A.1.b., *supra*.

favor one over the other, but this difference could influence the way the FCC perceives of different network management tools.

III. The Character of the Proposed Rules

When the FCC wants to make official rules, it must first publish a draft version of these rules and allow for comment in a document called a notice of proposed rulemaking.[24] The public and any interested parties have a defined time to comment on these rules.[25] The FCC has therefore published their Open Internet Notice. The comment period on this Notice is closed, and hundreds of significant comments have been filed on it. It is from these comments that most of the information on public perception and industry response to these rules comes.

The proposed rules within the Notice are divided into two sections: the substantive rules themselves and the definitions section, including the all-important definition of "reasonable network management."[26]

A. The Substantive Rules Themselves

The proposed rules in the Notice are based primarily on the Internet Policy Statement's four principles. In fact, one aim of the rules in the Notice is to formally codify these four principles.[27] These codified principles are essentially a list of activities that providers may not engage in, including blocking content, services, applications, devices, or other competitive options.[28]

Some aspects of the Notice are not very hotly debated. For instance, service providers recognize that noncompetitive behavior or wholesale blocking of legal content is illegitimate and therefore take no issue with the existence of the anti-discrimination rules. The rules also include

[24] 5 U.S.C. §553

[25] *Id.*

[26] Open Internet Notice, at Appendix A.

[27] *Id.* at Par. 88.

[28] *Id.*

a requirement that service providers maintain transparency about their network management practices.[29] Service providers disagree with the scope of this transparency rule, though all seem to agree on its importance. Therefore, since they do not greatly influence the debate within the FCC and among commenters, these provisions are not what I wish to discuss.

Perhaps the most contentious feature of the proposed rules is the actual extent of the nondiscrimination requirement. This requirement says that service providers cannot discriminate "against, or in favor of, any content, application, or service, subject to reasonable network management."[30] Essentially, it disallows treating traffic differently based on the source or type of that traffic. It is the contours of this nondiscrimination requirement that raise the most objections, and the nondiscrimination rule is the one under which the exceptions for "reasonable network management" seems to allow for the most leeway.

B. The Exceptions

The debate over the scope of the non-discrimination requirement is mostly a debate over how broadly the exceptions codified in the Notice apply; ISPs can discriminate up to the point of reasonable network management, so the definition of reasonable network management creates the actual standard being implemented by the nondiscrimination rule. The proposed rules include some general exceptions. These include an allowance for public safety measures (allowing emergency services and homeland security to distriminate), law enforcement (maintaining ISP obligations to prevent illegal behavior or comply with law enforcement officials), and construction with other laws (maintaining any other legal obligations).[31] But the real allowances

[29] *Id.* at Par. 118.
[30] *Id.* at Par. 104.
[31] *Id.* at Appendix A.

for otherwise discriminatory behavior are found within the multi-part definition of "reasonable network management."[32]

This definition is divided into four categories of allowed network management: limiting the transfer of unlawful content, the unlawful transfer of content, transfer of harmful or unwanted content, or any "other reasonable network management practices."[33] These four categories of network management, and the methods and standards by which these categories are applied, constitute the bulk of the debate over how far these proposed rules as codified do, and more importantly should, reach.

IV. The Debate Over Explicit Exceptions

Before embarking on any discussion of each of these exceptions as defined in the Notice, it is important to discuss the background factors that influence the form that these exceptions take. The debate among commentators over these exceptions starts with a discussion of whether the exceptions are necessary. The next question is whether the openness rules could be limited by an alternate case-by-case approach, as opposed to explicitly codified exceptions.

A. The Reason for Requiring Exceptions

The first debate among commenters is over the necessity of these exceptions in the first place. Many commenters agree that an unqualified rule or a strict policy is not a reasonable approach, granting that the proper functioning of connections to the Internet relies on ISPs' ability to manage their networks in some way.[34] Such strict rules would limit service providers' ability to experiment with solutions for congestion and developing new technologies for managing Internet traffic. As discussed below, Internet use is growing, but the amount of

[32] *Id.*

[33] *Id.*

[34] *See, eg.*, *Skype Comments*, GN Docket 09-191, January 14, 2010, at 14 (hereinafter "Skype Comment"); *Electronic Frontier Foundation Comments*, GN Docket 09-191, January 14, 2010, at 12; *Open Internet Coalition Comments*, GN Docket 09-191, January 14, 2010, at 41-42 (hereinafter "OIC Comment").

network connections can't grow forever. A strict non-discrimination standard would cut off ISPs' power to address this problem.

The FCC also seems to maintain that a strict, zero-tolerance approach to discrimination is unsound. The Notice does seem to recognize that the "proposed nondiscrimination and reasonable network management rule bears more resemblance to unqualified prohibitions on discrimination added to Title II in the 1996 Telecommunications Act than it does to the general prohibition on 'unjust or unreasonable discrimination' by common carriers in section 202(a) of the Act."[35] But the Commission still makes it clear that "a bright-line rule against discrimination, subject to reasonable network management and enumerated exceptions, may better fit the unique characteristics of the Internet."[36] This is a function of the FCC's desire to distinguish between kinds of discrimination that are "socially beneficial" and kinds that are not, as opposed to disallowing all types of discrimination.[37] The rules, with their exceptions, should allow flexibility for service providers to experiment with socially beneficial methods of discrimination in the future while maintaining their networks now.[38] So the consensus seems to be that unqualified rules would not meet the goals of the commission and would not allow for innovation in network management practices or for current management necessities, thereby undermining one of the core purposes of the rules.

B. Handling Exceptions Without Explicit Rules

While most agree that unqualified rules would be the wrong approach, some insist that a more case-by-case approach to defining allowable network management is a better approach than

[35] Open Internet Notice at Para. 109.
[36] *Id.*
[37] *Id.* at Par. 103.
[38] *Id.* at Par. 107.

trying to define this kind of allowable discrimination with codified rules.[39] Using adjudication for this process would allow the commission to build a record of different types of network management practices and discrimination.

There are obvious drawbacks to this approach, however. The FCC's history of regulation indicates that it will have difficulty with applying its rules actively, fairly, and in a timely and beneficial manner.[40] Issues of allowed network management would need to be addressed quickly. If the adjudication process takes too long, the decisions on what is allowed might not even be made by the time new problems arise. Also, the FCC has had some successes in the past with allowing market forces to set the standard in regulating an industry.[41]

The best approach seems to be a hybrid of the two: a set of rules outlining the framework in which types of discrimination are tested to provide clarity and predictable standards, then providing ongoing clarification using adjudication proceedings.

V. The Notice's Current Exceptions and Their Shortcomings

In the current version of the proposed rules, there is a general exception for reasonable network management, including four specific subcategories (congestion and quality of service, unlawful or harmful traffic, transfer of unlawful content, and unlawful transfer of content). Each category is a response to a status quo of current approaches to that problem. And each one raises its own questions, including objections to the form in which the exceptions appear in the Notice.

[39] *See, eg.,* OIC Comment at 47 ("OIC urges the Commission to adopt the proposed "Six Principle" framework, which can be enforced on a case-by-case basis as the Commission has done in other contexts."); *Sprint Nextel Comments,* GN Docket 09-191, January 14, 2010, at 25-26 ("In contrast, the "unjust and unreasonable" nondiscrimination standard - coupled with case-by-case adjudication - contains the flexibility needed to distinguish socially desirable discrimination from socially harmful discrimination."); *Telecommunications Industry Association Comments,* GN Docket 09-191, January 14, 2010, at 44 (hereinafter "TIA Comment").

[40] *See generally,* Daniel L. Brenner, *Creating Effective Broadband Network Regulation*, 993 PLI/Pat 69.

[41] *See Id.* (demonstrating that the laissez-faire regulatory approach to communications satellites, broadcast satellites, phone service competition, and broadband successfully allowed for innovation, while the ex-ante regulation of video dialtone, open video systems, and advanced instant messaging did not spur the desired innovation and competition).

A. Congestion and Quality of Service

The first of the four specifically delineated exceptions to the net neutrality principles is one allowing reasonable network management practices in order to prevent congestion and maintain quality of service (QOS).[42] The Notice lists a few examples of possible QOS management techniques. These include charging users more for higher bandwidth use, limiting bandwidth to some high-traffic users during times of extreme congestion, and prioritizing traffic that requires a more steady connection, also called latency-sensitive traffic.[43] I will discuss the current situation of congestion and quality of service management, then I will discuss the problems with the FCC's proposed response.

1. The State of Congestion and QOS Management Now

a. The Problem

The problem of quality of service is not an imagined one. Many ISPs report encounters with congestion, resulting from anything from usage hikes to malicious attacks.[44] Such congestion is not uncommon, and ISPs demand a way to mitigate its effects.

Data suggests that congestion will become a larger problem over time, though how large depends on whose data you look to. Some commenters say that we are headed for an "exaflood," or a time when the demand for bandwidth will grow beyond possible capacity.[45] This means that, no matter how quickly and rabidly we try to build capacity, our demand for badnwidth will always be growing faster.[46] The Free Press, however, in compiling data from a variety of

[42] Open Internet Notice at Appendix A.
[43] *Id.* at 137.
[44] *See, eg.,* TIA Comment at 25 (describing traffic spikes in traffic from individual connections that cannot be predicted); *Joint Declaration of Michael D. Poling and Thomas K. Sawanobori*, GN Docket 09-191, January 14, 2010, at 29 (discussing unpredictable spikes in traffic due to large events); *AT&T Comments*, GN Docket 09-191, January 14, 2010, at 76 (hereinafter "AT&T Comment") (describing worm attacks on their network).
[45] *See, e.g.*, *Alcatel-Lucent Comments*, GN Docket 09-191, January 14, 2010, at 5 (hereinafter "Alcatel Comment").
[46] This is a phenomenon that will be familiar to those that have any experience with the current spectrum crisis.

companies, reports that Internet growth hovers around 40% per year.[47] This sense of measurable, manageable growth is not echoed in comments from actual service providers and infrastructure companies. These commenters point to increasing broadband penetration, coupled with higher bandwidth applications, as evidence of a coming bandwidth shortage, unless action is taken.[48] While the dire predictions of the end of Internet capacity are alarming, the comments predicting the coming bandwidth drought are unfortunately rather short on actual data and prediction to support this point.

One contributing factor in the discussion of congestion is the directions in which this congestion actually moves. The average home user is likely going to notice if, for instance, a video loads drastically slower than usual. But the same user might not notice if a video being uploaded to YouTube is moving slower than usual. This is due to asymmetrical networks.[49] ISPs provide larger amounts of bandwidth for downloading than for uploading, because the average user is going to need more bandwidth for downloading.[50] But back-and-forth programs like Skype and BitTorrent, which require relatively steady upload *and* download speeds, are vulnerable to jitters and slowdowns as a result of low upstream bandwidth allowances.[51] The increasing availability of programs that require this kind of symmetry might contribute to congestion.

Another contributing factor in the ever-increasing congestion of the Internet is a common ISP practice called oversubscription.[52] Since ISPs connect users to the larger network using

[47] *Free Press Comments,* GN Docket 09-191, January 14, 2010, at 37 (hereinafter "Free Press Comment") (provides a roundup of studies that report Internet growth, from companies like AT&T, Comcast, and Verizon, with percentage growth rates ranging from 35 to 50).
[48] *See* Alcatel Comment at 8; *Google Comments*, GN Docket 09-191, January 14, 2010, at 28 (hereinafter "Google Comment").
[49] Lennett at 116.
[50] *Id.*
[51] *Id.*
[52] *Id.* at 104.

smaller loops and pathways that include only their immediate neighborhood, the availability of bandwidth for each node can be quite limited.[53] In fact, the amount of bandwidth necessary to supply the whole neighborhood often exceeds the maximum available bandwidth.[54] This means that a collection of homes, each demanding their share of bandwidth, might not have that share available to them. For instance, imagine a 50 Mbps connection to a neighborhood shared by ten homes with 15 Mbps Internet plans. If all ten access the network at the same time, their demand would be a total of 150 Mbps, or 3 times the available bandwidth. Such "oversubscription" is very common; Comcast, for instance, might operate some loops designed for 500 or more users with only enough bandwidth for 12 users at maximum connection speeds.[55] In short, capacity is not only limited by the availability of bandwidth, but by how ISPs allocate that bandwidth, something that will become important when discussing network management as a solution to congestion.

Either way, all commenters do recognize that Internet applications are quickly becoming more bandwidth intensive. Consumers are increasingly demanding higher quality and more reliable Internet for high definition streaming video, VoIP, and even new technologies still being developed. The net result of increased broadband penetration and higher bandwidth demand applications is the necessity of an increase in Internet capacity.

In all cases, the consensus is that the increased demand for bandwidth, no matter how severe or inflated, makes necessary some kind of response from ISPs to manage this increased traffic, one way or another.

[53] *Id.* at 103.
[54] *Id.* at 104.
[55] *Id.* at 105 (these subscription rates are estimated, as actual oversubscription ratios are carefully-guarded proprietary information).

b. Proposed Solutions

Proposed solutions to this increased Internet congestion are diverse. The most expensive option is the expansion of the Internet backbone, increasing network size, and therefore relieving congestion.[56] The Internet backbone is the actual large array of pipes that relay information from one network to another. ISPs are what connect users to this backbone. The option of expanding this backbone would relieve congestion, but it is only a temporary solution to a problem that will likely continue. It is also among the most expensive solutions. Alcatel-Lucent's model indicates that costs to expand internet pipelines would be more per user than such expansions would allow in revenue per user, i.e. the expense per user to expand backbone wouldn't be made up by the increased subscriptions from new users that the expanded backbone would allow.[57] Such a figure makes backbone expansion economically prohibitive. Plus, high oversubscription rates and constantly increasing bandwidth demands would mean that increased backbone would just get misallocated and overtaxed in the same ways as existing Internet backbone.

A more cost effective way of easing congestion and insuring quality of service is to more efficiently use existing bandwidth capabilities by managing Internet traffic.[58] Methods of accomplishing this task are equally diverse and mostly undefined in the comments. Some of these methods would likely be violative of the proposed rules without an exception for congestion and quality of service.

ISPs have already begun to experiment with some ways to counteract the congestion associated with this increased traffic. There is not a lot of hard data, and some of this is

[56] Alcatel Comment at 8.
[57] Alcatel Lucent, *Analysis of the impact of traffic growth on the evolution of Internet access*, GN Docket 09-191, at 5 ("Alcatel White Paper").
[58] Alcatel Comment at 8.

proprietary.[59] But some comments discuss current successful anti-congestion methods. Cox Communications discusses one specific trial that it conducted to determine the best network management scheme to put to use.[60] The trial was essentially a protocol specific management system, in that it determined which traffic required the most steady stream of bandwidth and then provided that traffic with priority.[61] Its findings were mostly specific and technical, but the lesson that the Cox trial teaches is that, when ISPs are left to experiment with innovative new management techniques, they can conserve bandwidth without passing infrastructure expansion costs onto consumers.[62]

CISCO also described three specific types of management that are currently used, and can be further implemented, to increase bandwidth capacity.[63] The first is IP routing, which involves applying labels to some packets that allow the network to make decisions about those packets without inspecting their actual contents.[64] This obviates the need for invasive packet inspection procedures and could likely be implemented in either protocol agnostic or protocol specific schemes. The second is packet differentiation, which is a method for applying priority levels to transmissions. These priority levels can be assigned based on source, type of traffic, or latency sensitivity.[65] Again, this could be either protocol agnostic or protocol specific. Finally, CISCO describes filtering, which is a simple mechanism of weeding out harmful traffic from harmless traffic. This method is more protocol specific, in that the filters detect the sources or destinations of traffic to make decisions. The appropriate use of successful network management

[59] *Comcast Corporation January 14, 2010 Comments*, GN Docket 09-191, at 46 (hereinafter “Comcast Comment”).
[60] *Cox Communications, Inc. January 14, 2010 Comments*, GN Docket 09-191, at 23 (hereinafter “Cox Comment”).
[61] *Id.*
[62] *Id.*
[63] *CISCO Comments*, GN Docket 09-191, January 14, 2010 at 12 (hereinafter “CISCO Comment”).
[64] *Id.*
[65] *Id.*

approaches like those listed here could increase Internet capacity by 2.5 times.[66] Large ISPs are also upgrading to new standards for data transmission that increase the efficiency of the way data is handled by the actual pipes that make up the Internet.[67] All of these network management practices and network changes could also greatly enhance Internet capacity.

Solutions, then, do exist for the problems of congestion and quality of service. But such solutions might require otherwise discriminatory practices, thus making this exception to the Notice's rules necessary.

2. *Objections and Suggested Changes to This Exception:*

In light of the above information about the quality of service and congestion problem, and considering the innovation necessary to solve this problem, it is clear that this exception is necessary to protect ISPs in its endeavors to experiment in this area. But it's also clear that too broad a version of this exception will allow many harmful types of discrimination to go unchecked.

One proposed solution to this threat of overbreadth is strict scrutiny of any network management asserted to be in the interest of easing congestion or improving quality of service.[68] Such strict scrutiny would require any such management to be temporary and closely tailored to the congestion problem that the ISP is experiencing. The different approaches to reasonableness and scrutiny are discussed below.

Another proposed solution to this possibly overbroad exception is to tie any allowable anti-congestion network management to objective standards or practices established by a third party group of experts. For instance, the FCC could immunize practices that arise out of

[66] *Id.*

[67] Lennett at 107.

[68] This standard will be discussed fully in IV.B.1, *supra.*

collaboration with respected Internet authorities, like the Internet Engineering Task Force.[69] Also, the FCC could presume reasonableness for measures adopted as best practices in compliance with trade alliances.[70] The Open Internet Notice does acknowledge the role of these kinds of organizations in defining reasonableness standards.[71] Such measures could streamline the analysis of management technologies for reasonableness.

While this kind of change would allow companies to proceed with experimental anti-congestion and improved quality of service measures under the protection and predictability of a known standard, they also slow innovation by confining it to the offices of large standards-making bodies, not to entrepreneurs and small-scale innovators. Allowing companies breadth to try new technologies outside of those currently imaginable is the best way to encourage newer and more effective network management tools

Others suggest that the transparency requirement could be used to require companies to disclose its network management methods prior to implementation, possibly for approval by the FCC.[72] Such a cache of approved practices would provide maximum clarity for those companies implementing new management tools. On the other hand, requiring pre-implementation disclosure of any management plans would interfere with the right of service providers to keep proprietary information within its own companies, thus chilling innovations in this area.[73]

The most measured approach seems to be the one adopted by the FCC: an exception allowing reasonable network management and a process by which the record of reasonable network management practices is built through adjudication. Pre-approval tips the balance too

[69] Comcast Comment at 56.
[70] *Id.*
[71] Open Internet Notice at Para. 141.
[72] *Sony Electronics Comments*, GN Docket 09-191, January 14, 2010, at 7 (proposing that practices that limit demand or access to Internet enabled devices be pre-approved by the FCC).
[73] *See* Comcast Comment at 46; *Fiber to the Home Council Comments*, GN Docket 09-191, January 14, 2010, at 28; AT&T Comment at 195.

far towards curtailing innovation, and presumptions of reasonableness tip the balance too far towards immunization of possibly harmful practices. Since network management mechanisms are so varied and change so quickly, the FCC should proceed with caution, with as few hard and fast specific allowances or prohibitions as possible.

Finally, quality of service and congestion could be managed through allowing users to do the discrimination instead. End user discrimination is not technically network management, but it does offer an alternative type of traffic handling that would be permissible under the FCC's proposed rules.[74] A narrowly tailored allowance for discrimination in the pursuit of quality of service would allow such discrimination to be initiated and decided upon by users, instead of within the network itself. Such allowable end-user discrimination would include choosing certain pricing models or specifically tailored service. Such end-user decisions on discrimination would not undermine the aims of the prohibition against discrimination, and indeed they would allow companies to innovate in what services they can offer users, instead of how they manage the network from within, a practice that is more likely to violate the proposed rules and lead to harmful discrimination. This option, like most instances of allowable discrimination, relies heavily on the transparency requirements, discussed in further detail below.

B. Unwanted or Harmful Traffic

The second specific type of excepted reasonable network management is to curtail unwanted or harmful traffic.[75] This includes blocking spam, malicious programs, or content that specific users have requested be blocked, such as adult content.[76]

[74] Skype Comment at 19.
[75] Notice at Appendix A.
[76] *Id.* at Para. 138.

1. *The State of Unwanted or Harmful Traffic Now*

While the FCC specifically allows ISPs to control traffic for the purposes of relieving congestion and maintaining quality of service, the prevention of unwanted or harmful traffic is one such way of relieving congestion.

There are a variety of examples of unwanted traffic. The simplest example is spam email, or unsolicited advertising email. Between 45 and 73% of email traffic is spam.[77] Such a large percentage of traffic, if eliminated, could ease congestion on the Internet. The Open Internet Notice also allows for more specifically user-defined standards for unwanted content, including blocking pornography.[78]

Traffic that is directly harmful to users and the network is just as common. According to the Government Accounting Office, reported cybersecurity incidents increased by 206% in the years between 2006 and 2008.[79] These kinds of threats include DDOS attacks and hacking. DDOS attacks are a carefully orchestrated barrage of requests to a site that is so high in volume that the server hosting the site crashes.[80] Hacking is the unauthorized accessing of protected files. AT&T reports finding 39 million indicators for instances of harmful attacks every month.[81] The Open Internet Notice indicates that the FCC is certainly well aware of these instances of unwanted and harmful traffic, saying that "the general usefulness of the Internet could suffer if spam floods the inboxes of users, if viruses affect their computers, or if network congestion impairs their access to the Internet."[82]

[77] SpamLaws, *SpamStatistics and Facts,* http://www.spamlaws.com/spam-stats.html.
[78] Open Internet Notice at 138.
[79] Statement of Gregory C. Wilshusen, Dir., GAO Info. Sec. Issues, & David A. Powner, Dir., Information Tech. Mgmt., *Cybersecurity: Continued Efforts Are Needed to Protect Information Systems from Evolving Threats*, at 5 (Nov. 19, 2009), http://www.gao.gov/new.items/d10230t.pdf.
[80] *Verizon Comments*, GN Docket 09-191, January 14, 2010, at 7.
[81]AT&T Comment at 184.
[82] Open Internet Notice at Para. 138.

2. Objections and Suggested Changes to This Exception

The most common objection to this exception is that defining "unwanted" traffic is not an easy task. This could be a way for ISPs to block content under the guise of presuming it is unwanted by its users. The easiest solution is to ensure that any blocked "unwanted" content is only so designated by users defining it and by subscribers opting in to have it blocked.

The definition of "harmful" is also not entirely clear. This certainly includes harm to the network itself, but it likely also refers to harm perpetrated against a user's computer. The Open Internet Notice seems to include this kind of harm as well,[83] but ISPs are traditionally not accountable for harmful viruses or malware that are transmitted over their network.

Another common complaint about the "harmful traffic" exception is that it could easily be covered in the rules themselves. The rule on devices specifically indicates that the right to connect a device only extends to content that does not harm the network.[84] This could be generalizable to content, i.e. that users only have the right to transmit or access content that does not harm the network. Any instances not covered by harm would be covered by quality of service and congestion.

The somewhat counterintuitive result is that ISPs would not be allowed to discriminate against traffic that harms users, that they would have to allow users equal access to traffic that harms their computers or that the users do not want. This strange result might be enough to create a necessity for this separate "harmful or unwanted traffic" provision.

The reality of the "unwanted or harmful" exception is that most commenters do not take issue with its probable modes of application. Commenters seem to agree that the blocking of

[83] *Id.*
[84] *Id.* at Appendix A.

unwanted or harmful traffic is generally socially beneficial, with not as many possibilities for misuse by ISPs.

C. Transfer of Unlawful Content

The third defined category of "reasonable network management" is the curtailing of the transfer of unlawful content.[85] The only example given in the Notice is the transfer of child pornography.[86]

1. Current State of Unlawful Internet Content

The transfer of unlawful content is prevalent on the Internet, but such content is usually policed and regulated by law enforcement agencies and other groups besides the FCC. The data on this particular area of network management is particularly scant. Nothing indicates that this kind of content has a particularly detrimental effect on the structure of the Internet, and nothing indicates that the actual volume of child pornography on the Internet has any measurable effects on bandwidth. So the impetus for protecting this kind of management is based entirely on policies external to the Internet.

2. Objections

The good news about an exception that is based mostly on external policy is that the grounds for the exception are fairly established already by traditional law enforcement regimes. There is not much debate as to whether or not this exception should exist; most commenters seem to agree that preventing the transfer of unlawful content is entirely reasonable network management.

The first major obstacle in allowing service providers to discriminate against the transfer of unlawful content is the mechanism by which such providers determine that content is

[85] *Id.*
[86] *Id.* at Para. 139.

unlawful. The Internet is designed in such a way that all content appears similar, as a stream of data between content providers and consumers. Anything different could pose privacy issues for consumers. However, deep-packet inspection technology[87] allows an intermediary (specifically the internet service provider) to investigate the contents of the transmitted packets to see what kinds of data or files are being transferred across its networks, sometimes down to the specific contents.[88] Such inspection is not new, and it's been in the kind of use anticipated by this exception for some time.[89] In fact, while deep packet inspection might sound ominous and like a particularly invasive technology, the Open Internet Notice only briefly mentions the technology, and the Commission takes no stance on it.[90]

The next concern under this exception is one of territoriality. What is unlawful content in some jurisdictions may not be unlawful content in others. This is less problematic from one United States jurisdiction to another, but the Internet is inherently global, making laws vary extremely form one end of an Internet transmission to the other. For example, Level 3 Communications was presented with a situation where the German government asked them to block the transmission to German citizens of national socialist propaganda hosted in the United States.[91] Such content is unlawful in Germany, but it is entirely lawful in the US.[92] In such a case, if a company can't block content on a country-by-country basis, the company must either violate the rules proposed in the Notice or violate German law. Another example is the blocking

[87] Described at II.A., *infra.*

[88] Notice at Para. 57.

[89] *Motorola, Inc. January 14, 2010 Comments,* GN Docket 09-191, at 6.

[90] Notice at Para. 57.

[91] *Level 3 Communications, LLC January 14, 2010 Comments*, GN Docket 09-191, at 8 (hereinafter "Level 3 Comment").

[92] *Id.*

of some foreign websites to comply with international trade laws.[93] The commission should therefore make its territoriality standards for unlawfulness of content clear.

While these particular examples seem easily fixable with minor variations to the formulation of the rules within the Notice, they illustrate the problem with this exception (and the exceptions in general): the Notice cannot anticipate the strange and new ways that the law might develop or conflict with the structure of the Internet in the future.

D. Unlawful Transfer of Content

The fourth delineated category of reasonable network management in the Notice is the prevention of unlawful transfer of content.[94] The Notice's definition of "unlawful transfer of content" is unclear, and the only clearly defined example in the Notice is the unlawful transmission of copyrighted works.[95]

1. The Current State of Unlawful Transfer of Content

The transfer of copyrighted content is possibly the largest problem facing the Internet as it exists today. Pressure to crack down on this kind of transfer is very strong, mostly coming from rights-holders' guilds and corporate lobbies.[96] This pressure might be with good reason; the unlawful trafficking of copyrighted material has measurable detrimental effects on hundreds of

[93] *Id.* at 9.

[94] Open Internet Notice at Appendix A

[95] *Id.* at Para. 139.

[96] *See, eg., Broadcast Music, Inc. Comments*, GN Docket 09-191, January 14, 2010, at 7 ("This is an area that must be protected. If piracy is permitted to flourish because ISPs cannot take action in fear o f being accused o f discrimination, the result will be a crippling o f the music industry."); *MLB Advanced Media Comments*, GN Docket 09-191, January 14, 2010, at 2 ("the Internet content piracy business is growing at an alarming rate, undermining the open Internet and threatening the U.S. economy."); *Motion Picture Association of America, Inc. Comments*, GN Docket 09-191, January 14, 2010, at ii ("MPAA Comment") ("inasmuch as illegal content currently clogs broadband pipes and undermines consumer confidence in the safety and security of the Internet, a reduction in unlawful online activity would help the Commission achieve its goal of widespread broadband deployment and adoption."); *Songwriters Guild of America Comments*, GN Docket 09-191, January 14, 2010, at 2 ("This illegal activity has serious financial consequences for all content owners. However, it has a disproportionately harmful effect on songwriters").

thousands of jobs and billions in government resources.[97] Moreover, reasonable estimates indicate that upwards of 50% of bandwidth usage is dedicated to the unlawful transfer of copyrighted material.[98] This would contribute heavily to the congestion problems that ISPs are dealing with.

The current methods for finding and controlling unlawful transfer of copyrighted material are loose and haphazard. The most common models are designed as partnerships between rights-holders and ISPs, where the rights-holding organizations provide ISPs with technology to seek out and block illegal transfers.[99] In other cases, rights enforcement is barely keeping up with the flow of this kind of content, since the standards to which ISPs must legally conform in blocking this content are as unclear as what kind of blocking methods are permissible under the Notice's rules.

2. *Objections and Suggested Changes to This Exception*

Since the most successful "unlawful transfer of content" management techniques are centered on the relationship between ISPs and rights-holders, commenters are concerned with what scrutiny of management decisions will do to these relationships.[100] Specific standards for what kinds of ISP-rights-holder relationships are permissible would be outdated quickly, but too nebulous a standard would discourage ISPs from cooperating with these rights-holder organizations, thus discouraging innovation in technologies and partnerships to solve the unlawful transfer problem.[101]

[97] *See* Stephen E. Siwek, "The True Cost of Copyright Industry Piracy to the U.S. Economy," Institute for Policy Innovation, IPI Policy Report No. 189 (2007), at i, 11-13.
[98] MPAA Comment at 9 (citing various studies on the prevalence of unlawfully transferred copyrighted material, finding that between 50 and 75% of internet traffic was this kind of content).
[99] MPAA Comment at 10.
[100] *Id.* at 3 ("Content creators, including MPAA studios, cannot tackle this problem alone. They depend on broadband Internet access service providers … to cooperate in combating content theft.")
[101] *Id.* at 14.

Another obvious problem with the enforcement of this exception dovetails with the exception discussed above: how to determine which transferred content is actually copyrighted. Deep packet inspection might offer some information about what kind of content is being transferred, but using this method would be subject to similar privacy concerns as those previously discussed. ISPs could also monitor from which servers downloads arrive, shutting off connections with servers known to be providing illegal content. This would likely lead to errors and unfair shutdowns of legitimate content providers, thus running contrary to the purpose of the Open Internet Notice in the first place.

E. Other Reasonable Network Management

The exception for "any other reasonable network management" is the most difficult to discuss. The Notice provides no real guidance as to what kind of management this could refer to that is not already covered in the previously delineated categories of allowable network management. The FCC does clarify that this general exception is included to cover instances of shortcomings in FCC knowledge and allow for innovation in management.[102]

The ill-defined nature and circularity of this exception is precisely what troubles most commenters.[103] Defining "reasonable network management" this way is circular and best and a license for possibly very discriminatory practices under its banner at worst.[104] "Other reasonable network management" might become the magic words that ISPs can utter to immunize themselves from scrutiny for all sorts of slyly anticompetitive discrimination. The FCC must be

[102] Open Internet Notice at 140.
[103] The definition of "reasonable network management" in the Open Internet Notice includes "other reasonable network management," a facially circular definition. Many commenters pointed this out and expressed their anger at this circularity. *See, eg.,* Free Press Comment at 82 ("The Commission's proposed definition is circular, ambiguous, and incomplete, and without further definition will create loopholes and result in future errors in policymaking.").
[104] Level 3 Comment at 6.

careful not to provide a magic phrase to ISPs that immunizes them from the rules they wish to enact.

The complaints about this last "other reasonable network management" cut both ways, however. Commenters on the other side of the issue insist that the rules should provide for a presumption of reasonableness for any network management undertaken by service providers. The specific mechanisms of this presumption, and the other methods of defining and testing reasonableness, are discussed below.

VI. Suggested Alternate Formulations

We've seen how the FCC has chosen to structure its rules and its exceptions in its Open Internet Notice. But throughout the discussion, it's become clear that alternate rule structures or differing standards for reasonableness could be implemented instead of, or in addition to, these rules and exceptions. I will discuss some of these alternate plans, highlighting an "anticompetitive behavior" approach, a focus on transparency, and some alternate definitions of "reasonable network management."

A. Anticompetitive Behavior

A rule based only on curtailing anticompetitive behavior would streamline the FCCs network management definitions, making any possibly quickly outdated exceptions unnecessary. The existing policy statements, which make the FCC's position clear, could, along with already in place antitrust law, fully control ISP practices.[105] Indeed, in Japan, the Ministry of Internal Affairs and Communications established "reasonable network management" standards that allowed discrimination, as long as the regulator placed in charge watched closely for anticompetitive behavior.[106]

[105] Alcatel Comment at footnote 24.

[106] *Id.* at 27.

The FCC could also enact this anticompetitive standard by conditioning the reasonable network management exception on proper purposes. For instance, the exception could require that any network management be implemented only for public purposes, not for anticompetitive purposes.[107]

The problem with using anticompetition restrictions is that anticompetitive behavior is not obviously so. It would take time and investigation to determine the purpose for which an ISP has undertaken a network management practice. Using anticompetitive indicators as a factor in analyzing the reasonableness of a measure is more efficient and allows for more flexibility in how reasonableness is determined.

B. The Importance of Transparency

Another alternative approach to defining "reasonable network management" would be to greatly enhance the transparency requirements. Transparency, as discussed throughout, is incredibly important, giving ISPs clarity about what network management would be permitted and giving consumers the information they need to make decisions. A robust transparency requirement might also diminish the necessity for specific reasonableness definitions. If an ISP's management practices are all fully disclosed, content providers and consumers would be fully informed as to what the available ISPs are doing to traffic and therefore able to take informed action in choosing their providers. The transparency requirement's effectiveness could be enhanced by standardizing the method by which these disclosures are made, including at what point they must be made and what kinds of information they must include.[108]

Allowing disclosures to aid consumer choice is limited by the actual competitiveness of the market, however. In choosing an ISP, most Americans face either a monopoly or duopoly at

[107] Free Press Comment at 93.
[108] Level 3 Comment at 13.

best (only one or only two choices for providers, respectively).[109] There is no shortage of commentary in support of this vision of limited choice in service providers.[110] The lack of choice would make any public disclosures much less effective, as a choice between only two competing companies that are both engaging in undesired network management is no choice at all, no matter how transparent those companies are.

C. The Standard by which Reasonableness is Judged

As discussed above, some commenters propose replacing the "nondiscrimination" rule with a "no unreasonable discrimination" rule, obviating the need for a clearly defined reasonable network management exception. Such an approach would require careful definition of the term reasonable, and the FCC would have to determine a level of scrutiny used on ISP practices.

The only clearly defined standard for reasonableness that the FCC has used is the one present in the Comcast decision. In that decision, the FCC declared that any practice that is in the service of a crucial end and is strictly tailored to serving that end would be deemed reasonable.[111] The FCC recognizes that the Comcast standard is too high for the purpose of encouraging innovative and effective management practices.[112] Most commenters agree that the standard as formulated for the Comcast decision is too strict to allow for the necessary flexibility in network management techniques.

The problem is that no one suggests actual terms for an alternate reasonableness standard. Commenters, as discussed above, suggest anticompetitive practice language, reasonableness presumptions, safe harbors, and a number of other ways to replace the reasonableness

[109] Google Comment at 19.
[110] Free Press Comment at 50-51 (compiling the Free Press's own findings on ISP choice duopoly, and also compiling findings from the National Telecommunications and Information Administration, the Department of Justice, the Federal Trade Commission, Chairman Genachowski, and the National Broadband Plan team).
[111] Comcast Order at Para. 47.
[112] Open Internet Notice at Para. 137.

determination, but no one suggests a form for the actual test of reasonableness, should the FCC adopt its rules unchanged.

I propose that the test is actually quite obvious. The FCC seeks to balance harms to speech and choice against the benefits of network management. A proper test would do just this: balance the interests on both sides, and decide whose interest is stronger. A list of factors, as was established for fair use in the copyright act,[113] would provide adjudicators with the guidance necessary to provide a reasoned ruling. The factors could draw from any of the above discussed areas, including harm to the user, anticompetitive nature of the policy, benefits to the network, etc. Maybe then reasonableness would come to have some meaning, beyond the conflicted, confused, and circular definition present in the Open Internet Notice.

VII. Conclusion

In discussing the scope of network neutrality regulations, the aim is to scaffold the next phase in the Internet's development. That means attempting to regulate future unknown harms while encouraging any future unknown positive developments. That bitter tension is what makes the exceptions to the general principles of nondiscrimination so important. Because the trouble is that practices, both harmful and positive, may not appear that different until they have already arrived.

It's the FCC's definition of "reasonable network management" that should give us guidance in trying to discern the difference between harmful and positive practices. When the definition of "reasonable network management" is finally debated and settled upon, it's also not clear whether the actual definition and implementation will be any clearer than what the FCC has proposed here. But the standard as presented here is as important as it is nebulous; its boundaries and scope are unclear, but its purpose is definite.

[113] 17 U.S.C. § 107.

That cloudiness is one of the FCC’s rules’ greatest benefits. It’s clear what the FCC wants and how it wishes to set out accomplishing it, but it leaves enough room for newly developed network management tools to flourish without squashing them in their infancy. The carefully tailored uncertainty of this Open Internet Notice’s exceptions might be just what makes them exactly the right tool for protecting the future of the Internet.