Implications of Internet Architecture on Net Neutrality

SCOTT JORDAN
University of California, Irvine

Net neutrality represents the idea that Internet users are entitled to service that does not discriminate on the basis of source, destination, or ownership of Internet traffic. The United States Congress is considering legislation on net neutrality, and debate over the issue has generated intense lobbying. Congressional action will substantially affect the evolution of the Internet and of future Internet research. In this article, we argue that neither the pro nor anti net neutrality positions are consistent with the philosophy of Internet architecture. We develop a net neutrality policy founded on a segmentation of Internet services into infrastructure services and application services, based on the Internet's layered architecture. Our net neutrality policy restricts an Internet service Provider's ability to engage in anticompetitive behavior while simultaneously ensuring that it can use desirable forms of network management. We illustrate the effect of this policy by discussing acceptable and unacceptable uses of network management.

Categories and Subject Descriptors: C.2.3 [**Computer-Communication Networks**]: Network Operations; K.4.1 [**Computers and Society**]: Public Policy Issues; K.5.2 [**Legal Aspects of Computing**]: Governmental Issues

General Terms: Economics, Legal Aspects, Management

ACM Reference Format:
Jordan, S. 2009. Implications of internet architecture on net neutrality. ACM Trans. Internet Technol., 9, 2, Article 5 (May 2009), 28 pages. DOI = 10.1145/1516539.1516540
http://doi.acm.org/10.1145/1516539.1516540.

1. INTRODUCTION

Net neutrality represents the idea that Internet users are entitled to service that does not discriminate on the basis of source, destination, or ownership of Internet traffic. The idea is rooted in the manner in which the Internet has historically operated, in which all traffic receives best-effort service, with

Portions of this work also appear in Jordan [2007]. Jordan [2007] is directed at the policy community and focuses on development of statue language; this article is directed at the networking community and focuses on the relationship of policy to Internet architecture.

Author's address: S. Jordan, Department of Computer Science, University of California, Irvine, CA 92697-3435; email: sjordan@uci.edu.

DOI 10.1145/1516539.1516540 http://doi.acm.org/10.1145/1516539.1516540

limited differentiation based on the application and without any performance guarantees. There is great disagreement, however, about the future implications of this relatively simple idea as the Internet progresses and as the economic communications landscape changes.

Proponents of net neutrality (generally, application providers and consumer groups) argue that without a prohibition on discrimination, Internet Service Providers[1](ISPs) may charge application providers discriminatory prices for access to dedicated bandwidth or for Quality of Service (QoS), or may outright block access to certain applications or Web sites, and that such activity will inhibit development of new Internet applications. Most proponents believe that ISPs should not be allowed to charge for priority access to the Internet portion of their service offerings.

Opponents of net neutrality (generally, ISPs) argue that there is no current problem, that competition is sufficient to ensure that commercially negotiated arrangements for bandwidth or QoS will not negatively impact consumers, and that any regulation will discourage investment in network infrastructure.

During 2005 and 2006, Congress worked on the most substantial rewrite of the nation's communications law in a decade. The issue initially motivating this task in both houses of Congress was video franchising, which attempted to streamline the laws regulating carriers offering broadcast video services, with the goal of enabling faster deployment of video products by telephone companies. The House bill included not only video franchising, but also a weak version of net neutrality, and a few other communications issues. The Senate bill included an even wider range of communications issues. Among all of these communication issues, net neutrality was easily the most contentious. Lack of resolution of net neutrality was generally given credit for derailing the communications bill.

Congressional action, or inaction, on net neutrality will greatly impact Internet architecture. If Congress adopts one of the more extreme versions of net neutrality, then the use of QoS could be prohibited. If Congress adopts a more moderate version of net neutrality, then QoS might be allowed but charging for QoS might be prohibited. If Congress does not act on net neutrality, then ISPs may block access to QoS mechanisms to competing application providers. Any of these actions would substantially affect the evolution of the Internet and of future Internet research.

The topic is timely for a combination of technological and policy reasons. On the technology side, convergence between telephone networks, video networks, cellular networks, and the Internet is resulting in all of these networks becoming capable of efficiently supporting a combination of voice, video, and data services. Competition between the carriers supporting the infrastructure of these various networks and application providers offering a wide range of services will thus intensify.

[1]In this article, we use the term ISP to represent Internet access providers, namely carriers who provide last-mile access to residential and business customers. ISPs can provide two types of service: access to Internet infrastructure, such as DSL, and application services, such as email. This distinction is made later in the article when we introduce definitions.

On the policy side, U.S. federal communications law was separately developed for telephone networks, cable video networks, and cellular networks. In the absence of any explicit statutes regarding Internet access or services, the Federal Communications Commission (FCC) recently declared that Internet access is not subject to the same common carrier regulation that addresses telephone networks, which had the effect of removing from Internet access several prohibitions on discrimination that were included in common carrier regulation. This lifting of discrimination constraints triggered the push for net neutrality.

Net neutrality was a heavily lobbied issue in 2006. Estimates of the amount spent on net neutrality advertising and lobbying range up to $100 million [Bloomberg News 2006]. The issue also attracted wide attention on the Internet and in the media. The term *net neutrality* appears on more than two million Web pages.[2] Net neutrality was the topic of editorials, commentaries, and news articles in dozens of national and local newspapers and magazines, and garnered some attention on television, ranging from *Moyers On America* to *The Daily Show*.

The academic literature on net neutrality has been split on the issue. Openists [Bar et al. 2000; Lemley and Lessig 2001] believe that the Internet is best served by maintaining a *dumb network* that does not differentiate among different types of traffic. They support a policy based on *open access*, in which Internet infrastructure and applications can not be bundled using either technical or business mechanisms. They argue that vertical integration harms consumers, that most innovation comes from application providers, and that open access will maximize social welfare. Openists support a strong version of network neutrality.

Deregulationists[3] [Owen and Rosston 2003; Yoo 2005] believe that ISPs are in the best position to determine the most beneficial evolution of the Internet. They expect that the Internet will become a *smart network* that uses traffic differentiation in order to increase the variety of the types of Internet access and applications to which customers subscribe. They support a policy that deregulates the Internet in a manner that allows ISPs to vertically integrate, bundle services, and use traffic differentiation as they see fit. They argue regulation will hinder investment by ISPs and that in the absence of regulation ISPs will only vertically integrate in ways that benefit consumers. Deregulationists oppose network neutrality.

There is a limited amount of academic literature that attempts to strike middle ground. Nondiscriminationists [Peha 2006; Wu 2004] believe that there are good and bad uses of traffic differentiation. They support a policy that allows vertical integration and traffic differentiation, but restricts their use to ensure that ISPs do not discriminate in a manner that allows them to increase prices because of the limited number of carriers offering service in a single market. They argue that such a balanced approach will allow development of a smart network in a manner that does not restrict development of applications. They

[2]Google search on 7/25/2007.

[3]The terms *openist* and *deregulationist* were introduced in Wu [2004].

would support a limited version of net neutrality that falls short of open access, but have not yet fully formulated such a policy.

Net neutrality and open access arguments are often related to consequences of a layered Internet architecture. Many of the Openist arguments are based on an interpretation of the end-to-end principle [Saltzer et al. 1984], which suggests that network functionality should be implemented in OSI layers 1 through 3, and hence in each router, only if it cannot be implemented effectively in higher layers. Openists believe that the end-to-end principle is responsible for the tremendous amount of innovation at the application layer, while Deregulationists believe there is a trend away from the end-to-end principle. One of the original authors of the end-to-end principle suggests that a new theory addressing network core functionality should be created to co-exist with the end-to-end principle [Blumenthal and Clark 2001], and that the end-to-end principle should be generalized with a new form of modularity that can gracefully accommodate tussle along various competing players [Clark et al. 2005].

Finally, there is an academic literature that addresses net neutrality and related topics in the context of broader revisions of communications policy [Bar and Sandvig 2000; Werbach 2002; Solum and Chung 2003], economic analysis [Speta 2000; Farrell and Weiser 2003; van Schewick 2007], and market analysis [Lehr et al. 2006].

In this article, we argue that net neutrality is, at its core, an attempt to address problems posed by a fragmented communications policy unable to deal with network convergence. We adopt an approach jointly grounded in Internet technology and communications policy. We argue that the evolving layered Internet architecture supports the model of a smart Internet that allows only certain types of discrimination. We accept the premise that vertical integration between infrastructure and applications poses potential threats to a level playing field. We suggest that an important tool in solving such problems is a proper delineation of Internet infrastructure and Internet applications. We illustrate how such a delineation can be used to restrict an ISPs ability to engage in anticompetitive behavior through discrimination, while simultaneously ensuring that ISPs can use desirable forms of network management. We further illustrate how this use of layering can appropriately limit the scope of regulation. Finally, we suggest that net neutrality can be addressed in a manner consistent with current Federal communications law.

Our analysis is heavily based on communications technology and historical communications policy. Our approach belongs in the Nondiscriminationist camp. Our solution essentially mandates the use of open interfaces, but does not mandate full open access.[4] Although we believe that there are many other interesting ways to address net neutrality, we hope that our convergence-inspired layered approach will illustrate a solution that is well grounded in both technology and policy, and that may even be somewhat agreeable to many of the opposing forces on this issue.

The article proceeds as follows. The first three sections are devoted to background information, presented in what we hope is a nonpartisan fashion. In

[4]We do not take a position here on open access, as this is outside the scope of the article.

Section 2, we briefly review relevant aspects of U.S. federal communications law. In Section 3, we briefly review evolving carriers' network architectures and the likely uses of discrimination. In Section 4, we review the relevant academic literature, present pro and anti net neutrality lobby arguments, and characterize the various congressional approaches to net neutrality. In the remaining three subsections, we develop our proposed net neutrality policy. In Section 5, we consider the challenges to telecom policy posed by convergence and how application of Internet layering can serve as a foundation for a new net neutrality policy. In Section 6, we combine elements of existing telecommunication policy with a proper delineation of Internet infrastructure and Internet applications to create a new net neutrality policy that restricts an ISP's ability to discriminate while allowing reasonable network management. Finally, in Section 7, we illustrate how this form of net neutrality can delineate acceptable and unacceptable uses of network management.

2. OVERVIEW OF RELEVANT COMMUNICATIONS LAW

2.1 History

U.S. federal communications law was separately developed for telephone networks, cable video networks, and cellular networks. Only recently has communications policy started to address the Internet.

Regulation of telephone carriers was formalized with the Communications Act of 1934, which codified a number of practices that were in place at the time. The Communications Act has been seen a number of updates over the years. However, before 1996, telephone common carrier policy was based on the assumption that local phone service was offered by a set of local monopoly carriers. The Telecommunications Act of 1996 [U.S. Congress 1996] removed barriers between local and long-distance telephone service, and allowed competition between local and long-distance carriers. However, communications law regarding common carriers has always principally focused only on voice service, and has not effectively addressed video or data service.

Regulation of cable carriers was largely laid down in the Cable Communications Policy Act of 1984 [U.S. Congress 1984] and the Cable Television Consumer Protection Act of 1992 [U.S. Congress 1992]. The policy was based on the likelihood of a set of local monopoly carriers; there was little video competition before the introduction of direct broadcast satellite video service. Communications law regarding cable carriers has always principally focused only on video service, and has not effectively addressed voice or data service.

In contrast, there is no similar body of federal communications law that uniquely address Internet access and services. The Federal Communications Commission (FCC) is tasked with creating regulations that interpret and implement laws passed by Congress relating to communications. In the absence of any explicit statutes regarding Internet access or services, it has been up to the FCC to determine the applicability of the Communications Act to the Internet. During the last thirty years, this issue has come up repeatedly. In 1970, the FCC wrestled with the convergence between communications and computation

in the consideration of the applicability of the Communications Act to data processing services. They implemented regulations [FCC 1971] that delineated between data processing services and communication services on the basis of the device, that is, whether a computer was used for communication or for the processing of information. The communications market continued to be subject to the common carrier provisions in the Communications Act, while the data processing services market was viewed as competitive. Telephone carriers who wanted to enter the data processing market were mandated to establish separate subsidiaries for this purpose.

By 1980, convergence between communication and computation had proceeded to the point where delineation on the basis of device had become unworkable. In response, the FCC replaced this framework with regulations [FCC 1980] that delineated on the basis of the service. *Basic services* were defined as pure transmission capabilities and the data processing capabilities required to support them, whereas *enhanced services* were defined as processing that changed the format of the information or provided additional information. In the Telecommunications Act of 1996, the terms changed from *basic services* to *telecommunication services* and from *enhanced services* to *information services*, but the delineation was similar. Telephone carriers who wanted to offer enhanced or information services were mandated to provide a certain degree of open access to competing providers of these services.

In a sense, therefore, the portion of the Internet that resembles (or is identical to) telephone networks has often been considered to be within the domain of telephone common carrier law. However, the portion of the Internet that resembles applications has been considered to be outside this domain. The delineation between these two portions remains a principal challenge to communications policymakers.

More broadly, the convergence between the networks supporting voice, video, and data services is a fundamental challenge to communications policy. Many of the topics in the 2006 communications bills stem from this network convergence, including video franchising and net neutrality. Despite this trend, Congress has yet to fundamentally consider creation of a communications policy that recognizes network convergence.

This fragmented approach to communications policy is a principal cause for net neutrality's timeliness. Recently, the FCC has declared that Internet access is an information service, and is thus not subject to common carrier regulation [FCC 2005a; 2002]. This decision removed from Internet access several prohibitions on discrimination included in common carrier regulation. This lifting of discrimination constraints triggered the push for net neutrality.

2.2 The Communications Act

The relevant Federal law is contained in the Communications Act of 1934, as updated by various Acts since its original introduction. Title I states general provisions applicable to all communications activities, and is perhaps the only title currently interpreted as applying to the Internet. Title II states

much more detailed provisions applicable to common carriers, namely telephone companies.

In addition, there are relevant ideas that are embedded in other titles; in particular Title VI states provisions applicable to cable communications, which includes some limits on vertical integration between carriers and application providers. We briefly summarize the relevant portions of each.

Title I is largely concerned with establishing the Federal Communications Commission (FCC) and its operation. The FCC has the authority to create regulations that interpret and implement laws passed by Congress relating to communications. Title I, however, has also been interpreted by the courts to give the FCC limited ancillary authority to create regulations on types of communications not covered by the other titles. As a result, whenever communications are not deemed to fall under Title II (telephone carriers), Title III (wireless carriers), or Title VI (cable carriers), the only authority the FCC has stems from Title I. In addition, Title I includes a provision which instructs the FCC to forbear from applying regulations unless they are in the public interest and required to ensure just and reasonable practices.

Title II regulates telephone carriers. Section 201 requires common carriers to offer service upon request and to interconnect with other carriers, and mandates that charges and practices be just and reasonable. Section 202 bars common carriers from unjust or unreasonable discrimination and from giving undue or unreasonable preference. Section 203 requires common carriers to post public lists of their charges. Section 205 gives the FCC authority to proscribe just and reasonable charges. Sections 251 and 252 regulate interconnection, and require incumbent local exchange carriers to interconnect with other carriers on rates, terms, and conditions that are just, reasonable, and nondiscriminatory and in a manner that is at least equal in quality to that provided to itself. In addition, there are many other provisions in Title II that are not pertinent to the discussion here.

Title VI regulates cable carriers. While most of the title is not pertinent to the discussion here, section 628 discusses the relationship between cable carriers and content providers. In particular, it addresses cases of vertical integration, in which a cable carrier has an attributable interest in a content provider. The section prohibits unfair methods of competition or unfair acts or practices, and certain types of discrimination in prices, terms, and conditions of sale.

3. DEPLOYMENT OF QOS

In this section, we briefly review evolving carriers' network architectures and the likely uses of discrimination. Most carrier networks are evolving to an architecture similar to that shown in Figure 1. Fiber, or a combination of fiber and copper, is deployed along the path from the customer premises to the carrier's IP network, from **A** to **D**. Consumer premises equipment at **A** often multiplex signals from separate phone, Ethernet, and coax cable tv connections in the customer premises. Remote electronics at **B** may multiplex lines from multiple customers. If any content (e.g., video) is provided in a circuit-switched manner, an additional multiplexor at **C** may connect to a dedicated circuit-switched

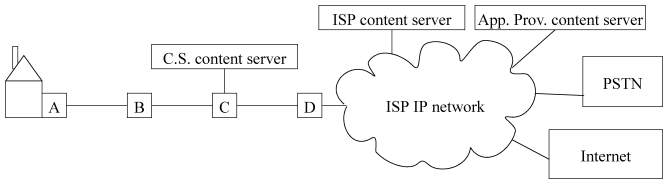

Fig. 1. Typical ISP network architecture.

content server. All packet-switched traffic typically progresses onto the ISP's IP network at **D**. Within the ISP's IP network, there are often content servers operated by the ISP (e.g., video over IP), content servers operated by other application providers that are business customers of the ISP, a gateway to the Public Switched Telephone Network (PSTN), and a gateway to the rest of the Internet.

Traffic management is typically used at several points. First, it is common to place a limit on the bandwidth of a broadband user's access link that is lower than the speed the link could accomplish; this is often implemented at **D**. Such limits, referred to as *tiering*, are often used to differentiate service offerings by a selection of access rates, and to charge different prices for these different service offerings. Second, some form of reservations (e.g., IntServ or dedicated bandwidth) or priorities (e.g., diffServ) are often used to provide higher performance for selected traffic. Third, although currently rare in commercial networks, many universities use traffic shaping to limit file sharing traffic to a small proportion of the organization's network capacity in order to protect the performance of other traffic viewed as more important to the organization's mission.

The decision of which traffic to enhance or degrade can be based on: (1) the type of application, (2) the source, (3) the destination, (4) consumer payment, or (5) application provider payment. For example, if priority service is given to all voice and video traffic, or if traffic shaping is applied to file sharing traffic, then application type is the basis for the decision. If capacity is reserved for all voice and video traffic to/from consumers for a fee, then consumer payment is the basis. If capacity is reserved for all traffic to/from application providers for a fee, then application provider payment is the basis.

Currently, ISPs often give priority service to VoIP and video over IP traffic preferred by the ISP. Other VoIP and video over IP traffic generally receives best-effort service. For instance, customers who subscribe to the ISP's VoIP service may receive preferential service from **A** to the PSTN, while customers who subscribe to a competitor's VoIP service may receive best-effort service from **A** to the Internet.

Most of the large carriers have announced plans to deploy QoS mechanisms, often in coordination with deployment of fiber and video service. In the near term, priorities or reservations will be used for selected traffic that both originates and terminates within the carriers network, or that transits onto the PSTN. In this case, a carrier's own network management can provide acceptable performance to limited traffic classes using QoS techniques. In the long term, we expect that QoS will also be applied to selected traffic that originates

or terminates (but not both) within the ISP's network and transits onto another carrier's portion of the Internet. In this case, acceptable performance may only be provided through cooperation with other carriers offering QoS.

We expect the initial use of QoS will be to support a carrier's own VoIP and video services. It is unclear whether carriers will offer QoS to competitors' applications, on the basis of application type, consumer payment, and/or application provider payment.

We believe these developments are at the core of the issue of net neutrality. Specifically, we are concerned about on the basis on which differentiated service may be used to support real-time applications and applications that require relatively large amounts of bandwidth. In contrast, we are *not* concerned about applications that are less interactive and require little bandwidth, such as most Web browsing; we do not foresee best-effort Internet transport degrading to the point where such applications do not receive acceptable performance.

In addition, as networks converge, it is becoming less clear what portion of the integrated network is considered to be the public Internet and what portion is considered to be a private network. Specifically, in the case of traffic that does not transit onto another carrier's portion of the Internet, the distinction between Internet traffic, VoIP traffic, and video traffic may be more a matter of user perception than of technical distinction. This causes a fundamental problem with separate regulation of Internet traffic, telephone traffic, and video traffic.

4. ARGUMENTS FOR AND AGAINST NET NEUTRALITY

In this section, we review the arguments for and against net neutrality. We start with the academic literature, proceed to the arguments presented by lobbyists on various sides of the issue, and then analyze how the various Congressional bills relate to net neutrality.

4.1 Academic Literature

The academic literature on net neutrality is rooted in an earlier literature on open access, which in turn partially relies on an older literature on Internet architecture.

A design paradigm related to Internet layering is called the *end-to-end principle* [Saltzer et al. 1984]. The principle suggests that network functionality should be implemented in OSI layers 1 through 3, and hence in each router, only if it cannot be implemented effectively in higher layers. This principle has been followed in much (but not all) of Internet design to-date. As a consequence, although there are a large number of protocols at the Internet LAN-link layer to accommodate transmission over various media (e.g., Ethernet, WiFi, DSL, cable modem, CDMA), and there are a very large number of protocols at the Internet application layer to support various applications (e.g., pop, smtp, http, p2p, IM), there are relatively few protocols at the intervening Internet network and transport layers (principally IP, TCP, and UDP).

The Openists believe that the end-to-end principle is responsible for the tremendous amount of innovation at the application layer. They see a danger in monopoly control over essential network infrastructure, and argue that

communications policy should mandate open access to the network infrastructure for application providers in order to assure continued innovation. The debate over open Internet access was precipitated by FCC consideration of whether to allow ISPs.[5] to bundle Internet access with applications such as email and Web hosting.[6] Bar et al. [2000] suggest that policy for Internet access should be developed on the model of open access mandated for the telephone network in the *Telecommunications Act of 1996*, by mandating availability on cost-effective terms of key network components to service providers who do not own Internet access infrastructure. Lemley and Lessig [2001] support a policy in which Internet infrastructure and applications can not be bundled using either technical or business mechanisms. They argue that an open network encourages greater innovation than a closed network and that vertical integration harms consumers. They also demonstrate how open access leads to net neutrality, by using the end-to-end principle as an argument for maintaining a dumb network that does not differentiate among different types of Internet traffic. Herman [2007] adds first amendment arguments to support net neutrality.

The Deregulationists believe that there is a trend away from the end-to-end principle. They believe that ISPs are in the best position to determine the most beneficial evolution of the Internet, and that this will likely entail a smart network that uses traffic discrimination in order to accomplish product differentiation. Speta [2000] argues against open Internet access. He postulates that vertical integration will not threaten competition among content providers, on the basis of an economic theory of network externalities that states that monopolists will only vertically integrate in ways that maximize consumer welfare. He argues that vertical integration may be necessary to justify investment in network infrastructure, and that ISPs will be motivated to allow competition to their own applications since this will add value to the broadband connection they sell. Owen and Rosston [2003] view net neutrality as a property rights issue. They claim that it is unknown whether welfare is maximized by assignment of property rights to ISPs or to application providers, and therefore that net neutrality regulation should not be imposed. Yoo [2005] extends the anti open Internet access argument put forward by Speta to net neutrality. He believes that the theory of network externalities undermines the need for mandated net neutrality, and furthermore that the lack of concentration of broadband access when measured at a national level predicts that ISPs will not have market power in the application space. Yoo also suggests that competition can be more effectively fostered between different physical networks that are matched to different types of applications. Yoo [2006] furthermore suggests that exclusive arrangements between ISPs and application providers can be justified on the basis of the efficiency of traffic shaping as a proxy for

[5]There is an unfortunate inconsistency in the use of the term *ISP*. In this article, we use the term ISP to mean carriers who provide last-mile access to residential and business customers. However, in the open Internet access debate, the term ISP was often used to mean providers of services such as email and Web hosting who do not own Internet access infrastructure.

[6]The FCC later allowed such bundling. We discuss this in more detail next.

usage-based pricing. Finally, McTaggart [2006] argues that the Internet is not neutral, using examples of Internet portals, caching, routing, and QoS.

Some economists counter that application of the theory of network externalities to open access or net neutrality is not so straightforward. Farrell and Weiser [2003] discuss potential problems with application of the theory to open access, and suggest that uncertainty dictates further study of the effect of various types of Internet regulation in other countries. van Schewick [2007] argues that a correct application of the theory demonstrates that ISPs do have an incentive to discriminate against competing application providers, and that the economic benefits of net neutrality outweigh the costs.

Nondiscriminationists see more complexity in the net neutrality issue than the Openists or Deregulationists. Wu and Negi [2003] supports net neutrality but opposes open access. He is not opposed to vertical integration, since he believes that QoS may require it. He argues that open access is inferior to direct rules on discrimination that ensure that bad types of discrimination are not used by ISPs. He suggests that such a policy might be constructed by allowing discrimination at the Internet LAN-link layer but not at the Internet network layer. As a result, he would allow bandwidth tiers and some forms of QoS, but discrimination on the basis of the type of application or on the source or destination would be prohibited. Peha [2006] surveys various technical methods of discrimination, and argues that some uses are beneficial and some are not. He discusses dangers from harmful uses of discrimination, focusing on oligopoly rents in upstream markets. He suggests that good policy might be based on lists of allowed and banned discrimination.

4.2 Lobbyist Arguments

The pro net neutrality lobby falls roughly into two categories: application providers (led by Google, Amazon, Yahoo, eBay, and Microsoft) and consumer groups (led by moveon.org, Free Press, Consumers Union, and Common Cause).

The application providers are roughly in the Nondiscriminationist camp. They are primarily concerned that ISPs will charge them discriminatory prices for access to reserved bandwidth. When they examine carrier plans for deployment of fiber, they worry that carriers will sell reserved bandwidth to selected application providers rather than adding capacity to their Internet access offerings for shared use by all Internet traffic. The application providers are large business customers of ISPs, and they consider it acceptable for an ISP to charge their *own customers* for access to the Internet based on bandwidth. However, the application providers do not want to be forced to make contractual arrangements with the ISPs of each of the *application provider's* customers. They argue that ISPs have always charged only their *own* customers, not customers of *peering ISPs*. They are amenable to ISPs charging their own customers different prices for different levels of bandwidth. Application providers are also concerned that ISPs will choose to strike exclusive deals for access to reserved bandwidth, tilting the playing field among application providers.

Application providers remain unconcerned that ISPs might refuse to offer access to reserved bandwidth to application providers who compete with the ISPs

own service offerings. Such behavior is likely for some ISPs who deploy their own video service, particularly if the available bandwidth is not high enough to support a large number of video streams. Application providers, however, typically believe that net neutrality need only apply to Internet access, and do not attempt to apply to it what they view as private network services.[7]

Some application providers are concerned that ISPs will not offer access to QoS mechanisms, or at least not without charging application providers. They worry that ISPs will use QoS to support their own real-time applications but refuse access to QoS to competing application providers' traffic. Alternately, they worry that ISPs will attempt to charge application providers, rather than the ISP's own customers, for QoS.

Application providers are concerned that lack of access to reserved bandwidth or QoS will hinder the development of new applications. They argue that small application providers will not be able to compete effectively.

The application providers want a form of net neutrality that prohibits discrimination on the basis of source, destination, or ownership of Internet traffic. Discrimination on the basis of application type or consumer payment, for example, enhanced performance for all voice and video traffic, would be allowed. This prohibition would only apply to what is thought of as Internet traffic, as opposed to private network services. Application providers generally believe that such a prohibition would not allow ISPs to charge application providers for reserved bandwidth or QoS.

The consumer groups in support of net neutrality are generally Openists. They share many of the same concerns as the application providers. In addition, they are concerned that ISPs may block access to certain applications or Web sites. Some groups are worried that an ISP may block access to Web sites based on the content of the Web site, for example, political beliefs, and interpret net neutrality as an issue of free speech. They are also worried about new charges that may increase the cost of Internet access.

The anti net neutrality lobby consists mostly of ISPs (led by AT&T, Verizon, and Comcast, but including many other wireline and wireless ISPs). They fall squarely into the Deregulationist camp. They argue that there is no current substantial problem, and hence no need for regulation. Furthermore, they argue that if substantial problems arise, then the FCC has the expertise and authority to effectively take action. The ISPs believe that net neutrality would reduce ISPs' incentives to invest in their network infrastructures by reducing the return they may earn on these investments.

The ISPs further believe that they should be allowed to choose whether to make available reserved bandwidth or QoS, to charge consumers for these services, and to make individual contractual arrangements with application providers. They argue that refusal to provide access to reserved bandwidth or QoS to application providers would not hinder the application provider's ability to deploy new services. Most ISPs expect to offer access to reserved bandwidth or QoS to application providers for a fee, but want the flexibility to do this

[7]However, as mentioned earlier, we believe it will be increasingly difficult to determine which services are Internet services and which are private services.

through individual contractual arrangements. Such ISPs are typically focused in the short term on arrangements in which the application provider becomes a ISP business customer, rather than providing access through peering ISPs. Some ISPs also claim that net neutrality impinges on the ISP's ability to perform network management tasks, including traffic shaping for p2p traffic.

Finally, many ISPs fear that net neutrality would reestablish many other provisions of common carrier regulation. In particular, they worry that net neutrality's occasional use of the term *nondiscriminatory* will lead to imposition of other common carrier rules such as unbundling and open access.

4.3 Bills

In 2005 through 2006, Congress considered several approaches to net neutrality. The House bill, the *Communications Opportunity, Promotion, and Enhancement Act of 2006* [U.S. Congress 2006b], and the Senate Commerce Committee bill, the *Advanced Telecommunications and Opportunities Reform Act* [U.S. Congress 2006a], take similar approaches to net neutrality and essentially present the anti net neutrality position. They give the FCC authority to enforce the FCC's four principles (discussed next), but do not allow the FCC to issue additional rulemaking to further define or extend these principles. Blocking of Web pages or applications is prohibited, but no further restrictions are placed on ISPs. ISPs remain free to offer or deny access to QoS mechanisms. Either degraded service or enhanced service can be based on application type, source, destination, consumer payment, and/or provider payment. Payments can be on the basis of commercial arrangements, meaning that there is no requirement that they be consistent or equitable between different application providers.

The *Internet Non-Discrimination Act of 2006* [U.S. Congress 2006e] (sponsored by Sen. Wyden) essentially presents the most extreme pro net neutrality position among the bills. It prohibits blocking of Web pages and applications. The bill includes a broad prohibition on discriminatory behavior; our interpretation is that any use of QoS is prohibited, since the bill prohibits discrimination in data transport, which presumably even includes differentiation based on application type. The bill explicitly prohibits charging application providers that are not customers of the ISP, although this seems redundant given the prohibition on discrimination. The bill prohibits degradation of traffic; our interpretation is that traffic shaping would fall into this prohibition. (These provisions are summarized in Table I.)

Three other bills, namely, the *Net neutrality Act of 2006* [U.S. Congress 2006d] (sponsored by Rep. Markey, defeated as an amendment to the House bill), the *Internet Freedom and Nondiscrimination Act of 2006* [U.S. Congress 2006c] (sponsored by Rep. Sensenbrenner), and the *Internet Freedom Preservation Act* [U.S. Congress 2006f] (sponsored by Sen. Snowe and Sen. Dorgan), take similar approaches to each other and present refined pro net neutrality positions. They each prohibit blocking of Web pages and applications, as do the other bills. They each also include a broad prohibition on discriminatory behavior; our interpretation is that use of QoS is allowed if based on the type of application, but prohibited if based on the source or destination of the traffic, for example,

Table I. Comparison of the Net Neutrality Provisions in Various Congressional Bills

	House & Senate bills	Wyden bill	Markey & Sensenbrenner & Snowe-Dorgan bills	Snowe-Dorgan amendment
Prohibit blocking	*Y*	*Y*	*Y*	*Y*
Degradation:				
allow traffic shaping	*Y*	*N*	*Y*	*Y*
prohibit if not based on type	*N*	*Y*	*Y*	*Y*
QOS:				
allow based on type	*Y*	*N*	*Y*	*Y*
prohibit exclusive use	*N*	*n/a*	*Y*	*Y*
QOS fees:				
allow consumer payment	*Y*	*n/a*	*N*	*Y*
allow provider payment	*Y*	*n/a*	*N*	?
prohibit discriminatory fees	*N*	*n/a*	*n/a*	?
prohibit unreasonable fees	*N*	*n/a*	*n/a*	?

based on the application provider. The bills explicitly prohibit charging for QoS, either in the form of customer or application provider payment. The bills also prohibit degradation of traffic on the basis of source or destination, but allow network management based on application type, such as traffic shaping.

Finally, Sen. Snowe and Sen. Dorgan proposed a streamlined version of their bill as an amendment to the Senate bill[8]; this amendment was defeated in the Committee markup on a tie vote. The amendment states a broad prohibition on behavior that discriminates on the basis of source, destination, or ownership of Internet traffic. This has the effect of prohibiting blocking of Web pages and applications, allowing QoS or degradation on the basis of the type of application, and prohibiting QoS or degradation on the basis of the source or destination of the traffic. The amendment is silent on whether an ISP can charge for QoS. Our interpretation is that it allows consumer payment, but it is less clear whether it allows provider payment and, if so, whether such payment would be mandated to be reasonable and nondiscriminatory.

5. A LAYERED APPROACH TO NETWORK NEUTRALITY

In the remainder of the article, we turn to development of our proposed net neutrality policy. In this section, we lay the foundation for a new net neutrality policy. We start by identifying outdated definitions that have caused the failure of current communications law and regulation to adequately address the issue of net neutrality. We then propose new definitions of Internet infrastructure and applications to remedy this failure. In the following section, these new definitions will serve as the foundation for a new net neutrality policy.

5.1 Challenges from Network Convergence

U.S. federal communications law was separately developed for telephone networks, cable video networks, and cellular networks. As others have noted,

[8]The key part read: "END USER ACCESS.–In addition to the principles contained in the broadband policy statement, end users shall be entitled to service from each broadband Internet access provider that does not discriminate in the carriage and treatment of Internet traffic based on the source, destination or ownership of such traffic."

convergence between the networks supporting voice, video, and data services is therefore a fundamental challenge to such *vertical regulation*. Common carrier regulation has focused on voice service, and does not effectively address video or data service. Similarly, cable services regulation has focused on video service, and does not effectively address voice or data service. Internet services are not effectively addressed anywhere in the Communications Act. Furthermore, when voice, video, and data services are offered over a integrated converged network infrastructure, it is unclear how to apply conflicting regulations from different titles.

As discussed previously, this vertical regulation approach started breaking down decades ago, as the telephone network progressed from a set of local monopolies to a somewhat competitive marketplace. The distinction between telecommunication services and information services has been an effective tool for communications policy. However, network convergence has now progressed to the point where this distinction is no longer sufficient. Indeed, many of the topics in the 2006 communications bills stem from convergence, including video franchising and net neutrality. Despite this trend, Congress has yet to consider creation of a communications policy that fundamentally recognizes network convergence.

In the absence of explicit statutes regarding Internet access or services, the FCC recently addressed the issue of whether Internet access is a telecommunications service or an information service. The decision by the FCC to classify Internet access (including DSL and cable modem service) solely as information services exempted these services from regulation under Title II of the Act, and left them only subject to the FCC's Title I ancillary authority [FCC 2005a, 2002]. This decision means that Internet access is not subject to the requirement that charges and practices be just and reasonable (section 201), the prohibition on unjust and unreasonable discrimination (section 202), or the requirement that interconnection rates, terms, and conditions be just, reasonable, and nondiscriminatory (section 251). It is this decision that inspired the pro net neutrality coalition to act now.

In 2005, the FCC issued a set of principles related to net neutrality [FCC 2005b]. They express the sentiment that consumers should be entitled to connect devices and to access content and applications of their choice. The principles also state that consumers are entitled to competition among network and application providers. However, the principle regarding competition is too vague to be useful, and the principles are not legally binding. The vagueness of this principle has left both the pro and anti net neutrality lobbies uncomfortable. The pro lobby is skeptical that it will have any impact on carrier behavior, and the carriers are nervous what might be expected of them.

5.2 Defining Appropriate Sets of Services

We begin by strongly disagreeing with the FCC's decision to classify Internet access solely as an information service. Telecommunication services can be thought of as the lower layers of the network, and information services can be thought of as the upper layers. The immediate consequence of this approach,

purely on a technical basis, is that the Internet clearly consists of both telecommunication services and information services. We believe that the FCC decision may have been based on a model that places Internet access solely at OSI layers 3 and 4, and classifies these layers as information services [Cannon 2003], whereas most networking researchers would describe Internet access as consisting of at least layers 1 through 3; we therefore believe that Internet access also consists of both telecommunication services and information services.

A number of papers have suggested using layers as a tool to formulate communications policy. Lessig [2001] considers a model consisting of physical, logical, applications, and content layers. He argues that the physical and logical layers of the Internet have historically been neutral. He believes it is acceptable for the physical layer to be closed, but proposes that the logical layer should be open and act as a commons. Werbach [2002] uses a similar set of four layers: The physical layer corresponds to OSI layer 1, the logical layer to OSI layers 2 through 6, and the application and content layers share OSI layer 7. He argues that communications policy should be formulated around these layers, with open interfaces between them. Solum and Chung [2003] propose a six-layer model, and argue that communications policy should attempt to respect the integrity of layers and to place regulation at or near the layer where the problem occurs. Whitt [2004] suggests a four-layer model similar to Werbach's, and presents principles concerning how layers should inform policy formulation.

Proper delineation of Internet infrastructure and Internet applications is an important tool in addressing net neutrality. However, we believe the relevant distinction is between functionality that must be provided *within the access network* and functionality that can be provided *elsewhere* in the Internet. As a starting point for this delineation, consider defining functions placed in OSI layers 1 through 3 (the Internet LAN-link and network layers) as network infrastructure, and functions placed in OSI layers 4 through 7 (the Internet transport and application layers) as network applications. As a consequence, functions that must be implemented at each IP hop are classified as infrastructure.

This distinction between network infrastructure and applications also helps with correct application of the end-to-end principle. Using the new terminology, we can view the principle as a suggestion that network functionality should be implemented in network infrastructure (at every IP hop) *only if* it cannot be implemented effectively in network applications (often solely at the endpoints).

We believe that these two sets—infrastructure and applications—are more appropriately suited to formation of a net neutrality policy than the layered models considered by others. This delineation can be used to distinguish between discrimination which is implemented in network infrastructure (e.g., QoS) and discrimination which is implemented in network applications (e.g., caching and search engines). In contrast, most of the other layered models discussed before combine OSI layers 3 and 4 (the Internet network and transport layers) into a single logical layer, and therefore fail to distinguish between network-layer functions provided in the access network and transport-layer functions provided only at endpoints.

However, a purely layer-based delineation of infrastructure and applications is not sufficient to capture the principle that infrastructure should contain

functionality that must be provided within the access network. While OSI layers 1 through 3 contain most of such protocols, for example, routing, addressing, and QoS, some protocols must be provided within the access network but not at every IP hop, such as DHCP. For this reason, we include in the definition of network infrastructure that functionality that is required to manage the network, using the following new definitions.

(1) (*Internet Infrastructure Services*). The term "Internet infrastructure services" means all services: (A) over a network that uses a public right-of-way; and (B) that reside at or below the network layer or are required to manage the network.

(2) (*Internet Application Services*). The term "Internet application services" means all services: (A) over a network that uses a public right-of-way; (B) that are not infrastructure services; and (C) that do not fall under Title VI of the Communications Act.

These terms agree with the intent of the older terms telecommunication services and information services. Internet infrastructure services do not change the content of information, similar to telecommunication services. In contrast, Internet application services create, store, or change the presentation of information, similar to information services.

Maintaining a distinction between the Internet infrastructure services and Internet application services is critical to formulation of good communications policy. Internet infrastructure services can only be provided by carriers, and must be provided by each carrier within their autonomous system. Such services include multi-user sharing of a wire or frequency (e.g., Ethernet or WiFi), routing (e.g., IP), and address assignment (e.g., DHCP). Internet infrastructure services require large investments into loops or wireless spectrum and switches or routers. These large initial fixed costs of the business are high relative to the costs per incremental customer served, and thus such Internet infrastructure services have a natural economy of scale that serves as a high barrier-to-entry. This high barrier-to-entry leads few carriers to offer service in any particular geographical region.

In contrast, Internet application services can be provided by carriers or by many other application providers on the Internet and can be placed at many locations within the Internet. Such services include email, Web hosting, caching, voicemail, and the portions of VoIP and IPtv that can be offered by independent application providers. Internet application services usually do have fixed costs that are small relative to incremental costs, and thus there is usually a low barrier-to-entry, which leads to a competitive market with a large number of application providers.

We recognize that some protocols that have usually been implemented in the access network (e.g., DNS[9], pop, and smtp) are thus classified as Internet application services. We believe this is appropriate, since these functions could be offered by an application provider other than the access carrier.

[9]DNS is a bit more complicated, as it relies on the coordination and cooperation between DNS servers among multiple autonomous systems.

6. A NEW NET NEUTRALITY POLICY

With the Internet's layered architecture and our new definitions of Internet infrastructure services and Internet application services in mind, we now turn to formulation of a new net neutrality policy. We start by examining the Communications Act for potential policy models. We then outline the basic tenets of a new net neutrality policy. Finally, we compare this new policy to the policies proposed in the academic literature discussed before.

6.1 Policy Models

The Communications Act broadly lays out three potential models to follow. First, if a communications market has sufficient competition, then minimal regulation is applied. This is captured in Title I of the Act, which applies forbearance from regulation unless required to ensure just and reasonable practices. The paradigm here is that a free market will regulate itself, and will promote investment in a fair manner that maximizes social welfare.

If sufficient competition is not present to ensure just and reasonable practices, Titles II and VI apply different models of regulation. Both models were initially formulated assuming a local monopoly over a portion of the corresponding network, but have since been applied to circumstances in which there is limited (but not sufficient) competition. Both models attempt to maximize social welfare by limiting monopoly behavior.

Title II (common carrier) uses the open access model. The network can be separated into two layers: infrastructure and applications. Local infrastructure, for example, the local loop, is usually deemed to be lacking sufficient competition, due to the high cost of building this infrastructure. In contrast, applications, such as telephone service, are often offered by a large number of competitors. Title II essentially dictates open access to infrastructure, and requires that carriers make critical network elements available to application providers without unjust or unreasonable discrimination and without undue or unreasonable preference. Title II also mandates reasonable and nondiscriminatory rates, terms, and conditions for interconnection with other infrastructure providers. As a result, telephone carriers do not limit who a customer can call and do not provide differentiated service based on who a customer calls.

Title VI (cable services), in contrast, uses a paradigm of a closed network. In this model, the network cannot be naturally segmented so that infrastructure and applications can be offered by separate entities. Vertically integrated carriers control both the infrastructure and the applications. Title VI recognizes this natural vertical integration, and applies limits to the cable carrier's behavior when it is also integrated with content providers. As a result, cable carriers do determine which content a customer can receive.

6.2 A Layered Net Neutrality Policy

In this subsection, we outline the basic tenets of a new net neutrality policy. We believe that the Internet's architecture supports the need for some type of openness. However, we do not believe that net neutrality by itself requires

open access. This intermediate approach places us loosely in the Nondiscriminationist camp. We suggest that net neutrality can be effectively achieved by a properly regulated *open interface* from applications to network infrastructure. The requirement of an open interface captures the central principle of a layered Internet architecture, and yet is less intrusive than the requirement of open access.

We also believe economic arguments support the need for some type of openness. We believe that carriers should offer an interface to network infrastructure to applications without unjust or unreasonable discrimination and without undue or unreasonable preference. We argue that an open interface mandate can be used to limit monopoly or oligopoly behavior. We adopt the free market model, based on Title I forbearance, for the application portion of the Internet. We mandate an open interface, based on selected elements of Title II, for the local infrastructure portion of the Internet. We believe this combination is most likely to encourage continued development of the Internet in a manner that maximizes social welfare.

To formulate this policy, we start by focusing on the goal that an open interface should attempt to provide a level playing field between application providers and ISPs who offer competing applications. Threats to a level playing field are potentially posed by vertical integration between Internet service providers and application providers. Certain Internet application services (e.g., video on demand, high-quality VoIP, video conferencing) require specialized Internet infrastructure services (e.g., dedicated bandwidth or QoS). Our first net neutrality tenets are therefore as follows.

(A) ISPs should be prohibited from refusing to provide enabling Internet infrastructure services to competing application providers in order to differentiate the ISP's own application offerings.
(B) ISPs should be prohibited from making exclusive deals to provide enabling Internet infrastructure services to certain application providers.

These tenets can be implemented by selectively applying current communications policy to Internet access. First, we require each ISP to allow each subscriber to receive service that does not discriminate in the carriage and treatment of Internet traffic based on the source, destination, or ownership of such traffic. This was the principal component of the Snowe-Dorgan amendment. Second, we apply section 202 of Title II of the Communications Act to ISPs with respect to Internet infrastructure services. This section will thus bar ISPs from unjust or unreasonable discrimination and from giving undue or unreasonable preference in use of Internet infrastructure services. In our context, these two provisions will give ISPs the right to implement bandwidth allocation and QoS and to sell such service, and give subscribers the right to purchase such service that discriminates only on application type, broadly supporting tenets A and B.

To guarantee a level playing field, we must not only ensure that Internet infrastructure services are available, but also that they are priced reasonably. We thus add a third net neutrality tenet.

(C) ISPs should be prohibited from providing Internet infrastructure services to competing application providers at inflated prices in order to favor the ISP's own application offerings.

This tenet can be implemented in two steps. First, consider the responsibilities of vertically integrated ISPs to *their own subscribers and peers*. When an ISP offers applications that rely on Internet infrastructure services, we require the ISP to make available to competitors the same Internet infrastructure services at the same prices. Second, consider the responsibilities of vertically integrated ISPs to *subscribers of other ISPs*. Without regulating peering agreements, we prohibit a vertically integrated ISP from using service-level agreements with its peers to favor its own applications. In both situations, section 628 of the Communications Act, which similarly addresses vertical integration between cable carriers and content providers, can serve as a model.

Such threats to a level playing field are limited in scope, and therefore any remedy should be similarly limited in scope. Vertical integration only affects applications which rely on specialized Internet infrastructure. A characteristic of such network services is that they must be provided directly in the *access network*. Networks services that can be provided outside the access network (e.g., caching, search engines, Web hosting) can be provided by competing ISPs or application providers, and therefore do not represent a threat to a level playing field. We therefore add a fourth net neutrality tenet.

(D) Limit tenets A through C by restricting their use to Internet infrastructure services.

This tenet has already been implemented by using the new definitions in the statement of tenets A through C.

When sufficient competition exists in the Internet access market, there is a lower danger that vertical integration would cause a nonlevel playing field. Our next net neutrality tenet is thus given.

(E) Assure forbearance where sufficient competition exists.

Section 10 in Title I of the Communications Act instructs the FCC to forbear from applying regulations unless they are required to ensure just and reasonable practices. This tenet can be therefore be elegantly implemented by applying this section to ISPs with respect to Internet infrastructure services. The result is that tenets A through C will not apply when there is a competitive market.

Finally, we add two net neutrality tenets to ensure that tenets A through C do not limit procompetitive behavior.

(F) Ensure that ISPs have the right to apply network management mechanisms that do not threaten a level playing field.

(G) Ensure that regulation does not impede an ISP from making arrangements with consumers, application providers, and peering ISPs for Internet infrastructure services in a manner that does not conflict with the aforesaid tenets.

Tenet F requires a delineation of allowed network management practices. The Senate Commerce bill describes allowed network management regarding network security, diagnostics and repair, authorized use, and parental control. We believe it would be of value to add a provision to explicitly give ISPs the right to alleviate congestion by treating all traffic similarly, or by treating all applications of the same type similarly.

Tenet G can be implemented be giving ISPs the right to sell reserved bandwidth and QoS to both their residential and business subscribers, and the right to discriminate in the carriage of Internet traffic based on peering arrangements with other ISPs. Such peering arrangements would allow for reserved bandwidth and QoS to be provided cooperatively by multiple ISPs.

Detailed proposed statue language implementing these net neutrality tenets can be found in Jordan [2007].

6.3 Comparison to Other Approaches

Openists believe that the Internet is best served by maintaining a dumb network that does not differentiate among different types of traffic. They use the end-to-end principle to support this position. We disagree with both the reasoning and the conclusion. We believe that the end-to-end principle suggests implementing functionality as Internet application services if feasible. However, as noted before, QoS mechanisms such as priority or reservations must be implemented at the Internet LAN-link and/or network layers in order to be effective. The end-to-end principle therefore does not ban their use, and therefore does not mandate a dumb network. Indeed, one of the original authors of the end-to-end principle was also one of the authors of RSVP [Braden et al. 1994], and has suggested that a new theory addressing network core functionality should be created to co-exist with the end-to-end principle [Blumenthal and Clark 2001] and that the end-to-end principle should be generalized with a new form of modularity that can gracefully accommodate tussle along various competing players [Clark et al. 2005]. We believe that the introduction of QoS, and hence the creation of a smart network, will *increase* the range of applications that can be efficiently run on the Internet.

Openists support a policy based on open access, including a broad ban on vertical integration. We agree with the need for an open interface between network infrastructure and applications, but do not believe that net neutrality by itself is threatened by vertical integration that crosses other layer boundaries.

Deregulationists present economic arguments that regulation will hinder investment by ISPs and that in the absence of regulation ISPs will vertically integrate only in ways that maximize consumer welfare. However, other economists argue that the same theory of network externalities indicates that ISPs do have an incentive to discriminate against competing application providers, and that the economic benefits of net neutrality outweigh the costs. We accept here this latter claim, and agree with the goal of providing a level playing field. Although we agree that development of applications drives network infrastructure investment, we do not believe that vertical integration violating an open interface is necessary to foster new applications.

Some deregulationists have also suggested that competition can be more effectively fostered between different physical networks that are matched to different types of applications. We strongly disagree, and believe that 20 years of network convergence provides strong evidence for the efficiency of integrated networks over the use of separate infrastructures.

Finally, many deregulationists equate QoS with other types of discriminatory network behavior (e.g., caching or priority listing of search results) and use this to argue that QoS should not be regulated. We strongly disagree. We are concerned with the use of prioritization when there is limited competition (e.g., QoS); we have faith that free market dynamics will adequately address the use of prioritization when there is sufficient competition (e.g., caching or priority listing of search results). Equating the two types of prioritization ignores the distinction between network infrastructure services (e.g., QoS) which must be implemented in the access network and network applications (e.g., caching) which can be implemented in other ISPs' networks.

We agree with the Nondiscriminationist camp's arguments that there are good and bad uses of traffic discrimination. Although we do not believe that QoS requires vertical integration, we support a policy that allows vertical integration and traffic differentiation, but restricts their use to ensure that ISPs do not discriminate in a manner that extracts oligopoly rents. However, whereas Wu and Negi [2003] would ban discrimination at the Internet network layer and hence ban QoS based on the type of application, we believe that an open interface is a cleaner and more effective solution. This difference will become clear as we analyze what types of ISP network management should be considered unacceptable.

7. ACCEPTABLE AND UNACCEPTABLE ISP NETWORK MANAGEMENT

In this section, we apply our new net neutrality policy to illustrate which types of network management should be considered acceptable and which unacceptable. We consider a wide range of network management actions related to net neutrality. For each, we decide whether the action should be allowed or prohibited by considering the tenets presented in the last section as well as the general intent of the Communications Act as applied to other technologies. We focus on the case in which sufficient competition does not exist in the Internet access market; where sufficient competition exists we choose to apply forbearance. When doubt remains as to the desirable policy goal, we choose the action most likely to encourage a level playing field.

7.1 Blocking

First, consider *blocking* of either a Web page or an application. Currently, blocking (when used) is usually done on the basis of the *type of application*. This type of blocking can be beneficial, for example, firewalls that block security intrusions. However, it can also be anticompetitive, for example, an ISP that blocks all VoIP traffic in order to limit competition to its own voice service. In some instances, blocking based on application type can be viewed as either beneficial

or harmful depending on user preferences, for example, a firewall that blocks all file sharing traffic.

Is it acceptable for an ISP to block applications on the basis of the application type? In the absence of sufficient competition in infrastructure, we believe such blocking is not acceptable because it may violate tenet A stated earlier. While some ISPs block certain applications to limit bandwidth use, we believe that network management should be transparent and that tiering is a superior solution. Even if sufficient competition exists, then the layered architecture of the Internet strongly suggests that such blocking should be rare. Our answer is that blocking based on application type should be a decision made by the consumer, not the ISP.

It is less common to block Web pages or applications based on the *source and/or destination*. However, sometimes this type of blocking is used for security purposes. Examples include parental control software and firewalls that block all traffic from unknown addresses. Is it acceptable for an ISP to block applications on the basis of the source and/or destination of the traffic? We believe this is fairly straightforward, and that blocking based on source and/or destination should also be a decision made by the consumer, not the ISP.

7.2 Degradation

Next, we consider *degradation* of Internet traffic. Currently, intentional degradation is common in a few scenarios. First, it is common to place a limit on the bandwidth of a broadband user's access link that is lower than the speed the link could accomplish. Such limits, referred to as *tiering*, are often used to differentiate service offerings by a selection of access rates, and to charge different prices for these different service offerings. Although tiering is a form of intentional degradation, neither the pro or anti net neutrality lobbies oppose its use. Second, it is increasingly common for an ISP to place bandwidth limits on certain classes of applications, such as traffic shaping. Currently, the most common such limits are on file sharing traffic, which otherwise may consume a large proportion of the available bandwidth and cause congestion for all other applications. The bandwidth limits degrade the performance of the applications to which they are applied, but do not block the applications entirely. Both tiering and traffic shaping are intentional degradation based on the type of application (and in the case of tiering on consumer payment). In contrast, degradation of traffic based on the source and/or destination is both technically difficult and uncommon.[10]

Is it acceptable for an ISP to intentionally degrade traffic? We believe that degradation based on application type, for example, tiering and traffic shaping, is acceptable since it does not bias the network toward certain providers. Indeed, we believe that such techniques can be crucial elements of network management, support tenet F, and should not be discouraged. However, we believe that any degradation based on source or destination should also be a

[10]We do not consider congestion control measures (e.g., TCP) which treat all users similarly as being intentional degradation based on source or destination.

decision made by the consumer, not the ISP, since it would violate tenet A in much as the same manner as blocking based on source or destination.

7.3 QoS

Finally, we consider QoS mechanisms such as reservations or priorities. QoS is often used by carriers to provide acceptable performance for their own VoIP or video conferencing traffic as part of business packages that offer combined voice, video, and data transport. Although use of QoS for residential service is currently less common, we expect its use to increase with the growth in residential VoIP, video conferencing, and video streaming.

A traffic stream receives enhanced performance through QoS if the stream requests it and the carrier grants the request. A carrier may grant QoS requests: (1) to only its *own applications* (e.g., only its own VoIP subscribers), (2) to *particular applications types* (e.g., VoIP packets from all providers), or (3) on the basis of *contractual arrangements* (e.g., consumer and/or application provider payment).

A central component of net neutrality policy is the decision of how to treat QoS mechanisms. Is it acceptable for an ISP to implement QoS in its network? Some net neutrality proponents argue that the Internet should only provide best-effort service so that all Internet streams are treated identically; this group opposes any use of QoS. Although some carriers may choose alternate methods to provide acceptable performance to voice and video applications, we believe that a broad ban on use of QoS is a Luddite instinct. Furthermore, we do not see any precedent in the Communications Act for a similar ban on a set of technologies. We therefore believe use of QoS can be acceptable. The real issue, we believe, is how QoS is used.

Is it acceptable for QoS to be based on application type? Currently, the Internet adopts different types of transport[11] for different types of applications. QoS would fall into the same category. We do not see a threat to tenets A through C through this use of QoS, and our open interface approach would allow it. As an example, if a carrier chooses to deploy QoS mechanisms in their network and to make these mechanisms available to all VoIP traffic, regardless of the application provider supporting the service, we believe this approach should be allowed. The anti net neutrality lobby supports this position. The pro net neutrality lobby, excepting those who believe QoS should not be allowed in any manner, also supports this position.

Is it acceptable if a carrier deploys QoS in its network to support some of its own applications, but does not offer these mechanisms to competing application providers? The anti net neutrality lobby argues that refusal to provide access for QoS to application providers would not hinder the application provider's ability to deploy new services. We disagree. QoS should be classified as infrastructure, not application, because it must be implemented in every autonomous system in which congestion occurs in order to be effective. As a result, there is no effective option available to application providers who wish to provide enhanced

[11]TCP versus UDP.

performance to voice and video applications; they must receive the same QoS as the carrier uses for its own similar applications in order to compete effectively. In support of tenets A and B, we conclude that it is not acceptable for a carrier to deploy QoS in its network to support some of its own applications, but not to offer these mechanisms to competing application providers. Therefore, we believe that QoS on the basis of source or destination is not acceptable.

The remaining issue with use of QoS is whether it is acceptable for QoS to be based on consumer and/or application provider payment. The anti net neutrality lobby argues that this approach should be allowed without constraint. Such contractual arrangements based on consumer payment are likely to be combined with tiering, for example, for a small additional fee all VoIP traffic to/from a particular residence will receive QoS. Alternately, contractual arrangements for QoS could be based on provider payment, for example, all VoIP traffic to/from a particular provider will receive QoS. Both approaches are technically feasible. The fee for such service could be a flat monthly fee, a per-minute charge, or a per-byte charge. Any consumer fee would almost certainly be based on a publicly available fee schedule, and therefore be nondiscriminatory. We see no threat to tenets A through C from QoS based on consumer payment, and believe allowing this approach supports tenet G. We also note that most pro net neutrality groups would agree.

Application provider fees for QoS, however, could potentially be subject to *commercially negotiated* agreements, as are commonly used when carriers offer virtual private networks to business customers. Such agreements open the door to potential discrimination between various application providers. Under the policy developed before, we have already chosen to prohibit an ISP from offering QoS to only selected application providers, so the issue here is whether an ISP can charge different fees to different application providers. We believe the open interface approach suggests that any such discrimination must be reasonable, for example, based on the differences in cost of providing QoS. We conclude that similar constraints should be placed on QoS based on provider payment.

7.4 Summary

The decisions made previously with regard to what types of network management should be allowed or prohibited result in an answer of yes for every row in Table I. All of the bills in Congress prohibit blocking, as would we. With regard to degradation, our decision to allow traffic shaping puts us at odds with the Wyden bill, whereas our decision to prohibit degradation on the basis on source or destination (unless elected by the consumer) puts us at odds with the House and Senate bills. With regard to use of QoS, our decision to allow QoS based on application type puts us at odds with the Wyden bill, whereas our decision to prohibit exclusive use of QoS puts us at odds with the House and Senate bills. Finally, our decision to allow either consumer or provider payment for QoS puts us at odds with the Markey, Sensenbrenner, and Snowe-Dorgan bills, whereas our mandate that such fees be reasonable and nondiscriminatory puts us at odds with the House and Senate bills.

Table II. Effect of Our Net Neutrality Policy upon an ISP Offering VoIP

Is it acceptable if . . .	Our answer
. . . a carrier *blocks* a competitor's VoIP traffic?	no
. . . a carrier doesn't block a competitor's VoIP traffic, but doesn't offer QoS to competitor's VoIP subscribers while using QoS for its own VoIP subscribers?	no
. . . a carrier gives a *broadband subscriber* who uses a competitor's VoIP service the choice of: (a) best-effort transport of their VoIP traffic as part of the basic broadband package, (b) enhanced performance for their VoIP traffic for an additional $0.01/min paid by the subscriber, or (c) enhanced performance for up to 500 minutes of their VoIP traffic for an additional $5/month paid by the subscriber?	yes
. . . a carrier gives a *VoIP competitor* the choice of (a) best-effort transport of their VoIP traffic as part of the subscriber's basic broadband package or (b) enhanced performance for their VoIP traffic for an additional $0.01/min paid by the VoIP provider?	yes
. . . a carrier charges different VoIP competitors *different prices* for QoS?	no
. . . a carrier charges VoIP competitors a uniform price for QoS, but a different price than charged to *its own affiliates*?	no

To illustrate these policy decisions, consider their effect upon an ISP who offers both broadband Internet service and a separate VoIP service. The effects of our policy tenets are shown in Table II.

8. CONCLUSION

Congressional action on net neutrality will substantially affect the evolution of the Internet and of future Internet research. We have argued that neither the pro nor anti net neutrality positions are consistent with the philosophy of Internet architecture, because both misconstrue the end-to-end principle. We have defined new terms Internet infrastructure services and Internet application services to distinguish between functionality that must be provided in the access network and functionality that can be provided by other application providers. Based on these new terms, we have proposed a new net neutrality policy that restricts an Internet Service Provider's ability to engage in anticompetitive behavior while simultaneously ensuring that it can use desirable forms of network management. We illustrated the effect of this policy by discussing acceptable and unacceptable uses of network management.

Our approach differs with many of the pro net neutrality bills in that it allows nondiscriminatory network management techniques and QoS and it allows consumer or provider payment for QoS. It differs with many of the anti net neutrality bills in that it prohibits discriminatory use of Internet infrastructure and discriminatory degradation and mandates reasonable and nondiscriminatory fees for QoS.

There are substantial issues remaining before this proposed net neutrality policy could be implemented. First, statues need to be drafted that faithfully implement the proposed tenets. (Proposed statue language can be found in Jordan [2007].) Second, the role of regulation and enforcement must be considered. In particular, we foresee substantial issues in regulation and enforcement of the prohibition on unreasonable price discrimination for fees charged to application providers for QoS.

REFERENCES

BAR, F., COHEN, S. S., COWHEY, P., DELONG, B. J., KLEEMAN, M., AND ZYSMAN, J. 2000. Access and innovation policy for the third-generation Internet. *Telecommun. Policy 24*, 489–518. http://repositories.cdlib.org/brie/BRIEWP137.

BAR, F. AND SANDVIG, C. 2000. Rules from truth: Communication policy after convergence. In *Proceedings of the 28th Telecommunications Policy Research Conference*.

BLOOMBERG NEWS. 2006. AT&T, Comcast Rout Google in 'Net Neutrality' Battle. *Bloomberg News*, July 20.

BLUMENTHAL, M. S. AND CLARK, D. D. 2001. Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world. *ACM Trans. Inter. Tech. 1*, 1, 70–109.

BRADEN, R., CLARK, D., AND SHENKER, S. 1994. Integrated services in the Internet architecture: An overview. Tech. rep., IETF.

CANNON, R. 2003. The legacy of the federal communications commission's computer inquiries. *Federal Commun. Law J.* http://www.cybertelecom.org/ci/index.htm.

CLARK, D., WROCLAWSKI, J., SOLLINS, K., AND BRADEN, R. 2005. Tussle in cyberspace: Defining tomorrow's Internet. *IEEE/ACM Trans. Netw. 13*, 462–475.

FARRELL, J. AND WEISER, P. 2003. Modularity, vertical integration, and open access policies: Towards a convergence of antitrust and regulation in the Internet age. *Harvard J. Law Technol. 17*, 85–134. http://ssrn.com/abstract=452220.

FCC. 1971. First computer inquiry, final decision, 28 fcc2d 267.

FCC. 1980. Second computer inquiry, final decision, 77 fcc2d 384.

FCC. 2002. FCC 02-77 cable modem service order. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-02-77A1.pdf.

FCC. 2005a. FCC 05-150 Internet over wireline facility order. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-150A1.pdf.

FCC. 2005b. FCC 05-151 Internet policy statement. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.pdf.

HERMAN, B. D. 2007. Opening bottlenecks: On behalf of mandated network neutrality. *Federal Commun. Law J. 59*. http://ssrn.com/abstract=902071.

JORDAN, S. 2007. A layered network approach to net neutrality. *Int. J. Commun.* (to appear).

LEHR, W. H., GILLETT, S. E., SIRBU, M. A., AND PEHA, J. M. 2006. Scenarios for the network neutrality arms race. In *Proceedings of the 34th Telecommunications Policy Research Conference*.

LEMLEY, M. A. AND LESSIG, L. 2001. The end of end-to-end: Preserving the architecture of the Internet in the broadband era. *UCLA Law Rev. 48*, 925–972. http://ssrn.com/abstract=259491.

LESSIG, L. 2001. The Internet under siege. *Foreign Policy 127*, 56–65.

MCTAGGART, C. 2006. Was the Internet ever neutral? In *Proceedings of the 34th Telecommunications Policy Research Conference*.

OWEN, B. M. AND ROSSTON, G. L. 2003. Local broadband access: Primum non nocere or primum processi?: A property rights approach. Stanford law and economics olin working paper no. 263. http://ssrn.com/abstract=431620.

PEHA, J. M. 2006. The benefits and risks of mandating network neutrality, and the quest for a balanced policy. In *Proceedings of the 34th Telecommunications Policy Research Conference*.

SALTZER, J. H., REED, D. P., AND CLARK, D. D. 1984. End-to-end arguments in system design. *ACM Trans. Comput. Syst. 2*, 4, 277–288.

SOLUM, L. B. AND CHUNG, M. 2003. The layers principle: Internet architecture and the law. U. San Diego public law research paper No. 55. http://ssrn.com/abstract=416263.

SPETA, J. B. 2000. Handicapping the race for the last mile?: A critique of open access rules for broadband platforms. *Yale J. Regulation 17*, 39–91.

U.S. CONGRESS. 1984. S.66.PP, Cable Communications Policy Act of 1984, 98th Congress, Public Law 98-549. http://thomas.loc.gov/cgi-bin/bdquery/z?d098:s.00066:.

U.S. CONGRESS. 1992. S.12.PP Cable Television Consumer Protection Act of 1992, 102nd Congress, Public Law 102-385. http://thomas.loc.gov/cgi-bin/bdquery/z?d102:s.00012:.

U.S. CONGRESS. 1996. S.652.PP Telecomunications Act of 1996, 104th Congress, Public Law 104-104. http://thomas.loc.gov/cgi-bin/bdquery/z?d104:s.00652:.

U.S. CONGRESS. 2006a. H.R.5252.EH Advanced Telecommunications and Opportunities Reform Act, 109th Congress. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:HR05252:.

U.S. CONGRESS. 2006b. H.R.5252.RS, Communications Opportunity, Promotion, and Enhancement Act of 2006, 109th Congress. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:HR05252:.

U.S. CONGRESS. 2006c. H.R.5417.RH Internet Freedom and Nondiscrimination Act of 2006, 109th Congress. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:h.r.05417:.

U.S. CONGRESS. 2006d. Net Neutrality Act of 2006, 109th Congress, Proposed Amendment to H.R.5252. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:HZ00987:.

U.S. CONGRESS. 2006e. S.2360.IS Internet Non-Discrimination Act of 2006, 109th Congress. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:s.02360:.

U.S. CONGRESS. 2006f. S.2917.IS Internet Freedom Preservation Act, 109th Congress. http://thomas.loc.gov/cgi-bin/bdquery/z?d109:s.02917:.

VAN SCHEWICK, B. 2007. Towards an economic framework for network neutrality regulation. *J. Telecommun. High Technol. Law 5*. http://ssrn.com/abstract=812991.

WERBACH, K. D. 2002. A layered model for Internet policy. *J. Telecommun. High-Tech Law 1*. http://ssrn.com/abstract=648581.

WHITT, R. S. 2004. A horizontal leap forward: Formulating a new Ccommunications public policy framework based on the network layers model. *Federal Commun. Law J. 56*, 587–672.

WU, D. AND NEGI, R. 2003. Effective capacity: A wireless link model for support of quality of service. *IEEE Trans. Wireless Commun. 2*, 4, 630–643.

WU, T. 2004. The broadband debate: A user's guide. *J. Telecommun. High Technol. Law 3*, 69–95. http://ssrn.com/abstract=557330.

YOO, C. S. 2005. Beyond network neutrality. *Harvard J. Law Technol. 19*, 1–77. http://ssrn.com/abstract=742404.

YOO, C. S. 2006. Network neutrality and the economics of congestion. *Georgetown Law J. 94*. http://ssrn.com/abstract=825669.

Received February 2007; accepted June 2007