REDEFINING NET NEUTRALITY AFTER *COMCAST V. FCC*

Alexander Reicher[†]

Critics sometimes describe James Joyce's modernist epic *Ulysses* as the most discussed, least read novel in the world.[1] Net neutrality may be the most discussed, least understood concept in the world of internet policy. Consequently, the term has so many definitions advancing so many different goals that the net neutrality debate seems at times only about *what* net neutrality is, not *why* it should (or should not) be. The debate was reopened this past year with the D.C. Circuit's decision in *Comcast Corporation v. Federal Communications Commission*, which invalidated the FCC's jurisdiction over broadband internet service providers (ISPs), including its jurisdiction to enforce a policy statement of net neutrality principles.[2] Although the court focused exclusively on Comcast's procedural challenge to the FCC's jurisdiction, the FCC and the policy community subsequently have engaged in a process of redrafting not only the jurisdictional basis but also the net neutrality principles themselves. In late December 2010, the FCC adopted a set of net neutrality rules for the first time through a formal rulemaking process—going beyond the general policy statement of net neutrality principles invalidated in *Comcast* by requiring transparency and forbidding most blocking and discrimination.[3] This Note analyzes and affirms the importance of mandating full ISP transparency, as the FCC has done in this recent regulation. Given that ISPs will now be required to disclose whether they discriminate among content, services, and applications, this Note also proposes a two-step analysis to determine whether a given practice should be considered reasonable or unreasonable network management.

[†] J.D. Candidate, 2012, University of California, Berkeley School of Law.

[1] JAMES JOYCE, ULYSSES (Hans Walter Gabler, ed., Vintage Books 1986) (1922); *see, e.g.*, Barbara Leckie, *"Short Cuts to Culture": Censorship and Modernism; or, Learning to Read* Ulysses, 17 European Joyce Studies 9, 25 (2006).

[2] Comcast Corp. v. FCC, 600 F.3d 642, 661 (D.C. Cir. 2010).

[3] Preserving the Open Internet Broadband Industry Practices, Report and Order, WC Docket No. 07-52 (Dec. 23, 2010), http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf [hereinafter *Open Internet Rules*].

The term "net neutrality" refers to a bundle of open access principles enforced in a variety of legal and technical ways. According to one common definition, "[n]et neutrality means simply that all like internet content must be treated alike and move at the same speed over the network."[4] As used by scholars, lawyers, and engineers, the term "net neutrality" can refer simultaneously to three different understandings. First, the term can address a collection of theoretical "net neutrality principles"—mainly, the principles that we should protect innovation, free speech, and competition on the Internet.[5] Second, it can encompass the set of legal rules and policies that the FCC enforces, first adopted in the "Internet Policy Statement" and, more recently, in the "Open Internet Rules."[6] Lastly, it can refer to the network protocols and internet architecture that can direct, on the technical level, how ISPs discriminate among content, services, or applications. Of course, the theoretical, legal, and technical definitions are related in that theoretical net neutrality principles often inform the legal codification and technical execution of net neutrality. This Note argues, however, that an operational legal definition of net neutrality must encompass not only the theoretical principles underlying the term but also the technical realities of the Internet, such as its physical architecture and interconnections. This Note will also suggest that the debate over the very definition of net neutrality and what constitutes reasonable network management may be resolved through the FCC's enforcement of a transparency principle. Requiring ISPs to disclose how they discriminate will force them to compete on how they define net neutrality and reasonable network management.

This argument proceeds in three parts. Part I, THEORETICAL NET NEUTRALITY, introduces the major net neutrality principles, which include protections for innovation, free speech, and competition. It also introduces various types of discrimination undertaken by ISPs. Not all forms of discrimination necessarily violate all of the net neutrality principles; the

[4] Lawrence Lessig & Robert W. McChesney, *No Tolls on the Internet*, WASH. POST, June 8, 2006, http://www.washingtonpost.com/wp-dyn/content/article/2006/06/07/AR2006060702108.html.

[5] *See*, *e.g.*, *Hearing on "Network Neutrality," Before the Senate Comm. on Commerce, Science and Trans.*, 110th Cong. 4 (2006) (statement of Lawrence Lessig, C. Wendell and Edith M. Carlsmith Professor of Law Stanford Law School) [hereinafter *Lessig Senate Hearing*]; Al Franken, *Net Neutrality Is Foremost Free Speech Issue of Our Time*, CNN.COM (Aug. 5, 2010), http://www.cnn.com/2010/OPINION/08/05/franken.net.neutrality/; Philip J. Weiser, *The Next Frontier for Network Neutrality*, 60 ADMIN. L. REV. 273, 277 (2008).

[6] *See* Appropriate Framework for Broadband Access to the Internet Over Wireline Facilities, 20 FCC Rcd. 14986 (2005) [hereinafter *Internet Policy Statement*]; *Open Internet Rules*, *supra* note 3, ¶¶ 43–115.

internet community accepts some discriminatory practices because they are technically necessary or because they do not violate the net neutrality principles in any substantial way.[7] Collectively, these discriminatory, but generally accepted, practices are known as "reasonable network management." Part II, LEGAL NET NEUTRALITY, surveys the history of the FCC's jurisdiction and enforcement of net neutrality through *Comcast v. FCC*. However, because all FCC actions have involved ISPs that completely block competitors' services and applications, these cases do not help distinguish reasonable network management practices from unreasonable ones in instances where ISPs only delay the delivery of certain content. To facilitate drawing this distinction, Part III, TECHNICAL NET NEUTRALITY, examines the technical realities of the Internet by reviewing the physical architecture, interconnection agreements among service providers, and protocol layers of the Internet. After considering the definition of net neutrality and reasonable network management from these three perspectives, this Note concludes that mandating ISP transparency is an essential part of an enforceable definition of net neutrality that accounts for the Internet's technical realities.

I. THEORETICAL NET NEUTRALITY

Articulating net neutrality principles serves the important purpose of envisioning the Internet as if it were, and has always been, a fully neutral network. This theoretical mode of discussion is important in forming a set of ideals for the Internet, which includes the principles of innovation, free speech, and competition, as well as the idea of reasonable network management.

A. NET NEUTRALITY PRINCIPLES

Net neutrality principles represent what we value most about the Internet: its ability to produce innovation, foster free speech, and promote competition. As such, these principles should always serve as a framework for understanding and enforcing legal and technical net neutrality. Although this Note ultimately concludes that these theoretical principles are inadequate as an enforceable definition of net neutrality, they are an essential starting point.

[7] Scarce network resources may force network administrators to violate certain net neutrality principles. At peak times of network congestion, for example, a network administrator may need to limit a highly innovative but bandwidth-intensive application to maintain a reliable network.

1. *Transparency*

Transparency is the idea that ISPs should disclose how they manage their networks. Mandating that ISPs disclose their network management practices is not itself a separate principle, since a network service provider may maintain a perfectly neutral network while (for whatever reason) failing to disclose how the network is managed. Transparency is rather a subservient concept to the other net neutrality principles, but it is nonetheless the most important component of a net neutrality definition. Encouraging ISPs to disclose "meaningful information" about their service plans, former FCC Chairman Michael Powell observed that the importance of such information is that it is "necessary to ensure that the market is working."[8]

Transparency is not only necessary to maintain honest competition in the market for the provision of broadband service, it is also essential to *create* new forms of competition among service providers on the basis of how they define net neutrality and reasonable network management. The transparency principle acknowledges that the theoretical net neutrality principles are ideals and that providers should be required to disclose how and when they deviate from those ideals—in essence, how they define net neutrality and reasonable network management. This empowers consumers with the opportunity to choose the form of net neutrality they value and the type of reasonable network management they can tolerate.

2. *Innovation*

It is now universally acknowledged that the Internet has become the platform for some of the most impressive innovations of the past several decades. According to one widely-accepted theory, this type of disruptive innovation occurs when users are able to adapt older technologies to entirely new purposes.[9] Based on this proposition, some conclude that the Internet's neutral design—its equal treatment of content, services, and applications—has allowed innovators to freely adapt it to entirely new uses with nearly no restrictions imposed by ISPs.[10] This argument that net neutrality protects innovation draws upon the engineering concept known as the end-to-end (e2e) principle, which provides that the middle, or "core," of the Internet

[8] Michael K. Powell, Chairman, FCC, Remarks at the Silicon Flatirons Symposium on "The Digital Broadband Migration: Toward a Regulatory Regime for the Internet Age," University of Colorado School of Law: Preserving Internet Freedom: Guiding Principles for the Industry 5 (Feb. 8, 2004).

[9] *See* JONATHAN L. ZITTRAIN, THE FUTURE OF THE INTERNET AND HOW TO STOP IT 86 (2008) (citing ERIC VON HIPPEL, DEMOCRATIZING INNOVATION 19 (2005)).

[10] *See, e.g.*, *Lessig Senate Hearing*, *supra* note 5, at 4.

should provide only general processing services so as not to favor one type of content, service, or application over another.[11] This principle is sometimes referred to as the "dumb pipe" argument, since an e2e network has little network "intelligence" between, for example, a user and a website.[12] Therefore, the network (and the service provider that controls it) cannot favor, disfavor, or otherwise disrupt the connection. Though the e2e principle originated as an engineering principle, it now stands for a theory that delegates the role of innovating new services, content, and applications to end-users rather than to ISPs.[13] This creates a competitive environment among the uncountable number of internet end-users, who develop applications that a smaller group of core ISPs could never have anticipated. Email, for example, was the "unintended by-product" of early internet users, rather than a central purpose envisioned by the original network service providers.[14] Net neutrality thus ensures that the Internet remains open to this kind of disruptive innovation from end-users.

3. Free Speech

As a net neutrality principle, protecting free speech on the Internet is related to, but conceptually separate from, protecting innovation. Both innovation and free speech are protected by a non-discriminating, e2e network, but the free speech principle is more concerned with censorship of perspectives than with barriers to entry for new companies. Senator Al Franken calls net neutrality "the most important First Amendment issue of our time."[15] He wrote in a guest column on CNN.com: "You're reading this op-ed online; it'll load just as fast as a blog post criticizing it. That's what we mean by net neutrality."[16] From this perspective, there is harm to free speech not only when content is censored entirely, but also when some points of view are prioritized over others. Thus, if one news source is "throttled" (slowed) by an ISP, over time users might migrate to other, faster-loading

[11] *See* BARBARA VAN SCHEWICK, INTERNET ARCHITECTURE AND INNOVATION 378 (2010). This engineering design principle was first articulated in J.H. Saltzer et al., *End-to-End Arguments in System Design,* 2 ACM TRANSACTIONS ON COMPUTER SYS. 277 (1984).

[12] *Cf.* David S. Isenberg, *The Rise of the Stupid Network*, COMPUTER TELEPHONY 16–26 (Aug. 1997) (calling the same phenomenon a "stupid network").

[13] *See* Tim Wu, *The Broadband Debate, A User's Guide*, 3 J. ON TELECOMM. & HIGH TECH. L. 69, 73–74 (2004).

[14] See Mark A. Lemley & Lawrence Lessig, *The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era*, 48 UCLA L. REV. 925, 932 (2001).

[15] Franken, *supra* note 5.

[16] *Id.*

sources. This form of discrimination, in addition to the wholesale blocking of content, violates the free speech net neutrality principle.

Courts have also recognized that the Internet is now *the* platform for both public and private speech. In *Reno v. American Civil Liberties Union*, the Supreme Court quoted Judge Dalzell of the Eastern District of Pennsylvania, who described the Internet as "the most participatory form of mass speech yet developed."[17] Indeed, some net neutrality supporters suggest that as social networks become fixtures of communication, the increasingly complicated human interactions that occur on those networks are becoming the central purpose of the Internet.[18] Net neutrality, based on this view, protects free speech on the Internet's blogs and social networks, some of which have become the new town square or the new Pruneyard Shopping Center.[19]

4. *Competition*

The net neutrality principle of maintaining competition concerns two separate but related markets: the market for the provision of internet service and the market for content, services, and applications. Under the competition principle, the call for net neutrality regulation responds to alleged failures in both of these markets.[20] Failure in the broadband services market means higher prices for subscribers. Failure in the content, services, and applications market means higher barriers for new (and potentially innovative) entrants. The latest data from the FCC Wireline Competition Bureau indicate that roughly half of households in the United States have access to just two choices of broadband ISPs.[21] The discussion about regulating this duopoly echoes debates over public utility regulation from the last one hundred years. According to this history, "a provider of basic infrastructure—a railroad or a telecommunications network—will often seek

[17] Reno v. ACLU, 521 U.S. 844, 863 (1997) (quoting ACLU v. Reno, 929 F. Supp. 824, 883 (E.D. Penn. 1996)).

[18] *See* Susan P. Crawford, *The Internet and the Project of Communications Law*, 55 UCLA L. Rev. 359, 362, 363 n.12 (2007).

[19] *See generally* Pruneyard Shopping Ctr. v. Robins, 447 U.S. 74 (1980) (affirming that a state, through its own constitutional free speech protections, can prohibit a privately-owned space from suppressing peaceful expressive activity).

[20] *See* J. Gregory Sidak, *What Is the Network Neutrality Debate Really About?*, 1 INT'L J. OF COMM. 377, 380 (2007).

[21] WIRELINE COMPETITION BUREAU: INTERNET ACCESS SERVICES: STATUS AS OF DECEMBER 31, 2009, FCC 7, *available at* http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db1208/DOC-303405A1.pdf (indicating that 44 percent of U.S. households have a choice of two broadband service providers and 7 percent have a choice of only one broadband service provider).

some share of the available rents from the goods or services carried on their platform."[22] Without regulatory oversight, the monopolist (or duopolist) will charge supra-competitive prices to end-users, raising the cost of internet service.[23] On the consumer end, paying supra-competitive prices for internet service is, on its face, more of an antitrust harm than a net neutrality violation. This becomes a neutrality issue, however, when ISPs, which have a de facto monopoly (a "terminating access monopoly") over each end-user, charge supra-competitive prices to *websites*, *services* and *applications*, particularly bandwidth-intensive ones.[24] This is a form of discrimination against certain content providers that may chill the growth of those products and services.

The second alleged market failure, in the content, services, and application market, concerns the vertical integration of these products with ISPs themselves, giving providers the incentive to prioritize their own integrated products over their competitors. This is known as the "next Google" argument, since it envisions a "pair of entrepreneurs who would make the next Google," but are deterred by the threat that the incumbent Google will join with a service provider to obtain prioritized service.[25] The net neutrality concerns in this respect are essentially the same as the concerns over protecting innovation.

B. TYPES OF DISCRIMINATION AND REASONABLE NETWORK MANAGEMENT

Though this Note has referred to network discrimination as any ISP practice that "violates" one of these net neutrality principles, there are in fact a number of categorically different discriminatory practices. Edward Felten, now Chief Technologist of the Federal Trade Commission, offers a framework that sorts network discrimination into four useful categories: minimal, non-minimal, minimal delay, and non-minimal delay discrimination.[26] These categories help to distinguish between more and less harmful practices on a theoretical level, and they will provide the basis for developing an operational definition of network neutrality. In particular, understanding how to categorize various forms of network discrimination is essential in determining what constitutes "reasonable network management."

[22] Weiser, *supra* note 5, at 302.

[23] *See id.*

[24] *See id.* at 307.

[25] Sidak, *supra* note 20, at 383.

[26] *See* Edward W. Felten, *Nuts and Bolts of Network Neutrality* 3–5 (July 6, 2006), http://itpolicy.princeton.edu/pub/neutrality.pdf.

Network administrators face the challenge of dealing with the "bursty" nature of internet traffic. Internet traffic patterns are characterized by periods of low activity followed by sudden "bursts" in transmissions.[27] During surges, internet servers may become overwhelmed and may be forced to drop a certain amount of network traffic because they reach their capacity to process incoming data. Discarding transmissions only when it is an absolute technical necessity is known as "minimal discrimination."[28] In contrast, discarding internet traffic for any other reason is known as "non-minimal discrimination."[29] When a server does not drop but merely delays the transmission, this is known as "delay discrimination," and delay discrimination can also be "minimal" (required by a server's capacity constraints) or "non-minimal" (delayed for any other reason).[30] To distinguish minimal from non-minimal discrimination, therefore, is to ask a purely technical question: "Is this discrimination a technical necessity?"

Practices that are technically necessary to prevent an ISP's network from failing during traffic surges ("minimal discrimination" and "minimal delay discrimination") should always be considered "reasonable network management." Even if such discrimination temporarily violates a net neutrality principle, it would be far worse if the network failed entirely during surges in traffic. Thus, the concept of reasonable network management is an important one because it bridges the theoretical definition of net neutrality with the technical reality that network discrimination is justified at certain times. Reasonable network management is not a part of a strictly theoretical definition of net neutrality that contemplates the Internet as a completely neutral, e2e network, because the exclusion allows discrimination that is either justified for technical reasons, imperceptible to the end-user, or sometimes even requested by the end-user.

Although there should be a bright-line rule defining minimal discrimination as reasonable network management, non-minimal discrimination is not so easily defined as reasonable or unreasonable. Some forms of non-minimal discrimination, particularly small amounts of delay discrimination, may not be noticeable to the end-user and therefore may not harm any of the net neutrality principles in any substantial way. Moreover, some users may want their ISPs to prioritize certain traffic. Consumers may prefer that their ISPs guarantee a higher quality of service (QoS) for certain

[27] *See id.* at 4.

[28] *Id.* at 3.

[29] *Id.*

[30] *Id.* at 3–4.

applications, such as online video, at the expense of slower speeds for other web content.[31] In allowing broadband service providers to deviate from a strict application of the nondiscrimination principle, reasonable network management accounts for a variety of acceptable discriminatory practices and is part of the FCC's net neutrality language discussed in the following section.

II. LEGAL NET NEUTRALITY

Because net neutrality principles fail to account for the technical realities of the Internet, it is important, as Tim Wu encourages, "to differentiate sharply between the *principle* of network neutrality and a network neutrality *law*."[32] The history of the FCC's enforcement of net neutrality will help to develop a rough outline of reasonable network management, as its two major enforcement actions both involved clear cases of unreasonable practices.

A. *MADISON RIVER*

The FCC first enforced net neutrality through a 2005 consent decree involving Madison River Communications, LLC, a North Carolina-based digital subscriber line (DSL) broadband ISP and telephone service provider.[33] Vonage, an early Voice over Internet Protocol (VoIP) provider, complained that Madison River was blocking Vonage's application, which allows users to place calls over the Internet.[34] At that time, Madison River served over 180,000 subscribers with telephone service, making Vonage's advance into the voice market a potential threat.[35] Vonage alleged that Madison River persistently blocked VoIP services not just during bursts in network traffic, but at all times.[36] If this allegation is accurate, Madison River's non-minimal blocking represents a clear case of unreasonable network management. It violated all of the net neutrality principles by chilling innovation and

[31] *See* Wu, *supra* note 13, at 76–77.

[32] NETWORK NEUTRALITY FAQ, http://timwu.org/network_neutrality.html (last visted Feb. 14, 2011).

[33] Madison River Commc'ns, LLC, 20 FCC Rcd. 4295 (2005).

[34] *Id.* at 4297; *see also* Ben Charny, *Vonage Says Broadband Provider Blocks Its Calls*, CNET.COM (Feb. 14, 2005), http://news.cnet.com/Vonage-says-broadband-provider-blocks-its-calls/2100-7352_3-5576234.html.

[35] *See* Declan McCullagh, *Telco Agrees to Stop Blocking VoIP Calls*, CNET.COM (Mar. 3, 2005), http://news.cnet.com/Telco-agrees-to-stop-blocking-VoIP-calls/2100-7352_3-5598633.html.

[36] *See* Charny, *supra* note 34.

restraining competition in the VoIP market, without being transparent about its practices.[37] The FCC's investigation ended early with a settlement in which Madison River agreed to cease blocking users from using VoIP applications and to pay a fine.[38]

In the *Madison River* settlement, the FCC enforced net neutrality principles during a time when DSL broadband ISPs were regulated as "telecommunications services" under Title II of the Communications Act of 1934 (as amended by the Telecommunications Act of 1996).[39] Title II imposes a number of common carrier duties on telecommunications services, such as reasonable rates (§ 201), non discrimination (§ 202), and unbundling and interconnection obligations (§§ 251, 252).[40] Title I of the Communications Act, by contrast, applies to "information services" and contains no specific duties for carriers.[41] Rather, it grants the FCC the authority to "perform any and all acts, make such rules and regulations, and issue such orders, not inconsistent with this chapter, as may be necessary in the execution of its functions."[42] Although the FCC reclassified *cable* broadband internet providers under Title I three years prior to *Madison River*, the agency left DSL providers under Title II as common carriers.[43] Shortly after *Madison River*, however, the FCC reclassified DSL broadband ISPs under Title I.[44]

B. *BRAND X*

In *National Cable & Telecommunications Ass'n v. Brand X Internet Services*, the Supreme Court decided a challenge to the reclassification of cable internet providers under Title I.[45] This case is important because it contains the dicta

[37] Since Madison River blocked essentially the use of a VoIP application, its discriminatory practice was less aimed at suppressing a particular perspective, though it certainly blocked the free transmission of speech generally.

[38] *Madison River*, 20 FCC Rcd. at 4297.

[39] 47 U.S.C. § 201 (2006).

[40] 47 U.S.C. §§ 201, 202, 251, 252 (2006).

[41] 47 U.S.C. § 154(i) (2006).

[42] *Id.*

[43] Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities, 17 FCC Rcd. 4798, 4802–03 (2002) [hereinafter *Cable Order*].

[44] Appropriate Framework for Broadband Access to the Internet Over Wireline Facilities, 20 FCC Rcd. 14853 (2005) [hereinafter *DSL Order*].

[45] Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967 (2005).

language upon which the FCC built its jurisdictional foundation to enforce net neutrality after cable and DSL broadband deregulation.[46]

In the initial administrative action, the FCC issued an order that re-categorized cable broadband Internet as an "information service" (one that transforms or processes the communication) instead of a "telecommunications service" (one that does not change the form or content of the communication).[47] The FCC's re-categorization effectively deregulated cable broadband. Applying the deferential test developed in *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.* to evaluate an agency's authority to interpret statutory ambiguities (here, surrounding the terms "telecommunications service" and "information service"), the Supreme Court held that the reclassification was within the FCC's jurisdiction.[48]

While affirming the FCC's decision to move cable broadband Internet out of Title II regulation, Justice Thomas, writing the majority opinion, also commented on the FCC's Title I authority. Comparing "telecommunications services" to "information services," Justice Thomas wrote: "Information-service providers . . . are not subject to mandatory common-carrier regulation under Title II, *though the Commission has jurisdiction to impose additional regulatory obligations under its Title I ancillary jurisdiction to regulate interstate and foreign communications*."[49] This language goes beyond the holding in *Brand X*, as the Court was only reviewing whether the FCC had the authority to resolve the cable broadband classification ambiguity; the Court was not interpreting Title I. However, this language became the jurisdictional foundation of the FCC's authority to enforce net neutrality after it deregulated both cable and DSL broadband.[50]

C. "INTERNET POLICY STATEMENT"

With the encouragement of *Brand X*, the FCC embarked upon the enforcement of net neutrality principles with the publication of its "Internet Policy Statement" in 2005.[51] Citing the key dicta language from *Brand X*, the

[46] *See id.* at 976 ("[T]he Commission has jurisdiction to impose additional regulatory obligations under its Title I ancillary jurisdiction to regulate interstate and foreign communications.").

[47] *Cable Order*, *supra* note 43, at 4802–03.

[48] *Brand X*, 545 U.S. at 1002–03 (citing Chevron, U.S.A. v. Natural Res. Def. Council, Inc., 467 U.S. 837, 865–66 (1984)).

[49] *Id.* at 976 (emphasis added).

[50] *See* Appropriate Framework for Broadband Access to the Internet Over Wireline Facilities, 20 FCC Rcd. 14986, 14988 (2005) [hereinafter *Internet Policy Statement*].

[51] *See id.*

FCC concluded that it had the "jurisdiction necessary to ensure that providers of telecommunications for Internet access or Internet Protocol-enabled (IP-enabled) services are operated in a neutral manner."[52] The FCC adopted four principles to encourage broadband deployment and preserve the open and interconnected nature of the public Internet. Under these principles, consumers are entitled to:

> [1] . . . access the lawful Internet content of their choice;
>
> [2] . . . run applications and use services of their choice, subject to the needs of law enforcement;
>
> [3] . . . connect their choice of legal devices that do not harm the network; [and]
>
> [4] . . . competition among network providers, application and service providers, and content providers.[53]

Importantly, the FCC made clear that these principles are also subject to "reasonable network management."[54] Principles 1, 2, and 3—each a form of nondiscrimination rule—embody the theoretical net neutrality principles of protecting innovation and free speech in the respective markets of internet content, services, applications, and devices. Principle 4 articulates the net neutrality competition principle, and it notably reaches both the market for broadband service providers and the market for applications, services, and content. Ultimately, however, the adoption of these principles in a policy statement rather than through a rule-making or through a grant of authority by Congress undermined the FCC's ability to enforce net neutrality.

D. *COMCAST V. FCC*

In April 2010, five years after the adoption of the "Internet Policy Statement," the D.C. Circuit decided *Comcast Corp. v. Federal Communications Commission*, which held that the FCC did not have jurisdiction over broadband service providers to enforce neutrality principles.[55] The case involved Comcast's non-minimal blocking of peer-2-peer (p2p) file networking applications. The holding, however, did not reach the FCC's technical argument against Comcast's unreasonable network management practice. Rather, *Comcast* reflects the application of the D.C. Circuit's jurisdictional doctrine developed in earlier cases to determine the boundaries of the FCC's Title I authority.

[52] *Id.*

[53] *Id.*

[54] *See id.* at 14988 n.15.

[55] Comcast Corp. v. FCC, 600 F.3d 642, 661 (D.C. Cir. 2010).

1. *Facts and Procedural History*

In 2007, several subscribers to Comcast's high-speed internet service noticed that the company was slowing or blocking traffic through peer-to-peer networking applications, including those relying on BitTorrent.[56] That same year, the Associated Press conducted nationwide tests confirming that Comcast "actively interfere[d] with attempts by some of its high-speed Internet subscribers to share files online."[57] In response, two non-profit organizations, Free Press and Public Knowledge, filed a complaint with the FCC alleging that Comcast violated the FCC's "Internet Policy Statement" by interfering with users' internet access.[58] After first denying any responsibility for the disrupted peer-to-peer access,[59] Comcast later acknowledged and defended its practice as necessary for reasonable management of its network's limited capacity.[60]

2. *Comcast's Network Management Practices*

After a period of public comment, the FCC issued an order finding that Comcast's practice "unduly squelches the dynamic benefits of an open and accessible Internet and does not constitute reasonable network management."[61] When Comcast detected that BitTorrent users were attempting to share files, Comcast issued a "reset packet" that would terminate the connection.[62] Because the packet looked like it came from the other user's computer, Comcast was "falsifying network traffic" through a process that was very difficult to circumvent.[63] The FCC observed that Comcast was determining how to route its connections (or, more precisely, whether to *terminate* some of its connections) based "not on their destinations but on their contents."[64] Thus, as the FCC noted, Comcast was "open[ing] its customers' mail because it want[ed] to deliver mail not based on the address

[56] *Id.* at 644.

[57] Peter Svensson, *Comcast Blocks Some Internet Traffic*, WASH. POST (Oct. 19, 2007), http://www.washingtonpost.com/wp-dyn/content/article/2007/10/19/AR2007101900842.html.

[58] Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028 (2008) [hereinafter *Complaint Against Comcast*].

[59] Marguerite Reardon, *Comcast Denies Monkeying with BitTorrent Traffic*, CNET.COM (Aug. 21, 2007), http://www.news.com/8301-10784_3-9763901-7.html.

[60] *Comcast*, 600 F.3d at 645.

[61] *Complaint Against Comcast*, *supra* note 58, at 13028.

[62] *Id.* at 13031.

[63] *Id.*

[64] *Id.* at 13051.

or type of stamp on the envelope but on the type of letter contained therein."[65] Moreover, the majority of experts the FCC consulted found that inserting a "reset packet" into consumer traffic did not constitute reasonable network management and did not conform to any standard practice in network engineering.[66] As a result, the order required Comcast to disclose its network management practices, construct a plan to amend its discriminatory practice, and disclose its new practices to the public.[67]

3. D.C. Circuit's Analysis

After complying with the order, Comcast appealed the FCC's decision on jurisdictional, procedural, and Due Process grounds. In April 2010, the D.C. Circuit ruled that the FCC lacked sufficient statutorily-mandated responsibility and vacated the FCC's order on jurisdictional grounds alone.[68]

As the FCC had no express statutory authority to regulate Comcast's purportedly unreasonable network management, it relied on Title I of the Communications Act, which states in relevant part that the FCC may "perform any and all acts, make such rules and regulations, and issue such orders, not inconsistent with this chapter, as may be necessary in the executions of its functions."[69] This section has come to be known as the FCC's "ancillary jurisdiction."[70] Prior to *Comcast*, the D.C. Circuit held that the FCC "may exercise ancillary jurisdiction only when two conditions are satisfied: (1) the Commission's general jurisdictional grant under Title I [of the Communications Act] covers the regulated subject and (2) the regulations are reasonably ancillary to the Commission's effective performance of its statutorily mandated responsibilities."[71]

Before applying this test to the facts of *Comcast*, the D.C. Circuit first addressed the FCC's two threshold arguments, both of which asserted that the normal jurisdictional test should not apply. First, the FCC argued that Comcast should be judicially estopped from challenging the FCC's jurisdiction since Comcast had acknowledged the FCC's jurisdiction over

[65] *Id.*

[66] *Id.* at 13055.

[67] *Id.* at 13060.

[68] Comcast Corp. v. FCC, 600 F.3d 642, 661 (D.C. Cir. 2010).

[69] 47 U.S.C. § 154(i) (2006).

[70] *Comcast*, 600 F.3d at 644.

[71] *Id.* at 646 (quoting Am. Library Ass'n v. FCC, 406 F.3d 689, 691–92 (D.C. Cir. 2005)).

peer-to-peer services in a district court case two years earlier.[72] The D.C. Circuit disagreed, finding that Comcast's admission in the prior case applied only to the first part of the jurisdictional test—the "regulated subject" element of the two-part test—and did not preclude Comcast from disputing the FCC's jurisdiction for other reasons.[73] Second, the FCC argued that the Supreme Court had already decided the jurisdictional question in *Brand X*.[74] Acknowledging that this language from *Brand X* is technically dicta, the D.C. Circuit also dismissed this argument by examining a line of Supreme Court decisions directly defining the FCC's ancillary jurisdiction.[75] Based upon those cases, the D.C. Circuit concluded that *Brand X* does nothing to eliminate the requirement that ancillary authority must be independently justified.[76]

Comcast conceded, and the D.C. Circuit accepted, that the FCC's action satisfied the first element of the two-part jurisdictional test because Comcast's internet service qualified as "interstate and foreign communication by wire" as that term is used in Title I.[77] Turning to the second element—the "statutorily mandated responsibilities" element—the D.C. Circuit found that none of the FCC's cited provisions of the Communications Act delegated sufficient regulatory authority over broadband Internet.[78] The court divided these provisions into two general categories: those that articulate only congressional policy and those that potentially delegate regulatory authority.[79] Congressional policy statements alone, the court said, "cannot provide the basis for the Commission's exercise of ancillary authority," since it is an "axiomatic principle" that "administrative agencies may [act] only pursuant to authority delegated to them by Congress."[80] Thus, the sections of the Communications Act relied upon by the FCC that express only policy could not support the FCC's jurisdiction to regulate Comcast's network

[72] *Id.* at 647 (citing Hart v. Comcast of Alameda, No. 07-6350, 2008 WL 2610787 (N.D. Cal. June 25, 2008)).

[73] *Id.* at 648.

[74] *Comcast*, 600 F.3d at 649; *see* Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967, 996–97 (2005).

[75] *Comcast*, 600 F.3d at 650–51 (citing United States v. Sw. Cable Co., 392 U.S. 157 (1968); United States v. Midwest Video Corp., 406 U.S. 649 (1972)).

[76] *Id.* at 651.

[77] *Id.* at 646.

[78] *Id.* at 661.

[79] *Id.* at 651. The court designated sections 230(b) and 1 of the Communications Act statements of policy, and sections 706, 256, 257, 201, and 623 plausible delegations of regulatory authority. *Id.* at 651, 658–61.

[80] *Id.* at 654 (quoting Am. Library Ass'n v. FCC, 406 F.3d 689, 691 (D.C. Cir. 2005)).

management practices.[81] Although the remaining provisions upon which the FCC relied could have "arguably delegate[d] regulatory authority to the Commission," the court found that each failed to deliver a specific delegation of jurisdiction over broadband Internet.[82] None of the provisions, therefore, could provide the FCC with the appropriate, independently justified authority required by the two-part jurisdictional test. Thus, the court overturned the FCC's order.[83]

E. "OPEN INTERNET RULES"

About eight months after *Comcast* invalidated the FCC's jurisdiction over broadband Internet, the FCC responded with a reassertion of authority and a new set of net neutrality rules in the "Open Internet Rules."[84] In these new rules, the FCC adopted a new jurisdictional theory by relying heavily on section 706 of the Telecommunications Act of 1996, which directs the FCC to "encourage the deployment on a reasonable and timely basis of advanced telecommunications capability to all Americans."[85] "Advanced telecommunications capability," as defined in the 1996 Act, includes broadband Internet.[86] Thus, the FCC argued that section 706 provides authority for the net neutrality regulation because the "Open Internet Rules" will encourage broadband Internet deployment.[87]

In order to adopt this theory, the FCC had to reconcile it with the *Comcast* court's earlier interpretation of section 706. In *Comcast*, the D.C. Circuit considered section 706 to be a provision that could "at least arguably be read to delegate regulatory authority."[88] However, because the FCC had acknowledged that section 706 "does not constitute an independent grant of authority" in a separate, earlier order (the "Advanced Services Order"), the FCC could no longer use section 706 as a basis for their jurisdiction.[89] Responding to this holding in *Comcast*, the FCC asserted a different reading of section 706 and the Advanced Services Order in the "Open Internet Rules." Specifically, the FCC clarified that the Advanced Services Order only meant that section 706 conferred no authority upon the FCC "over and

[81] *Id.*

[82] *Id.* at 659–61.

[83] *Id.* at 661.

[84] *See generally Open Internet Rules*, *supra* note 3.

[85] 47 U.S.C. § 1302(a) (2009); *see Open Internet Rules*, *supra* note 3, ¶ 116.

[86] § 1302(d)(1); *see Open Internet Rules*, *supra* note 3, ¶ 117.

[87] § 1302(d)(1); *see Open Internet Rules*, *supra* note 3, ¶ 117.

[88] Comcast Corp. v. FCC, 600 F.3d 642, 658 (D.C. Cir. 2010).

[89] *Id.* (quoting *Deployment of Wireline Servs. Offering Advanced Telecomms. Capability*, 13 FCC Rcd. 24012, 24047 (1998) [hereinafter *Advanced Services Order*]).

above what it otherwise possessed" (in other words, "independent" of what it already had).[90] Consequently, the FCC argued, section 706 still "authorizes the [FCC] to address practices, such as blocking VoIP communications, degrading or raising the cost of online video, or denying end users material information about their broadband service, that have the potential to stifle overall investment in Internet infrastructure and limit competition in telecommunications markets."[91] Two wireless providers, Verizon and Metro PCS, have already filed complaints challenging the "Open Internet Rules."[92]

Aside from a new jurisdictional basis, the FCC also adopted three net neutrality rules. These include a rule for ISP transparency and rules against blocking and discrimination as follows:

> [Transparency rule:] A person engaged in the provision of broadband Internet access service shall publicly disclose accurate information regarding the network management practices, performance, and commercial terms of its broadband Internet access services sufficient for consumers to make informed choices regarding use of such services and for content, application, service, and device providers to develop, market, and maintain Internet offerings.[93]
>
> [No blocking rule:] A person engaged in the provision of fixed broadband Internet access service, insofar as such person is so engaged, shall not block lawful content, applications, services, or non-harmful devices, subject to reasonable network management.[94]
>
> [Non discrimination rule:] A person engaged in the provision of fixed broadband Internet access service, insofar as such person is so engaged, shall not unreasonably discriminate in transmitting lawful network traffic over a consumer's broadband Internet access service. Reasonable network management shall not constitute unreasonable discrimination.[95]

By incorporating a transparency principle, these rules represent a significant advancement over the "Internet Policy Statement" for reasons discussed in Section I.C, *supra*. The "Open Internet Rules" further define a network management practice as reasonable "if it is appropriate and tailored to achieving a legitimate network management purpose, taking into account the

[90] *Open Internet Rules*, *supra* note 3, ¶ 118.

[91] *Id.* ¶ 120.

[92] *See FCC Seeks to Dismiss Net Neutrality Challenges*, HUFFINGTON POST, Jan. 28, 2011, http://www.huffingtonpost.com/2011/01/28/fcc-net-neutrality-news_n_815626.html.

[93] *Open Internet Rules*, *supra* note 3, ¶ 54.

[94] *Id.* ¶ 63.

[95] *Id.* ¶ 68.

particular network architecture and technology of the broadband Internet access service."[96] The FCC offered a few examples of legitimate network management practices, which include "ensuring network security and integrity . . ., addressing traffic that is unwanted by end users . . ., and reducing or mitigating the effects of congestion on the network."[97] To the extent that these examples clarify the concept of "reasonable network management," they do so only in a generalized way. The FCC acknowledged that they will "develop the scope of reasonable network management on a case-by-case basis, as complaints about broadband providers' actual practices arise."[98] In the end, therefore, while the adoption of the "Open Internet Rules" introduces an important transparency rule, it does little to develop the concept of reasonable network management.

F. LEGAL NET NEUTRALITY AND REASONABLE NETWORK MANAGEMENT

By failing to adequately elaborate criteria for reasonable network management in the "Open Internet Rules," the FCC left the concept wide-open to interpretation by future litigants. This is particularly so given the *Madison River* and *Comcast* decisions, which define reasonable network management in only the bluntest way: both cases involved the persistent, non-minimal blocking of internet applications, which could not be justified by a continuing technical necessity. The FCC and the D.C. Circuit, therefore, offer little guidance in analyzing more subtle forms of discrimination, such as the delay discrimination that occurs when an ISP does not block but merely delays a transmission.[99] The next Part argues that reasonable network management is best defined through a technical analysis of the Internet because determining whether a discriminatory practice is "minimal" (and thus reasonable) should be rooted in whether the practice is a "technical necessity" for broadband network administrators.

III. TECHNICAL NET NEUTRALITY

The decentralized architecture of the Internet—a network of networks—requires ISPs to enter into service provider agreements for exchanging traffic. These agreements dictate the cost of sending traffic. Using them, ISPs

[96] *Id.* ¶ 82.

[97] *Id.*

[98] *Id.* ¶ 83.

[99] *See* Types of Discrimination and Reasonable Network Management, *supra* Section I.B.

often manipulate protocols to route transmissions along the lowest-cost paths. The following Part examines the physical architecture of the Internet, interconnection agreements among ISPs, and the technical protocols that define connections and routing—revealing a number of potential ways that network owners can discriminate among content, services, and applications. Understanding how network administrators discriminate on a technical level and why they would decide to deviate from full neutrality will help in classifying ISP discrimination practices as reasonable or unreasonable.

A. PHYSICAL ARCHITECTURE AND SERVICE PROVIDER AGREEMENTS

At its most basic level, the Internet is divided into a three-level hierarchy of "last mile" ISPs, regional ISPs, and internet backbones. This tripartite structure tracks the original hierarchy of the early Internet, which went online as the NSFNET backbone in 1986 to provide universities nationwide access to federally funded supercomputers located at a small number of universities.[100] Originally, ISPs entered into two general forms of interconnection agreements: *transit agreements* and *peering agreements*. In a transit agreement, one ISP agrees to deliver internet traffic from another ISP for a fee, often because there is an unequal exchange of traffic. In peering agreements, by contrast, ISPs agree to exchange roughly equal traffic free of charge.[101] Internet backbones originally entered into settlement-free peering agreements based on an approximate determination that their packet exchange was symmetrical. Because of the high transaction costs associated with precise measurement of the exchange, service providers during the early days of the Internet still favored free peering relationships even when the exchange was not completely equal.[102]

Today, there are still the three levels of service providers. However, these providers no longer connect exclusively through one-to-one relationships. This is because a hierarchical Internet consisting of one-to-one relationships among the three levels of service providers made each network participant completely dependent upon the level above them—providing internet

[100] Christopher S. Yoo, *Innovations in the Internet's Architecture That Challenge the Status Quo*, 8 J. ON TELECOMM. & HIGH TECH L. 79, 81 (2010). For a discussion of the parallels between the divestiture arrangement with the long distance telephone companies and the three-level hierarchy of the Internet, see Juan D. Rogers, *Internetworking and the Politics of Science: NSFNET in Internet History*, 14 INFO. SOC'Y 213, 219 (1998).

[101] Stanley Besen et al., *Advances in Routing Technologies and Internet Peering Agreements*, 91 AM. ECON. REV. 292, 292 (2001).

[102] Peyman Faratin et al., *The Growing Complexity of Internet Interconnection*, 72 COMMC'NS & STRATEGIES 51, 52–57 (2008).

backbones at the top of the hierarchy with the potential power to charge monopoly rents.[103] As a result, service providers entered into new arrangements, through *secondary peering* and *multihoming*, in which lower-level ISPs could connect to more than just the ISP directly above them. Regional ISPs, for example, no longer needed to connect to an internet backbone through a transit agreement; they could also connect to another regional ISP for free on the basis of roughly equal exchange. This process is known as *secondary peering*.[104] Regional ISPs could also connect to more than one internet backbone, which is known as *multihoming*.[105] As a result, while service providers still enter into peering and transit agreements, those arrangements now represent just two among a variety of contractual arrangements.[106]

In addition, ISPs now draft increasingly sophisticated peering and transit agreements. *Paid peering*, for example, resembles normal peering in almost every respect, except that one network pays the other network even when the exchange of traffic is roughly the same. These more sophisticated agreements reflect the fact that while the traffic exchange may be equal, the cost of maintaining the networks' respective infrastructures may be unequal.[107] ISPs serving a smaller number of large internet content websites (known as "content networks") have lower costs in maintaining their infrastructure than ISPs serving home users ("eyeball networks"), since residential neighborhoods require more equipment investment (such as wiring) and maintenance than commercial areas.[108] These interconnection agreements create the economic incentives for ISPs to route internet traffic along the lowest-cost paths, which can sometimes have a discriminatory effect on certain types of content, applications, and services.

B. THE PROTOCOL LAYERS OF THE INTERNET

Interconnection agreements are realized on a technical level through network protocols. As service provider agreements provide strong economic incentives for ISPs to discriminate in ways that keep transit costs low, network administrators can discriminate by manipulating certain protocols in a variety of minimal and non-minimal ways. For example, a network administrator can send a signal to both ends of a connection that has the effect of resetting the connection and effectively blocking traffic between

[103] *See* Yoo, *supra* note 100, at 83.

[104] *Id.* at 86.

[105] *Id.*

[106] *See id.* at 61.

[107] *See id.* at 96.

[108] *See id.*

two end users.[109] Administrators can also prioritize traffic based on traffic class designations, adjust routing tables to send traffic along faster or slower routes, adjust routes based on cost, and block sending or receiving traffic from certain networks altogether.[110] These practices represent some (though certainly not all) of the network administrator's "tools" for network discrimination. This Article proposes, *infra* Section III.C, that determining whether these practices constitute reasonable or unreasonable network management should involve two inquiries. First, are the practices technical necessities? If they are not, then second, do they violate any of the theoretical net neutrality principles of innovation, free speech, and competition in any serious way?

Unlike the Internet's physical infrastructure, which is largely privately owned, protocols are, for the most part, community assets. As a network comprised of smaller networks, the Internet is not governed by any one entity. Rather, it is advised by a voluntary group of users known as the Internet Engineering Task Force (IETF). Through online working groups, the IETF produces technical and engineering documents to "influence the way people design, use, and manage the Internet."[111] Among the different types of documents it produces, the IETF circulates memoranda describing protocol standards known as Requests for Comments (RFCs). In this way, the Internet is "governed" by individual networks' voluntary adherence to a complex set of protocols defining the format and order of messages sent and received by devices on the network.[112]

The Internet's complex protocols can be understood as a system of *layers*—a conceptual aid that allows engineers to envision the transmission of a message from one computer to another as a series of wrappings and unwrappings of the message. In a typical exchange between two end-users, a message is sent from an application, such as an email program, using a protocol in the *application layer*. It is then wrapped according to a protocol that defines how it will be transported in the *transport layer*. The message is then further encapsulated according to a protocol that will determine how the

[109] *See*, *e.g.,* Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028, 13031 (2008).

[110] *See, e.g.,* Scott Bradner & Allison Mankin, *The Recommendation for the IP Next Generation Protocol*, RFC 1752, at 3 (Jan. 1995), http://www.rfc-editor.org/rfc/pdfrfc/rfc1752.txt.pdf.

[111] Harald Tveit Alvestrand, *A Mission Statement for the IETF*, RFC 3935, at 1 (Oct. 2004), http://www.ietf.org/rfc/rfc3935.txt.

[112] JAMES F. KUROSE & KEITH W. ROSS, COMPUTER NETWORKING: A TOP-DOWN APPROACH FEATURING THE INTERNET 8 (3d ed. 2005).

message will move from one host to another in the *network layer*. Finally, the message is wrapped according to a protocol in the *link layer* based on whether it is traveling on an Ethernet network or through some other system. The first two protocols developed on the Internet were the *Transmission Control Protocol (TCP)*, which is a transport layer protocol, and the *Internet Protocol (IP)*, which is a network layer protocol. These two protocols, along with a growing set of other protocols, form the TCP/IP protocol suite on the Internet today, which determines how a message will be transported from one end of a network to another.

This Section will focus on the transport and network layers because ISPs have used these layers to implement discrimination practices. ISPs can, for example, interfere with traffic using transport layer protocols by "posing" as an end-user. The network layer is also critically important because it determines the input and output decisions of individual routers and the global coordination of internet routing.

1. *Transport Layer*

TCP is one of two common transport layer protocols.[113] With TCP, the transport layer establishes a logical connection between, for example, a user's computer and an internet email provider's server.[114] Over a logical connection, an email application running on a server makes a direct connection to software on the user's computer (a web browser, for example), even if the application and software are actually separated by thousands of miles.[115] A digital "handshake" between the two ends establishes the connection and creates a reliable transfer in which TCP ensures that all data is delivered correctly and in order.[116]

The transport layer is generally implemented only at the ends of the network. According to the layered approach to protocols, this means that transport layer protocols (such as TCP) are packaged inside network layer protocols (such as IP) when traveling through the Internet's core. This raises the question: how can an ISP interfere using a protocol layer with which it does not communicate? Revisiting the facts of *Comcast* may be helpful here. In *Comcast*, the ISP blocked peer-to-peer networking applications by sending a message to both sides of a connection such that the message looked like it

[113] The other is the User Datagram Protocol (UDP), defined in Jonathan Postel, *User Datagram Protocol*, RFC 768 (Aug. 1980), http://www.rfc-editor.org/rfc/pdfrfc/rfc768.txt.pdf.

[114] *See* KUROSE & ROSS, *supra* note 112, at 184.

[115] *See id.*

[116] *See id.* at 188.

was being sent by the other end-user to reset the connection.[117] TCP reserves a field, the RST flag bit, in the header of every message to allow one side of the communication to reset the connection.[118] By posing as an end-user and repeatedly sending reset messages, an ISP can effectively block a connection. Jon Peha, former Chief Technologist for the FCC, condemned Comcast's practice, stating that he was "unaware of any technical literature that has proposed that ISPs adopt this particular practice as a way of dealing with congestion."[119] This observation does not preclude the possibility that RST blocking could be used in some minimal way to control congestion during surges in activity, but it certainly suggests that it is unconventional and thus more likely to be indicative of non-minimal discrimination.

2. *Network Layer*

Like reset packet blocking in the transport layer, discriminatory practices in the network layer can be used in minimal and non-minimal ways. Network layer protocols control routing and forwarding on the Internet. Forwarding refers to transfers that take place *within* a router, from the input to the output link. Routing, on the other hand, refers to the process of determining the network-wide path for the data.[120] Every router contains a forwarding table, which tells the router where to output its data based on the address assigned to the incoming data. Routing protocols compute these forwarding tables.[121] Though there are many forwarding and routing protocols, there is one dominant forwarding protocol—Internet Protocol (IP)—and there are three dominant routing protocols.

As a forwarding protocol, IP describes how a single internet router should deal with data inputs and outputs. IP directs a server to attach a "header" to the data it receives from the layer above it (the transport layer). This can be roughly understood as taking a letter, folding it, and putting it in the envelope with a stamp, destination, and return address. The format of the IP protocol header (the "envelope") requires certain categories of information. There are two IP versions—the older IPv4 and the newer IPv6—each with slightly different header formats containing different required categories. Both IPv4 and IPv6, however, have required bits

[117] *See* Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028, 13031 (2008).

[118] *See* KUROSE & ROSS, *supra* note 112, at 254.

[119] Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028, 13055 (2008).

[120] KUROSE & ROSS, *supra* note 112, at 302.

[121] *Id.* at 324.

designating the type of service (TOS) (in IPv4 and IPv6) or "traffic class" (in IPv6), which allow routers to distinguish among different types of data.

The design of this forwarding protocol suggests that priority designation—and thus the ability for ISPs to discriminate—has been a part of the Internet since the beginning.[122] The TOS field has been a part of the IP since RFC 791, which first defined the protocol in 1981.[123] Cisco, for example, uses the first three TOS bits to define different levels of service within its routers.[124] The TOS bits are significant in that they are mandatory and would be an inefficient use of network resources if they were not used.[125] As a result, ISPs have the potential to implement non-minimal discrimination practices into their network management.

The newer version of IP (IPv6) has also increased the potential for ISP discrimination by expanding rather than eliminating the fields dedicated to flagging priority levels for different types of data. IPv6 allocates a separate field, known as "flow," to allow applications to designate data that require special handling, higher quality, or real-time service.[126] Like TOS bits, "flow" designations might be employed for minimal discrimination if used to select higher priority traffic and drop or delay lower priority traffic during congested periods. They might equally be employed for non-minimal discrimination if used to favor a certain application, content, or service at all times.

The network layer's second function is data routing. Broadly speaking, a routing protocol finds a "good" path from the origin of the data to the destination.[127] But what is a "good" path? RFC 1058 describes how routing protocols calculate (and tabulate) the best paths, which are recorded as a set of "hops" between routers:

> [I]n order to define which route is best, we have to have some way of measuring goodness. This is referred to as the "metric." . . . In more complex networks, a metric is chosen to represent the total amount of delay that the message suffers, the cost of sending it, or

[122] *See* Kai Zhu, *Bringing Neutrality to Network Neutrality*, 22 BERKELEY TECH. L.J. 615, 634 (2007) (quoting Information Sciences Institute, *Internet Protocol: Darpa Internet Program Protocol Specification*, RFC 791, at 11 (Sept. 1981), http://www.ietf.org/rfc/rfc0791.txt).

[123] *See generally* Information Sciences Institute, *Internet Protocol: Darpa Internet Program Protocol Specification*, RFC 791 (Sept. 1981), http://www.ietf.org/rfc/rfc0791.txt (defining the Internet Protocol).

[124] KUROSE & ROSS, *supra* note 112, at 326.

[125] *See* Zhu, *supra* note 122, at 634–35, n.135.

[126] *See* Bradner & Mankin, *supra* note 110, at 3.

[127] *See* KUROSE & ROSS, *supra* note 112, at 351.

> some other quantity which may be minimized. The main requirement is that it must be possible to represent the metric as a sum of "costs" for the individual hops.[128]

On the Internet, there are three main routing protocols that roughly correspond to the three hierarchical levels of the Internet's architecture, discussed in Section III.A, *supra*. The *Routing Information Protocol (RIP)* coordinates routing within the networks of the "last mile" providers; the *Open Shortest Path First (OSPF)* protocol manages routing within the regional ISPs; and the *Border Gateway Protocol (BGP)* coordinates routes between regional and "last mile" providers, which often includes routes across internet backbones.[129]

All three routing protocols decide how to route data as a function of a route's cost, but they all compute cost in different ways. In a network using RIP, every "hop" between intermediate devices on the way to a destination costs the same amount by default.[130] This default can be changed to account for differences in cost between individual "hops," but RIP does not allow much freedom to customize the cost metric, nor does it allow for real-time metric updating to account for delays further down the path.[131]

In a network using OSPF, by contrast, the network administrator can configure the individual costs per hop so that the protocol will automatically choose the minimum-cost hop route or avoid certain paths.[132] Quite predictably, the lower a hop costs, the more likely that the network administrator will use that hop to send traffic.[133] OSPF does not generate its routing tables exclusively from its own cost-based algorithms; OSPF derives some of its routing data from external sources, including route calculations by BGP, which, as discussed *infra*, can be set by network administrators.[134]

Finally, in a network using BGP, network administrators' discretion plays an even larger role. As RFC 1164 explains, BGP can be used in response to

[128] C. Hedrick, *Routing Information Protocol*, RFC 1058, at 7 (June 1988), http://www.rfc-editor.org/rfc/pdfrfc/rfc1058.txt.pdf.

[129] RIP, OSPF, and BGP are technically implemented at the application layer, but because they control routing on the Internet, are often associated with the network layer. *See, e.g.*, KUROSE & ROSS, *supra* note 112, at 370–83.

[130] *See* Hedrick, *supra* note 128, at 4; *see also* KUROSE & ROSS, *supra* note 112, at 371.

[131] *See* Hedrick, *supra* note 128, at 4.

[132] *See* John Moy, *OSPF Version 2*, RFC 2178, at 18 (July 1997), http://www.rfc-editor.org/rfc/pdfrfc/rfc2178.txt.pdf; *see also* KUROSE & ROSS, *supra* note 112, at 384.

[133] *See* Moy, *supra* note 132, at 18.

[134] *See id.*

"non-technical" concerns.[135] This is because BGP policies are set by the administrator of the network running BGP (usually an internet backbone), and these administrators can manipulate the selection of paths based on cost, for example, when multiple paths are available.[136] This can result in a wholesale refusal to carry traffic from a particular regional network or simply "favoring" or "disfavoring" traffic from certain networks.[137] Christopher Yoo provides the following illustration of the effects of lower transit costs on routing:

> [A]ssume that an end user is downloading content from both CNN.com and MSNBC.com. Assume further that the end user's regional ISP has a secondary peering relationship with the regional ISP serving CNN.com, but does not have a secondary peering relationship with the regional ISP serving MSNBC.com. The absence of a secondary peering relationship means that traffic from MSNBC.com will have to pay transit charges, while traffic from CNN.com will not. The result is that traffic that is functionally identical will end up paying different amounts.[138]

The fact that traffic to these two functionally identical websites (both fall into the category of "mainstream news") can cost ISPs different amounts incentivizes ISP administrators to employ non-minimal discrimination by slowing traffic going to content or services for which the transit costs are greater. Correspondingly, network administrators may encourage traffic going to content or services for which the transit costs are lower due to the free peering agreement between ISPs. As such, ISPs can encourage users to switch websites by slowing traffic to websites involving more expensive transit costs.

To take this example one step further, as Yoo does, we may also consider a situation in which the same end-user's regional ISP connects to CNN.com *both* through a slower, often-congested secondary peering arrangement and a faster, higher capacity transit agreement.[139] Once again, the end user's regional ISP would have every economic incentive to route traffic through the slower (but free) secondary peering connection. In this scenario, the end-user is provided with a slower connection to CNN.com that costs the regional ISP nothing in transit fees. This end user also retains a connection to

[135] Jeffrey C. Honig et al., *Application of the Border Gateway Protocol in the Internet*, RFC 1164, at 6 (June 1990), http://www.rfc-editor.org/rfc/pdfrfc/rfc1164.txt.pdf.

[136] *See id.*

[137] *Id.*

[138] Yoo, *supra* note 100, at 87.

[139] *See id.*

MSNBC.com, but since the regional ISP has to pay transit fees, it remains in the ISP's interest to encourage the user to choose CNN.com for his news.

By manipulating routing protocols, network administrators can also route traffic to *overlay networks*, which are physical additions to the Internet in the form of servers deployed widely across the Internet.[140] Content Distribution Networks (CDNs) are some of the most popular overlays on the Internet today. They consist of servers distributed geographically across the Internet that retain a cache of the most frequently demanded content and services from publishers and providers. CDNs work by shortening the physical distance between the end-user and the content, enabling CDNs to optimize content delivery based on different criteria, including faster response time or optimal bandwidth costs.[141] In 2007, Akami, one of the world's largest CDNs, was estimated to manage approximately 20,000 servers in 70 countries and to deliver approximately 15 percent of the world's internet content.[142] Because CDNs are networks separate from the three-tier system, they are outside the minimal versus non-minimal classification of discrimination that this Note adopts to analyze net neutrality. However, because CDNs can also have the effect of prioritizing certain routing, they also constitute a potentially discriminatory routing practice.

[140] *See* Dave Clark et al., *Overlay Networks and the Future of the Internet*, 63 COMMC'NS & STRATEGIES 1, 3–4 (2006).

[141] *See* KUROSE & ROSS, *supra* note 112, at 610.

[142] Peyman Faratin, *Economics of Overlay Networks: An Industrial Organization Perspective on Network Economics* 2, http://netecon.seas.harvard.edu/NetEcon07/Papers/faratin_07.pdf (last visited Dec. 21, 2010).

C. TECHNICAL NET NEUTRALITY AND REASONABLE NETWORK MANAGEMENT

Figure 1: Defining Reasonable Network Management

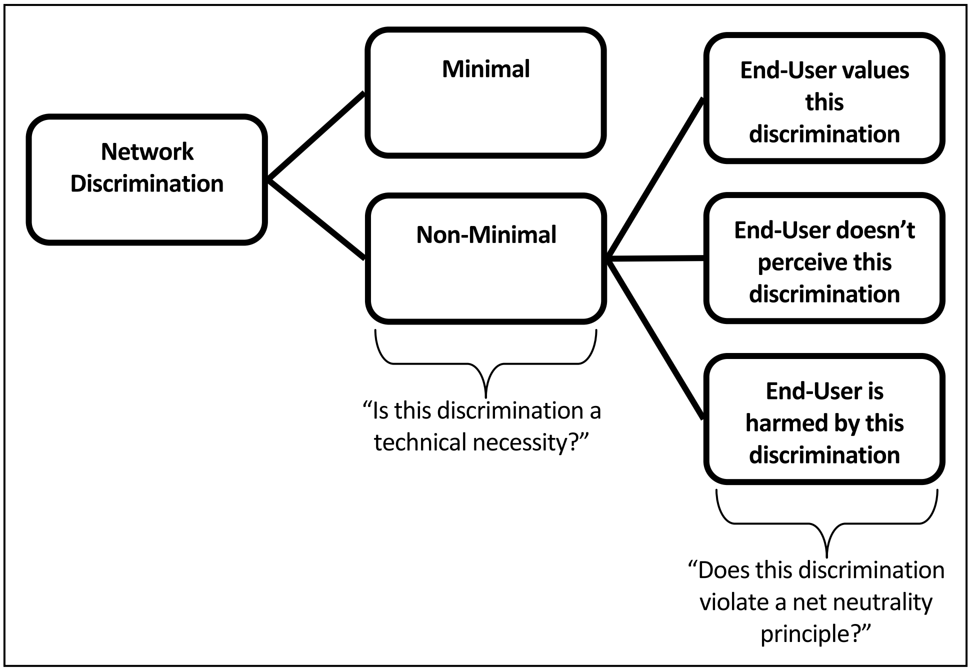
Network Discrimination
Minimal
Non-Minimal
End-User values this discrimination
End-User doesn't perceive this discrimination
End-User is harmed by this discrimination
"Is this discrimination a technical necessity?"
"Does this discrimination violate a net neutrality principle?"

In order to evaluate whether these potentially discriminatory practices fall into the category of reasonable or unreasonable network management, one must answer two questions (shown in Figure 1): is this discrimination a technical necessity? If not, does this discrimination violate a net neutrality principle? As explained, Section III.B, *supra*, the same protocol-level tools available to network administrators can be used for both minimal and non-minimal discrimination. These questions are impossible to answer, therefore, if ISPs are not transparent about when and why they discriminate on a technical level. Indeed, as Kevin Martin, then Chairman of the FCC, observed in his order reviewing Comcast's network management practices: "A hallmark of whether something is reasonable is whether a provider is willing to disclose to its customers what it is doing."[143]

The first inquiry classifies the discrimination as either "minimal" or "non-minimal." Minimal discrimination should always be considered reasonable network management, because it is an absolute technical necessity

[143] Formal Complaint of Free Press and Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028, 13059 (2008).

to prevent network failures during bursts in internet traffic. If minimal discrimination happens to temporarily violate one of the net neutrality principles, this is comparatively less harmful to consumers than total network failure.

Non-minimal discrimination, by contrast, requires a second inquiry into how the particular form of discrimination affects net neutrality principles. This inquiry can conclude in three possible ways:

First, end-users may value the discrimination. Using IP TOS or "flow" designations, an ISP could guarantee that certain types of traffic—such as video content or VoIP calls—will be prioritized on the network at all times. An ISP might offer a similar guarantee by routing certain content through more expensive, but less congested, paths using a manipulation of the BGP (undoubtedly passing along the premium cost to the consumer). This kind of quality of service (QoS) guarantee at first appears to violate net neutrality principles, but consumers may value QoS at the expense of innovation, free speech, and competition on the Internet. Moreover, depending on how QoS is implemented at the technical level, it may actually promote net neutrality values by fostering innovation, free speech, and competition in products and services that would otherwise not function without internet service guarantees. If ISPs perceive this change in consumer priorities, they should not be prohibited from offering QoS guarantees, provided they are fully transparent about their network discrimination. Through this disclosure, therefore, ISPs would essentially compete based on how they define reasonable network management.

Alternatively, some forms of non-minimal discrimination may be imperceptible to the end-user. In this case, slight delays because of small amounts of discrimination through either TOS/flow designations or inferior routing may not significantly affect access to content, services, and applications. This kind of non-minimal discrimination would have essentially no effect on net neutrality principles. This discrimination, however, may be very important to ISPs in reducing transit costs by routing traffic along lower cost (or free) paths through peering and secondary peering relationships. Given that this kind of non-minimal discrimination does not violate any of the net neutrality principles in any perceptible way, it should be included within the definition of reasonable network management.

Finally, some forms of non-minimal discrimination may harm the user by violating one (or more) net neutrality principle(s) with no compensating QoS benefit. Non-minimal blocking, through TCP reset packet blocking or through the manipulation of routing tables to avoid interconnection with certain networks, will frequently fall into this category.[144] Both *Madison River* and *Comcast* were clear cases of non-minimal blocking that were held to be unreasonable network management.[145] Similarly, non-minimal delay discrimination that prioritizes one application (violating the innovation and competition principles) or one perspective (violating the free speech principle) should be considered unreasonable. Mandating that ISPs disclose all discriminatory practices, as the FCC requires in the recent Open Internet Rules, discussed in Section II.E, *supra*, will be particularly effective in reducing unreasonable delay discrimination, since it is unlikely that ISPs will continue chilling innovation, free speech, and competition if such practices are publicized.

IV. CONCLUSION

The concept of reasonable network management calibrates net neutrality principles to the technical realities of the Internet. Reasonable network management, in turn, should be defined first by whether or not the discriminatory practice is technically necessary, and, if not technically necessary, by the discrimination's effect on net neutrality principles. As theoretical ideals, net neutrality principles articulate what we value most in the Internet: its ability to foster innovation, free speech, and competition. This list of values, however, should remain open to new additions. With a strongly enforced requirement that ISPs disclose all discriminatory practices, some forms of non-minimal discrimination could be considered reasonable network management. This could include certain types of QoS guarantees, provided that disclosure makes consumers fully aware of the network discrimination. With this transparency, ISPs would then compete to define QoS in a way that conforms to consumers' preferences. Through this kind of development, demand for QoS internet service would show either that consumers value QoS guarantees higher than the other net neutrality principles or that QoS guarantees actually facilitate the net neutrality principles by supporting otherwise impossible innovations that demand a

[144] It should be noted that discrimination for blocking certain types of illegal content such as child pornography and for security purposes should still be permissible. For more on these exceptions, see Jon M. Peha, *The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy*, 1 INT'L J. COMM. 644, 648–49 (2007).

[145] *See* Madison River Commc'ns, LLC, 20 FCC Rcd. 4295, 4297 (2005); Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13028, 13060 (2008).

high quality connection to the Internet. Conversely, demand for neutral (non-QoS) internet service would confirm that consumers value the ideals protected by the current set of net neutrality principles. In either case, mandating transparency represents a significant step forward from the current state of competition in the provision of broadband internet service. Unlike the opacity of *Ulysses*, in which James Joyce's literary challenges define his style, ISPs should be open books.