Network Security and the Need to Consider Provider Coordination in Network Access Policy*

Aaron J. Burstein
Samuelson Law, Technology &
Public Policy Clinic
Berkeley Center for Law & Technology
University of California, Berkeley
School of Law (Boalt Hall)
Berkeley, CA 94720

Fred B. Schneider
Computer Science Department
Upson Hall
Cornell University
Ithaca, New York 14853

August 17, 2007

Abstract

The policy debate over how to govern access to broadband networks has largely ignored the objective of network trustworthiness—a set of properties (including security) that guarantee that a network will behave as expected. Instead, the terms of the network access debate have focused on whether imposing a nondiscrimination, or network neutrality, obligation on service providers is justified by the condition of competition among last-mile providers. Some argue that, in the absence of a nondiscrimination obligation, service providers will discriminate against content, applications, and services that they (or their affiliates) do not provide. Others argue that this kind of discrimination is unlikely and that a nondiscrimination obligation would reduce incentives to invest in improving networks and developing new applications and services.

One point of agreement is that any nondiscrimination obligation must allow network providers to take measures to protect network security. This agreement, however, is rather abstract. Legislative, regulatory, and scholarly proposals have set forth substantially different security exceptions to nondiscrimination rules; but there has been little analysis of how these exceptions would affect the corresponding rule. Just as importantly, there has been little analysis of whether various exceptions allow sufficient room to defend against modern-day attacks. Moreover,

*The authors acknowledge support for this work from TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422). FBS also acknowledges support from AFOSR grant F9550-06-0019, NSF grant 0430161, and funding from Microsoft Corporation.

the question of how network access policy affects other elements of trustworthiness, such as privacy, have gone unexamined. Put simply, network trustworthiness and network neutrality are closely related technologically and through network access policy. Decisions about technology or policy that are based on either trustworthiness or network neutrality principles in isolation pose the risk of affecting the other area in unexpected and undesirable ways.

This paper seeks to expand the network access policy debate to include both trustworthiness and neutrality. Our analysis leads to three principal conclusions. First, network providers need leeway to block or degrade traffic within their own subnets, as well as traffic exchanged between providers' subnets, in order to offer guarantees against certain kinds of attacks. Some currently proposed security exceptions to network neutrality requirements fail to allow such blocking. Second, some trustworthiness guarantees that are within technical reach, such as routing guarantees, would require service providers *not* to refuse to interconnect. The potential competitive effects of service provider coordination—which is critical in establishing these guarantees—warrant further study. Finally, individual providers are well situated to provide stronger privacy and confidentiality guarantees, without either coordinating with other providers or awaiting new technology. Drawing greater attention to the competitive dimensions of these elements of trustworthiness would likely help induce service providers to strengthen these guarantees.

Introduction

Network trustworthiness—a concept that encompasses not only security but also safety, survivability, and other properties that guarantee a network will behave as expected—is becoming crucial to the operation of national infrastructures and day-to-day business. This paper discusses the dependencies between network competition policy and trustworthiness with an eye toward establishing a framework that will better inform the network access policy debate. Using network neutrality proposals from policymakers and legal scholars, we provide a foundation for relating neutrality to trustworthiness. Improved network trustworthiness will ultimately require cooperation among network providers, and we expect that our framework could extend to this more highly coordinated world.

Until now, network trustworthiness has played a peripheral role in analyses of network competition policy. Advocates of a deregulatory access policy—which would reject nondiscrimination obligations for network access providers—cite improved trustworthiness as a potential benefit of this policy.[1] Scholars in this camp, however, have failed to demonstrate in detail how a deregulatory policy would help improve trustworthiness. Advocates of network neutrality, on the other hand, tend to accept that the current lack of discrimination against applications and protocols is at least partly responsible for certain kinds of security threats.[2] This group expresses hope that trustworthiness measures

[1] *See, e.g.,* Christopher S. Yoo, *Beyond Network Neutrality*, 15 Harv. J.L. & Tech. 1, 9 (2005).

[2] *See, e.g.,* Tim Wu, *Wireless Net Neutrality: Cellular* Carterfone *on Mobile Networks*, New America Foundation Wireless Future Program Working Paper #17 39-40 (Feb. 2007), *at* `http://ssrn.com/abstract=962027` (stating that "any allowance of open entry and competition is likely to lead to greater

implemented at the edges of the network will fix the problem, but it is unclear that such measures would mitigate current threats, let alone foster the Internet's evolution toward providing greater trustworthiness.[3]

The scant attention given to trustworthiness reflects the centrality of competition among network operators and innovation—based on a more highly connected global information system—in the debate over network access policy. Specifically, this debate has revolved around the question of how network providers might unilaterally exploit their market power through terms imposed on users, content providers, or both. The concern among many is that a lack of competition may lead to discrimination against—or even blocking access to—content or applications that are not provided by a network access provider or its affiliates.[4] Preventing such discrimination from taking root, some argue, requires a network neutrality law or regulation, the essence of which would be to prohibit service providers from degrading or blocking traffic, applications, or services from unaffiliated sources.

But a full consideration of network trustworthiness reveals a need to develop a more nuanced picture of discrimination, as well as a broader picture of competition among network access providers. Service providers' defenses against several kinds of trustworthiness compromises would likely run afoul of the nondiscrimination principle that is central to most conceptions of network neutrality. Moreover, many proposed defenses against these threats, as well as some kinds of service guarantees that are on the horizon, would require cooperation among multiple providers, because many network-based threats span organizational boundaries. The increasing need for coordinated defenses and guarantees raises the possibility that improving network trustworthiness could threaten competition in ways that participants in the network neutrality debate—who have focused on single-firm conduct—have not addressed. In addition, some elements of trustworthiness, such as privacy, have substantial competitive dimensions and are crucial to ongoing discussions about network architecture;[5] yet they have been largely ignored within the network access policy debate.

Given that the process of network access competition is unfolding at the same time as private firms, researchers, and the government are seeking ways to improve network trustworthiness, relating these two areas will lead to more fully informed network access

abuses") [hereinafter Wu, *Wireless Net Neutrality]*; Tim Wu and Lawrence Lessig, *Ex Parte Submission in CS Docket No. 02-52*, Aug. 22, 2003, at 4 (noting that some network operators' technical and contractual restrictions reflect "legitimate security concerns").

[3] *See, e.g.,* Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. Telecomm. & High Tech. L. 141, 141-44 (2003) (relating network neutrality to the general problem of "promoting fair evolutionary competition in any privately owned environment") [hereinafter Wu, *Broadband Discrimination.*

[4] *See* John Windhausen, Jr., Public Knowledge, *Good Fences Make Bad Broadband: Preserving an Open Internet Through Net Neutrality* 6, 16-23, Feb. 6, 2006, *at* `http://www.publicknowledge.org/pdf/pk-net-neutrality-whitep-20060206.pdf` (formulating the network neutrality problem as concerning "network owners . . . discriminating against web sites, applications, services or equipment that are not affiliated with the network operator" and providing examples of content and service blocking by service providers) [hereinafter Windhausen, *Good Fences Make Bad Broadband*].

[5] See Section 1.3 for a full definition of trustworthiness and Section 2.3 for an explanation of how privacy fits into a broader picture of network access competition than the network neutrality debate has thus far provided.

policy. Our approach to establishing this relationship is not to advocate a position in favor of or against network neutrality. Nor do we argue that improving network trustworthiness should be a trump card in the network neutrality debate. Instead, we begin, in Part 1, by defining trustworthiness and examining its increasingly important role in the development of network competition policy and scholarship. We find that advocates of network neutrality recognize the need to provide a trustworthiness exception to any neutrality obligation, but they differ in their prescriptions for the scope of this exception. We define a spectrum of these exceptions to guide the rest of our discussion. In Part 2 we examine several categories of trustworthiness improvements through the lens of neutrality and competition policy. We argue that a narrow trustworthiness exception would prevent service providers from implementing trustworthiness improvements that are likely to be important in future networks; but an overly broad exception would effectively swallow a neutrality rule. Finally, in Part 3, we suggest ways to mitigate some of the tensions that we uncover.

1 Tracing Trustworthiness Through the Network Access Debate

Trustworthiness and network neutrality both bring research and policy agendas that could shape the Internet (or its successor). Our focus in this paper is on how efforts to build a more trustworthy network might affect competition in access markets, although competition is not the only societal interest that these efforts might affect. A trustworthy system has been described as one that "does what people expect it to do—and not something else—despite environmental disruption, human user and operator errors, and attacks by hostile parties."[6] Trustworthiness is a "multidimensional" concept encompassing "correctness, reliability, security . . . privacy, safety, and survivability."[7] Security, in turn, means resistance to attacks that "can compromise the secrecy, integrity, or availability of data and services."[8] We provide examples of many of these elements of trustworthiness in Part 1.3.

Where the Internet is concerned, trustworthiness is important for a number of reasons. Computer networks have become elements of infrastructure. Indeed, networks have become the "nervous system" of infrastructure in the United States and throughout much of the world, connecting transportation, energy, water, and food distribution systems.[9] Network-based attacks can last for days and have major effects on a national economy. For example, in May 2007, Estonia suffered a distributed denial of service attack that

[6] National Research Council, Computer Science & Telecommunications Board, *Trust in Cyberspace* 13 (ed. Fred B. Schneider, 1999), `http://www.nap.edu/readingroom/books/trust/trustapk.htm` [hereinafter CSTB, *Trust in Cyberspace*].

[7] CSTB, *Trust in Cyberspace*, *supra* note 6, at 14.

[8] CSTB, *Trust in Cyberspace*, *supra* note 6, at 14.

[9] President's Critical Infrastructure Protection Bd., *National Strategy to Secure Cyberspace* vii (Feb. 2003).

brought banking and other services to a halt for several days.[10] Vulnerabilities in a network can also lead to leaks of personal information, potentially leading to a loss of privacy as well as personal financial losses.

1.1 Policymakers' Views

1.1.1 The FCC

Trustworthiness has been lurking, in some form, in network competition policy ever since the D.C. Circuit decided *Hush-a-Phone Corp. v. United States* in 1956.[11] Hush-a-Phone sold a telephone receiver attachment that reduced background noise present at the speaker's location and also prevented the speaker's voice from being heard by others around him or her. AT&T the Bell companies sought to ban the use of the Hush-a-Phone device under a rule that forbade the "attachment to the telephone of any device not furnished by the telephone company."[12] At the end of a lengthy proceeding to hear Hush-a-Phone's complaint against AT&T's application of this "foreign attachment" rule, the FCC found that the lower volume and distorted sound of a Hush-a-Phone user's voice effected a "public detriment" to the phone system and, on this ground, upheld the Hush-a-Phone ban.[13] The *Hush-a-Phone* court, however, found that the FCC's own findings did not support its conclusion and ordered the Commission to reverse the ban of Hush-a-Phone devices.[14] In doing so, the D.C. Circuit announced a broader principle that forms part of the intellectual foundation of network competition policy: the device prohibition was an "unwarranted interference with the telephone subscriber's right reasonably to use his telephone in ways which are privately beneficial without being publicly detrimental."[15] The court did not specify what a "public detriment" might be, but it clearly recognized the possibility that one user's attaching the wrong type of device to the phone network, or using a device in the wrong way, could degrade or disrupt phone service for others. That is, new devices must not threaten the trustworthiness of the phone system as a whole. The device at issue in *Hush-a-Phone* did not pose such a threat. Nevertheless, preserving the trustworthiness of the phone network was integral to the *Hush-a-Phone* principle.

More than a decade later, the FCC considered whether the Carterfone device, which allowed a mobile radio user to connect to a party on the phone network, had a "material adverse effect upon use of the telephone system" when deciding whether to prohibit it.[16]

[10] John Schwartz, *Bit Wars: When Computers Attack*, N.Y. Times, June 24, 2007.

[11] Hush-a-Phone Corp. v. United States, 238 F.2d 266 (D.C. Cir. 1956).

[12] *Hush-a-Phone*, 238 F.2d at 267 (internal quotation omitted).

[13] *Id.*

[14] *Id.* at 269.

[15] *Hush-a-Phone*, 238 F.2d at 269. *See also In re* Use of the Carterfone Device, 13 F.C.C.2d 420, 423-24 (1968) (referring to the statement in the main text as "the principle of *Hush-A-Phone*").

[16] *In re* Use of the Carterfone Device, 13 FCC 2d 420 (1968). AT&T argued in the *Carter* proceeding that allowing the device to connect to AT&T's network would "divide the responsibility for assuring that each part of the system is able to function effectively"—a duty that AT&T asserted it should be solely responsible for bearing.

The FCC found that a device that provided "nonharmful interconnection" of a telephone system user to a user off the grid did not prevent AT&T from "carry[ing] out its system responsibilities," including maintaining a reliable phone system. Again, trustworthiness appears as a limitation on the scope of permissible innovations.

The FCC followed the *Hush-a-Phone* principle when computer connections to the phone network became common. In the *Second Computer Inquiry*, the FCC again affirmed *Hush-a-Phone*'s and *Carterfone*'s articulation of a "consumer right" to use the network "in ways [that] are privately beneficial without being publicly detrimental."[17]

More recently, as the FCC and the federal courts have removed broadband service providers from the common carrier regulations that applied to the telephone system, the Commission has begun to revisit the relationship between network access and network trustworthiness.[18] In the midst of these regulatory shifts, former FCC Chairman Michael Powell articulated four "Internet Freedoms":[19]

1. Freedom to Access Content.
2. Freedom to Use Applications.
3. Freedom to Attach Personal Devices.
4. Freedom to Obtain Service Plan Information.

Consistent with prior network access regulations, Chairman Powell bounded some of these freedoms with trustworthiness considerations. Specifically, the "Freedom to Access" content was subject to network providers' "legitimate needs to manage their networks," and the "Freedom to Use Applications" was subject to the qualification that they "will not disrupt the network."[20]

Though the FCC has used trustworthiness in a simple and consistent way, it has not articulated in greater detail how to distinguish a genuinely trust-related use or device restriction from a spurious one. Perhaps this elaboration was not necessary; when the FCC complied with *Hush-a-Phone*, it might have been plausible to think of a single entity

[17] *See In re* Second Computer Inquiry, 77 F.C.C. 2d 384, ¶ 142 (1980) (quoting and citing *Hush-a-Phone* and *Carter*) [hereinafter *Computer II*].

[18] Much of the complicated history of these developments is recounted in *National Cable & Telecommunications Association v. Brand X Internet Services*, which held that broadband service delivered via cable modem is an "information service," and hence not subject to the common carrier regulations that apply to a "telecommunications service." *See* 545 U.S. 967, 974-80 (2005) (describing the history of FCC regulations concerning access to communications as well as the particular proceeding that led to *Brand X*); *id.* 985-1000 (explaining the Court's decision to uphold the FCC's classification of cable modem services). Shortly after *Brand X* was decided, the FCC classified broadband Internet service via DSL as an information service. *See In re* Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, 20 F.C.C. Rcd. 14853 (Sept. 23, 2005) [hereinafter FCC, *Wireline Order*]. For a brief, readable history of all of these proceedings, see Windhausen, *Good Fences Make Bad Broadband*, *supra* note 4, at 8-12.

[19] Michael K. Powell, *Preserving Internet Freedom: Guiding Principles for the Industry*, 3 J. Telecomm. & High Tech. L. 5, 11-12 (2004) [hereinafter Powell, *Preserving Internet Freedom*.

[20] *Id.* at 11.

as owning a communications network and defending it against threats arising from the ends of that network. By the time Chairman Powell described the "Internet Freedoms," however, the diversity of network ownership, the level of network interconnections, and diversity of devices connected to networks might have made the notion of providers managing "their" networks somewhat simplistic.

1.1.2 Congress

Still, this framework persists not only at the FCC but also in legislative proposals concerning network neutrality. The assumptions of many of these proposals are that network providers can protect "their" networks alone and can do so without violating the central tenet of network neutrality: not degrading connectivity based on the source of content or the application or service in use. In the current Congress, for example, a bill introduced by Senators Dorgan and Snowe to mandate a form of broadband neutrality would provide an exception for "protecting the security of a user's computer on the network *of such broadband service provider*, or managing such network in a manner that does not distinguish based on the source or ownership of content, application, or service."[21] Thus, the Dorgan-Snowe bill's security exception would extend only to protection of a service provider's "own" network; it would essentially require providers to act unilaterally to improve security. State-level proposals in New York and Maine have taken the same approach.[22] The essential features of these trustworthiness exceptions are that they generally prohibit a service provider from discriminating on the basis of content source

[21] *See* Internet Freedom Preservation Act § 2, S. 215, 110th Cong., `http://thomas.loc.gov/cgi-bin/query/z?c110:S.215.IS:` (emphasis added). [hereinafter "S. 215" or "the Dorgan-Snowe bill" The previous version of this bill contained an identical exception. *See* S. 2917, 109th Cong., `http://thomas.loc.gov/cgi-bin/query/z?c109:S.2917.IS:`. Along similar lines, the Internet Freedom and Nondiscrimination Act of 2006 would have allowed prioritization of certain types of data, so long broadband service providers treated all providers of such data equally. This bill did not explicitly mention security. Instead, it contained a number of exceptions that might encompass network security. For example, § 3(c)(1) would have allowed a service provider "to manage the functioning of its network, on a systemwide basis, provided that any such management function does not result in discrimination"; and § 3(c)(4) explicitly allows a provider to "offer consumer protection services (such as parental controls), provided that a user may refuse or disable such services." *See* H.R. 5417, 109th Cong., `http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5417.RH:`.

[22] The New York State Assembly is considering a network neutrality resolution, which provides this security exception:

> Nothing in this section shall be construed to prevent a broadband or Internet network provider from taking reasonable and nondiscriminatory measures . . . to manage the functioning of *its* network to protect the security and to offer parental controls and other consumer protection measures of such network and broadband or internet network services if such management does not result in discrimination among the content, applications, or services on the network.

A. 3980-B § 243(2)(A), `http://assembly.state.ny.us/leg/?bn=A03980&sh=t` (emphasis added).

Similarly, a bill introduced in the Maine legislature would have mandated "nondiscriminatory access" but permitted a service provider to "[p]rotect the security of a user's computer or provide services in a manner that does not distinguish the source of ownership of content, application or service." *See* LD 1675, `http://www.mainelegislature.org/legis/bills/billtexts/LD167501.asp`.

or application, though discrimination might be permissible for services that subscribers may refuse or disable.[23]

An example of a broader trustworthiness exception, which would would relax the nondiscrimination requirement without eliminating it entirely, comes from the Network Neutrality Act of 2006 ("H.R. 5273," or "the Markey bill").[24] The Markey bill's security exception differs in two important ways from the exception in the Dorgan-Snowe bill. First, though the security exception in the Markey bill would require providers to use "reasonable and nondiscriminatory measures" to protect security,[25] the overall structure of the bill suggests that not all forms of discrimination are prohibited. Specifically, the line between permissible and impermissible discrimination appears to be whether a service provider takes into account the distinction between content or services that it (or an affiliate) provides, versus an unaffiliated provider.[26] In addition, like the Dorgan-Snowe bill, the Markey bill would allow a service provider to offer "consumer protection services" that might include trustworthiness guarantees, so long as subscribers may opt out of them.[27]

An even broader security exception to a network neutrality proposal would have provided removed any nondiscrimination requirement, though the single-firm view of network security remained in place. The Internet Consumer Bill of Rights Act would have allowed an ISP to "protect the security, privacy, or integrity of the network or facilities of such provider, the computer of any subscriber, or any service, including by (A) blocking worms or viruses; or (B) preventing denial of service attacks."[28]

Despite their differences on the issue of discrimination for the purposes of improving network trustworthiness, these legislative proposals (as well as the four Internet Freedoms[29]) share a common approach to the increasing need for coordination among service providers: they ignore it. All of these proposals reflect a single-firm outlook on trustworthiness—service providers may decide when to act in the interests of securing the subnets that they operate (or their subscribers' computers), albeit with varying levels of immunity from the broader nondiscrimination requirements. Whether this silence precludes providers from coordinating on matters of trustworthiness, or what may be sensible guidelines for determining whether a provider's actions are sufficiently protec-

[23] *See supra* note 21.

[24] To our knowledge, a successor to H.R. 5273 has not been introduced in the current Congress. Still, it is worth discussing because it provides a distinct approach to a trustworthiness exception to a network neutrality mandate.

[25] H.R. 5273 § 4(b)(3), 109th Cong., 2d Sess., `http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5273.IH:`.

[26] *See id.* § 4(b)(1) (creating an exception to allow a service provider to "manage the functioning of its network, on a systemwide basis, provided that any such management function does not result in discrimination between content, applications, or services offered *by the provider and unaffiliated providers*") (emphasis added).

[27] *Id.* § 4(b)(4).

[28] This act was Title IX of the Communications Opportunity, Promotion, and Enhancement Act of 2006, H.R. 5252 § 906, 109th Cong., `http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5252.RS:`. Like the Markey bill, the Internet Consumer Bill of Rights Act does not appear to have been reintroduced in the 110th Congress. *But see* note 24.

[29] *See* Part 1.1.1, *supra*.

tive of its subnet in the case of coordinated defenses, are questions that we do not settle here.[30] Still, this silence is worth noting, given the importance that coordinated defenses will play in improving network trustworthiness.[31]

Finally, Congressional forbearance from imposing a nondiscrimination obligation would likely leave service providers with broad power to block or degrade communications for security purposes without regard to their source or contents.[32] In the absence of a Congressional network neutrality mandate, the regulatory levers would remain to address discrimination by service providers include conditions on telecommunications provider mergers and FCC rulemakings.[33] to address abuses of market power, including discrimination.

These trustworthiness exceptions are summarized in Table 1.

1.2 Legal Scholars' Views

1.2.1 Trustworthiness in Tension with Innovation

One view of the relationship between trustworthiness and network neutrality is that they are in tension, if not outright conflict. Like the legislative proposals surveyed above, scholars holding this view focus on security as the element of trustworthiness that provides the most compelling reason to allow deviations from neutrality. As a leading advocate of network neutrality writes: "Spam, viruses, junk mail and telemarketing are different

[30] Whatever these boundaries may be, the antitrust laws would provide *some* limit on the kinds of information that providers may share, as well as the purposes for which they may share it. Specifically, Sherman Act § 1, 15 U.S.C. § 1, prohibits agreements that unreasonably restrain trade; and sharing information about industry practices may sometimes run afoul of this law. *See* Complaint, United States v. Professional Insurance Consultants Ins. Co., Civil No. 1:05CV01272 (D.D.C. June 24, 2005), *available at* `http://www.usdoj.gov/atr/cases/f209700/209728.htm` (alleging that actuarial consulting firms moved toward an industry standard of limitations on liability clauses by sharing competitively sensitive information about such clauses and their efforts to implement them independently).

Note that, in the past, concerns about potential § 1 liability for sharing security-related information have prompted Congress to propose an antitrust exception for sharing such information. *See* Cyber Security Information Act of 2000, H.R. 4246, 106th Cong., 2d Sess., Apr. 12, 2000, *at* `http://thomas.loc.gov/cgi-bin/query/z?c106:H.R.4246:`. *See also* Center for Democracy & Technology, *Davis-Moran Cyber Security Information Act—H.R. 4246*, May 5, 2000, *at* `http://www.cdt.org/security/000504davismoran.shtml` (criticizing the antitrust exemption as "unnecessary").

[31] See our discussion of this point in the Introduction.

[32] *See, e.g.,* Yoo, *Beyond Network Neutrality*, *supra* note 1, at 9, 22, 31, 71.

[33] For discussions of the possibility of FCC intervention outside of the merger context, see *Wireline Order*, *infra* note 18, ¶ 102 (reserving possibility that the FCC will use Title I ancillary jurisdiction to regulate broadband Internet access) *and* Harold Feld, *DSL Item Released—Coulda Been Worse*, WETMACHINE, Aug. 5, 2005, `http://www.wetmachine.com/totsf/item/333`. In the merger context, the FCC imposed a condition of "maintain[ing] a neutral network and neutral routing" on the merger of AT&T and BellSouth, effective for 30 months after closing. *See* Press Release, *FCC Approves Merger of AT&T and BellSouth Corporation* 8, Dec. 29, 2006, *at* `http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-269275A1.pdf`. For a proposal to give the FCC "antitrust-like" authority to adjudicate complaints about service providers abusing their market power, see generally Robert D. Atkinson & Philip J. Weiser, *A "Third Way" on Network Neutrality*, May 30, 2006.

Table 1: Summary of network neutrality trustworthiness exceptions

Exception breadth	Restrictions	Permits security as service?	Source
Narrow	No discrimination on basis of communication source, application, or service	Yes	S. 215
Medium	No discrimination based on whether a communication source, application, or service is from a provider's affiliate	Yes	H.R. 5273
Broad	Discrimination permitted	Yes	H.R. 5252
No neutrality mandate	Constraints: FCC Title I ancillary jurisdiction; merger conditions		

names for problems that every information network faces. What this suggests is that network security must be taken seriously, but also cannot become a blanket answer to any scrutiny of carrier practices."[34] This view appears to fit the pattern set by *Hush-a-Phone* of trustworthiness as a limiting principle on innovation.

1.2.2 Trustworthiness as a Buffer for Innovation

Another scholar cautions that framing a choice between security and "generativity"—a combination of competition and innovation in networked computer systems—creates a "false dichotomy."[35] As a practical matter, this scholar argues, "complete fidelity" to the principle of "placing control and intelligence at the edges of a network" could drive consumer demand for closed platforms as well as increased government regulation of the Internet.[36] Thus, moving some trustworthiness functions away from the edges of the network—i.e., end-users and their computers—may help preserve the values of openness and innovation that are central terms in the network access debate.[37] Bringing security threats under control could, in turn, reduce market-based and regulatory forces leaning in favor of more tightly controlled platforms.[38]

[34] Wu, *Wireless Net Neutrality*, *supra* note 2, at 39-40.

[35] *See* Jonathan Zittrain, *The Generative Internet*, 119 Harv. L. Rev. 1974, 2026 (2006).

[36] *See id.* at 2030 (discussing "end-to-end theory").

[37] An example of such a shift in responsibility for security is protection against computer viruses: "it may be preferable in the medium term to screen out viruses through ISP-operated network gateways rather than through constantly updated PCs." *Id.* at 2031.

[38] *See id.* at 1977-78 (arguing that "we should establish the principles that will blunt the most unappealing features of a more locked-down technological future while acknowledging that unprecedented and, to many who work with in- formation technology, genuinely unthinkable boundaries could likely become the rules from which we must negotiate exceptions").

1.2.3 Trustworthiness as Innovation

Scholars who argue against mandating neutrality view increased trustworthiness as a type of innovation that a mandate might preclude. For example, a proponent of "network diversity" argues that, in the absence of a neutrality mandate, we might see the emergence of multiple last-mile networks, including one that "incorporat[es] security features to facilitate e-commerce and to guard against viruses, spam, and other undesirable aspects of life on the Internet."[39] Thus, this theory postulates that improving security in the last mile is both possible and potentially attractive to consumers.

1.3 Technologists' Views

To make progress in relating network trustworthiness to network competition, it is helpful to have some concrete examples of trustworthiness properties. By focusing on guarantees, we avoid limiting the discussion to the known, specific attacks of today. Attacks evolve, but the kinds of trustworthiness properties one might expect from a network are independent of the threats and the attacks they might employ.

To start, we introduce a refinement of the model typically used to describe relationships on the Internet. The traditional model focuses on the vertical flow of information from content providers to network operators to end users. When considering trustworthiness, however, it is important to recognize that individual end users are not the only consumers of data services that networks carry; the subnets that comprise the Internet also exchange traffic with one another. These interconnections depend on peering and transport agreements, whose significance will become evident in the discussion below.

With our notion of network customer expanded to include subnets as well as individual users and computers, we can list examples of network properties that are useful for building trustworthy networked information systems. For each such property, we discuss the extent to which the current Internet architecture provides support for a corresponding guarantee (only some of which are available to today's Internet customers).

Confidentiality. A sender might want a guarantee that the data she sends are neither intercepted nor stored and later accessed by unauthorized third parties. Such unauthorized access can be prevented by encrypting data, and the current Internet protocols will handle the data in the same manner as unencrypted data.[40]

Privacy. In addition to preventing third parties from gaining access to the contents of a communication, a user might wish to prevent others from learning about the very existence of a communication. Guarding against the disclosure of this kind of information

[39] Yoo, *Beyond Network Neutrality*, *supra* note 1, at 9.

[40] In practice, the strength of the guarantee against a confidentiality breach will depend on a number of other factors: the strength of the encryption algorithm, the sender's and recipient's key management practices, the trustworthiness of any certificate authority involved, and whether the encrypted data are dumped and decrypted offline. These factors are related to cryptography rather than network design. The point of the example in the main text is that the current Internet handles encrypted data in the same manner as unencrypted data.

would require a guarantee concerning the dissemination of traffic logs and restricting access to packets in transit. Currently, network operators decide whether to keep logs of the traffic they carry; the Internet architecture does not provide users with a means to direct a network provider not to log traffic.

Integrity. One of the Internet's core networking protocols, the Transmission Control Protocol (TCP), provides a guarantee that data accepted by a receiver have not been corrupted while in transit. Each TCP header contains a field for a checksum, which is a (more or less) unique numerical coding of the bit strings comprising the header and data in a TCP packet.[41] A receiver may calculate the checksum of incoming data and compare it to the checksum that was calculated on the sender's end and carried by the packet. A difference in these two checksums indicates the data were corrupted during transmission and causes the packet to be discarded by the receiver; the sender will then retransmit that packet. Thus, packets that are not discarded are identical on the sending and receiving ends of a communication.

Availability. The current Internet architecture offers some limited guarantees concerning availability. Specifically, the Internet architecture provides guarantees that users who persist for long enough in attempting to communicate will be able to do so, aided (in part) by the multiplicity of routes that packets may take from sender to recipient. TCP enforces the availability guarantee by requiring the sender to repeatedly retransmit a packet until an acknowledgment packet has been returned to the receiving computer.[42] This particular guarantee of delivery, however, does not imply that the delivery is timely, and TCP delivers data on a best-effort, first-in-first-out basis. This means that network providers can shape traffic based on its source, destination, and application type. Because traffic shaping decisions lie with network providers, they are beyond the control of most users.

Note that outages, such as those caused by earthquakes or accidental severing of network cables, in one subnet might cause traffic to take suboptimal routes and leave destinations on the affected subnet unreachable; but the Internet's current routing architecture renders other hosts usable during such outages. Though network design might help to mitigate some environmental threats, it is unlikely to defend against all of them. Moreover, the current Internet does not provide guarantees of *negative* availability.[43] That is, the Internet does not provide a way to ensure that a user will not receive communications from a specific set of hosts. We discuss how proposed negative availability guarantees would relate to network neutrality in Part 2.

Correctness. The Internet currently employs a primitive service—the domain name system (DNS)—for translating between names that are easy to use and remember, such as www.whitehouse.gov, and the numerical IP addresses used for routing packets. The DNS is vulnerable to a variety of attacks that undermine the trustworthiness of the network. For example, compromising DNS allows attackers to send traffic from malicious

[41] W. Richard Stevens, UNIX Network Programming, Vol. 1, 32 (2d ed., Prentice Hall PTR, 1998).

[42] Stevens, UNIX Network Programming, *supra* note 41, at 32.

[43] *See* Clark, *infra* note 45.

hosts that impersonate legitimate ones, which allows attackers to collect usernames and passwords. This form of attack, known as "phishing," facilitates identity theft and the fraudulent use of personal information to commit financial crimes.[44] The Internet itself (or successor networks) may provide facilities for higher-level queries, such as the search engine queries that have become many users' primary means of navigating the Internet, as well as queries that allow programs to find services. Compromises to these services could severely harm the trustworthiness of those networks.

2 Network Trustworthiness in a Network Neutrality Framework

In this section we consider whether proposals to enhance network trustworthiness would be permissible under the security exceptions found in various network neutrality proposals. The range of network neutrality trustworthiness exceptions identified above suggests two questions to guide further analysis of the relationship between these two principles:

1. What trustworthiness improvements are available without discriminating against traffic based on its source?

2. What is left of network neutrality's general nondiscrimination principle if network operators may discriminate against communications sources, applications, or services in order to enhance network trustworthiness?

Our discussion answers these questions in the context of three guarantees that would help to improve network trustworthiness. Though by no means exhaustive, these example guarantees provide a diverse set of test cases for the narrow, medium, and broad exceptions identified in Part 1.1. Specifically, section 2.1 examines a trustworthiness guarantee that might require service providers to agree not to exchange traffic. By contrast, guarantee discussed in Part 2.2 would require providers to relinquish their right not to exchange traffic with each other. Finally, the privacy guarantee given Part 2.3 could be implemented effectively by a provider acting unilaterally. Examining this range of trustworthiness guarantees permits us to evaluate whether network neutrality trustworthiness exceptions accommodate the range of defenses that are available today and that appear to be promising for the near future.

[44] According to FBI estimates, phishing attacks cause approximately $1 billion in damage per year, identity theft costs $49.3 billion annually, and computer crime overall costs the United States $67.2 billion per year. U.S. Gov't Accountability Office, *Cybercrime: Public and Private Entities Face Challenges in Addressing Cyber Threats* 2

2.1 Negative Availability

The current Internet does not support a guarantee of *negative* availability,[45] which would allow a user to employ the network to block traffic from a specific set of hosts. A negative availability guarantee is useful for defending against distributed denial of service attacks. Blocking traffic from certain hosts could also prevent the spread of viruses or worms from one host to another. Limiting the spread of these programs, in turn, could interrupt the formation of "botnets"—networks of compromised computers under the control of a remote attacker—which can then be used to launch distributed denial of service attacks, send spam, or store data that are useful in committing financial crimes.

Current defenses against these attacks, however, are implemented predominantly at the edge of the network. Firewalls, for example, block traffic with specific characteristics; and anti-virus programs installed on individual PCs reduce the end-user's risk of executing malicious software. These defenses, though helpful, have significant limitations. Authors of worms and viruses have become adept at crafting programs that evade detection by anti-virus programs. Furthermore, firewalls are usually ineffective against denial of service attacks because the attacks saturate network resources near the edge or on the target host; so even if a firewall prevents traffic from reaching the intended target, that host nevertheless remains unavailable if its link to the Internet is saturated by attack traffic.

The questions of whether networks should support and will support negative availability guarantees remain under active debate by technologists and others who are considering future Internet designs.[46] Still, the basic contours are clear enough to discuss within the context of network access competition policy. Basically, achieving negative availability would require automatic detection of malicious traffic and the quarantine of infected hosts.[47] Detecting malicious traffic, in turn, might require the exchange of network data among multiple service providers,[48] as well as agreements among them *not* to exchange traffic. This is due to the fact that certain kinds of attacks, such as distributed denial of service attacks, might be perpetrated using traffic whose packet-level characteristics are indistinguishable from legitimate traffic—only when traffic observations from many points on the network are correlated could a picture of an attack emerge.[49]

2.1.1 Negative Availability as a Consumer Service

Suppose that an ISP offers to its subscribers a package of trustworthiness services relating to negative availability, e.g., filtering traffic from botnets, worms, and viruses and

[45] David D. Clark, *Requirements for a Future Internet: Security as a Case Study*, ver. 2.0, Dec. 3, 2005, *at* `http://find.isi.edu/presentation_files/David_Clark-Security-Requirements-2.pdf`.

[46] *See* Clark, *Requirements for a Future Internet*, *supra* note 45, at 7.

[47] *See id.*

[48] *See, e.g.,* Yinglian Xie, Vyas Sekar, Michael K. Reiter & Hui Zhang, *Forensic Analysis for Epidemic Attacks in Federated Networks* [hereinafter Xie et al., *Forensics in Federated Networks*].

[49] *See id.*; Mark Allman, Ethan Blanton, Vern Paxson & Scott Shenker, *Fighting Coordinated Attackers with Cross-Organizational Information Sharing*, ACM SIGCOMM HotNets V, Nov. 2006.

blocking traffic believed to be part of a distributed denial of service attack. It is only under the narrowest trustworthiness exception—which would allow no blocking based on the source of network traffic—that this service might be impermissible. As noted above, the successful identification of certain kinds of attacks depends upon finding patterns in the source and timing of traffic; without the ability to discriminate on the basis of network traffic source, this type of mitigation would be ineffective. Still, the Dorgan-Snowe bill, which represents the narrow pole of the security exception spectrum, permits service providers to offer "consumer protection services" so long as each user may refuse or disable the service.[50] The broader security exceptions, which permit at least some discrimination based on source for network security purposes, would allow this service.

2.1.2 Negative Availability as Provider Policy

But two widely repeated observations about computer and network security might make this service inadequate, and serve as a basis to examine a second negative availability scenario. The first observation is that end-users are reluctant to invest much in improving security. The second observation is that the insecurity of one host on a network can harm end-users at another host. These observations are related: end-users do not fully internalize the benefits of their investment in security; and, conversely, any given user is subject to attacks launched from the "weakest link" in the network. A possible response from network access providers is to block suspected worm, virus, and botnet traffic for *all* of its subscribers. That is, instead of offering negative availability guarantees as a separate service, the service provider imposes them by default.[51] From the perspective of the range of security exceptions to network neutrality mandate, this approach by the service provider would not look much different from offering a separate service. So long as a security exception allows blocking on the basis of a communication's source, as all but the narrowest exceptions do, blocking worm, virus or botnet traffic by default at the service provider level would not raise concerns under proposed network neutrality mandates.

An alternative approach is the agreement of multiple ISPs to form a federation for

[50] *See* S. 215, *supra* note 21, § 12(b)(3). We discuss the difficulties in this approach later in this section.

[51] This scenario make assumptions about service provider behavior that are unrealistic, at least at present. ISPs are developing managed security services that are aimed primarily at large enterprise customers; thus, at least some service providers see managed security services as a potential new source of revenue. *See* Sarah D. Scarlet, *Pipe Cleaners*, CSOonline.com, July 1, 2007, *at* `http://csoonline.com/read/070107/fea_pipecleaners.html`.

In addition, service providers are reluctant to take aggressive, blanket action to block traffic. *See id.* ("For now, and maybe for the long run, companies like AT&T will have to continue to make careful decisions about what traffic they can safely delete without violating their service-level agreements with customers or overstepping their bounds as common carriers that just pass bits from left to right."). Relying on broadband providers' sense of fidelity to the principles of common carriage, however, may be misguided. At minimum, it assumes away the concern that is fundamental to the network neutrality debate, namely, the classification of broadband service providers as "information services" that are not subject to common carrier regulations..

exchanging data about possible attacks.[52] The rationale for this federation is that smaller service providers administer smaller slices of the Internet's address space; unlike backbone providers or large ISPs, these providers might not command a sufficiently wide view of the Internet to identify subtle threats.[53] A last-mile ISP might also agree to share information with a backbone provider. The backbone provider, which handles a higher volume of traffic and is likely to have a more comprehensive view of Internet traffic than a last-mile ISP, would be able to provide the ISP with a broader view than the ISP could obtain on its own. Finally, two or more backbone providers might agree to exchange information about malicious traffic in order to provide their respective downstream customers—last-mile ISPs or large enterprise networks—with guarantees that they will not forward malicious traffic.

Neither coordination among last-mile ISPs nor coordination between an ISP and one or more backbone providers is addressed in network neutrality security exceptions or in the network neutrality debate more generally. The network neutrality security exceptions are silent about the prospect of coordination among network access providers to implement negative availability guarantees. As was the case with vertically integrated operations—whether performed as a service that a subscriber requests, or as a default policy of the service provider—the key from an implementation perspective is being able to block traffic based on its source.

2.1.3 Negative Availability Without a Nondiscrimination Constraint

A final consideration raised by the examples in this section is whether a security exception without a nondiscrimination requirement would swallow a network neutrality rule, irrespective of the level of coordination that a service provider uses to implement a negative availability guarantee.

Attackers have methods to remotely install malicious software that evades both firewalls and anti-virus software. For example, users risk unwittingly downloading malicious software simply by viewing Web pages that have been corrupted by attackers.[54] These threats pervade the Internet; accordingly, a service provider might be able to find justification for degrading the performance of an application or to degrade or block connections to specific hosts on the Internet. Moreover, a service provider would not have to coordinate with other providers to handle traffic in this manner; a provider could degrade the performance of applications or protocols using today's Internet architecture. As we stated in section 2.1, the uncoordinated security responses of service providers face increasing challenges from increasingly distributed and coordinated security threats. Thus,

[52] *See* Xie et al., *Forensics in Federated Networks*, *supra* note 48, for a discussion of how this might work in practice.

[53] *See* Scalet, *Pipe Cleaners*, *supra* note 51 (quoting Gartner vice president John Pescatore: "[I]t's not just economies of scale . . . It's that the carriers have access to information that the individual enterprise doesn't.").

[54] *See* Niels Provos et al., *The Ghost In The Browser: Analysis of Web-based Malware*, in *Proceedings of the First Workshop on Hot Topics in Botnets (HotBots)* (2007) (demonstrating how malicious HTML and JavaScript can be used to cause a browser to download malicious software automatically to an end-user's computer—a so-called "drive-by download").

the result of a security exception without a nondiscrimination might well be perverse: the exception would shelter provider conduct that holds an attenuated relationship to Internet trustworthiness—but which may be motivated by reasons that are not related to any aspect of trustworthiness—just as strongly as it shelters provider conduct that is specifically intended to improve trustworthiness.

2.2 Availability and Integrity: Attribution of Path

The routing of communications over the Internet is currently beyond individual users' control. Once Internet communications leave a sender's last-mile ISP's network, they are carried by backbone providers until they arrive at the receiver's ISP.[55] These backbone providers exchange traffic under barter agreements in an unregulated market. As others have noted, peering agreements are responsible for a number of problems, including sub-optimal routing and a lack of investment in innovations to the Internet's core.[56] Though an indifference to the route between a sender and a receiver makes connections between end points resilient to failures of some subnets (by giving service providers license to update routes as needed), this also requires users effectively to trust the routing infrastructure for the entire Internet. Two examples will illustrate how routing guarantees would be useful in improving network trustworthiness.

First, consider a user who trusts routers only in certain countries. For instance, this user might be a defense industry consultant who is traveling abroad and needs to communicate confidentially with her colleagues in the United States. But she surmises that her communications are likely to pass through countries that monitor the contents of Internet communications and would be highly motivated to try to break the encryption on communications relating to the U.S. defense industry.[57] If this user can control the routes that her communications take, she will be able to ensure that those communications travel only through countries whose routers she trusts; she would no longer have to trust the entire Internet.

A second example is a guarantee of *disjoint* paths, i.e., paths that do not rely on any of the same routers. The use of such paths increases the probability of delivering

[55] *See* FTC, Staff Report, Broadband Connectivity Competition Policy 25-26, June 2007; Paul Laskowski & John Chuang, *Network Monitors and Contracting Systems: Competition and Innovation* 183, in *Proceedings of ACM SIGCOMM* (2006) [hereinafter Laskowski & Chuang, *Network Monitors*]; Syliva Ratnasamy, Scott Shenker & Steven McCanne, *Towards an Evolvable Internet Architecture* 315, in *Proceedings of ACM SIGCOMM* (2005); Ramesh Johari & John Tsitsilis, *Routing and Peering in a Competitive Internet*, Jan. 30, 2003 [hereinafter Johari & Tsitsilis, *Competitive Internet*]; David D. Clark, Karen R. Sollins, John Wroclawski & Robert Braden, *Tussle in Cyberspace: Defining Tomorrow's Internet*, in *Proceedings of ACM SIGCOMM* (2002)

[56] *See* Johari & Tsitsilis, *Competitive Internet*, *supra* note 55 (discussing "hot potato" routing under backbone provider peering agreements); Laskowski & Chuang, *Network Monitors*, *supra* note 55 (analyzing how peering agreements diminish incentives to invest in core Internet innovation).

[57] Information about the likely route of an Internet communication can be obtained by using the `traceroute` command on Unix and Mac systems, or the `tracert` command on a Microsoft Windows system.

any given packet, because the probability of failure (or compromise) of a machine on any given path is independent of the other paths.[58]

Providing stronger routing guarantees—whether a guarantee to follow a route specified by an end-user or a service provider's guarantee of diverse routing—requires coordination among network access providers. Specifically, to implement these guarantees, network providers would have to: (1) implement a technical mechanism to express and communicate preferred routes; (2) agree to follow route specifications, and (3) provide some means for others to verify that a given provider had followed its promise to route traffic in the specific manner.[59]

We set aside the considerable change in economic relationships among last-mile and backbone providers that would be necessary to achieve such guarantees, in order to examine how they fit in to the spectrum of network neutrality security exceptions that we identified in section 2.[60] In both of the examples that we presented, end-users sought guarantees concerning the paths that their communications would take. The service provider did not draw distinctions among the end-points to which these users wanted to connect. In other words, a service provider's ability to offer attributions of path does not necessarily imply that the provider would use control over routing to degrade performance based on the end-user's choice of application or the identity of the other party to the communication. So long as the end-user may control this choice, these guarantees would fall within the scope of even the narrowest of network neutrality security exceptions.

A more difficult question would arise if a service provider were to select routes based on its own security considerations. As a practical matter, a network architecture that provided routing guarantees would open several possibilities for providers to discriminate against traffic based on its source or the application in use. A provider might decide, for example, that a particular Web browser leaves its users unacceptably vulnerable to the installation of malicious software by remote attackers. This vulnerability, the service provider might conclude, poses a security threat to the provider's network by opening it to further propagation of malicious software, or by enlisting the network's participation in distribute denial of service attacks.[61] Suppose that the provider further reasons that alternatives to this browser with the same functionality are available at no cost. A

[58] An alternative to full user control over the routes for their communications is to provide guarantees of diverse routing. The current Internet architecture does not support these guarantees, either.

[59] For a proposal for how to implement these requirements in practice, see Karthik Lakshminarayanan et al., *Achieving Convergence-Free Routing Using Failure-Carrying Packets*, ACM SIGCOMM (Aug. 2007, Kyoto, Japan).

[60] Recall that the distinction between the narrow and broad security exceptions is that the broad exception would permit a network access provider to take the source of communication into account when acting out of concern for security, while this consideration is not permitted under the narrow exception.

[61] We are aware that an ISP may have incentives to be disingenuous, tacking a security rationale onto a service or application degradation whose primary motivation may be a financial agreement with the provider of another, similar service. Indeed, the existence of such an agreement would create some suspicion about the service provider's motives. To keep this example simple, however, we assume that the service provider acts solely to impose a penalty for using a highly vulnerable browser. In Part 3, however, we explore the implications for this change in Internet architecture if we allow for the possibility that a service provider might (mis)use security to evade a network neutrality obligation.

network architecture that supports path attribution would allow the provider to choose relatively slow routes for requests from that browser, thus degrading the service based on the application that the subscriber has chosen.

In this case, the service provider would likely run afoul of the two relatively narrow network neutrality security exceptions that we examined. The service provider has clearly decided in this example to degrade the performance of a particular application, something that the relatively narrow security exceptions flatly prohibit.[62]

The broadest of the security exceptions that we examined, however, probably offers some cover for the service provider's decision to degrade the performance of the browser in question through route manipulation. The service provider in this example acted to preempt remote threats to the security of subscribers' computers by penalizing users who used a relatively vulnerable browser. On the one hand, this exception would allow a provider to block traffic from worms or viruses, or to "prevent[] denial of service attacks."[63] There is no requirement that the service provider act only to prevent or counter a denial of service attack once it is underway; a set of logically connected considerations—discouraging vulnerable browser use by degrading its performance might prevent malicious software installation, and thus prevent the use of such software to carry out denial of service attacks—might be sufficient to bring the provider's conduct within the scope of this security exception. In addition, this security exception would allow a provider to prevent "unauthorized" uses of its network, without any restrictions on the means employed to achieve that goal.[64]

2.3 Privacy and Confidentiality: Guarantees Against Logging

As the definition of trustworthiness in Part 1.3 suggests, the conditions of network access encompass more than whether service providers will degrade or block communications involving certain hosts or applications. To take one example, service providers play an essential in setting guarantees of end-user privacy.[65] In contrast to the trustworthiness guarantees discussed above, which individual service providers have relatively limited power to make, providers exert significant control over privacy guarantees. Competition among service providers shows promise to strengthen privacy guarantees, yet this dimension of competition is one that the network access policy debate has largely ignored.

[62] *See* S. 215 (requiring a service provider to manage security in a manner that "does not distinguish based on the source or ownership of content, *application*, or service") (emphasis added); H.R. 5273 (requiring a provider to protect of the security of its network or a subscriber's computer using "reasonable and *nondiscriminatory*" measures) (emphasis added).

[63] *See* H.R. 5252, *supra* note 28, § 906(1).

[64] *Id.* § 906(3). Note that the other provisions of H.R. 5252's security exception do not limit this exception. Users would have the right to run any application "without interference from an Internet service provider, *except as otherwise provided by law*" (emphasis added). *Id.* §§ 903(a)(7), (b)(1).

[65] As stated in Part 1.3, we focus on the aspect of privacy that pertains to preventing third parties from learning of the *existence* of a communication. We emphasize that this aspect of privacy—communications privacy—is but one element of a far more complicated concept. This focus is appropriate to keep our discussion focused on the intersection between trustworthiness and network competition policy.

This section expands the framework of network access competition to include end-user privacy.

Privacy fits naturally into the framework that we have established for relating network trustworthiness to network access competition. One reason is that individual privacy protection affects end-users' decisions about Internet use. For example, a user who is concerned about breaches of privacy might avoid visiting certain websites out of fear that her use will be revealed (or used in public or private surveillance).[66] Thus, privacy guarantees could help to promote the goal of openness on the Internet that network neutrality advocates seek to promote. As the Computer Science and Telecommunications Board of the National Research Council wrote in a recent report that makes "confidentiality of stored information and information exchange" part of a "Cybersecurity Bill of Rights":[67]

> One central function of information technology is the communication and storage of information. Just as most people engage in telephone conversations and store paper files with some reasonable assurance that the content will remain private even without their taking explicit action, users should expect electronic systems to communicate and store information in accordance with clear confidentiality policies and with reasonable and comprehensible default behavior.
> . . . As a particularly important way of ensuring confidentiality, responsible parties should have the technical capability to delete or expunge selected information that should not be permanently stored.

Another reason to view privacy as standing on equal footing with availability and integrity guarantees is that individual users—and technical approaches that focus on the edge of the network—are limited in what they can do to improve privacy. Anonymizers provide some measure of privacy by making traffic analysis more difficult, but these technical measures can be cumbersome to use and do not address the more fundamental problem of logging by ISPs.[68] Thus, like the other trustworthiness guarantees discussed in this Part, privacy guarantees could provide a basis for network service provider differentiation and competition in the near term and technical improvements in the long term. In other words, technical and policy decisions about privacy will be made alongside the decisions that affect other elements of trustworthiness as well as the Internet's support for innovation and openness.

As a starting point, the current Internet architecture does not provide technical guarantees to protect individual privacy. Last-mile ISPs, backbone providers, and Internet hosts (such as e-commerce sites) set their own network traffic logging policies. United

[66] *See* Helen Nissenbaum, *Privacy as Contextual Integrity*, 79 Wash. L. Rev. 119, 121 (2004) (discussing how end-users' online activities are recorded, stored, and analyzed into individual profiles for commercial use).

[67] Computer Science and Telecommunications Board, *Toward a Safer and More Secure Cyberspace* 3-4 (June 26, 2007 draft).

[68] *See, e.g.,* Tor: Anonymity Online, Aug. 10, 2007, *at* `http://tor.eff.org/`. Tor uses a distributed network of servers to route communications in a manner that makes them resistant to traffic analysis by parties with access to network traffic logs.

States law does not require network access providers to retain data, but, on the other hand, it does not impose any limits on the amount of data that these providers may retain.[69] Though details about the data retention practices of specific network service providers are scarce, some prominent providers appear to retain significant amounts of data about their subscribers.[70]

Thus, a straightforward and potentially far-reaching means of compromising individual privacy on the Internet is for a last-mile provider to link a user's personal identifying information to his or her IP address and a list of addresses that that subscriber visited. Whether a provider makes this link voluntarily or under compulsion,[71] last-mile providers occupy a central role in setting communications privacy protections because they control subscriber information, IP address assignments, and may retain logs about their subscribers' Internet use.[72] Backbone providers may log Internet communications records but typically do not have the information necessary to link these records to individuals. Individual websites, on the other hand, may collect information about individuals but typically do not control the same breadth and volume of data that a last-mile a last-mile provider does. Thus, privacy guarantees from a last-mile provider, such as a policy limiting the scope and duration of data retention, could significantly reduce threats to privacy, though it would not eliminate them. This guarantee would gain little strength from coordination among different providers; this trustworthiness measure is one that lends itself to unilateral implementation by a single provider.

Raising the profile of privacy guarantees as a dimension of service provider competition would begin with seeking more information about current practices.[73] This, in turn, would provide end-users with sufficient information to discipline service providers in the marketplace, either by registering complaints with their providers or switching to a different one. Thus, improving privacy guarantees could follow naturally from greater disclosure of service plan information—which is a pillar of the current network neutrality regulatory environment—provided that policymakers, market participants, and advo-

[69] The Electronic Communications Privacy Act (ECPA) does establish a data *preservation* requirement; under specific circumstances, a service provider must preserve data that it has in its possession, but the ECPA has no provision to require retention prospectively. *See* 18 U.S.C. § 2703(f).

[70] *See* Recording Indus. Ass'n of Am. v. Verizon Internet Servs., 240 F. Supp. 2d 24, 28 (D.D.C. 2003), *remanded by* 351 F.3d 1229 (D.C. Cir. 2003); Charter Comms., Inc., Memorandum in Support of Motion to Quash Subpoena Served by Recording Indus. Ass'n of Am., No. 4:03MC00273CEJ (E.D. Mo., Oct. 3, 2003) (not arguing that Charter did not have the information necessary to comply with the RIAA's subpoena for personal identifying information linked to an IP address).

[71] One of the three titles in the Electronic Communications Privacy Act (ECPA) regulates the circumstances under which a service provider may disclose such data voluntarily as well as in response to subpoenas or other compulsory process. *See* 18 U.S.C. §§ 2701-2710.

[72] Though the details of specific network access providers' data retention practices are not publicly known, several sources of evidence suggest that they log considerable amounts of information about their customers' Internet use. First, last-mile providers have a strong incentives—protecting against fraud, abuse, and bandwidth hogs—to keep information that will allow them to identify an IP address with a particular subscriber.

[73] Obtaining this information is likely to pose a significant challenge. *See* Ryan Singel, *Why ISP Data Survey Matters: One Smart Lawyer's Take*, Threat Level, Mar. 29, 2007, *at* `http://blog.wired.com/27bstroke6/2007/03/why_isp_data_su.html` (discussing the difficulties involved in "ferret[ing] out how ISPs store and share user Internet usage histories").

cates recognize that privacy is an element of trustworthiness that competition could help to improve.[74]

3 Toward Reconciling Improved Trustworthiness and Network Neutrality

Section 2's examples of trustworthiness improvements present a complex picture of the relationship between trustworthiness and network access competition. Service providers are well positioned to unilaterally provide stronger privacy and confidentiality guarantees; greater attention to providers' privacy policies could spur competition along this dimension of trustworthiness. However, effective defenses against worms, viruses, and distributed denial of service attacks—which are guarantees of negative availability—depend on better exchanges of information among service providers and, potentially, agreements among providers not to exchange traffic. Routing guarantees and attribution for the paths of Internet communications, on the other hand, would require service providers not to refuse traffic from other subnets.

We found in both of these latter cases that the broad trustworthiness exception—which would allow discrimination based on source, application, or service so long as the discrimination has some plausible tie to protecting the provider's network or subscribers' computers—could effectively swallow a corresponding neutrality rule. That is, the broad exception could serve as cover for service provider practices that are not actually related to improving trustworthiness. In this section we consider how to mitigate this tension.[75]

Our principal suggestion is based on a consideration that appears repeatedly in legislative and regulatory proposals for network neutrality: transparency in service provider practices. Specifically, requiring service providers to report practices that they undertake under the trustworthiness exception would likely be effective in limiting the scope of the broad trustworthinessexception.[76]

This reporting requirement would have to balance a number of factors. Informing consumers of service provider practices would allow them to seek changes in provider practices, either through direct complaints to the provider or by switching providers. Making data about uses of the trustworthiness exception publicly available would also

[74] *See infra* note 76 (discussing the "Freedom to Obtain Service Plan Information" as one of the four Internet Freedoms articulated by former FCC Chairman Michael Powell).

[75] In doing so, we do not intend to advance a position in favor of or against network neutrality in general.

[76] For specific evidence of the role of reporting in policy proposals relating network neutrality, see S. 215 § 3 (requiring the FCC to file an annual report with Congress stating, among other things, the terms and conditions for transmitting information over broadband networks); H.R. 5273 § 4(a)(3) (requiring service providers to "clearly and conspicuously disclose to users, in plain language, accurate information about the speed, nature, and limitations of their broadband service"); Michael K. Powell, *Preserving Internet Freedom: Guiding Principles for the Industry*, 3 J. Telecom. & High Tech. L. 5, 12 (2004) (arguing that the "Freedom to Obtain Service Plan Information" is one of the "Four Freedoms" of the Internet).

facilitate enforcement of any public or private right of action that is part of a legislatively enacted neutrality obligation. On the other hand, service providers have an interest in protecting the confidentiality measures that they develop and deploy unilaterally to improve trustworthiness. Similarly, providers will likely want to maintain the confidentiality of any commercial relationships that they have formed with other providers to facilitate coordinated defenses against security threats. A possible way to balance these interests is to require providers to disclose (1) the number and purposes of trustworthiness-related agreements that they have with other service providers (but not their identities); (2) general descriptions of measures that the service provider uses to detect and respond to security threats; and (3) the occasions on which the provider availed itself of the the trustworthiness exception.

Conclusion and Directions for Future Work

This paper offers a few conclusions that, we hope, will serve as a basis for establishing better discourse between the technical and policy debates over network neutrality and network trustworthiness. Cyberthreats are an increasingly urgent matter for network operators and end-users. A trustworthiness exception that does not allow a provider to discriminate based on the source of Internet communications is unlikely to give service providers sufficient latitude to respond to modern-day threats. And even if a trustworthiness exception allows service providers to offer security services, based on discriminating against traffic sources or application, separately from basic Internet service, this provision might leave providers incapable of protecting users from the "public detriments" that have set limits on the extent of network openness ever since *Hush-a-Phone* was decided.

Still, a trustworthiness exception that does not keep any limits on discrimination could swallow the neutrality rule. The threats that currently face the Internet are far more varied and complex than those facing the telephone system in *Hush-a-Phone* and *Carterfone*, but we argued in Part 2 that the broad exception would allow at least some spurious claims of protecting security to serve as cover for practices that have, at most, a tenuous connection to network trustworthiness. To the extent that policymakers are concerned about service providers using a trustworthiness exception to evade a neutrality obligation, they should consider the reporting requirements that we discussed in Part 3.

We have also identified the possibility that information sharing among providers to improve network trustworthiness could present a threat to competition that the network neutrality debate has not considered thus far. Determining whether these agreements could affect competition among network providers, is an important area for future work that will require combining the findings of technical research with a more detailed empirical picture of the economic relationships among network providers and economic and legal theories for evaluating competition under these conditions.