Columbia Law School

Scholarship Archive

Faculty Scholarship

Faculty Publications

2003

Network Neutrality, Broadband Discrimination

Tim Wu

Columbia Law School, twu@law.columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship

Part of the [Communications Law Commons](#), [Intellectual Property Law Commons](#), and the [Internet Law Commons](#)

Recommended Citation

Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. ON TELECOMM. & HIGH TECH. L. 141 (2003).
Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1281

This Article is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu, rwitt@law.columbia.edu.

NETWORK NEUTRALITY, BROADBAND DISCRIMINATION

TIM WU*

INTRODUCTION

Communications regulators over the next decade will spend increasing time on conflicts between the private interests of broadband providers and the public's interest in a competitive innovation environment centered on the Internet. As the policy questions this conflict raises are basic to communications policy, they are likely to reappear in many different forms. So far, the first major appearance has come in the "open access" (or "multiple access") debate, over the desirability of allowing vertical integration between Internet Service Providers and cable operators.¹ Proponents of open access see it as a structural remedy to guard against an erosion of the "neutrality" of the network as between competing content and applications. Critics, meanwhile, have taken open-access regulation as unnecessary and likely to slow the pace of broadband deployment.

This paper takes a more general perspective. The questions raised in discussions of open access and network neutrality are basic to both

* Associate Professor of Law, University of Virginia Law School. I am grateful for comments on this paper from Tom Nachbar, Lawrence Lessig, Mark Lemley, Glen Robinson along with participants at the 2003 Silicon Flatirons Conference and the 2003 University of Ottawa Tory Law Speaker Series. The ideas in this paper were aided by discussions of network neutrality questions with individuals at the Federal Communications Commission and Congress, including Jordan Goldstein, James Assey, Jessica Rosenworcel and Commissioner Michael Copps.

1. See generally Joseph Farrell & Philip J. Weiser, *Modularity, Vertical Integration, and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internet Age*, 17 HARV. J.L. & TECH. (forthcoming 2003), available at <http://repositories.cdlib.org/iber/cpc/CPC02-035> (last visited Sept. 24, 2003); Glenn A. Woroch, *Open Access Rules and the Broadband Race*, 2002 L. REV. MICH. ST. U. DET. C.L. 719 (2002); Glen O. Robinson, *On Refusing to Deal with Rivals*, 87 CORNELL L. REV. 1177, 1224-27 (2002); 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 (2001); Phil Weiser, *Paradigm Changes in Telecommunications Regulation*, 71 U. COLO. L. REV. 819 (2000); James B. Speta, *Handicapping the Race for the Last Mile? A Critique of Open Access Rules for Broadband Platforms*, 17 YALE J. ON REG. 39, 77-90 (2000).

telecommunications and innovation policy. The promotion of network neutrality is no different than the challenge of promoting fair evolutionary competition in any privately owned environment, whether a telephone network, operating system, or even a retail store. Government regulation in such contexts invariably tries to help ensure that the short-term interests of the owner do not prevent the best products or applications becoming available to end-users. The same interest animates the promotion of network neutrality: preserving a Darwinian competition among every conceivable use of the Internet so that the only the best survive.

Given the likely recurrence of these kinds of questions, this paper compares three general approaches to the regulation of broadband providers: structural remedies, a non-discrimination regime, and self- or non-regulation. It questions, first, the merits of structural remedies like open access as a means for promoting network innovation in favor of less intrusive models. While structural restrictions like open access may serve other interests, as a remedy to promote the neutrality of the network they are potentially counterproductive. Proponents of open access have generally overlooked the fact that, to the extent an open access rule inhibits vertical relationships, it can help maintain the Internet's greatest deviation from network neutrality. That deviation is favoritism of data applications, as a class, over latency-sensitive applications involving voice or video. There is also reason to believe that open access alone can be an insufficient remedy for many of the likely instances of network discrimination.

The preferable framework for ensuring network neutrality, I argue, forgoes structural remedies for a direct scrutiny of broadband discrimination. The link between anti-discrimination regulations and network innovation are as old as the *Hush-a-Phone*² and *Carterfone*³ decisions, which controlled AT&T's efforts to destroy innovative network attachments. The basic principle behind a network anti-discrimination regime is to give *users* the right to use non-harmful network attachments or applications, and give innovators the corresponding freedom to supply them. Such a regime avoids some of the costs of structural regulation by allowing for efficient vertical integration so long as the rights granted to the users of the network are not compromised.

But might network neutrality be accomplished without any regulation at all? Basic economic theory suggests that operators have a long-term interest coincident with the public: both should want a neutral platform that supports the emergence of the very best applications.

2. *Hush-A-Phone Corp. v. United States*, 238 F.2d 266 (D.C. Cir. 1956).

3. *Use of the Carterfone Device in Message Toll Tel. Serv.*, 31 F.C.C.2d 420 (1968).

However the evidence suggests the operators may have paid less attention to their long-term interests than might be ideal. A 2002 survey of operator practices conducted for this paper suggests a tendency to favor short-term results.⁴ In that year, evidence of a discrimination problem became clear from several sources, including consumer complaints about operators who ban classes of applications or equipment, like servers, Virtual Private Networks, or WiFi devices,⁵ and in filings at the Federal Communications Commission by application developers.⁶ The survey in this paper shows that operators indeed had implemented significant contractual and architectural limits on certain classes of applications. Operators showed an unfortunate tendency to want to ban new or emerging applications or network attachments, like WiFi devices or Virtual Private Networks, perhaps out of suspicion or an (often futile) interest in price-discrimination. On the whole the evidence suggests that the operators were often pursuing legitimate goals, such as price discrimination and bandwidth management. The problem was the use of methods, like bans on certain forms of applications, which are likely to distort the market and the future of application development. In short, the recent historical record gives good reason to question the efficacy of self-regulation in this area.

I don't want to suggest that operators are somehow incapable of understanding their long-term interests. Yet, when we return to the open access debate, one account of the utility of the debate is that it played an important informational role—the debate itself helped cable operators evaluate their long-term self-interests, and many have chosen to allow rival ISPs access to their networks, for a variety of reasons.⁷ Even strong believers in deregulation and the advantages of vertical integration recognize that incumbents may occasionally become set in their ways.⁸ In this respect, one of the functions of raising issues of broadband discrimination is to challenge broadband operators to ask whether applications restrictions are a good long-term policy. Indeed many of the improvements in operator behavior in the year 2003 may be

4. See *infra* Appendix.

5. Complaints about restrictions on broadband applications like filesharing applications or VPNs are common on discussion forums like DSL Reports. See, e.g., BROADBAND REPORTS, at <http://www.dslreports.com/forum/remark,3775421;mode=flat;root=sware> (July, 2002).

6. See Comments of the High Tech Broadband Coalition, In re: Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities (filed June 18, 2002), available at http://www.itic.org/policy/fcc_020618.pdf; see also FCC Ex Parte Letter, Aug. 22 2003, available at http://faculty.virginia.edu/timwu/wu_lessig_fcc.pdf.

7. For example, AT&T Broadband has recently begun to open parts of its network to ISP competition. See Peter J. Howe, *Earthlink Debuts On AT&T Networks Offers High-Speed Internet Service*, BOSTON GLOBE, Oct. 17, 2002, at C4.

8. See, e.g., Farrell & Weiser, *supra* note 1, at 33-36.

linked to the Federal Communications Commission's increased oversight of this area.

This paper encompasses a mixture of empirical and theoretical sections. The first part of five is an effort to explain the relationship between several related concepts in this area: open access, broadband discrimination, and network neutrality. Network neutrality, as shorthand for a system of belief about innovation policy, is the end, while open access and broadband discrimination are the means. I suggest that open access regulation, as a structural remedy to ensure network neutrality, is not ideally suited to that task. A direct analysis premised on normative principle of network neutrality may provide a better means to discuss the harm in question.

The second part develops the theoretical framework for a broadband discrimination regime. It asks whether we can differentiate between justified and unjustified restrictions on user behavior, with particular reference to the restrictions seen in the survey in the third part. The use of restrictions on classes of application to pursue bandwidth management and price discrimination is troubling when those restrictions might be pursued through less restrictive means. The section also asks whether self-regulation is likely, and concludes that the threat of regulation might serve useful.

The third part is a survey of the degree to which broadband operators restrict certain applications and favor others. The study surveys the nation's 10 largest cable operators and six largest DSL providers. The results are mixed. First, cable operators tend to employ far more contractual restrictions than do DSL operators. The contractual restrictions and network designs tend to favor, as a class, one-to-many applications development. Second, there is a tendency to use restrictions on application classes to pursue goals such as price discrimination and bandwidth management.

The fourth part shows what a workable principle of network neutrality would look like and what it would mean for the conduct of broadband providers. It suggests that operators should have the freedom to "police what they own," or act reasonably to control the local broadband network. On the other hand, it suggests that that the Internet community (and, at some point, regulators) should view with suspicion restrictions premised on inter-network criteria. A sample text of an anti-discrimination law is included to show how such a principle could be implemented. Finally, the fifth and final part of this paper addresses several possible counterarguments to the network neutrality regime discussed in this article.

I. NETWORK NEUTRALITY & OPEN ACCESS

The relationship between concepts like open-access, network neutrality, and broadband discrimination may be unclear to the reader. It is best to understand network neutrality as an end, and open access and broadband discrimination as different means to that end. In this section we will examine both why network neutrality might be an attractive goal, and, how an open-access and broadband discrimination regime differ as means toward that end.

A. The Case for Network Neutrality

So what is attractive about a neutral network—that is, an Internet that does not favor one application (say, the world wide web), over others (say, email)? Who cares if the Internet is better for some things than others?⁹

The argument for network neutrality must be understood as a concrete expression of a system of belief about innovation, one that has gained significant popularity over last two decades. The belief system goes by many names.¹⁰ Here we can refer to it generally as the evolutionary model.¹¹ Speaking very generally, adherents view the innovation process as a survival-of-the-fittest competition among developers of new technologies. They are suspicious of models of development that might vest control in any initial prospect-holder, private or public, who is expected to direct the optimal path of innovation, minimizing the excesses of innovative competition.¹² The suspicion arises from the belief that the most promising path of development is difficult to predict in advance, and the argument that any single prospect holder will suffer from cognitive biases (such as a predisposition to continue with current ways of doing business) that make it unlikely to come to the right decisions, despite best intentions.

9. More general arguments in favor of a network neutrality regime can be found in Lawrence Lessig & Tim Wu, FCC Ex Parte Letter, Aug. 22, 2003, available at http://faculty.virginia.edu/timwu/wu_lessig_fcc.pdf.

10. A full treatment of the names given to evolutionary theories of innovation is beyond the scope of this paper. Some adherents would ascribe such theories to economist Joseph Schumpeter, while in recent legal work the argument is stated as an argument over what should be owned and what should be free. See generally LAWRENCE LESSIG, *THE FUTURE OF IDEAS* 3-17 (2001).

11. See, e.g., John Ziman, *Evolutionary Models for Technological Change*, in *TECHNOLOGICAL INNOVATION AS AN EVOLUTIONARY PROCESS* 3 (John Ziman ed., 2000); RICHARD NELSON, *UNDERSTANDING TECHNICAL CHANGE AS AN EVOLUTIONARY PROCESS* (1987).

12. In the legal field, Edmund W. Kitch's *The Nature and Function of the Patent System*, 20 J.L. & ECON. 265 (1977) is often taken to exemplify this approach.

This account is simplistic; of interest is what the theory says for network design. A communications network like the Internet can be seen as a platform for a competition among application developers. Email, the web, and streaming applications are in a battle for the attention and interest of end-users. It is therefore important that the platform be neutral to ensure the competition remains meritocratic.

For these reasons, Internet Darwinians argue that their innovation theory is embodied in the "end-to-end" design argument, which in essence suggests that networks should be neutral as among applications.¹³ As network theorist Jerome Saltzer puts it: "The End-to-End argument says 'don't force any service, feature, or restriction on the customer; his application knows best what features it needs, and whether or not to provide those features itself.'"¹⁴ The Internet Protocol suite (IP) was designed to follow the end-to-end principle, and is famously indifferent both to the physical communications medium "below" it, and the applications running "above" it.¹⁵ Packets on the Internet run over glass and copper, ATM and Ethernet, carrying .mp3 files, bits of web pages, and snippets of chat. Backers of an evolutionary approach to innovation take the Internet, the fastest growing communications network in history, as evidence of the superiority of a network designed along evolutionary principles.¹⁶

There is much to this debate, and I do not want to suggest that the discussion about the general merits of evolutionary innovation models are settled, nor are the debates over whether a neutral platform best stimulates competition among applications.¹⁷ But sentiments like those I have just expressed have come to enjoy a broad normative following. From this we can understand why preserving a neutral network might be taken as a suitable goal of Internet communications policy.

13. See J.H. Saltzer et al., *End-to-End Arguments in System Design*, 2 ACM TRANSACTIONS COMPUTER SYS. 277 (1984), available at <http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf> (last visited Oct. 9, 2003).

14. *Id.* at 3.

15. The metaphors of "above" and "below" come from the fact that in a layered model of the Internet's design, the application layers are "above" the TCP/IP layers, while the physical layers are "below." See ANDREW S. TANENBAUM, *COMPUTER NETWORKS* 39 (4th ed. 2002).

16. LESSIG, *supra* note 10, at 14 ("No modern phenomenon better demonstrates the importance of free resources to innovation and creativity than the internet.").

17. For a recent work doubting the merits of open platform designs under some circumstances, see, e.g., Douglas Lichtman, *Property Rights In Emerging Platform Technologies*, 29 J. LEGAL STUD. 615 (2000).

B. The Open Access Remedy and its Limitations

Taking network neutrality as the goal, we can understand open access as one kind of remedy. The term open-access is used in many different ways; it generally refers to a structural requirement that would prevent broadband operators from bundling broadband service with Internet access from in-house Internet service providers.¹⁸ Certain proponents, like Jerome Saltzer, Larry Lessig and Mark Lemley, make the logical link between open-access regulation and the preservation of a neutral Internet. They argue that if cable operators are allowed to bundle ISP services with cable services, cable operators would be in a position to destroy the neutrality of the network by foreclosing competition among Internet applications. As Lemley and Lessig put it,

[T]here is, in principle, no limit to what a cable company could bundle into its control of the network. As ISPs expand beyond the functions they have traditionally performed, AT&T or Time Warner might be in a position to foreclose all competition in an increasing range of services provided over broadband lines. The services available to broadband cable users would then be determined by the captive ISPs owned by each local cable company. This design would contradict the principle that the network should remain neutral and empower users. It further could constitute the first step in a return to the failed architecture of the old AT&T monopoly.¹⁹

Critics of this argument, like Phil Weiser, Jim Speta, and Glen Robinson, have, in the main, cast doubt on the claim that regulation is needed to prevent cable operators from foreclosing competition when it would be efficient, or ask whether network neutrality is an appropriate goal.²⁰ But I want to raise a slightly different question. If we agree with the normative goal of network neutrality, to what degree does the structural remedy of open-access actually serve its interest? Might we do better by targeting network neutrality directly with questions of broadband discrimination?

I believe there are several reasons to question the fit between open-access remedies and network neutrality. First, the concept of network neutrality is not as simple as some IP partisans have suggested. Neutrality, as a concept, is finicky, and depends entirely on what set of

18. The FCC, for example, has outlined three forms of open access remedy in ongoing open access rulemaking. See Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities, *Declaratory Ruling and Notice of Proposed Rule Making*, 17 F.C.C.R. 4798, ¶ 74 (2002) (discussing various models of open access regulation).

19. See Lemley & Lessig, *supra* note 1, at 942-43.

20. See Speta, *supra* note 1, at 76; Farrell & Weiser, *supra* note 1, at 4-6; Robinson, *supra* note 1, at 1216-17.

subjects you choose to be neutral among.²¹ A policy that appears neutral in a certain time period, like “all men may vote”, may lose its neutrality in a later time period, when the range of subjects is enlarged.

This problem afflicts the network neutrality embodied in the IP protocols. As the universe of applications has grown, the original conception of IP neutrality has dated: for IP was only neutral among *data* applications. Internet networks tend to favor, as a class, applications insensitive to latency (delay) or jitter (signal distortion). Consider that it doesn't matter whether an email arrives now or a few milliseconds later. But it certainly matters for applications that want to carry voice or video. In a universe of applications, that includes both latency-sensitive and insensitive applications, it is difficult to regard the IP suite as truly neutral as among all applications.

This point is closely linked to questions of structural separation. The technical reason IP favors data applications is that it lacks any universal mechanism to offer a quality of service (QoS) guarantee.²² It doesn't insist that data arrive at any time or place. Instead, IP generally adopts a “best-effort” approach: it says, deliver the packets as fast as you can, which over a typical end-to-end connection may range from a basic 56K connection at the ends, to the precisely timed gigabits of bandwidth available on backbone SONET links. IP doesn't care: it runs over everything. But as a consequence, it implicitly disfavors applications that do care.

Network design is an exercise in tradeoffs, and IP's designers would point out that the approach of avoiding QoS had important advantages. Primarily, it helped IP be “downwardly” neutral as to the underlying physical media. But this requires us to be more circumspect in our discussions of network neutrality. IP's neutrality is actually a tradeoff between upward (application) and downward (connection) neutrality. If it is upward, or application neutrality that consumers care about, principles of downward neutrality may be a necessary sacrifice.

This returns us to the question of structural separation. We have a public network that is indeed a great creative commons for data applications, but it is less so for any application that requires a minimum quality of service. True application neutrality may, in fact, sometimes require a close vertical relationship between a broadband operator and Internet service provider. The reason is that the operator is ultimately

21. Cf. *Lamb's Chapel v. Ctr. Moriches Union Free Sch. Dist.*, 508 U.S. 384, 397-400 (1993) (Scalia, J., concurring) (on the meaning of neutrality in the context of church and state).

22. Efforts to add quality of service functionality to the Internet protocol, such as the IETF's DiffServ and IntServ's approaches, have never been implemented to provide end-to-end quality of service on an IP network.

the gatekeeper of quality of service for a given user, because only the broadband operator is in a position to offer service guarantees that extend to the end-user's computer (or network). Delivering the full possible range of applications either requires an impracticable upgrade of the entire network, or some tolerance of close vertical relationships.

This point indicts a strict open-access requirement. To the extent open access regulation prevents broadband operators from architectural cooperation with ISPs for the purpose of providing QoS dependent applications, it could hurt the cause of network neutrality.²³ By threatening the vertical relationship required for certain application types, it could maintain IP's discrimination in favor of data applications. More broadly, this argument shows that the concept of network neutrality cannot be taken as counsel against all vertical integration.²⁴

A second, and simpler, problem with open access from a neutrality perspective is that the structural remedy may also be an underinclusive means of ensuring network neutrality. Competition among ISPs does not necessarily mean that broadband operators will simply retreat to acting as passive carriers in the last mile. As the survey in this study shows, operators continue to have reasons to want to control usage of the Internet based on their status as broadband operators, regardless of ISP competition. Hence, open-access does not end the debate over whether broadband operators are capable of engaging in undesirable behavior from the perspective of the public network.

For these reasons, this paper seeks to see if we might do better to address questions of network neutrality directly, through the remedial concept of "broadband discrimination," rather than through structural solutions like open-access.

II. THE CONCEPT OF BROADBAND DISCRIMINATION

The question of controlling what people do with their network services is hardly new to communications regulation. It is as least as old as *Hush-A-Phone*, and the D.C. Circuit's interpretation of the 1934 Communications Act to find that the subscriber has a "right reasonably to use his telephone in ways which are privately beneficial without being publicly detrimental."²⁵

23. This might happen, for example, if an open-access regulation slowed the development of vertically integrated layer 2 / layer 3 architectures.

24. Ultimately, this line of argument echoes the economists' point that efficiencies exist from vertical integration. The point here is to show that principles of network neutrality lead to the same conclusion.

25. *Hush-A-Phone Corp. v. United States*, 238 F.2d 266, 269 (D.C. Cir. 1956).

Nor is the prevention of discrimination a new topic in communications regulation. Over the history of communications regulation, the Government has employed both common carriage requirements (similar to the neutrality regime discussed here) and limits on vertical integration as means of preventing unwanted discrimination. The goal of this section is to further explain how a common carriage or anti-discrimination model might be better developed to address the current Internet environment:

Why might thinking in discrimination terms be useful? Only because it borrows from what is familiar to achieve new goals. What is critical to the study of discrimination regimes is the existence of both justified and suspect bases of discrimination. For example, in the employment context, where discrimination norms are most developed, employers are generally permitted to fire or refuse to hire individuals for a range of reasons, such as education-level, intelligence, and demeanor.²⁶ The law implicitly recognizes that it is essential that the employer retain the freedom to fire incompetents and hire only those with necessary skills. On the other hand, criteria such as race, sex, or national origin are suspect criteria of discrimination, but can only be justified by a *bona fide* rationale.²⁷

While discrimination among Internet applications is a different context, the framework of analysis can be usefully retained. As the proposal in Part IV develops, it is possible to distinguish between classes of restrictions that should generally be allowable, and those that might raise suspicion. Overall, there is a need to strike a balance between legitimate interests in discriminating against certain uses, and reasons that are suspect either due to irrationality or because of costs not internalized by the broadband operator.

To get a better feeling for what a discrimination approach entails, it is helpful to map out some of the extremes of clearly permissible and clearly troublesome discrimination in the broadband context. At one extreme, many of the usage or application bans surveyed are clearly justified. For example, operators usually ban users from using applications or conduct that are meant to hurt the network or other users, like network viruses.²⁸ It is true that this is a departure from

26. See, e.g., 42 U.S.C. § 2000e *et seq.* (2002) (codification of Title VII of the Civil Rights Act of 1964).

27. See *id.*

28. An example from the Cox Acceptable Use Policy:

You are prohibited from posting, transmitting or disseminating any information or software that contains a virus, Trojan horse, worm or other harmful program or that generates levels of traffic sufficient to impede others' ability to send or retrieve information. Prohibited conduct of this type includes denial of service attacks or

network neutrality, because it disfavors a class of applications—those that are disruptive to the network. Yet, it is clear that the operator has acted to solve a problem of a negative externality—the costs imposed by one user on others. Few could or would argue that this is a bad thing.

At the opposite extreme, the harm from totally unjustified discrimination is equally clear. Leaving aside whether operators would actually act in this way, imagine that the nation's broadband operators came to feel that IP "chat" programs were just a waste of time, and were able to use their control over the last mile to ban their use.²⁹ Such discrimination has both a direct harm as well as several negative externalities. The direct harm is obvious: existing broadband consumers who like chat programs lose the opportunity to use a valued application, while creators of chat programs lose whatever revenue opportunity chat programs create. But the more interesting costs are the various losses of positive externalities. Three stand out. First, if chat programs have positive externalities for other network applications—say, if the chat program is middle-ware for a file-exchange program, as in the case of Aimster—dependent applications are hurt as well. Second, to the degree other applications depend on a critical mass of high-bandwidth users, they are hurt by potential subscribers who at the margin are not willing to pay for broadband without the chat programs. Finally, to the extent chat programs have positive social externalities, like helping people to plan meetings or meet new boyfriends, the public suffers too.³⁰ Thus, there are considerable potential costs from an irrational or unjustified ban on certain application types.

These are the easy cases. We next consider whether reasons like price discrimination and bandwidth management should justify discrimination among applications.

A. Price Discrimination & Restrictions on Commercial Use

As detailed in the survey below, nearly every operator places limits on "commercial" use, sometimes including limits on Virtual Private

similarly disruptive transmissions, as well as transmissions containing other harmful or malicious features.

Cox Communications Policies, *Acceptable Use Policy*, Cox Communications, Inc., at <http://support.cox.net/custsup/policies/acceptableuse.shtml> (revised Feb. 3, 2003).

29. For example, by screening chat program activity by TCP port number. Such a restriction could be avoided, but it suffices for the example.

30. Conversely, as we will see in a second, if chat programs have negative externalities because they actually do waste everyone's time, the operators may have done the world a big favor.

Networks, as well as limits on acting as a server.³¹ Why might an operator put such a restriction on usage? Doing so obviously makes the service less attractive to consumers who might want to act in a commercial way, even in a fairly casual manner.³²

The simple answer is price discrimination. That this is the case is not just intuition, but can be confirmed by company policy. As evidence we can consider Comcast's reply in 2001 to a user who had complained about the ban on VPN usage on Comcast's network:

Thank you for your message.

High traffic telecommuting while utilizing a VPN can adversely affect the condition of the network while disrupting the connection of our regular residential subscribers.

To accommodate the needs of our customers who do choose to operate VPN, Comcast offers the Comcast @Home Professional product. @Home Pro is designed to meet the needs of the ever growing population of small office/home office customers and telecommuters that need to take advantage of protocols such as VPN. This product will cost \$95 per month, and afford you with standards which differ from the standard residential product.

If you're interested in upgrading³³

As the letter shows, Cable and DSL operators typically offer commercial packages at a considerable markup from basic broadband service. For example, phone companies like Verizon or BellSouth offer T-1 lines at prices far higher than basic DSL or cable service.³⁴ The goal is to exact a premium price from the customers who most desire the commercial service. Allowing subscribers to basic service to operate hosting services might erode such profits.

It is true that mainstream antitrust analysis has come to see price discrimination as generally uncontentious, or at least ambiguous.³⁵ As

31. See, e.g., *Cable Modem Service Subscription Agreement*, Time Warner Cable, at http://help.twcable.com/html/twc_sub_agreement.html (last visited Mar. 12, 2003) [hereinafter *Time Warner Usage Agreement*].

32. Network design already discourages hosting activity, because most broadband services give asymmetric bandwidth (more downstream than upstream) and a dynamic, as opposed to fixed, IP address. These design features preclude serious commercial website operation, but leave room for casual hosting operations, such as participating in a peer-to-peer network.

33. See *Comcast VPN letter*, Practically Networked, at <http://www.practicallynetworked.com/news/comcast.htm> (last visited Mar. 12, 2003).

34. A T-1 line, providing 1.5 mbps of symmetric data, is usually priced at over \$1000 per month.

35. See, e.g., RICHARD POSNER, *ANTITRUST LAW* 203-06 (2d ed. 2001).

between consumers and producers, it hurts some consumers and helps others, while raising the producers' profits. Yet this analysis can, and should, change as in the broadband context, because the practice of price discrimination may have external effects on the process of innovation and competition among applications. That is to say, while price discrimination among applications may not be troubling from a static perspective (as between existing consumers and producers), it may have dynamic consequences, for the competitive development of new applications.

We can see this in the present example of a ban on commercial operations. The goal, as we've seen, is to maintain a customary markup on business services. But the restrictions on the market for what can be termed commercial applications used on home connections come at a cost. The direct effect of a ban on hosting is to make the connection slightly less valuable to the basic consumer, which presumably the operator takes into account in her pricing scheme. But there are other costs that the operator may not internalize. The bans on commercial use or acting as a server constrain the competitive development of applications that might rely on such a function. In the Comcast letter example the problem was VPN applications, which typically can rely on end-users functioning both as clients and servers, and which can be classified as a commercial use.³⁶ And it is also the case that hosting services may have positive social externalities not taken into account by the operator's decision. For example, VPNs may facilitate greater productivity among employees, a benefit that may be lost in their prohibition.

Another major restriction that interests broadband operators is barring users from providing content to the public or running servers. Why do broadband operators act in this way, if, again, it might lower the value of its service to its users? One reason may be the price discrimination rationale discussed above. Yet from the reports of cable operators themselves, a major goal is bandwidth management.³⁷ The restrictions appear to be efforts to manage how users consume bandwidth by discriminating against types of usage. As the survey showed, such restrictions are more common on cable networks, which operate shared connections and tend to lack technological means for restricting individual bandwidth consumption.³⁸ Hence, the restrictions, for

36. "Servents" in Gnutella terminology.

37. See, e.g., JUSTIN PEARSE, *UK shrugs off American broadband troubles*, ZDNET NEWS.COM, at <http://news.zdnet.co.uk/story/0,,t269-s2077792,00.html> (Mar. 20, 2000).

38. More recent incarnations of the DOCSIS protocol attempt to add better QoS functionality, but implementation at this date seems to be scarce. See *Cable Modem/DOCSIS™*, CABLELABS, at <http://www.cablemodem.com/faq> (last visited Mar. 13, 2003) [hereinafter *CABLELABS, DOCSIS*].

example, on running "game" or "ftp" programs are most likely efforts to eliminate a potential source of bandwidth consumption.

The goal of bandwidth management poses an even more difficult question than does price discrimination. The goal of bandwidth management is, at a general level, aligned with network neutrality. As discussed above, certain classes of applications will never function properly unless bandwidth and quality of service are guaranteed. Hence, the absence of bandwidth management can interfere with application development and competition.

There are good reasons to question whether price-discrimination, without more, should be permissible grounds for allowing discrimination among applications. As we have seen, such usage restrictions may harm consumer welfare without offering a public benefit. This is particularly the case when there are less-restrictive means for engaging in price discrimination. Selling different tiers of service (low, medium, and high bandwidth) does not favor or discriminate against particular application types. In the presence of a means for differentiating among customers in a way that does not distort the process of competitive innovation, we should view discrimination on the basis of application with suspicion.

Similarly, while managing bandwidth is a laudable goal, its achievement through restricting certain application types is an unfortunate solution. The result is obviously a selective disadvantage for certain application markets. The less restrictive means is, as above, the technological management of bandwidth. Application-restrictions should, at best, be a stopgap solution to the problem of competing bandwidth demands.

B. Self-Regulation and the Educational Properties of Regulation

The previous sections show that broadband operators may want to discriminate amongst the uses of its network for various reasons. We have also seen that there are a variety of justifications—some good and some not—for such restrictions. Even if the goal itself is legitimate, the method of achieving that goal may be suspect. The question, then, is whether cable operators will self-regulate and come up with the best policies on their own, or whether regulation may be necessary.

In this section I argue that while cable operators may come to understand that broadband discrimination is not in their best interest, both the threat of, or actual implementation of, anti-discrimination regulation may otherwise serve a useful informational or educational function. Like anti-discrimination legislation in other contexts, it may serve an educational function, forcing operators to ask whether the

restrictions they draft are actually serving their interest in maximizing the value of their services.

As a baseline, the attractiveness of broadband service is a function of the applications it offers the consumer. Hence, any restriction on use will lower the value of the service, and consequently either the price the operator can charge or the number of customers who will sign up (assuming a negative demand curve). To make this clear: if an operator operated a service that screened all uses except web-access alone it might be worth \$30 to the average consumer, while a service that offered access to every kind of Internet application—including, say, the opportunity to get copyrighted music for free—might be worth \$50. The difference in the value to the consumer will affect the price the operator can charge.

This basic point is captured by Joseph Farrell and Philip Weiser's argument that a "platform monopolist has a powerful incentive to be a good steward of the applications sector for its platform."³⁹ The point reflects, as the authors stress, classic arguments from antitrust. A monopolist may still want competition in its input markets, to maximize profit in the monopoly market.

But it is easy for a steward to recognize that the platform should support as many applications as possible. The more difficult challenge has always been the dynamic aspect: recognizing that serving a tangible goal—like controlling bandwidth usage—may affect the intangible status of the Internet as an application development platform. Some of the restrictions, such as those on running various types of server, are applications that are now likely to be used by only a small minority of broadband users. Their sacrifice may appear like a good cost-saving measure.

More generally, the idea that discrimination may not always be rational is a well-understood phenomenon. In the employment context, the various discrimination laws have an explicitly educational function. For example, an express purpose of age discrimination legislation is to force employers to reconsider stereotyped perceptions of the competency of the elderly in the workforce.⁴⁰ Broadband operators may simply disfavor certain uses of their network for irrational reasons, such as hypothetic security concerns or exaggerated fears of legal liability. Additionally, a restriction may become obsolete: adopted at a certain time for a certain reason that no longer matters. Practical experience suggests that such things happen.

39. Farrell & Weiser, *supra* note 1, at 21. This they describe as the "internalization of complementary efficiencies, or ICE."

40. See *Gilmer v. Interstate/Johnson Lane Corp.*, 500 U.S. 20, 27 (1991) ("the ADEA is designed not only to address individual grievances, but also to further important social policies").

For these reasons, anti-discrimination regulation or the threat thereof can also serve a useful educational function. It can force broadband operators to consider whether their restrictions are in their long-term best interests. And in the absence of law it can establish norms around discrimination that may preserve network neutrality over the long term.

The events of the year 2003 provide evidence to support the utility of a regulatory threat in promoting desirable conduct. Both Comcast and Cox Communications openly disavowed their old practices of placing bans on Virtual Private Networks, and filed documents with the FCC to that respect.⁴¹ The cable industry has furthermore begun to publicly insist that it wants to avoid broadband discrimination in the future, stating, for example, that "Cable Believes in Open Connectivity for the Internet."⁴²

There is the possibility that the current regulatory process has forced cable operators to rethink their practices and conclude that discrimination is not in their long term self-interest. The process demonstrates the continuing utility of communications regulators in remaining appraised on potential problems of anti-competitive practices.

III. A SURVEY OF BROADBAND USAGE RESTRICTIONS

Have broadband operators tended to favor certain uses of the Internet? To what extent? The goal of this section is to answer these questions, to the extent possible, for broadband networks during the year 2002.⁴³

The study divides measures of favoritism and discrimination into two categories: contractual and architectural. The study surveyed the network designs (to the extent that the information was available) and usage restrictions in subscriber agreements and incorporated acceptable

41. See Comcast Corp., FCC Ex Parte Letter, May 9, 2002 ("the 'VPN restriction' about which certain parties have complained has been eliminated from and is no longer part of Comcast's subscriber agreements and terms of service for its high-speed Internet customers."); Cox Enterprises Inc., FCC Ex Parte Letter, May 1, 2003 ("Cox hereby informs the Commission that the language of that [VPN] provision has been changed. . .").

42. NTCA, *Cable Believes in Open Connectivity for the Internet*, at <http://www.ntca.com/legislative/legAffairs.cfm?legRegID=20>; see also NTCA, Ex Parte Letter, Sept. 8, 2003 (arguing that network neutrality legislation is unnecessary because of cable's commitment to non-discrimination.).

43. Unfortunately, nearly any feature of network design or policy can be described as a deviation from a "purely" neutral design. Something as innocuous as the length of the IP packet header could, potentially, help or hurt certain applications. To avoid an exercise in the esoteric, the goal of this section is to study major, intentional deviations from neutrality that clearly favor certain application types over others.

use policies from the 10 largest cable operators (AT&T,⁴⁴ Time Warner, Comcast, Cox Communications, Adelphia, Mediacom, Charter Communications, CableOne, Insight, and Cablevision), and 6 major DSL operators (Verizon, SBC, Qwest, BellSouth, Sprint and WorldCom). A chart containing full results can be found in the appendix.

The survey showed the following general results. On the whole, broadband operators' networks and usage restrictions favored the applications of the late 1990s (primarily the World Wide Web and other client-server applications), and disfavored more recent applications and usage, like home networking, peer-to-peer applications, and home telecommuting.

There are differences between cable and DSL operators. On the contractual side, cable operators tended to impose far more restrictions on usage than do DSL operators. Major differences exist with respect to the extent of restrictions on home networking, operation of servers, commercial use, and overuse of bandwidth.

An illustrative example is the difference in attitudes toward home networking.⁴⁵ At the extremes, then-Cable operator AT&T Broadband defined home networking as "theft of services" and threatened subscribers with civil and criminal penalties.⁴⁶ In contrast, DSL provider Verizon makes it clear in its service contract that home networking is permissible, as does Sprint.⁴⁷

There existed variation between individual cable operators and DSL operators on some of the restrictions. On the cable side, AT&T Broadband and Comcast (later combined to form the nation's largest cable operator), stood out for having the strictest usage restrictions. AOL Time-Warner, Charter Communications and smaller operators CableOne and Insight Broadband had the least restrictions. Among DSL operators, BellSouth stood out with the most restrictions, similar in extent to a cable operator. Overall, perhaps the most "liberal" broadband provider was DSL provider Sprint. Sprint has very few usage restrictions, tells subscribers in FAQs that they may run home networks,

44. At the time the survey was conducted, AT&T and Comcast were still operating independently.

45. Home networking refers to the practice of sharing a broadband connection amongst all of the computers in a home, as opposed to the single computer attached to the cable modem. This usually requires the purchase of additional equipment, such as a home router.

46. *AT&T Broadband Internet Subscriber Agreement*, § 6(g), available at http://help.broadband.att.com/listfaqs.jsp?category_id=973&category-id=34 (last revised Dec. 5, 2001).

47. Verizon Online Internet Access, *Terms of Service*, available at <http://www.verizon.net/policies/internetaa.asp> (2003).

web servers, and promises users that they "will have complete unrestricted access to all content available on the Internet."⁴⁸

On the architectural side, the outstanding deviation from neutrality in broadband networks today is the asymmetric bandwidth common across networks. Other, future controls may include application specific controls, as the survey of equipment vendors' offerings shows.

A. Contractual Restrictions

We first consider the individual types of restrictions found in usage agreements, focusing attention on restrictions that are likely to influence the development of certain application-types. The following chart shows the 13 main types of restrictions along with the percentage of major cable operators and DSL operators who stated such restrictions:

TABLE 1. MAJOR USAGE RESTRICTIONS

RESTRICTION	CABLE	DSL
Using a Virtual Private Network	10%	0%
Attaching WiFi Equipment	10%	0%
Making the Connection a Network End Point	10%	0%
Using Home Networking	40%	0%
Misusing IP Addresses	60%	0%
Any Commercial or Business Use	100%	33%
Operating a Server or Providing Public Information	100%	33%
Overusing Bandwidth	100%	33%
Reselling Bandwidth or Acting as an ISP	100%	33%
Conducting Spam or Consumer Fraud	100%	100%
Hacking or Causing Security Breaches	100%	100%
Any Unlawful Purpose	100%	100%
Any Offensive or Immoral Purpose	100%	100%

The appendix indicates which operators in the survey implemented the restrictions above. The following pages provide further details on the language of the most controversial restrictions: (1) providing information to the public or operating a server, (2) commercial uses, (3) Home Networking, and (4) WiFi network operation.

48. Sprint FastConnect DSL, *Frequently Asked Questions*, available at <http://csb.sprint.com/home/local/dslhelp/faq.html#gen16> (2003).

1. Restrictions on Providing Content

Nearly every cable operator and one third of DSL operators restricted operating a server and/or providing content to the public.⁴⁹ This restriction has the greatest potential significance because it affects the broadest class of applications—those where the end-user shares content, as opposed to simply downloading content. The potential breadth of server restriction can be seen from AT&T Broadband's acceptable use agreement:

[Subscriber may not] run programs, equipment or servers from the Premises which provide network content or any other services to anyone outside of the your home Examples of prohibited programs and equipment include, but are not limited to, mail, ftp, http, file sharing, game, newsgroup, proxy, IRC servers, multi-user interactive forums and Wi-Fi devices.⁵⁰

Again, this restriction can be understood as favoring a “one-to-many” or vertical model of application over a “many-to-many” or “horizontal” model. In application design terms, the restriction favors client-server applications over peer-to-peer designs.⁵¹ If taken seriously, the inability to provide content or act as a server would serve to restrict a major class of network applications.

Not all the restrictions are as broad as AT&T Broadband's. More typical is a simple ban on servers, as seen in this example from Cox Systems: “Servers You may not operate, or allow others to operate, servers of any type or any other device, equipment, and/or software providing server-like functionality in connection with the Service, unless expressly authorized by Cox.”⁵² Others, like Charter Communications, name banned applications: “Customer will not use, nor allow others to use, Customer's home computer as a web server, FTP server, file server

49. The exception is Time Warner. See *infra* Appendix.

50. AT&T Broadband Internet Acceptable Use Policy, ¶ xiv, available at http://help.broadband.att.com/faq.jsp?content_id=1107&category_id=34 (revised July 25, 2002).

51. The Internet's most popular application of the early 1990s—the world wide web—followed a client-server design, where a single specialized, centralized server provides services to a large number of clients. However, today an increasing number of applications use fully or partially decentralized designs. Email was always partially decentralized, for example, and the many popular “chat” programs embody a design that technically requires the user to act as a server as well as a client. Similarly, users who want to access a home computer from work (using, for example, rlogin) need to set up the home computer to act as a server. Peer-to-peer application designs also ask home users to act both as a client and server.

52. Cox Systems, Acceptable Use Policy § 6, available at <http://www.cox.com/iNetIncludes/policy/acceptable.asp> (updated Apr. 28, 2003). See also AT&T Broadband Internet Acceptable Use Policy, *supra* note 50.

or game server or to run any other server applications.”⁵³ The narrowest form of server restriction is seen in the Verizon terms of service agreement: “You may not use the Service to host a dedicated or commercial server.”⁵⁴ Finally, contrary to others, DSL provider Sprint suggests that consumers may, in fact, run a web server, based on the following excerpt from Sprint’s FAQ site:

Q: Can I run a web server?

A: Yes it is possible to set-up a web server using your Sprint FastConnect DSL service.⁵⁵

2. Bans on Commercial Use

A second restriction with potential implications for application development is a limit on “commercial” or “enterprise” use of residential broadband connections. Every cable operator and most DSL operators surveyed had some ban on using a basic residential broadband connection for commercial use.

The broadest and most controversial of such restrictions barred home users from using “Virtual Private Network” (VPN) services, which are used by telecommuters to connect to their work network through a secure connection. Cox Systems provides an example of a ban on Virtual Private Networks: “You agree not to use the Service for operation as an Internet service provider, or for any other business enterprise, including, without limitation, virtual private network usage, IP address translation, or similar facilities intended to provide additional access.”⁵⁶ More typical bans on commercial use came in the following form, as seen in the Time Warner Subscriber Conduct provision in its acceptable use agreement:

The ISP Service as offered and provided under this Agreement is a residential service offered for personal, non-commercial use only. Subscriber will not resell or redistribute (whether for a fee or otherwise) the ISP Service, or any portion thereof, or otherwise charge others to use the ISP Service, or any portion thereof. Subscriber agrees not to use the ISP Service for operation as an internet service provider, for the hosting of websites (other than as expressly permitted as part of the ISP Service) or for any enterprise

53. Charter Communications Pipeline, *Acceptable Use Policy* § 1(A), available at <http://www.chartercom.com/site/rules.asp#aup> (last checked Oct. 8, 2003).

54. Verizon Online Internet Access, *Terms of Service*, *supra* note 47, at § 2.4(C).

55. Sprint FastConnect DSL, *Questions & Answers*, available at http://csb.sprint.com/servlet/Faq/faq_category?category=DSLGenQuestions (2003).

56. Cox Systems, *Acceptable Use Policy*, *supra* note 52, at § 5.

purpose whether or not the enterprise is directed toward making a profit.⁵⁷

Again, the limitations found in DSL restrictions were far less extensive. For example, the BellSouth subscriber agreement mixed the restrictions on providing content and acting commercially as follows: "Subscribers may not provide public or commercial information over such [residential DSL] connections."⁵⁸

3. Home Networking

When home networking first became widespread in 2002, four of ten of the nation's largest cable operators contractually limited the deployment of home networks.⁵⁹ They did so by stating restrictions on the number of computers that could be attached to a single connection. The strongest example of such a usage restriction in 2002 came from AT&T Broadband:

THEFT OF SERVICE. Customer shall not connect the Service or any AT&T Broadband Equipment to more computers, either on or outside of the Premises, than are reflected in Customer's account with AT&T Broadband. Customer acknowledges that any unauthorized receipt of the Service constitutes theft of service, which is a violation of federal law and can result in both civil and criminal penalties. In addition, if the violations are willful and for commercial advantage or private financial gain, the penalties may be increased.⁶⁰

A milder approach was taken by Adelphia's online FAQ:

Can I network more than one computer?

Yes. Please check with a reputable computer electronics retailer for home networking solutions that are right for you. Adelphia will support a cable modem that is connected to a hub or router to the gateway or host computer. Adelphia does not install or support the

57. Time Warner, *Cable Modem Service Subscription Agreement* § 5(a), available at http://help.twcable.com/html/twc_sub_agreement.html (last visited Oct. 8, 2003).

58. BellSouth Internet Service, *Acceptable Use Policies*, available at http://home.bellsouth.net/csbellsouth/s/editorial.dll?fromspage=cg/legal/legal_homepage.htm&categoryid=&bfmind=354&ceid=376138&ceetype=article&xrender=y5ck= (last visited Oct. 8, 2003).

59. MediaOne, Comcast, AT&T and Adelphia. Due to enforcement difficulties and the ongoing regulatory proceedings at the Federal Communications Commission, most of these restrictions have been rescinded.

60. *AT&T Broadband Internet Subscriber Agreement*, § 6(g), at http://www.attbi.com/general-info/bb_terms.html (last visited Mar. 13, 2003).

network. Adelphia Power Link may not be connected to a broadcast server of any kind.⁶¹

In contrast, some DSL operators in their agreements explicitly acknowledged that multiple computers could be connected to the DSL connection. As Verizon's agreement stated: "You may connect multiple computers/devices within a single home or office location to your DSL modem and/or router to access the Service, but only through a single DSL account and a single IP address obtained from Verizon Online."⁶² Other DSL providers were vague. For example, in BellSouth's terms of service: "Unless otherwise specified in the BellSouth Internet Service subscriber's pricing plan agreement, sharing of accounts and/or connections on unlimited usage plans with anyone other than immediate family members in the same dwelling is strictly prohibited."⁶³

4. Restrictions on Wireless (WiFi) Networks

In addition to restrictions on home networking, several cable operators signaled a particular interest in controlling the deployment of home wireless networks. This is clearest with AT&T Broadband: They explicitly banned the connection of "Wi-Fi" equipment.⁶⁴ The provider also made it a breach of the subscriber's agreement to maintain a WiFi service that is available to outsiders. "[It is a breach of the agreement to] resell the Service or otherwise make available to anyone outside the Premises the ability to use the Service (i.e. WiFi, or other methods of networking)."⁶⁵

B. Architectural Controls, Present & Future

1. Present

Today, the principal deviation from network neutrality through architecture is, and continues to be, asymmetric bandwidth: that is, the practice of designing networks to provide more "downstream" bandwidth than "upstream." It is difficult to obtain a full set of data on the extent of

61. *Adelphia FAQ, Home Networking*, at http://www.adelphia.com/high_speed_internet/faqs.cfm (last visited Mar. 13, 2003).

62. *Verizon Online's Terms of Service*, § 2.5B, at <http://www.verizon.net/policies/internetaa.asp>.

63. See BellSouth, *Acceptable Use Policies*, *supra* note 58.

64. *AT&T Broadband Internet Acceptable Use Agreement*, *supra* note 50, at ¶ 14 ("Examples of prohibited . . . equipment include . . . Wi-Fi.").

65. *Id.* at ¶ ix. Cox Systems, *Acceptable Use Policy*, *supra* note 52, at 17, has a similar restriction.

asymmetry, because many cable operators do not make public the maximum bandwidth permitted by their networks. However, from the few sources of data that are available, we find that there is greater asymmetry in cable networks than DSL—though the shared architecture of cable networks makes the significance of this fact unclear. Published DSL rates included residential bandwidth with as low as 1:1 ratios, while the modal ratio is 6:1 ratios.⁶⁶ The few cable networks with public data promised maximum bandwidth ratios ranging from 5.3:1 (Time Warner / Earthlink) to as much as 12:1 (Cox Communications).⁶⁷

As others have recognized, allowing more downstream than upstream bandwidth obviously favors the development of applications that are one-to-many, or client-server in design. Applications that would demand residential accounts to deliver content as quickly as they receive it will do less well under conditions of asymmetric bandwidth.

2. Future – Better Bandwidth Management or Application Layer Controls?

It is difficult to predict what application controls broadband operators might implement in the future. Yet future possibilities can be gleaned from the marketing efforts of equipment vendors who target the cable and DSL market. Two trends can be briefly noted, though the full topic is well beyond the scope of this paper.

First, over the last several years, several companies have begun to market equipment described to facilitate application-based screening and control for broadband networks. Two prominent examples are Allot Communications and Packeteer Communications. The former markets a product named “NetEnforcer” to cable and DSL operators,⁶⁸ promising to control problems from both peer-to-peer traffic and unauthorized WiFi connections.⁶⁹ Allot’s competitor, Packeteer, markets a similar product, named “PacketShaper,” described as “an application intelligent traffic management appliance providing visibility into and control over network utilization and application performance.”⁷⁰ The company claims that the product is used on hundreds of University campuses, primarily to

66. See *infra* Appendix.

67. *Id.*

68. Allot Communications Netenforcer® Data Sheet, at http://www.allot.com/html/products_netenforcer_sp.shtml (last visited Mar. 13, 2003).

69. Jim Barthold, *Allot looks to help servers with bandwidth congestion problems*, TELEPHONY.ONLINE.COM, available at http://telephonyonline.com/ar/telecom_allot_looks_help/index.htm (Dec. 3, 2002).

70. Packeteer, at <http://www.packeteer.com/products/packetshaper.com> (last visited Mar. 13, 2003).

control peer-to-peer traffic.⁷¹ When this survey was conducted, despite the marketing efforts of both companies, there was no evidence of deployment by cable or DSL operators. It is therefore impossible to conclude whether broadband operators will begin using technological means to facilitate restrictions on usage.

Second, vendors of cable data equipment promise improved bandwidth management capabilities as between individual customers on cable networks.⁷² This is the promise of the DOCSIS⁷³ 1.1 and 2.0 standards, which are an update to the current DOCSIS 1.0 standard in use today.⁷⁴ As the new equipment is not yet widely deployed, these claims or their impact cannot be verified.

C. Conclusions & Evidence of Enforcement

What, generally, can be concluded from this survey? On the one hand, there is no broad effort to ban everything that might be said to threaten the interests of cable and DSL operators. For example, cable operators have not now barred streaming video, despite the potential to compete with cable television, and despite Dan Somers' famous comment that "AT&T didn't spend \$56 billion to get into the cable business to have the blood sucked out of [its] veins."⁷⁵ This conclusion is reinforced by the general perception that broadband access is not substantially limited.

To what degree are these usage restrictions enforced? While there is little formal data on enforcement patterns, there exists anecdotal evidence of enforcement on websites like DSL Reports,⁷⁶ which are dedicated to users complaining about broadband service and usage restrictions. Some examples of enforcement include the enforcement of monthly or daily bandwidth limits through a threatening to terminate or restrict the accounts of users who use too much bandwidth in a single month. For example, Cox Cable in November 2002 sent letters to users

71. Gwendolyn Mariano, *Schools declare file-swapping truce*, CNET NEWS.COM, at <http://news.com.com/2100-1023-859705.html?tag=rn> (Mar. 14, 2002).

72. See, e.g., <http://www.cisco.com/warp/public/779/servpro/solutions/cable> (last visited Mar. 13, 2003).

73. DOCSIS stands for Data Over Cable Service Interface Specifications. See *Seven Cable Modem Manufacturers Seek DOCSIS Certification*, CABLELABS, at <http://www.cablelabs.com/news/newsletter/SPECS/specnewsaug/news/pgs/story2.html> (last visited Mar. 13, 2003).

74. For an explication of the claims of DOCSIS 1.1 and 2.0, see CABLELABS, DOCSIS, *supra* note 38.

75. See David Lieberman, *Media Giants' Net Change Establish Strong Foothold Online*, USA TODAY, Dec. 14, 1999, at B3 (Dan Somers was CEO of AT&T Broadband at the time the comment was reported).

76. See BROADBAND REPORTS.COM, at <http://www.dslreports.com> (Mar. 2002).

who downloaded more than 2 gigabytes of bandwidth per day, or 30 gigabytes of bandwidth per month.⁷⁷ Other cable operators, though no DSL providers, have suggested similar policies may be on their way.⁷⁸ In addition, broadband consumers have complained of efforts to enforce specific bans on applications, such as threats to enforce contractual limits on VPN operations⁷⁹ and users who run file-sharing applications.⁸⁰

IV. A PROPOSAL FOR NETWORK NEUTRALITY

Recognizing that discrimination in broadband service is a potential problem is one thing; constructing an approach to dealing with it, is another. The open-access proposal, as we saw earlier, advocated structural separation between Internet service providers and broadband operators. This approach has the advantage of simplicity, but it has the disadvantage of retarding potential efficiencies of integration. This approach also may fail to deter other forms of discrimination.

What follows is a proposed antidiscrimination principle (a rule, only if necessary). The effort is to strike a balance: to forbid broadband operators, absent a showing of harm, from restricting what users do with their Internet connection, while giving the operator general freedom to manage bandwidth consumption and other matters of local concern. The principle achieves this by adopting the basic principle that broadband operators should have full freedom to “police what they own” (the local network) while restrictions based on inter-network indicia should be viewed with suspicion.

This non-discrimination principle works by recognizing a distinction between *local network* restrictions, which are generally allowable, and *inter-network* restrictions, which should be viewed as suspect. The principle represents ultimately an effort to develop *forbidden* and *permissible* grounds for discrimination in broadband usage restrictions.

77. See Karl Bode, *Defining Gluttony: Cox Cable Gets Specific*, at <http://www.dslreports.com/shownews/23465> (Nov. 12, 2002).

78. John Borland, *ISP download caps to slow swapping?* CNET NEWS.COM, at <http://news.com.com/2100-1023-975320.html> (Nov. 26, 2002).

79. Practically Networked Earthweb, VPN Comcast Letter, at <http://www.practicallynetworked.com/news/comcast.htm>. (last visited Mar. 10, 2003).

80. Many users have accused cable operators of blocking specific file-sharing applications like KaZaa, through port blocking, though the reports are unverified. See, e.g., *RoadRunner Blocking kaZaA*, ZEROPAID.COM, at <http://www.zeropaid.com/news/articles/auto/07142002a> (July 13, 2002).

A. Let Operators Police What They Own

Broadband carriers are members of two networks. They are each members of a local network, which they own and manage individually. They are also members of the inter-network, which they collectively manage with other service providers.

FIGURE 1: BROADBAND CARRIERS, MEMBERS OF TWO NETWORKS

Once we recognize that carriers are engaged in a collective management scheme, the origin of the externalized cost problem described above becomes clear. The effects of local network restrictions will, usually, affect only the network run by a single service provider. Such restrictions moreover, are necessary for good network management. In contrast, by definition, restrictions at the inter-network layer or above will always affect the entire network, and can create externality problems.

B. The Neutrality Principle

What follows is an example of a network neutrality law:

§__ Forbidding Broadband Discrimination

- (a) Broadband Users have the right reasonably to use their Internet connection in ways which are privately beneficial without being publicly detrimental. Accordingly, Broadband Operators shall impose no restrictions on the use of an Internet connection except as necessary to:
 - (1) Comply with any legal duty created by federal, state or local laws, or as necessary to comply with any executive order,

- warrant, legal injunction, subpoena, or other duly authorized governmental directive;
- (2) Prevent physical harm to the local Broadband Network caused by any network attachment or network usage;
 - (3) Prevent Broadband users from interfering with other Broadband or Internet Users' use of their Internet connections, including but not limited to neutral limits on bandwidth usage, limits on mass transmission of unsolicited email, and limits on the distribution of computer viruses, worms, and limits on denial-of service-or other attacks on others;
 - (4) Ensure the quality of the Broadband service, by eliminating delay, jitter or other technical aberrations;
 - (5) Prevent violations of the security of the Broadband network, including all efforts to gain unauthorized access to computers on the Broadband network or Internet;
 - (6) Serve any other purpose specifically authorized by the Federal Communications Commission, based on a weighing of the specific costs and benefit of the restriction.
- (b) As used in this section,
- (1) "Broadband Operators" means a service provider that provides high-speed connections to the Internet using whatever technology, including but not limited to cable networks, telephone networks, fiber optic connections, and wireless transmission;
 - (2) "Broadband Users" means residential and business customers of a Broadband Operator;
 - (3) "Broadband Network" means the physical network owned and operated by the Broadband Operator;
 - (4) "Restrictions on the Use of an Internet Connection" means any contractual, technical, or other limits placed with or without notice on the Broadband user's Internet Connection.

This law expressed the inter-network neutrality principle, operationally, as a non-discrimination rule. As the analysis above recognized, the concept of a total ban on network discrimination is counterproductive. Rather, we need distinguish between forbidden grounds of discrimination, those that distort secondary markets, and permissible grounds, those necessary to network administration and harm to the network.

Reflecting the dual-network membership just described, it will be inter-network criteria of discrimination that cause concern. In technical terms, this means discrimination based on IP addresses, domain name, cookie information, TCP port, and others as we will describe in greater

detail below. Hence, the general principle can be stated as follows: absent evidence of harm to the local network or the interests of other users, broadband carriers should not discriminate in how they treat traffic on their broadband network on the basis of inter-network criteria.

The negative inference (expressed most clearly in exceptions (a)(3) and (4)) is that operators generally may discriminate in their treatment of traffic on the basis of local network criteria. In technical terms, this means imposing restrictions on the basis of what network engineers call "link" or "layer 2" information, like bandwidth, jitter, or other local Quality of Service indicia.

C. In Practice: Online Gaming

Popular online gaming applications⁸¹ like Everquest, Asheron's Call, or Online Quake tend to be bandwidth intensive, particularly compared with episodic applications like email. As seen above, concerned broadband carriers have therefore been inclined to restrict the usage of such applications. However, with the neutrality principle in mind, we can distinguish between a "better" and a "worse" way for this to happen.

First, in today's environment, a broadband carrier could block traffic from gaming sites. It could do it either by enforcing a contractual provision in a usage agreement, or in the future, using its control of the local network to block traffic from gaming sites based on either application information, or the IP address of the application provider.⁸² Some carriers might elect, for a given supplemental fee, to remove the filter for specified users.

Under the neutrality principle here proposed, this approach would be frowned upon. Instead, a carrier concerned about bandwidth consumption would need to invest in policing bandwidth usage, not blocking individual applications. Users interested in a better gaming experience would then need to buy more bandwidth—not permission to use a given application.

The neutrality of such control would prevent the distortion in the market for Internet applications. If carriers choose to block online games in particular, this gives a market advantage to competing applications that have not been blocked. But if broadband carriers only police bandwidth, the result is an even-playing field. It may be that the expense

81. Also commonly referred to as "Massively Multiple Online Games," or MMOGs.

82. For an explanation of how a broadband carrier would do so, see, e.g., *The Cisco Content Delivery Network Solution for the Enterprise*, Cisco White Paper (Apr. 2002), available at http://www.cisco.com/warp/public/cc/so/neso/ienesv/cxne/cdnen_wp.htm; See also Cosine Communications., *Digital Subscriber Lines and Managed Network-based Services: A Perfect—and Profitable—Marriage*, White Paper, available at <http://cnscenter.future.co.kr/resource/rsc-center/vendor-wp/cosine/dslwp.pdf>.

of more bandwidth leads people to choose different ways to spend their money. But if so, that represents a market choice, not a choice dictated by the filtering policy of the broadband carrier.

D. Borrowing from Well-Established Categories

One advantage of the proposal is that it relies on well-established legal and technological criteria to achieve its consumer-welfare goals. Respectively, it borrows from principles of harm requirements and non-discrimination familiar to lawyers, along with a local/inter-network distinction that is fundamental to Datacom networks.

1. The Harm Requirement

In the telephony context, the “foreign attachment” problem discussed above was addressed by a “harm” rule; that is, a rule barring the Bells from preventing attachment of equipment unless harm to the network could be shown. Its origins are found in the Hush-a-Phone case, where the FCC ordered Bell to allow telephone customers to attach devices that “[do] not injure . . . the public in its use of [Bell’s] services, or impair the operation of the telephone system.”⁸³

In the broadband context, it is discrimination against certain content and applications that is the major problem. But the practice of requiring public harm to justify restrictions can be usefully employed.

2. Local/Inter-Networking

Finally, on the technological side, the distinction between inter-networking and local networking is very well established in the Datacom industry. While the distinction is best reflected, and usually discussed, in the context of the OSI network reference model (as the difference between layer 2 and layer 3 networks),⁸⁴ it is in fact independent of OSI. As a practical matter, different physical equipment and different protocols run the different networks. In a given network, “switches” run local networks, while “routers” collectively manage the layer 3 network. Services can be offered at both levels—for example, VPNs and telephony can be offered either as a layer 2 service or as a layer 3 service.

In addition, other schema used to describe network layers embody the same, fundamental, local / inter-network distinction. For example,

83. Hush-A-Phone Corp. v. AT&T, 22 FCC 112, 114 (1957). This led in turn to the broader *Carterfone* decision, 13 F.C.C.2d 420 (1968), and finally Part 68, which adopted a protective circuitry approach to protecting the telephone network, see 47 CFR §68 *et seq.*

84. Cf. ANDREW TANENBAUM, COMPUTER NETWORKS 10-18 (4th ed. 2002).

the TCP/IP network model maintains a distinction between the "link" layer and the "network" layer. This is exactly the same distinction as the layer 2/layer 3 distinction in the OSI model, and the local/inter-network distinction more generally. Again, this is no surprise, because virtual description simply reflects the physical network design. The existence and pervasiveness of the local / inter-network distinction makes it a natural dividing line for reasonable restrictions on use.

V. OBJECTIONS TO THE PROPOSED NETWORK NEUTRALITY REGIME

Before concluding, it is useful to consider some objections and challenges to the proposed network neutrality regime. We consider (1) whether it overly interferes with broadband carriers' ability to earn a return on their infrastructure investment, (2) whether local restrictions can be used to achieve the same problems as inter-network control, and (3) whether the principle interferes with administration of Internet addressing.

A. Return on Investment

First, does the neutrality principle restriction overly impinge on the ability of broadband carriers to earn a return from their infrastructure investments? While a full analysis of broadband economics is beyond the scope of this proposal, we can nonetheless suggest that the neutrality principle is unlikely to interfere with the special advantages that a carrier gains from building its own infrastructure.

The simple answer is that investing in a local network infrastructure creates its own rewards, as it creates particular advantages in the offering of network services. We can see this clearly by considering the particular example of Virtual Private Networks under the neutrality principle. A broadband operator who owns the local infrastructure has a natural advantage in offering local VPN services. The advantage comes from the fact that they can offer service level guarantees that cannot be provided on a shared network. Nothing in the neutrality principle would prevent a broadband operator from being in the unique position to sell such services.

But the principle would prevent operators from blocking use of Internet VPNs – that is, VPNs that used the Internet to reach sites that no single local network can encompass. For example, a home user on the East Coast will almost certainly need to use an Internet VPN to connect to his business on the West Coast. In offering this service, a broadband operator is in the exact position as any other Internet VPN

provider. Restricting use of Internet VPNs should therefore not be allowed, to preserve undistorted competition for this application.

B. Can Local Control Disrupt Application Markets?

Some might observe that the local and inter-network are interdependent in certain ways. Won't broadband operators simply use their control over the local network to achieve the same distortion of application markets?

No rule can perfectly stamp out all undesirable behavior. The point of the network neutrality principle is to make interference with the application markets much harder. Without the ability to discriminate on the basis of the origin of a packet or the application being used, the broadband carrier is left with the far blunter tools of local restrictions.

It might be argued that the address resolution protocol (ARP)⁸⁵ could be used to achieve the same goals as IP-address filtering, since the job of ARP on a typical network is to convert IP addresses into Ethernet MAC addresses. But, in fact, a broadband carrier manipulating ARP could only succeed in making his own users unreachable. The ARP-cache only holds the information to match up local physical addresses with local IP addresses. ARP has no idea how to stop a user from reaching a specific IP address, other than making that user unreachable. The example shows, in fact, the power of limiting a broadband carrier to local control.

C. The Need to Administer IP

Finally, some might point out that broadband carriers need some control over the Internet Protocol side of their network. They must, for example, be able to allocate static and dynamic IP addresses, maintain routing tables, and so on. Does the network neutrality principle interfere with this?

The point of the neutrality principle is not to interfere with the administration of the Internet Protocol side of a broadband carrier's network. It is, rather, to prevent discrimination in that administration. Since it is phrased as a non-discrimination principle, a negative inference is that most aspects of IP administration can be conducted without concern. For example, the allocation and administration of IP

85. Described in IETF RFC 826, available at <http://www.ietf.org/rfc/rfc1027.txt>.

addressing should not pose any discrimination problems, so long as the administration of such addresses is in an even-handed manner.⁸⁶

VI. CONCLUSION

The goal of this paper was to make an initial case for a broadband discrimination regime as an alternative to the structural remedy of open access to achieve the goal of network neutrality. At this point, the newness of the concept means much unavoidable vagueness as to its operation. It is easier to point out examples of application discrimination that seem unjustified than to elucidate a standard that nearly separates the legitimate from the suspect. For example, there remains much work to better define what the concepts of network neutrality and discrimination would fully entail as a regulatory matter, or even as a regulatory threat. Should neutrality be defined by IETF standards? The intuitions of network theorists? Government definition? Any workable regime designed to achieve network neutrality will need a more precise conception of this and other matters. Nonetheless, the hope is that the general framework described here might serve to begin the effort to discourage the most blatant or thoughtless disfavoring of certain application types through network design.

86. In today's environment, the scarcity of IPv4 addresses does appear to justify a form of discrimination: charging more for static addresses, than dynamic addresses. This forms a good example of "permissible" discrimination.

DSL OPERATORS:

RESTRICTION	VERIZON	SBC	QWEST	BELLS	SPRNT	WLDCM	FREQ
Home Networking	OK	OK			OK		0%
Operating a Server	R			R	OK		40%
Commercial / Enterprise / Business Use	R			R			40%
Overuse of Bandwidth	R			R			40%
Resell Bandwidth	R			R			40%
Spam / Consumer Fraud	R	R	R	R	R	R	100%
Hacking / Security Breaches	R	R	R	R	R	R	100%
Any Offensive or Immoral Purpose	R	R	R	R	R	R	100%

Legend:

R = Contractually Restricted

OK = Explicitly Permitted

CmCst = ComCast Communications

Cox = Cox Communications

CableV = CableVision, Inc.

Insight = Insight Communications

BellS = BellSouth

AT&T BB = AT&T Broadband

TW = Time Warner

Chartr = Charter Communications

Adphia = Adelphia Communications

MediaCM = MediaCom

Cable1 = CableOne

Sprnt = Sprint

UPSTREAM / DOWNSTREAM BANDWIDTH RATIOS

PROVIDER NAME	BANDWIDTH DOWN (K)	BANDWIDTH UP (K)	RATIO
QWEST	256	256	1:01
	640	256	2.5:1
SPRINT	256	96	2.66:1
	512	128	4:01
VERIZON	1.5M	256	6:01
	768	128	6:01
	1.5M	128	12:01
SBC	384	128	3:01
BELLSOUTH	1.5M	256	6:01
WORLD COM	1.5M	256	6:01
	384	128	3:01
AT&T BB	1.5M	256	6:01
	3M	384	8:01
TIME WARNER	2	384	5.33:1
COX	3M	256	12:01