Barbara van Schewick
Assistant Professor of Law and (by Courtesy) Electrical Engineering
Co-Director, Center for Internet and Society
Stanford Law School

**Official Testimony at the
Federal Communications Commission's Second Public En Banc Hearing
on Broadband Network Management Practices
at Stanford University, Stanford, CA on April 17, 2008,
Docket No. 07-52**

Mr. Chairman, Commissioners,

Thank you for giving me the opportunity to testify today. I would like to make four points:

1. The FCC should require broadband providers to disclose their broadband management practices.
2. The current level of disclosure is not sufficient.
3. Disclosure alone is not enough.
4. The FCC needs to establish some ground rules for reasonable network management. In particular, it should ban network management practices that single out specific applications or classes of applications in order to manage bandwidth consumption on broadband networks.

1. The FCC should require broadband providers to disclose their broadband management practices.

Full disclosure is a necessary condition to enable competition to work. Disclosure improves competition by enabling customers to make informed decisions when choosing providers. Disclosure also enables competitors to differentiate themselves along these dimensions.

Today, network providers in the US compete based on maximum upload and download speed and price. If, however, customers are unable to note the differences between the offerings along other dimensions (e.g. how oversubscribed is the network, how often is traffic management used, how is traffic prioritized), they cannot take these factors into account when making a decision, and network providers do not have an incentive to compete on these factors.

Compare this with the situation in Europe: A quick look at network provider offerings shows that ISPs compete on much more than just price and maximum upload or download bandwidth. For example, BT offers customers three different options with differing monthly bandwidth allowances that indicate the maximum amount of bandwidth you are allowed to use in a month. To protect customers from unexpected increases in their broadband bills, BT does not charge customers if they exceed their monthly usage allowance in one month.[1] Laribu, a provider in Belgium, offers three service packets which offer different combinations of upload/download speed and monthly bandwidth allowances. Usage that occurs during 2 am and 8 am only counts with half of the actual bandwidth used, an offering that clearly targets customers interested in peer-to-peer file sharing.[2] PlusNet, an ISP in the UK that is an independent subsidiary of BT, offers different combinations of maximum upload/download speed, monthly volume caps and traffic prioritizations; traffic between midnight and 8 am is not counted towards the monthly volume cap.[3] This type of differentiation enables someone who knows she wants to do a lot of file-sharing or who is an avid gamer to choose the Internet service offering that best supports her needs.

The UK providers have clearer descriptions of their acceptable use policies, and of the type of network management they use in their network.[4] For example, BT's policy usage allowance and fair use policy informs customers that "[…] we restrict P2P speeds if it's having a negative impact on the online experience of the majority of our customers. We normally place restrictions in the evenings at peak time, but we do apply them during the day if a lot of customers are using P2P at the same time."[5] BT's broadband management practices are very different from Comcast, and their disclosure is clearer. PlusNet explicitly states how different traffic is prioritized on their network, and what bandwidth rates customers can expect for different applications at different times of day.[6] As a side effect, such disclosure may help alleviate congestion by enabling customers to adjust their behavior.

Disclosure provides visibility to regulators, competitors and industry observers. It avoids the waste of resources spend when users or application or content providers try to figure out what is going on on a particular network, as when the Associated Press, EFF or private parties like Robert Topolski and David Reed were running tests on Comcast's network to understand what Comcast was doing. Additionally, testing by end users may be unable to detect the "next generation" of network management tools, which may be completely hidden from end users.

---

[1] BT (2008).
[2] LaTribu (2008).
[3] PlusNet (2008a); PlusNet (2008c).
[4] Not all these practices would be considered reasonable network management under the standard advanced below. The point here is that the disclosure is clearer than the disclosure of practices in the US.
[5] BT (2008).
[6] PlusNet (2008b); PlusNet (2008c).

Finally, under effective competition, disclosure can discipline provider's behavior.[7] Customers who do not like how a provider manages its traffic can switch to another provider.

2. The current level of disclosure is not sufficient.

To realize these goals, disclosed information must provide enough detail to enable customers to make an informed decision and to enable them to adjust their behavior. Comcast's current acceptable use policy falls short of these goals. Customers shopping for Internet service on Comcast's website are not directed to the acceptable use policy (if you click on terms and conditions, the pop-up window states: "Service is subject to terms and conditions of Comcast High-Speed Internet Subscriber Agreement and Home Networking Amendment if applicable. For restrictions, minimum requirements and details about service and prices, call 1-800-Comcast.")[8] While Comcast new acceptable use policy mentions that Comcast employs traffic management during periods of heavy usage, it does not give any indication when that may be, or how often it occurs. There also is an unresolved tension between the prohibition on running servers, and the allowing of peer-to-peer file-sharing applications. Compare "Examples of prohibited equipment and servers include, but are not limited to, e-mail, Web hosting, file sharing, and proxy services and servers"[9] with "Common activities that may cause excessive bandwidth consumption in violation of this Policy include, but are not limited to, numerous or continuous bulk transfers of files and other high capacity traffic using (i) file transfer protocol ("FTP"), (ii) peer-to-peer applications, and (iii) newsgroups, whether provided by Comcast or a third party."[10] and "Does Comcast block peer-to-peer ("P2P") traffic or applications like BitTorrent, Gnutella, or others? No. Comcast does not block P2P traffic or applications."[11] The "Frequently Asked Questions on Network Management and Excessive Use" are buried on the website under "Customers -> FAQs / Product Information / Comcast High-Speed Internet" under the heading "hot".[12] Comcast also reserves the right to change this policy at any time without giving notice to the consumer, forcing their customers to constantly monitor the acceptable use site, if they want to understand what is going on.

To support the effectiveness of disclosure, the FCC could set up a website where network providers would have to post their network management practices in a standardized format that would enable customers to compare providers' network management practices easily.

---

[7] van Schewick (2007), p. 376-377 (explaining how network providers' ability to secretly degrade the performance of applications limits the effectiveness of competition in the market for broadband Internet services); van Schewick (forthcoming 2009), chapter 5 (noting how disclosure requirements may alleviate the market power provided resulting from secret degradation).

[8] Comcast (2008a).

[9] Comcast (2008b), under "Technical restrictions".

[10] Ibid., under "Are there restrictions on bandwidth consumption that apply to the Service?".

[11] Comcast (2008d).

[12] Comcast (2008c).

3. Disclosure alone is not enough.

Disclosure can only facilitate competition and discipline providers if there is effective competition. In order for disclosure to have a disciplining effect, customers need to be able to switch to another provider that does not impose a similar restriction, and they need to be able to do so at low costs. In the US, none of these conditions is currently satisfied:

First, according to the GAO, the median number of broadband providers accessible to broadband Internet service customers in the US is two.[13] This market structure is often characterized as "duopoly +/-".[14] In some parts of the country, customers are facing a monopoly. According to the FCC's own data, 34 % of ZIP codes have one or less cable or ADSL provider who serves at least one subscriber living within the ZIP code.[15] As the GAO has pointed out, this measure (i.e. the number of providers reporting at least one subscriber in a certain ZIP code) overstates the level of competition to individual households.[16] While a duopoly is often better than a monopoly, duopolists enjoy a degree of market power that enables them to impose restrictions on their customers that they would not be able to impose in a competitive market.[17]

Second, as I have described in detail in my written work, the market for Internet services is characterized by significant switching costs that limit the effectiveness of competition. Although rules that require network providers to disclose whether and how they interfere with applications and content on their networks may reduce the problem of incomplete customer information, they cannot reduce the switching costs in the market for Internet services, and still leave the network provider with a substantial degree of market power over its customers that enables it to restrict some applications and content on its network without loosing too many Internet service customers. [18]

3. The FCC needs to establish some ground rules for reasonable network management. In particular, it should ban network management practices that single out specific applications or classes of applications in order to manage bandwidth consumption on broadband networks.

Without such a rule, "reasonable network management" becomes the back door that enables network providers to undermine the non-discriminatory nature of the Internet that the FCC's Internet Policy Statement is designed to protect. The Internet Policy Statement promotes user choice and fosters application-level innovation by providing application developers with the certainty that they will not be discriminated against. If network providers can single out specific applications in order to manage bandwidth on their network, application developers face a fundamental uncertainty. The network may turn against them at any time. This risk of being cut off from access to end users at any time

---

[13] United States Government Accountability Office (2006), p. 18.
[14] See, e.g., Farrell (2006), p. 202.
[15] Federal Communications Commission (2008), Table 16.
[16] United States Government Accountability Office (2006).
[17] See, e.g., Farrell (2006), pp. 202-205.
[18] van Schewick (2007), pp. 374-377; van Schewick (forthcoming 2009), chapter 5.

and at the sole discretion of the network provider constitutes a substantial risk that will make it much more difficult to get funding. For application developers, the impact of the threat of being discriminated against on their incentives to innovate[19] is the same, regardless of the motivation of the network provider, e.g. whether it is technically anticompetitive or not. Given the nature of the Internet as a general purpose technology, this reduction in application-level innovation is bad for society as a whole.[20]

Given that there is no real competition and significant switching costs, disclosure alone will not be sufficient to constrain network providers' incentives to impose restrictions that are not in the public interest.

As long as the current pricing structure continues to be based on flat-rate pricing, network providers have the perverse incentive to block or degrade applications on their network that use more bandwidth or deviate from the usage patterns for which the network was designed:[21] Applications that use more bandwidth than expected or exhibit usage patterns that deviate from the expected ones may put strain on the existing access network or the links to other providers, threatening congestion if the network provider does not upgrade the capacity of the network, or increase network providers' operating costs by raising their interconnection fees. From the network provider's point of view, blocking or degrading selected applications is a quick fix that requires less investment than upgrading the network or coming up with a non-discriminatory solution. As Brett Frischmann and I have shown in a recent paper, singling out specific applications to control bandwidth on a network has significant social costs that are not internalized by network providers. It harms application-level innovation by distorting the playing field between applications, and reduces consumer welfare by preventing users from using the applications of their choice.[22]

The original architecture of the Internet was based on a design principle called the end-to-end arguments.[23] As a result of this design, the network is general and can support a large variety of applications with different requirements. The network is not optimized in favor of specific applications. While this may increase the performance of particular applications, it also constitutes an unnecessary and therefore inefficient feature for applications that do not need this function and may even rule out the implementation of

---

[19] van Schewick (2007), pp. 378-380 (describing the impact of a threat of discrimination on application developers' incentives to innovate).

[20] Ibid., pp. 382-386; Frischmann and van Schewick (2007), pp. 423-425 (describing the benefits for social welfare associated with application-level innovation).

[21] MIT Communications Futures Program and Cambridge University Communications Research Network Broadband Working Group (2005); van Schewick (forthcoming 2009), chapter 5.

[22] Frischmann and van Schewick (2007).

[23] There are two versions of the end-to-end arguments: a narrow version, which was first identified, named and described in a seminal paper by Saltzer, Clark and Reed in 1981 (Saltzer, Reed and Clark (1981)), and a broad version which was the focus of later papers by the authors (e.g., Reed, Saltzer and Clark (1998); Blumenthal and Clark (2001)). While both versions have shaped the original architecture of the Internet, only the broad version is responsible for the application-blindness of the network.). For a detailed analysis of the two versions and their relationship to the architecture of the Internet, see van Schewick (2004), pp. 87-129.

applications that are not foreseen at the time of the design. Instead, all application-specific functionality is implemented at the end hosts, i.e. the computers at the edge of the network. As I show in my forthcoming book, this design results in an economic environment for innovation that is much more conducive to application-level innovation than network architectures that deviate from the broad version of the end-to-end arguments.[24]

However, while maintaining the openness for new applications, general solutions are sometimes more complicated to design and more costly to implement. As my forthcoming book shows, there is a market failure regarding the evolution of the core of the Internet.[25] While society as a whole has an interest in keeping the architecture non-discriminatory and general in order to preserve its openness for new applications, network providers' incentives are not necessarily aligned with this goal. Recent years have seen a myriad of technical solutions driven by the short-term interests of particular actors which do not take account of the impact on the long-term evolvability of the Internet.[26] The deployment of asymmetric DSL and cable, network address translators and now the Comcast tactics for managing bandwidth on their portion of the Internet are examples of such solutions. They solve the problem at hand, but at the same time reduce the generality of the Internet, with significant costs for application-level innovation. Network providers' preference for short-term solutions is not surprising:[27] While the benefits of deviating from the broad version such as the reduction in cost or the improvement in performance of certain applications are immediately apparent, the associated costs for the evolvability of the system are almost impossible to determine: The applications that may suffer from the deviation are yet unknown. This makes it impossible to determine whether and to what extent some of these future applications would be harmed. As the application is not known yet, the value of the application to society cannot be known either; this makes it impossible to determine what the ultimate costs to society will be. As research in behavioral economics has shown, humans tend to assign disproportionately more weight to present benefits than to future costs that are yet uncertain, making it more likely that network designers would deviate from the broad version of the end-to-end arguments. The fact that network providers are unable to capture all of the gains associated with a non-discriminatory Internet,[28] further exacerbates the problem.

By prudently defining limits to what constitutes reasonable network management, the FCC can provide network providers with incentives to work towards general solutions that preserve the openness of the Internet for new applications and consumer choice.

---

[24] van Schewick (forthcoming 2009).

[25] Ibid..

[26] See, e.g., Braden, Clark, Shenker and Wroclawski (2000), pp. 5-9.

[27] The following paragraph is based on van Schewick (forthcoming 2009), chapter 2.

[28] For example, they will not be able to capture all the productivity gains that users get from introducing new applications, or the positive spillovers associated with users' productive uses of the Internet (such as the improvement in democratic discourse resulting from widespread blogging on the Internet). For a long version of the argument, see Frischmann (2005); For a short summary, see Frischmann and van Schewick (2007), pp. 424-425.

If the openness of the Internet for new applications is to be preserved, the network providers need to manage their networks in a non-discriminatory way – without singling out specific applications or classes of applications. Of course, there needs to be an exemption for malicious applications that are engaged in an attack on the network.

Network providers would be able to enforce fairness among users, but how a user decides to use its "share" of bandwidth, both in general and at a particular point in time should be decided by the user. Network management solutions that enable network providers to police the amount of bandwidth used by a particular user are available today.

Some may claim that forcing network providers to treat a user's traffic in a non-discriminatory way may ultimately harm the user, as some of the user's applications may be more sensitive to the delay than others. They would suggest that enabling network providers to differentiate between classes of applications according to their needs (i.e. prioritize real-time VoIP traffic over file-sharing) would be in the interest of the user and should therefore be allowed.

While intuitively appealing, this view overviews that the value a user attaches to a particular application is not necessarily fixed:[29] my priorities may differ considerably depending on the circumstances: if I'm using BitTorrent to download a movie that I want to watch tomorrow, I do not care if the download is delayed a bit. If, however, I'm using BitTorrent to download a critical security patch that I need to get quickly or if I'm a programmer and want to download the source code for the project that I want to work on now, this may have priority over all the other applications I'm currently using.

Similarly, if I'm doing a VoIP call to my friend and just want to chat, I may not care as much about the quality of the call as when I'm doing a job interview or a project phone conference using VoIP and want to hear or be heard in a crystal clear way.

This implies that having users signal their priority instead of having network providers determine prioritization within the network based on classes of applications may be preferable.[30] For users who do not want to set their own priorities, network providers could offer different sets of potential prioritizations, among which users could choose. Thus, network providers could offer and users could opt-in to prioritization based on application class, but without enforcing this prioritization on users whose needs differ. (To maintain the non-discriminatory nature of the Internet, prioritization choices offered by the network provider would have to treat applications belonging to the same class in the same way).

Why the emphasis on user choice? First, user choice is fundamental if the Internet is to create the maximum value to society. The Internet is a general purpose technology. It does not create value through its existence alone.[31] It creates value by enabling users to

---

[29] Briscoe, Moncaster and Burness (2007), section 3.3.
[30] David Clark made the same point when he testified at the first FCC En Banc Hearing on Broadband Network Management Practices in Harvard.
[31] See, e.g., van Schewick (2007), pp. 385-386; van Schewick (forthcoming 2009).

do the things they want or need to do. Users know best what this is. As a result, users, not network providers should be able to decide how they would like to use the network, and what is important to them. Of course, in order for users to behave efficiently, they also need to bear (at least some of) the costs of their actions, something which the current system does not sufficiently provide.

User choice is also a fundamental component of the mechanism that enables application-level innovation to function effectively.[32] In the current Internet, it is impossible to predict what future successful applications will be. Enabling widespread experimentation at the application-level and enabling users to choose the applications they prefer is at the heart of the mechanism that enables innovation under uncertainty to be successful.

By singling out specific applications, network providers start picking winners and losers on the Internet. As we have seen, whom they pick may be driven by a number of motivations that are not necessarily identical with what users would prefer, leading to applications that users would not have chosen and forcing users to engage in an Internet usage that does not create the value it could. Consumers, not network providers, should continue to choose winners and losers on the Internet.

While there are mechanisms available that enable network providers to manage their networks in non-discriminatory ways now, the ways in which the Internet bandwidth on the Internet will be allocated between users and applications and in which it will deal with congestion is still evolving. Constraining the range of possible solutions in a way that provides network providers with incentives to evolve the Internet in a way that is good for society, not just good for network providers, while maintaining enough freedom to come up with good technical solutions, is the challenge in front of the FCC.

Thank you again for the opportunity to testify. I look forward to your questions.

Literature

Blumenthal, M. S., & Clark, D. D. (2001). Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World. *ACM Transactions on Internet Technology, 1*(1), 70-109.

Braden, R., Clark, D., Shenker, S., & Wroclawski, J. (2000). *Developing a Next-Generation Internet Architecture*: USC Information Sciences Institute Computer Networks Division, MIT Laboratory for Computer Science, International Computer Science Institute.

Briscoe, B., Moncaster, T., & Burness, L. (2007). *Problem Statement: We Don't Have To Do Fairness Ourselves*   (Internet Draft). IEFT.

---

[32] van Schewick (2004), chapter 11.

BT. (2008). BT Total Broadband Usage Allowance Guidelines and Fair Usage Policy. Retrieved April 15, 2008, from http://bt.custhelp.com/cgi-bin/bt.cfg/php/enduser/cci/bt_adp.php?p_faqid=10495&cat_lvl1=346&p_cv=1.346&p_cats=346

Comcast. (2008a). Comcast. Product Terms and Conditions. Performance with PowerBoost® (Speeds up to 12 Mbps). Retrieved April 15, 2008, from http://www.comcast.com/Shop/Buyflow/default.ashx?Popup=true&RenderedBy=Products&FormName=ProductTermsAndConditions&ProductID=20571

Comcast. (2008b). Comcast.net Terms of Service. Comcast Acceptable Use Policy For High-Speed Internet Services. Retrieved April 15, 2008, from http://www6.comcast.net/terms/use/

Comcast. (2008c). FAQs / Product Information / Comcast High-Speed Internet. Retrieved April 15, 2008, from http://www.comcast.com/Customers/FAQ/FaqCategory.ashx?CatId=193

Comcast. (2008d). Frequently Asked Questions about Network Management. Retrieved April 15, 2008, from http://www.comcast.com/Customers/FAQ/FaqDetails.ashx?Id=4567

Farrell, J. (2006). Open Access Arguments: Why Confidence is Misplaced. In T. M. Lenard & R. J. May (Eds.), *Net Neutrality or Net Neutering: Should Broadband Internet Services Be Regulated?* (pp. 195-214). New York, NY: Springer.

Federal Communications Commission. (2008). *High-Speed Services for Internet Access: Status as of June 30, 2007*.

Frischmann, B. M. (2005). An Economic Theory of Infrastructure and Commons Management. *Minnesota Law Review, 89* (April), 917-1030.

Frischmann, B. M., & van Schewick, B. (2007). Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo. *Jurimetrics Journal, 47*(Summer), 383–428.

LaTribu. (2008). Tableau Comparatif. Retrieved April 15, 2008, from http://www.latribu.be/internet.html

MIT Communications Futures Program and Cambridge University Communications Research Network Broadband Working Group. (2005). *The Broadband Incentive Problem (White Paper)*. Cambridge, MA.

PlusNet. (2008a). How Broadband Your Way Works FAQ. Retrieved April 15, 2008, from http://www.plus.net/support/broadband/products/broadband_your_way/info.shtml

PlusNet. (2008b). How PlusNet Guarantees Quality Broadband. Expected Speeds. Retrieved April 15, 2008, from http://www.plus.net/support/broadband/quality_broadband/speed.shtml

PlusNet. (2008c). How PlusNet Guarantees Quality Broadband. Traffic Prioritisation. Retrieved April 15, 2008, from http://www.plus.net/support/broadband/quality_broadband/traffic_prioritisation.shtml

Reed, D. P., Saltzer, J. H., & Clark, D. D. (1998). Commentaries on "Active Networking and End-to-End Arguments". *IEEE Network, 12*(3), 69-71.

Saltzer, J. H., Reed, D. P., & Clark, D. D. (1981, April 8-10). *End-to-End Arguments in System Design.* Paper presented at the 2nd International Conference on Distributed Systems, Paris, France.

United States Government Accountability Office. (2006). *Report to Congressional Committees; Telecommunications; Broadband Deployment is Extensive throughout the United States, but It Is Difficult to Assess the Extent of Deployment Gaps in Rural Areas.*

van Schewick, B. (2004). *Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet.* Unpublished PhD Dissertation. Technical University Berlin, Berlin, Germany.

van Schewick, B. (2007). Towards an Economic Framework for Network Neutrality Regulation. *Journal on Telecommunications and High Technology Law, 5*(2), 329-391.

van Schewick, B. (forthcoming 2009). *Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet*. Cambridge, MA: MIT Press.