

Barbara van Schewick
Professor of Law
Helen L. Crocker Faculty Scholar

Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610
Tel 650 723.8340
Fax 650 725.0253
schewick@stanford.edu

September 19, 2014

ELECTRONIC FILING

Marlene H. Dortch, Secretary
Federal Communications Commission
445 12th Street, SW
Washington, DC 20554

Re: Notice of *Ex Parte* Meeting, GN Docket No. 09-191, GN Docket No. 14-28

Dear Ms. Dortch:

On September 17, I, Barbara van Schewick, met with Scott Jordan, the FCC's Chief Technology Officer.

Network neutrality and Quality of Service, including user-controlled Quality of Service

We discussed my views on the relationship between network neutrality and Quality of Service as put forth in my forthcoming article "Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like."¹ The article is attached to this *ex parte*. Page v of the article provides a guide to the Quality-of-Service-related discussions in the paper.

The network neutrality debate is often framed as a debate for or against Quality of Service.² As the paper shows, the reality is much more nuanced. Some proposals take an all-or-nothing approach to discrimination. They ban or allow all forms of discrimination and, consequently, Quality of Service. Most proposals take a more nuanced position. They allow some, but not all forms of Quality of Service, with different proposals drawing the line between acceptable and unacceptable forms of Quality of Service in different ways.

¹ Barbara van Schewick, "Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like", Forthcoming Stanford Law Review, Volume 67, Issue 1; Stanford Law and Economics Olin Working Paper No. 462; Stanford Public Law Working Paper No. 2459568. The article is available online at: <http://ssrn.com/abstract=2459568>.

² The discussion in this subsection is adopted from the article.

Often, it is not immediately apparent how a specific non-discrimination rule affects network providers' ability to offer Quality of Service. To address this problem, the article explores the effect of the various proposals on the different forms of Quality of Service. The results of this analysis are summarized in *Table 4: The Impact of Different Non-Discrimination Rules on Quality of Service*.³

Underlying the differences between the proposals are disagreements over the social benefits and costs of the different forms of Quality of Service. In this respect, the article offers interesting insights, which are summarized in the article's conclusion.

The non-discrimination rule I propose would prohibit all application-specific discrimination and would allow all application-agnostic discrimination.⁴

The rule allows network providers to offer certain (though not all) forms of Quality of Service. In particular, it allows network providers to offer different classes of service if they meet the following conditions:

- (1) the different classes of service are available equally to all applications and classes of applications;
- (2) the user is able to choose whether, when and for which application to use which class of service;
- (3) the network provider is allowed to charge only its own Internet service customers for the use of the different classes of service.⁵

For example, a network provider could offer a low-delay service, a best-efforts service, a less-than-best-efforts service, and a guaranteed-bandwidth service. The decision of whether and when to use which service would be left to the user. For example, one user could use the low-delay service for Internet telephony, another may use it for online gaming, and a third user may use it for e-mail, if that is what that user wants. This type of user-controlled Quality of Service is technically feasible. (The technical feasibility of this type of Quality of Service and other questions regarding the impact of the proposed rule on Quality of Service are discussed in *Box 23: Frequently Asked Questions on Quality of Service under the Proposed Non-Discrimination Rule* in the paper.)

While the first two conditions directly flow from the proposed non-discrimination rule,⁶ the third condition is based on additional considerations and would need to be encoded separately.⁷

³ Ibid., p. 128.

⁴ For a more detailed discussion of the rule and for a definition of the terms application-specific and application-agnostic, see *ibid.*, pp. 90-115.

⁵ I explain the rationale for this criterion in van Schewick (2010b); van Schewick (2014b), Section "3. Allowing access fees is bad policy"; van Schewick (2014a), Section "Tough Lessons From Mobile and Music."

⁶ Deviating from the first condition by making a specific type of service available only to some applications or classes of applications (e.g., only to the provider's own online video application, or only to online gaming, but not Internet telephony) would make distinctions among applications and classes of applications based on application-specific criteria (here: application or application type) and would thus violate the requirement that

A network provider who is allowed to charge for Quality of Service has an incentive to degrade the quality of the baseline, best-effort service to motivate users to pay for an enhanced type of service. The existence of this incentive is well-documented in the economic literature on price discrimination and one of the main motivations behind proposals to ban Quality of Service.⁸ To mitigate this problem, the rules should require the regulatory agency in charge of enforcing the network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below appropriate levels.^{9, 10}

This type of user-controlled Quality of Service offers the same potential social benefits as other, discriminatory or provider-controlled forms of Quality of Service without the social costs. In particular, it does not raise any of the problems associated with “like treatment.” Contrary to like treatment, it preserves the application-agnosticism of the network, the principle of user choice, and the principle of innovation without permission:

First, this type of Quality of Service preserve the *application-agnosticism* of the network: The provision of Quality of Service is not dependent on which applications users are using, but on the Quality-of-Service-related choices that users make; thus, the network providers does not need to know anything about which applications are using its network in order for this scheme to work. The network provider only makes different classes of service available, but does not have any role in deciding which application gets which Quality of Service; this choice is for users to make. As a result, network providers cannot use the provision of Quality of Service as a mechanism to distort competition among applications or classes of applications. *Second*, since users choose when and for which applications to use which type of service (in line with the principle of *user choice*), they can get exactly the Quality of Service that meets their preferences, even if these preferences differ across users or (for a single user) over time. *Third*, in line with the principle of *innovation without permission*, an innovator does not need support from the network provider in order for his application to get the Quality of Service it needs. The only actors who need to be convinced that the application needs Quality of Service are the innovator, who needs to communicate

differential treatment must be application-agnostic. The second condition ensures that the differential treatment associated with the actual provision of the different types of services in the network happens based on an application-agnostic criterion (here: the type of service chosen by the user for that particular packet).

⁷ See footnote 5 above.

⁸ See footnotes 104 and 105 in the attached paper and accompanying text.

⁹ The incentive to degrade the quality of the baseline service arises only if network providers are allowed to charge for Quality of Service. If they are not allowed to charge for it, they do not benefit from users’ increased use of better-than-best-effort services and, therefore, do not have an incentive to degrade the quality of the baseline, best-effort service to motivate users to use more enhanced services. Thus, instead of adopting the solution proposed in the text, regulators could mitigate this problem by prohibiting network providers from charging for the provision of Quality of Service. Such a ban creates its own social costs, though.

¹⁰ The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. See Article 22(3) of the Universal Service Directive; European Commission (2007), pp. 92, 95-97, 101.

this to the user, and the user, who wants to use the application. This greatly increases the chance that an application can get the type of service it needs.

As I explained at the meeting and in the paper, I would not *allow network providers to offer different types of service to different provider-defined classes of applications, even if the network provider treats like traffic alike*. In other words, I would not allow network providers to provide different types of service to different provider-defined classes of applications that are not alike, as long as they do not discriminate among classes of applications that are alike or among applications within a class of like applications. This requirement is often called “like treatment.”¹¹ Under this approach, a network provider would be allowed to offer low-delay service to Internet telephony, but not to e-mail, as long as it does not treat Vonage differently from Skype, or Gmail differently from Hotmail.¹² In the US, the AT&T BellSouth Merger conditions and various draft bills in Congress allowed this form of Quality of Service.

Those who would allow forms of Quality of Service that provide like treatment assume that discriminating among classes of applications that are not alike is socially harmless and should therefore be allowed. As the article shows, this assumption is not correct. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike. For example, some Internet applications such as Internet telephony applications, Internet messaging applications or Internet video offerings compete with network-provider services that are sold separately from Internet access and do not run over the Internet-access portion of the network provider’s access network. In these cases, discriminating against all applications in that class allows the network provider to favor its own offering without discriminating among applications within the class. Moreover, applications in a class can be harmed by differential treatment even if they do not compete directly with applications in other classes that are treated more favorably.

In addition, like treatment negatively affects several of the factors that have fostered application innovation in the past. *First*, like treatment removes the *application-agnosticism* of the network. Allowing network providers to treat classes of applications differently requires the network provider to identify the different applications on its network in order to decide which class they belong to and determine the appropriate type of service. Thus, like treatment requires network providers to treat data packets differently based on information about the applications on the network. Since the concept of “like applications” is not well defined, network providers have broad discretion to decide which applications are alike, which allows them to deliberately or inadvertently distort competition among applications or classes of applications. *Second*, like treatment violates the principle of *user choice*. Under like treatment, network providers, not users, choose which application should get which Quality of Service. Since users’ preferences for Quality of Service are not necessarily the same across users and may even vary for the same user over time, letting network providers determine which applications gets which Quality of Service will result in levels of Quality of Service

¹¹ On this form of Quality of Service, see Section “Allow Discrimination Among Classes of Applications That Are Not Alike” in the paper.

¹² Internet telephony is sensitive to delay, but e-mail is not, so the two classes of applications are not alike.

that do not meet users' needs. *Third*, like treatment harms application innovation by requiring innovators to convince network providers that their application belongs to a certain class. Requiring network providers to take action before an application can get the Quality of Service it needs violates the principle of *innovation without permission* and reduces the chance that new applications actually get the type of service they need. *Finally*, disputes over which classes of applications are alike, or whether a certain application belongs to a certain class, are likely to be frequent and difficult to resolve, creating high costs of regulation.

Thus, forms of Quality of Service that respect the principle of like treatment do not adequately protect the values that network neutrality is designed to protect and should not be allowed under a network neutrality regime.

We also discussed the treatment of Quality of Service by the Open Internet Order. The discussion was in line with pp. 120-121 and *Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule* on p. 123 of the attached paper.

Use of the term “paid prioritization” by network neutrality proponents

We discussed the use of the term “paid prioritization” by network neutrality proponents. When network neutrality proponents call for a ban on paid prioritization, they use the term in a way that is both narrower and broader than the term suggest.

First, while the term “paid prioritization” does not specify who is paying, network neutrality proponents use it to refer to fees that an Internet service provider imposes on application and content providers who are not its Internet service customers (“access fees”).

Second, when network-neutrality proponents call for a ban on “paid prioritization,” they are calling for a ban on all forms of access fees, not just for a ban on charging applications and content providers for prioritized treatment on the access network.

Access fees come in two variants:

In the first variant, a network provider charges application providers who are not its Internet service customers¹³ a fee for the right to access the network providers' Internet service customers. Applications whose providers do not pay the access fee cannot be used on the network provider's access network.

In the second variant, a network provider charges application providers for prioritized or otherwise enhanced access to the network provider's Internet service customers. For example, if an application provider has paid such an access fee, the application's data packets may receive a better type of service (e.g., priority, or a guaranteed amount of bandwidth) on the network provider's access network or may not count against a user's monthly bandwidth cap (“zero-rating”).

¹³ Any Internet service provider can charge fees to customers of its Internet access service, regardless of whether these customers are providers of applications or “normal” end users. In the past, Internet users directly paid fees for Internet service only to their own Internet access provider.

Treatment of access fees by the Open Internet Order

We also discussed the treatment of access fees by the Open Internet Order. The Open Internet rules themselves do not address access fees. The text of the order discusses the two types of access fees separately.

Fees for access to end users

The text of the order **clearly prohibits** network providers from charging application and content providers for access to the network providers' Internet service customers (i.e. from just charging for access, without offering anything in return).¹⁴

The order discusses this question in the context of the rule against blocking on the fixed Internet. To the extent that the rules prohibit blocking of a specific application on the mobile Internet, the no-blocking rule also prevents network providers from charging this application an access fee.¹⁵

Fees for prioritized or otherwise enhanced access to end users (“third-party-paid prioritization”)

While the text of the order stops short of an outright ban of “third-party-paid prioritization” arrangements, it seems to **get as close to explicitly banning these arrangements as one can get without explicitly banning them**. The order explicitly endorses the concerns against these arrangements,¹⁶ unequivocally rejects the main arguments in favor of them,¹⁷ and concludes that “as a general matter,” arrangements of this kind are “unlikely” to be considered reasonable.¹⁸

The Open Internet order discusses the limits on access fees for prioritized or otherwise enhanced access to end users in the context of the non-discrimination rule. Conceptually, however, the rule as clarified by the text of the order is more accurately characterized as a limit or ban on charging. If it was a non-discrimination rule, the rule would allow Internet service providers to charge this type of access fees, but require Internet service providers to offer and charge for enhanced access in non-discriminatory ways.

In addition, limits on access fees rest on different considerations than rules against blocking or discrimination, and are therefore best treated separately – both in the text of eventual rules and in their justification.¹⁹

¹⁴ Federal Communications Commission (2010), para 67.

¹⁵ See the explicit reference to para 67, which contains the access fee discussion, in the discussion of the rule against blocking on mobile networks on p. 56, note 306 of the order.

¹⁶ Federal Communications Commission (2010), paras 76 and 24-34.

¹⁷ Federal Communications Commission (2010), paras 40 and 28.

¹⁸ Federal Communications Commission (2010), para 76.

¹⁹ For an explanation of the policy concerns underlying the ban on access fees, see van Schewick (2010a), pp. 278-280; van Schewick (2010b); van Schewick (2014b), Section “3. Allowing access fees is bad policy”; van Schewick (2014a), Section “Tough Lessons From Mobile and Music.”

Banning access fees under Section 706

We discussed the limits that *Verizon v. FCC* and relevant precedent impose on the FCC's ability to ban access fees under Section 706 as outlined in the attached blog post.²⁰ In particular, *Verizon v. FCC* explicitly requires the FCC to allow access fees. Moreover, the Court's decision suggests that the FCC needs to allow the worst kind of access fees if it wants the Court to uphold a general ban on blocking.

Should you have any questions, please do not hesitate to contact me.

Sincerely,

/s/ Barbara van Schewick

Barbara van Schewick
Professor of Law and (by courtesy) Electrical Engineering
Helen Crocker Faculty Scholar
Faculty Director, Center for Internet and Society, Stanford Law School

cc:

Scott Jordan

References

- European Commission. 2007. "Commission Staff Working Document. Impact Assessment". SEC(2007) 1472. http://ec.europa.eu/information_society/policy/ecomm/doc/library/proposals/1472/comm_pdf_sec_2007_1472_1_en_documentdetravail.pdf.pdf
- Federal Communications Commission. 2010. "Preserving the Open Internet. Report and Order". FCC 10-201. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf
- van Schewick, Barbara. 2010a. *Internet Architecture and Innovation*. Cambridge, MA: MIT Press.
- van Schewick, Barbara. 2010b. *Opening Statement at the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet*. Federal Communications Commission. http://www.law.stanford.edu/display/images/dynamic/publications_pdf/schewick-statement-20100428.pdf
- van Schewick, Barbara. 2014a. "The Case for Rebooting the Network-Neutrality Debate." *The Atlantic*. May 6. <http://www.theatlantic.com/technology/archive/2014/05/the-case-for-rebooting-the-network-neutrality-debate/361809/>
- van Schewick, Barbara. 2014b. "The FCC Changed Course on Network Neutrality. Here is Why You Should Care." *Stanford Center for Internet and Society Blog*. April 25. <http://cyberlaw.stanford.edu/blog/2014/04/fcc-changed-course-network-neutrality-here-why-you-should-care>

²⁰ van Schewick (2014b), Section "2. Section 706 of the Telecommunications Act requires the FCC to allow access fees."