Stanford Law School

Barbara van Schewick
M. Elizabeth Magill Professor of Law
Director, Center for Internet and Society
Professor, by courtesy, Electrical Engineering

Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305
Tel. 650 723.8340
Fax. 650 725.0256
schewick@stanford.edu

April 22, 2024

ELECTRONIC FILING

Marlene H. Dortch
Secretary
Federal Communications Commission
45 L Street NE
Washington, DC 20554

Re: Notice of *Ex Parte* Meeting, Safeguarding and Securing the Open Internet, Docket No. 23-320

Dear Ms. Dortch:

Pursuant to Section 1.1206(b)(2)(iv) of the Commission's rules, I am filing this reply to a written ex parte filed by CTIA, dated April 18, 2024 and posted in the docket on April 19,[1] which references and attempts to rebut my recent filings concerning the bright line rule on throttling.[2]

In its filing, CTIA argues that clarifying that the no-throttling rule prohibits ISPs from speeding up or otherwise provide quality of service only to some applications or categories of applications "depart[s] from the core principle" that "consumers should be able to choose the Internet experience they want," claiming that the requested clarification would "curb consumers' ability to choose offers and plans that they find best satisfy their preferences"[3] and would "take choice *away* from consumers."[4]

Allowing ISPs to limit the provision of quality of service to applications or categories of applications limits rather than enhances consumer choice and is unlikely to meet user needs.

[1] CTIA, 2024, Written Ex Parte Letter dated April 18, 2024 ("CTIA Ex Parte").

[2] CTIA Ex Parte, pp. 2-4.

[3] CTIA Ex Parte, pp. 1-2.

[4] CTIA Ex Parte, p. 3.

As I have explained in my 2015 Stanford Law Review article on network neutrality and Quality of Service, letting ISPs, not users, choose which applications should get which Quality of Service or differential treatment violates the principle of user choice.[5]

In those plans, only the applications or categories of applications that the ISP has chosen get a better quality of service. That's a problem. User preferences vary widely across users, and even for a single user over time. An ISP has no way of knowing which applications are most important to a specific user at a particular moment in time. As a result, letting ISPs chose which applications or categories of applications should receive which quality of service is unlikely to meet users' needs:

> "As I have explained elsewhere, the incentives of network providers and users are not necessarily aligned.[6] Network providers' incentive to offer low-delay service only to online gaming, but not to Internet telephony, or to reduce the performance of applications that may reduce their revenue from applications that are offered and provided separately from Internet access are examples of this phenomenon.[7] Thus, network providers do not always want to meet users' preferences. But even when they do, they may not be able to do so.[8] For example, if a network provider decides whether and when to offer Quality of Service, it is forced to guess what the average user's priorities may look like, but these priorities may differ among users, and, for the same user, over time. In particular, a specific user's needs with respect to a particular application are not necessarily fixed.[9] A user's desire for Quality of Service may differ considerably depending on the circumstances. For example, I may not care as much about the quality of my VoIP call when I am chatting with a friend as when I am doing a job interview. If I am playing a quick game at night, I may be willing to tolerate a level of latency that I would not be willing to tolerate during an online gaming tournament. Normally, I may want file uploads to happen in the background and may want them to yield to other applications that are more important to me right now. But if I am uploading a large paper to a conference website just before the submission deadline, finishing this upload as quickly as possible will have the highest

---

[5] Barbara van Schewick, 2015, Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like, Stanford Law Review, Volume 67, pp. 121-122, available at https://law.stanford.edu/publications/network-neutrality-and-quality-of-service-what-a-non-discrimination-rule-should-look-like/ (attached) ("NN & QoS").

6. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 350-51.

7. *See supra* notes 424-27 and accompanying text; *see also* 388-91 and accompanying text.

8. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 351.

9. van Schewick, Official Testimony, *supra* note 68, at 7; *see also* Free Press Open Internet Comments, *supra* note 3, at 102-03; RILEY & SCOTT, *supra* note 412, at 8; Lennett, *supra* note 357, at 143-45; Yiannis Yiakoumis et al., *Putting Home Users in Charge of Their Network*, 2012 PROC. 2012 ACM CONF. ON UBIQUITOUS COMPUTING 1114, 1115; B. Briscoe et al., Internet Eng'g Task Force, Internet-Draft, Problem Statement: Transport Protocols Don't Have to Do Fairness 13-14 (July 14, 2008), https://tools.ietf.org/html/draft-briscoe-tsvwg-relax-fairness-01; *supra* note 147 and accompanying text; *infra* notes 484-85 and accompanying text.

> priority.[10] Thus, any Quality of Service system that lets network providers determine whether and when to provide Quality of Service may not be well aligned with users' needs. Network providers' attempts to determine which applications are time-sensitive and should receive special treatment during times of congestion will fail to meet users' needs for the same reasons.[11]"

The fact that users' preferences for Quality of Service vary widely has also been shown empirically. For example, when Stanford researchers implemented a user-controlled fast lane that gave users a choice which website traffic to prioritize, they found that users choices differed widely, leading researchers to conclude that the study "demonstrate[s] that user preferences have a heavy tail, and suggest that a one-size-fits-all approach is unlikely to work for most users."[12]

> "We prototyped a service—a user-defined fast lane called Boost and deployed it in 161 homes, during an internal "dogfood" test of OnHub, a commercial home WiFi router built by Google. With Boost, users can decide which traffic gets higher priority (or decide to not give higher priority to any traffic). They can express their preferences through a web browser extension, either by prioritizing a specific tab, or by always prioritizing the traffic from a specific website. Figure 1 summarizes the behavior of a fairly homogeneous group of users: 43% of expressed preferences were unique, i.e., the preferred website was picked by only one user, while the median popularity index of prioritized websites was 223. While our sample is small, it demonstrates the diverse and heavy-tailed nature of user preferences, and that users are willing to express them if it is easy to do so."[13]

10. This example is taken from Yiakoumis et al., *supra* note 434, at 1115.

11. *See, e.g.*, RILEY & SCOTT, *supra* note 412, at 8; Lennett, *supra* note 357, at 143-44.

[12] Yiannis Yiakoumis, Sachin Katti, Nick McKeown, 2016, Neutral Net Neutrality, SIGCOMM '16: Proceedings of the 2016 ACM SIGCOMM Conference, August 2016, pp. 483–496, available at https://doi.org/10.1145/2934872.2934896 (attached) ("Neutral NN")

[13] Neutral NN, p. 2.

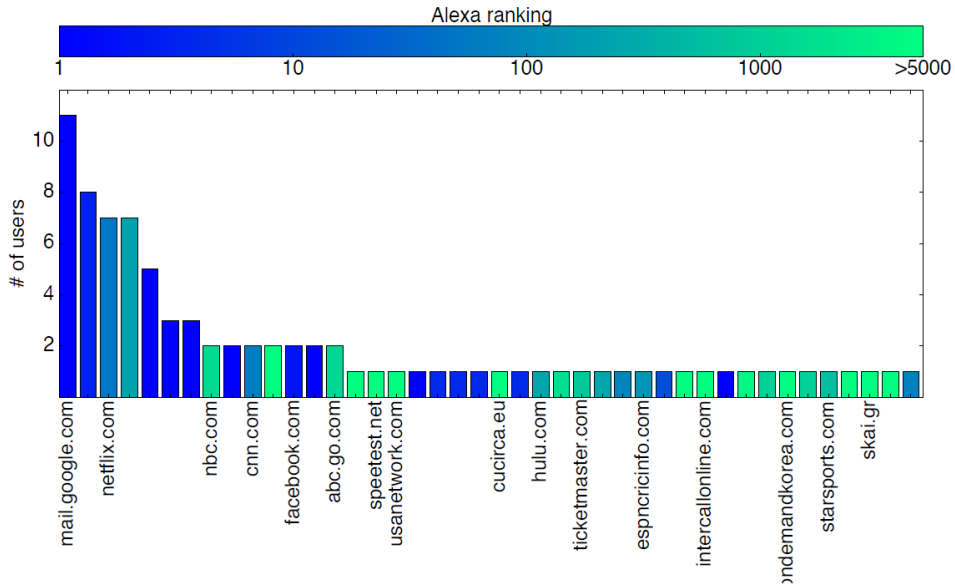

Figure 1: If given the choice, which websites would home users prioritize? The preferences have a heavy tail: 43% of the preferences are unique, with a median popularity index of 223.

More generally, letting ISPs chose which applications or categories of applications should get Quality of Service or receive preferential treatment is unlikely to meet user needs because ISPs have no way of knowing the quality of service requirements of each application. Even if they knew which applications would benefit from which quality of service, they are unable to accurately and reliably identify applications. That's both because applications are now widely encrypted, and Deep Packet Inspection technology is famously unreliable and prone to misclassifications.[14]

Comcast's Chief Technology Officer Jason Livingood explains this in a recent Internet Engineering Task Force Draft:[15]

[14] See, e.g., the remarkable case study by Alissa Cooper and Ian Brown. Cooper and Brown, Net Neutrality: Discrimination, Competition, and Innovation in the UK and USACM Transactions on Internet Technology, Volume 15, Issue 1 Article No.: 2 pp 1–21 https://doi.org/10.1145/2700055.

[15] Jason Livingood, 2024, ISP Dual Queue Networking Deployment Recommendations: draft-livingood-low-latency-deployment-05, Internet Draft, April 18, 2024, available at https://datatracker.ietf.org/doc/draft-livingood-low-latency-deployment/ ("Livingood").

> "* Application developers know best whether their application is compatible with low latency networking and which aspects of their traffic flows will or will not benefit.
>
> * Application traffic is almost entirely encrypted, which makes it very difficult for networks to accurately determine application protocols and to further infer which flows will benefit from low latency and which flows may be harmed because they need to build a queue.
>
> * Network operators and equipment vendors attempting to infer application type and application need will sometimes make mistakes, incorrectly classifying traffic [Lotus], and potentially negatively affecting certain flows.
>
> * The pace of innovation and iteration is necessarily faster-moving in the application edge at layer 7, rather than in the network at layer 3 (and below) - where there is greater standards stability and a lower rate of major changes. As a result, the application layer is best suited to rapid experimentation and iteration. Network operators and equipment vendors trying to infer application needs will in comparison always be in a reactive mode, one step behind changes made in applications."

ISPs are likely to target application categories for low-latency lanes that are used by a large number of people that have a high willingness to pay for premium performance. That's why trials have focused on categories like online gaming (targeting passionate gamers), online video conferencing (targeting people who need high-quality video conferencing for work), or online video (people who like high-quality entertainment). They are less likely to offer special fast lanes for categories of app that are used by smaller groups of users or users that don't have a high willingness or ability to pay. For example, many designers use Figma, a web-based service that lets designers collaborate on design projects in real time. Such real-time collaboration apps are unlikely to be used by enough people to motivate ISPs to offer a "fast lane for real-time collaboration tools," even though they would benefit immensely from the lower latency that a low-latency lane can offer.

More generally, being excluded from low-latency lanes because the app does not belong to a category for which the ISP created a fast lane will negatively affect a much wider range of applications (and, in turn, these applications' users) than one might assume. That's because almost all applications benefit from lower latency, including websites, as the letter by startups Outpost and Quiet explained in detail.[16]

By contrast, our proposals do not deprive consumers of the ability to improve the performance of applications that would benefit from quality of service. As we have explained before, there are ways to offer quality of service in a way that is consistent with Open

[16] Outpost/Quiet Startup Letter filed April 18, 2024.

Internet principles. [17]These kinds of application-agnostic, user-controlled, user-paid Quality of Service offerings violate neither the no-throttling rule nor the no-paid prioritization rule.

Allowing ISPs to offer a different quality of service as part of broadband internet access that is application-agnostic (i.e. open to all applications and categories of applications) and lets users decide whether, when, and for which applications to use the different quality of service, does not have the problems outlined above. Since ISPs are not allowed to limit which applications and categories of applications can receive quality of service, users can choose to use the different type of service in the way which best meets their needs.

As other commenters have explained, Comcast's low latency trial shows that these kinds of application-agnostic, user-controlled, and user-paid quality of service are feasible with today's technology, and mobile equipment vendors have implemented the same technology that Comcast is using as well.[18]

More generally, CTIA's filing creates the impression as if prohibiting ISPs from speeding up apps or classes of apps is a revolutionary idea. That's wrong.

Letting ISPs determine which apps or categories of apps get special treatment violates core net neutrality principles.

Net neutrality means people using the internet get to decide what they do online. ISPs don't get to interfere with their choices by blocking, speeding up or slowing down apps or kinds of apps. Apps compete on a level playing field, and users, not ISPs, determine which apps are successful. In other words, net neutrality means that ISPs don't get to pick winners and losers online.

Letting ISPs decide which apps get to be in a fast lane violates these principles. Apps in a fast lane work better than those that aren't, especially when the network is busy. That gives apps in the fast lane a competitive advantage over apps that are not.

Apps in the premium-slice fast lane continue to work well during congestion. Most apps, sites, and services in the regular lane suffer – not just the apps that are particularly sensitive to delay: websites load more slowly, video buffers, AIs lags, calls break up.

---

[17] See, e.g., van Schewick February 8 Ex Parte, pp. 3-4; Written Ex Parte of Open Technology Institute at New America, Public Knowledge, Professor Barbara van Schewick and Professor Scott Jordan, filed March 11, 2024 ("Joint Written Ex Parte"), pp. 4-5, available at https://www.fcc.gov/ecfs/document/103120890811342/1. See also CCIA & INCOMPAS, 2024, Written Ex Parte Letter dated March 14, 2024 ("CCIA & INCOMPAS Written Ex Parte") (supporting the same framework and proposals), p.2, https://www.fcc.gov/ecfs/search/search-filings/filing/10314012196846.

[18] See, e.g., Willars, P., Wittenmark, E., Ronkainen, H., Johansson, I., Strand, J., Ledl, D., and D. Schnieders, "Enabling time-critical applications over 5G with rate adaptation", May 2021, https://www.ericsson.com/49bc82/assets/local/reports-papers/white-papers/26052021-enabling-timecritical-applications-over-5g-with-rate-adaptationwhitepaper.pdf>.

Most people have no idea that their ISP can make their applications work better or worse. So if Google Meet works better than video services in the regular lane, people will think it's because it's is a better product.

This works in ways that people don't even notice. Even small differences in load times affect how long people stay on a site, how much they buy, and whether they come back.

As a result, apps in a fast lane will have a competitive advantage over apps that are not. Letting ISPs select the apps/classes of apps in a fast lane lets them distort the market.

The distortion in competition is easy to see if the ISP only includes some of the applications in a category in the fast lane. If Google Meet and Cisco‘s WebEx are in a video-conferencing fast lane, they will continue to work well even if the network is busy, while video calls on Signal and Discord are buffering. This directly distorts competition among video conferencing apps.

Some people think that letting ISPs decide which applications should go in a fast lane is less problematic if the fast lane is open to all apps in a category.

However, the distortion in competition resulting from ISP-controlled fast lanes is not limited to apps that are in the same category. That’s because all apps constantly compete for users’ time and attention. If an app doesn’t belong to a category that the ISP deems worthy of being in the low-latency lane, it will have higher latency and perform less well during times of congestion, so users will use it less. As explained above, that affects a much wider range of applications than those that are deemed particularly latency sensitive.

In other cases, an ISPs decision to limit the fast lane to a specific category might be driving by its desire to protect non-internet services from competition. Offering a 5G fast lane only to online games, but not to online telephony apps, makes it harder for Signal, WhatsApp, or Vonage to compete with carriers’ traditional telephony services.

Most importantly, as we have seen with zero-rating programs, opening a fast lane to all applications in a category is unlikely to remove the distortion of competition among apps in the category. Even when fast lanes are supposedly open to all apps in a category and no apps are paying the ISP, these fast lanes favor the most popular apps. That’s because it takes time and effort to work with an ISP to become part of its fast lane – from the paperwork to apply to be in the fast lane to the technical engineering work that’s required to ensure that the application works with the ISPs specific technology. As we saw with zero-rating, smaller players do not have the resources to do this work, and even when they make the effort to try to be included in a program, ISPs often don’t get back to them or prioritize working with larger players first. That means the biggest apps will end up in all the fast lanes, while most others would be left out: startups, small businesses, sites serving marginalized communities or niche audiences, millions of other apps & sites in the long tail.

As a result, ISP-controlled fast lanes hamper startups & small businesses and help cement platform dominance.

More generally, allowing ISPs to pick and choose which apps get quality of service negatively affects several of the factors that have fostered application innovation in the past even if ISPs have to treat like traffic alike. *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 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.

Finally, CTIA points to the general conduct rule as providing sufficient protections if an ISP limits quality of service or other forms of preferential treatment to select apps or categories of apps.

This has a certain irony. In the past, ISPs and their trade associations have complained that the general conduct rule does not provide sufficient certainty for them.

More generally, case-by-case standards do not sufficiently protect innovators and users.

We saw this with zero rating. In the 2015 Open Internet Order, the FCC declined to create bright line rules banning harmful zero-rating, instead opting for leaving zero-rating to future case-by-case evaluations under the general conduct standard.

In response, mobile ISPs blanketed the market with various zero-rating offerings, including clearly harmful ones that zero-rated the ISPs' own apps and charged apps to be zero-rated.

Even though more than a hundred public interest groups, companies, Senators, and more than 100,000 individuals repeatedly urged the FCC in letters and complaints to open a public proceeding, enforce the general conduct rule, and prohibit harmful zero-rating offerings, the FCC investigated existing offerings behind closed doors for more than two years and only published a staff report finding that some of these plans "likely" violated the general conduct rule in January 2017, right before Chairman Pai took office.

By contrast, California's net neutrality law included bright line rules prohibiting the most harmful forms of zero-rating.

The day California's net neutrality law became enforceable, AT&T and Verizon announced they were stopping zero-rating their own video apps in California – without waiting for enforcement action.

CTIA's argument also ignores the substantial other costs associated with the proposed vague, case-by-case standard. A vague, unbounded standard creates unnecessary confusion, requires costly adjudication and litigation, provides opportunities for both FCC overreach and underenforcement, and tilts the playing field in favor of large, established companies that can afford long and costly proceedings at the FCC. And in the meantime, apps not chosen to be in a fast lane suffer the competitive harms that net neutrality is supposed to prevent.

By contrast, bright line rules provide certainty to the market, keep the costs of regulation low, and allow small players to navigate the complaint process.

Pursuant to Section 1.1206 of the Commission's rules, this reply letter is being filed in ECFS. Should you have any questions, please do not hesitate to contact me.

Sincerely,

Barbara van Schewick

Barbara van Schewick
M. Elizabeth Magill Professor of Law and Professor, by Courtesy, of Electrical Engineering
Director, Stanford Law School Center for Internet and Society

Attachments:

Jason Livingood, 2024, ISP Dual Queue Networking Deployment Recommendations: draft-livingood-low-latency-deployment-05, Internet Draft, April 18, 2024, available at https://datatracker.ietf.org/doc/draft-livingood-low-latency-deployment/

Barbara van Schewick, 2015, Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like, Stanford Law Review, Volume 67, pp. 1-166, available at https://law.stanford.edu/publications/network-neutrality-and-quality-of-service-what-a-non-discrimination-rule-should-look-like/

Yiannis Yiakoumis, Sachin Katti, Nick McKeown, 2016, Neutral Net Neutrality, SIGCOMM '16: Proceedings of the 2016 ACM SIGCOMM Conference, August 2016, pp. 483–496, available at https://doi.org/10.1145/2934872.2934896