Independent Stream
Internet-Draft
Intended status: Informational
Expires: 26 October 2025

J. Livingood
Comcast

24 April 2025

ISP Dual Queue Networking Deployment Recommendations
draft-livingood-low-latency-deployment-09

Abstract

The IETF’s Transport and Services Working Group (TSVWG) has finalized experimental RFCs for Low Latency, Low Loss, Scalable Throughput (L4S) and new Non-Queue-Building (NQB) per hop behavior. These documents describe a new architecture and protocol for deploying low latency networking. Since deployment decisions are left to implementers, this document explores the potential implications of those decisions and makes recommendations that can help drive adoption and acceptance of L4S and NQB.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 26 October 2025.

Copyright Notice

Table of Contents

1. Introduction

The IETF’s Transport and Services Working Group (TSVWG) has finalized RFCs for Low Latency, Low Loss, Scalable Throughput (L4S) and Non-Queue-Building (NQB) per hop behavior [RFC9330] [RFC9331] [RFC9332] [RFC9435] [I-D.ietf-tsvwg-l4sops] [I-D.ietf-tsvwg-nqb]. These documents do a good job of describing a new architecture and protocol for deploying low latency networking. But as is normal for many such standards, especially new or experimental ones, certain deployment decisions are ultimately left to implementers.

This document explores the potential implications of key deployment decisions and makes recommendations for those decisions that may help drive adoption by network operators and application developers. That is a key issue for low latency networking, because the more applications developers and edge platforms that adopt new packet

marking for low latency traffic, then the greater the value to end users, so ensuring it is received well is key to driving strong initial adoption.

It is worth stating though that these decisions are not embedded in or inherent to L4S and NQB per se, but are decisions that can change depending upon differing technical, regulatory, business or other requirements. Even two network operators with the same type of access technology and in the same market area may choose to implement in different ways. Nevertheless, this document suggests that certain specific deployment decisions can help maximize the value of low latency networking to end users, network operators, and application developers.

In addition, the design of the protocols also make clear that applications developers are best positioned to understand the needs of their applications and to, by extension, express any such low latency needs via appropriate L4S or NQB packet marking.

For additional background on latency and why latency matters to the Internet, please read [BITAG].

1.1. A Different Understanding of Application Needs

In the course of working to improve the responsiveness of network protocols, the IETF concluded with their L4S and NQB work that there were two main types of traffic and that these two traffic types could benefit from having separate network processing queues in order to improve the way the performance of delay-sensitive and/or interactive applications. In addition, introducing a new queue better supports incremental development of a new standard rather than changing existing congestion control algorithms - which would be complex.

One of the two major traffic types is mostly file download or upload, such as downloading an operating system update or uploading files to a cloud backup. This type of traffic tends not to be particularly delay-sensitive, at least on a millisecond level basis. The other type of traffic is real-time, interactive traffic that is typically latency-sensitive, such as video conferencing and gaming.

The value of dual queue networking (simply "low latency networking" hereafter) seems potentially good, and at least one major ISP has deployed it [Comcast]. It seems possible that this new capability might enable entirely new classes of applications to become possible, driving a wave of new Internet innovation, while also improving the latency-sensitive applications that people use today.

1.2. New Thinking on Low Latency Packet Processing

L4S does *not* provide low latency in the same way as previous technologies like DiffServ Quality of Service (QoS). That prior QoS approach used packet prioritization, where it was possible to assign a higher relative priority to certain application traffic, such as Voice over IP (VoIP) telephony. This approach could provide consistent and relatively low latency by assigning high priority to a partition of the capacity of a link, and then policing the rate of packets using that partition. This traditional approach to QoS is hierarchical in nature.

That QoS approach is to some extent predicated on the idea that network capacity is very limited and that links are often highly utilized. But on today’s Internet, many users have experienced poor application performance, such as video conferencing, despite having sufficient bandwidth. In many of these scenarios, prioritization will not improve a flow. But finding a way to reduce latency has proven beneficial. This new low latency networking approach is not based on hierarchical QoS prioritization. Rather, it is built upon conditional priority scheduling between two queues that operate at best effort QoS priority.

2. Key Low Latency Networking Concepts

In the past, many thought that the only way to improve application quality was via more bandwidth or by using QoS priority. The advent of low latency networking enables a re-examination of those approaches.

2.1. Best Effort Priority

Low latency traffic to is not prioritized over other (best effort priority) "classic" Internet traffic. That is the case over the ISP network and the broader internet, though it may not not necessarily be the case for a user’s in-home Wi-Fi network due to the particulars of how the IEEE 802.11 wireless protocol [IEEE] functions at the current time - see [RFC8325]). In addition, some user access points may prioritize certain traffic (such as gaming) and some traffic such as NQB may use the AC_VI Wi-Fi link layer queue [I-D.ietf-tsvwg-nqb]. This best effort approach stands in contrast to prior differential quality of service (QoS) approaches or to what has been discussed for 5G network slicing [CDT-NN] [van-Schewick-1A] [van-Schewick-1B] [van-Schewick-2] [van-Schewick-3].

2.2. Shared Throughput

Low latency networking flows do not get access to greater throughput than "classic" flows. Thus, a user's total provisioned or permitted throughput on an ISP access network link is shared between both classic and low latency queues.

2.3. Access-Agnostic

Low latency networking can be implemented in a variety of network technologies. For example in access network technologies this could be implemented in DOCSIS [LLD], 5G [Ericsson], PON [CTI], and many other types of networks.

3. Application Developer Recommendations

Application developers need to add L4S or NQB packet marking to their application, which will often depend upon the capabilities of a device's operation system (OS) or a software development kit (SDK) [Apple] that the OS developer makes available. In addition, the application server will also need to support the appropriate marking and, when L4S is used, to implement a responsive congestion controller.

3.1. Delivery Infrastructure for L4S

Since L4S uses the Explicit Congestion Notification (ECN) field of the packet header, to ensure ECN works end-to-end, application developers need to be certain that their servers, datacenter routers, and any transit, cloud provider, or content delivery network (CDN) server involved in their application IS NOT altering or bleaching the ECN field. For an application to use the L4S queue, they must mark their packets with the ECT(1) code point to signal L4S-capability or with the Congestion Experienced (CE) code point when appropriate. Coupled with client marking, if an application client or server detects CE marks, it should respond accordingly (e.g., by reducing the send rate), which typically means that the server must be running a "responsive" congestion controller (i.e., is able to adjust rate based the presence or absence of CE marks for L4S traffic - such as DCTCP, TCP Prague, SCReAM, and BBRv2). See Section 4.3 of [RFC9330] and Section 4.3 of [RFC9331] for more information about this.

3.2. Delivery Infrastructure for NQB

Since NQB uses the DSCP-45 code point in the DiffServ part of the packet header, to ensure NQB works end-to-end, application developers need to be certain that their servers, datacenter routers, and any transit, cloud provider, or content delivery network (CDN) server involved in their application IS NOT altering or bleaching a DSCP-45 mark. The server DOES NOT need to run a special responsive congestion controller. However, it is common for networks to bleach or modify DSCP marks on ingress today, so networks will need to change that policy for NQB to work end-to-end (in contrast, ECN is rarely bleached).

3.3. Only Mark Delay-Sensitive Traffic for L4S or NQB

It may seem tempting to mark all traffic for L4S or NQB handling, but it may not help in all cases. For example, a video gaming service may benefit from using L4S or NQB for real-time controller inputs and gameplay, while major game software updates would best be left in the classic queue.

3.4. Consider Application Needs in Choosing L4S vs. NQB

Determine whether your application needs "sparse" flows or "congestion-controlled" (higher capacity) flows. Sparse flows that are latency senstive should be marked as NQB (thus DSCP-45). This may be things like DNS queries or VoIP media flows, where maximizing the bandwidth of the flow is not necesary.

Latency-sensitive flows that need more bandwidth are congestion controlled, and identified via ECN marking. These types of applications are less limited by the application protocol itself (i.e., a small DNS query), which means the application quality can improve as more bandwidth is available - such as shifting a video stream or a video conference session from Standard Definition (SD) to 4K quality.

4. ISP Recommendations

Like any network or system, good deployment design decisions matter. In the context of deploying low latency networking in an ISP network, these recommendations should help ensure that a deployment is resilient, well-accepted, and creates the environment for generating strong network effects.

4.1. Allow ECN Across Network Boundaries

Traffic sent TO a peer network marked with ECT(1) or CE in the ECN header MUST pass to that peer without altering or removing the ECT(1) or CE marking (see exception below). Traffic FROM peers marked with ECT(1) or CE in the ECN header MUST be allowed to enter the network without altering or removing the ECT(1) or CE marking (see exception below). The only exception would be when a network element is CE-aware and able to add a CE mark to signal that it is experiencing congestion at that hop.

This part - allowing unmodified ECN across the network - is likely to be easier than DSCP-45 for NQB (see next section), since it appears rare that networks modify the ECN header of packet flows.

4.2. Allow DSCP-45 Across Network Boundaries

Traffic sent TO a peer network marked with DSCP value 45 MUST pass to that peer without altering or removing the DSCP 45 marking (see exception below). Traffic FROM peers marked with DSCP value 45 MUST be allowed to enter the network without altering or removing the DSCP 45 marking (see exception below).

However, some networks may use DSCP 45 for internal purposes other than NQB within their network. In these cases, the peer using DSCP 45 for other purposes is responsible for remarking as appropriate. If possible, for operational simplicity, a network should try to maintain the use of DSCP 45 on an end-to-end basis without remarking in their interior network hops.

4.3. Last Mile Network (Access Network)

There are two hops of interest in the last mile access network. One will be a point of user aggregation, such as a Cable Modem Termination System (CMTS) or Optical Line Terminal (OLT). The second is at the user location, such as a Cable Modem (CM) or Optical Network Unit (ONU), both of which are example of CPE.

In theses two queues, ISPs should consider using the optional Queue Protection function [I-D.ietf-tsvwg-nqb] [I-D.briscoe-docsis-q-protection]. This can potentially detect mismarking and take corrective action as needed.

4.4. Customer Premise Equipment (Customer Edge)

In most residential Internet services, there are typically two equipment modes. One is very simple CPE that hands off from the ISP's access network (i.e., DSL, 5G, DOCSIS, PON) and provides the customer with an Ethernet interface and IP address(es). The customer then connects their own router and wireless access point (often integrated into the router, typically referred to as a "wireless gateway" or "wireless router"). The other model is more typical, which is that the CPE integrates a link layer termination function (i.e., Cable Modem, 5G radio, or Optical Network Unit) as well as a wireless gateway.

Not all ISP networks support both of these models; sometimes only a wireless gateway is available. Even in this case, some users "double NAT" and install their own router and wireless access point(s) to get whatever functionality and control over their home network that they desire. The cases explored below are commonplace but may not apply to all networks.

In some cases, dual queue networking and associated packet marking is supported up to the ISP's demarcation point - such as in a cable modem. It is recommended that packet markings should pass from such a demarcation point to any attached customer-administered CPE, such as a router or wireless access point. That enables a customer-administered router to implement dual queue networking, rather that it only being possible with ISP-administered CPE.

4.5. Inside the Home - Customer Local Area Network (LAN)

As noted above with the mention of an integrated wireless gateway, the CPE and router/wireless network gear is integrated into a single CPE device. Even though these are functionally in one piece of hardware, we can think of the wide area network interface and local area network as functionally separate for purposes of this analysis.

4.5.1. 802.11 WiFi Queuing

As noted above with respect to prioritization of packets in the ISP network, all packets should be handled with the same best effort priority in the ISP access network and on the internet. However, in a user's home Wi-Fi (wireless) local area network (WLAN), this is more complicated, because there is not a precise mapping between IETF packet marking and IEEE 802.11 marking, explored in [RFC8325]. In short, today's 802.11 specifications enable a Wi-Fi network to have multiple queues, using different "User Priority" and "Access Category" values. At the current time, these queues are AC_BK (Background), AC_BE (Best Effort), AC_VI (Video), and AC_VO (Voice).

As explored in [I-D.ietf-tsvwg-nqb], packets in the low latency queue may be expected to be marked for the best effort (AC_BE) or video (AC_VI) wireless queue. For additional context, please refer to Section 8.1 of [I-D.ietf-tsvwg-nqb]. In some situations, such as a user-owned wireless access point or CPE, it may not be possible for the user to select which wireless queue is used. In cases where the CPE is ISP-administered, selecting a specific wireless queue may be possible - though it is not yet clear what the best practice may be for this selection until ISPs and application developers have more experience with low latency networking. As of the writing of this document, it appears that the AC_VI queue may be used for the low latency queue in some networks - and that many latency-sensitive applications are already marking their upstream wireless traffic for AC_VI and AC_VO.

4.5.2. Use Permissive Upstream NQB Queue Admission

Since the IETF's NQB specification is only recently completed, many applications that have been using other DSCP marks for their latency-sensitive flows have not yet shifted to adopt DSCP-45. One example is the Microsoft Xbox platform [Microsoft], which is using DSCP-46. So in the relatively short-term, ISPs may find it beneficial to their customers to use a more permissive upstream NQB admission policy, allowing DSCP-40, 45, 46, and 56 admission into the low latency queue. It may take a year or more after the NQB DSCP assignment is made by IANA for developers to shift to DSCP-45, given other items in their development backlog and their software release schedule.

4.6. Do Not Use Middleboxes

As noted in [Tussle] there has always been a tension in the end-to-end model between how much intelligence and processing takes place along the end-to-end path inside of a network and how much takes place at the end of the network in servers and/or end user client devices and software. In this new approach to low latency networking, entry into a low latency queue depends upon marks in the packet header of a particular application flow. In practice, this marking is best left to the application edge of the network, rather than it being a function of a so-called middlebox in the ISP network. As explored below, this is the most efficient, least prone to mis-classification, and is most acceptable from the standpoint of network neutrality.

The best approach is for applications to mark traffic to indicate their preference for the low latency queue, not the network making such a decision on its own. This is for several reasons:

* According to the end-to-end principle, this function is best delegated to the edge of the network as an architectural best practice (the edge being the application in this case).

* Application marking maintains the loose coupling between the application and network layers, eliminating the need for close coordination between networks and application developers.

* 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.

* Only the application (not the network) knows whether a scalable congestion control algorithm congestion control is being used on the application server. Thus, only the developer and server administrator know if they are correctly responding to Congestion Experienced (CE) markings for L4S (see Section 4.1 of [RFC9331]).

* 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. It is likely that false positives [Lotus] and false negatives will occur if network-based inference is used; all of which can be avoided simply by relying solely on application marking.

* 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.

* This avoids issues arising from mis-classification of application flows [Lotus].

* Any application provider should be able to mark their traffic for the low latency queue, with no restrictions other than standards compliance or other reasonable and openly documented technical guidelines. This maintains the loose cross-layer coupling that is a key tenet of the Internet’s architecture by eliminating the need for application providers and networks to coordinate and creates an environment of so-called "permissionless innovation".

5. Acknowledgements

Thanks to Bob Briscoe, Gorry Fairhust, Mat Ford, Vidhi Goel, Mirja Kuhlewind, Eliot Lear, Sebastian Moeller, Sebnem Ozer, Jim Rampley, Dan Rice, Greg Skinner, Joe Touch, Greg White, and Yiannis Yiakoumis for their review and feedback on this document.

6. IANA Considerations

RFC Editor: Please remove this section before publication.

This memo includes no requests to or actions for IANA.

7. Security Considerations

The key security consideration pertains to Queue Protection. As the current time, it is recommended that implementers utilize Queue Protection, to ensure that any traffic that is incorrectly marked for low latency can be detected and remarked for the classic queue. The necessity of Queue Protection remains something of a debate, with some firmly believing it is necessary but others believing that it is not needed. The latter view is that application developers have a natural incentive to correctly mark their traffic, because to do otherwise would worsen the quality of experience (QoE) for their users. In that line of thinking, if a developer mismarks, they and/or their users will notice and they will fix that error. However, it is also conceivable that malicious software could be operating on a user’s device or home network and that malicious software could try to send some much traffic to the low latency queue that the queue or both queues become unusable. This is quite similar to other "traditional" denial of servce (DoS) attacks, so it does not necessarily seems unique to low latency networking. But due to the possibility of this occuring, and low latency networking being such a new approach, it seems prudent to implement Queue Protection.

8. Network Neutrality Considerations

Network Neutrality (a.k.a. Net Neutrality) can mean a variety of things within a country, as well as between different countries, based on different opinions, market structures, business practices, laws, and regulations. Generally speaking, In the context of the United States’ market, it has come to mean that Internet Service Providers (ISPs) should not block, throttle, or deprioritize lawful application traffic, and should not engage in paid prioritization, among other things. Net Neutrality concerns can sometimes affect the deployment of new technologies by ISPs, so they should carefully consider regulatory issues when making deployment decisions.

As it is envisioned in the design of the IETF's new low latency networking protocols, the addition of a low latency queue at a network link is merely a second packet queue and does not mean that this queue is hierarchically prioritized or that it has more capacity. As a result, low latency networking appears to pose NO new Net Neutrality issues.

One key aspect of low latencty networking is that it operates, from the perspective of an ISP's deployment, is application-agnostic. The ISP creates a second network queue on key network links, but does not decide on their own what applications can use this queue. Rather, they add the queue and packet flows are sent to that queue based on packet marking set by application developers. This approach is far superior to older approaches, which caused significant Net Neutrality risks [Lotus], that used middleboxes to attempt to infer applications based on observing packet flows on ISP network links.

9. Revision History

RFC Editor: Please remove this section before publication.

v00: First draft

v01: Incorporate comments from 1st version after IETF-115

v02: Incorporate feedback from the TSVWG mailing list

v03: Final feedback from TSVWG and prep for sending to ISE

v04: Refresh expiration before major revision

v05: Changes from Greg Skinner and Eliot Lear

v06: More changes from Eliot Lear

v07: More changes from Eliot Lear

v08: Misc updates from IETF review

v09: Additional updates during review

10. Open Issues

RFC Editor: Please remove this section before publication.

- Open issues are being tracked in a GitHub repository for this document at https://github.com/jlivingood/IETF-L4S-Deployment/issues

11. Informative References

[RFC8325] Szigeti, T., Henry, J., and F. Baker, "Mapping Diffserv to IEEE 802.11", RFC 8325, DOI 10.17487/RFC8325, February 2018, <https://www.rfc-editor.org/info/rfc8325>.

[RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <https://www.rfc-editor.org/info/rfc9330>.

[RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <https://www.rfc-editor.org/info/rfc9331>.

[RFC9332] De Schepper, K., Briscoe, B., Ed., and G. White, "Dual-Queue Coupled Active Queue Management (AQM) for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9332, DOI 10.17487/RFC9332, January 2023, <https://www.rfc-editor.org/info/rfc9332>.

[RFC9435] Custura, A., Fairhurst, G., and R. Secchi, "Considerations for Assigning a New Recommended Differentiated Services Code Point (DSCP)", RFC 9435, DOI 10.17487/RFC9435, July 2023, <https://www.rfc-editor.org/info/rfc9435>.

[I-D.ietf-tsvwg-l4sops] White, G., "Operational Guidance on Coexistence with Classic ECN during L4S Deployment", Work in Progress, Internet-Draft, draft-ietf-tsvwg-l4sops-07, 17 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-l4sops-07>.

[I-D.ietf-tsvwg-nqb] White, G., Fossati, T., and R. Geib, "A Non-Queue-Building Per-Hop Behavior (NQB PHB) for Differentiated Services", Work in Progress, Internet-Draft, draft-ietf-tsvwg-nqb-27, 8 November 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-nqb-27>.

[I-D.briscoe-docsis-q-protection]
Briscoe, B. and G. White, "The DOCSIS(r) Queue Protection Algorithm to Preserve Low Latency", Work in Progress, Internet-Draft, draft-briscoe-docsis-q-protection-07, 23 November 2023, <https://datatracker.ietf.org/doc/html/draft-briscoe-docsis-q-protection-07>.

[BITAG] Broadband Internet Technical Advisory Group, "Latency Explained", 10 January 2022, <https://bitag.org/documents/BITAG_latency_explained.pdf>.

[Lotus] Eckerseley, P., von Lohmann, F., and S. Schoen, "Packet Forgery By ISPs: A Report on the Comcast Affair", 28 November 2007, <https://www.eff.org/wp/packet-forgery-isps-report-comcast-affair>.

[IETF-114-Slides]
White, G., "First L4S Interop Event @ IETF Hackathon", 25 July 2022, <https://datatracker.ietf.org/meeting/114/materials/slides-114-tsvwg-update-on-l4s-work-in-ietf-114-hackathon-00.pdf>.

[LLD] White, G., Sundaresan, K., and B. Briscoe, "Low Latency DOCSIS: Technology Overview", February 2019, <https://cablela.bs/low-latency-docsis-technology-overview-february-2019>.

[Ericsson] 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-time-critical-applications-over-5g-with-rate-adaptation-whitepaper.pdf>.

[CTI] International Telecommunications Union - Telecommunication Standardization Sector (ITU-T), "Optical line termination capabilities for supporting cooperative dynamic bandwidth assignment", Series G: Transmission Systems and Media, Digital Systems and Networks Supplement 71, April 2021, <https://www.itu.int/rec/T-REC-G.Sup71-202104-I>.

[IEEE] IEEE Computer Society (IEEE), "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", DOI 10.1109/IEEESTD.2021.9363693, IEEE Standard for Information Technology--Telecommunications and Information Exchange between Systems - Local and

Metropolitan Area Networks--Specific Requirements 802.11-2020, 26 February 2021, <https://ieeexplore.ieee.org/document/9363693>.

[Microsoft] Microsoft, "Quality of service (QoS) packet tagging on Xbox consoles", 19 August 2022, <https://learn.microsoft.com/en-us/gaming/gdk/_content/gc/networking/overviews/qos-packet-tagging>.

[Comcast] Comcast, "Comcast Introduces Nation’s First Ultra-Low Lag Xfinity Internet Experience with Meta, NVIDIA, and Valve", 29 January 2025, <https://corporate.comcast.com/press/releases/comcast-introduces-nations-first-ultra-low-lag-xfinity-internet-experience-with-meta-nvidia-and-valve>.

[CDT-NN] Doty, N. and M. Knodel, "Slicing the Network: Maintaining Neutrality, Protecting Privacy, and Promoting Competition. A technical and policy overview with recommendations for operators and regulators.", April 2023, <https://arxiv.org/pdf/2308.05829>.

[van-Schewick-1A] van Schewick, B., Jordan, S., Open Technology Institute at New America, and Public Knowledge, "FCC Ex Parte In the matter of Safeguarding and Securing the Open Internet, WC Docket No. 23-320", 20 March 2024, <https://www.fcc.gov/ecfs/document/103120890811342/1>.

[van-Schewick-1B] van Schewick, B., Jordan, S., Open Technology Institute at New America, and Public Knowledge, "Net Neutrality & Non-BIAS Data Services", 20 March 2024, <https://www.fcc.gov/ecfs/document/10323701322790/2>.

[van-Schewick-2] van Schewick, B., "Net Neutrality & 5G Network Slicing", 3 April 2024, <https://law.stanford.edu/wp-content/uploads/2024/08/van-Schewick-2024-5G-Network-Slicing-and-Net-Neutrality-Shetler-Steffen1.pdf>.

[van-Schewick-3]
van Schewick, B., "Network Slicing and Net Neutrality: No Throttling Rule", 18 April 2024, <https://law.stanford.edu/wp-content/uploads/2024/08/van-Schewick-2024-5G-Network-Slicing-and-No-Throttling-Rule-20240418.pdf>.

[Apple] Apple, "Testing and Debugging L4S in Your App", <https://developer.apple.com/documentation/network/testing-and-debugging-l4s-in-your-app>.

[Tussle] Clark, D., Wroclawski, J., Sollins, K., and R. Braden, "Tussle in Cyberspace: Defining Tomorrow’s Internets", 19 August 2002, <https://dl.acm.org/doi/10.1145/633025.633059>.

Author’s Address

Jason Livingood
Comcast
Philadelphia, PA
United States of America
Email: jason_livingood@comcast.com