

Independent Stream
Internet-Draft
Intended status: Informational
Expires: 3 September 2026

J. Livingood
Comcast
2 March 2026

ISP Dual Queue Networking Deployment Observations
draft-livingood-low-latency-deployment-15

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 some of the implications of those decisions and makes suggestions that can help drive adoption and acceptance of L4S and NQB based on observations from the world's first large scale deployment.

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 3 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	2
1.1. A Different Understanding of Application Needs	3
1.2. New Thinking on Low Latency Packet Processing	4
1.3. Application Performance Benefits	5
2. Key Low Latency Networking Concepts	5
2.1. Throughput is Bounded by Congestion Control Algorithms	6
2.2. Best Effort Priority	7
2.3. Shared Capacity	8
2.4. Access-Agnostic	8
2.5. End-to-End Diagram	8
3. Application Developer Implementation Suggestions	9
3.1. Delivery Infrastructure for L4S	9
3.2. Delivery Infrastructure for NQB	9
3.3. Only Mark Delay-Sensitive Traffic for L4S or NQB	10
3.4. Consider Application Needs in Choosing L4S vs. NQB	10
3.5. Example: iOS Application for L4S	10
4. ISP Implementation Suggestions	10
4.1. Allow ECN Across Network Boundaries	11
4.2. Allow DSCP-45 Across Network Boundaries	11
4.3. Last Mile Network (Access Network)	12
4.4. Customer Premise Equipment (Customer Edge)	12
4.5. Inside the Home - Customer Local Area Network (LAN)	13
4.5.1. 802.11 WiFi Queuing	13
4.6. Packet Loss Effects	13
4.7. Avoid Middleboxes	14
5. IANA Considerations	16
6. Security Considerations	16
7. Acknowledgements	16
8. Revision History	16
9. Open Issues	17
10. Informative References	17
Author's Address	22

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 some of the key deployment decisions and makes suggestions that may help drive adoption by network operators and application developers. These decisions are not inherent to L4S and NQB per se, but are decisions that can change depending upon differing technical, regulatory, business or other requirements. This document suggests that certain specific deployment decisions - based on observations from the world's first large scale deployment - that can help maximize the value of low latency networking to end users, network operators, and application developers.

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. Introducing a new queue enables incremental development of a new standard rather than changing existing congestion control algorithms for existing traffic on all networks, hosts, and clients - 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 is not delay-sensitive, at least on a millisecond level basis. The other type of traffic is real-time, interactive traffic that is latency-sensitive, such as video conferencing, gaming, and artificial intelligence (AI) interactions.

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.

Primary Goal	Examples	Flow Properties	IETF Classification
Low Latency	VoIP, DNS, game sync packets, some machine-to-machine and IoT communications	Low data rate, application-limited flows	Non-Queue-Building (NQB)
Low Latency with High Throughput	Cloud gaming, video conferencing, adaptive video streaming, instant messaging, web browsing, VoIP, cloud-based interactive video, cloud-based virtual and augmented reality	Low data rate, application-limited flows and high data rate, throughput-limited flows	Low Latency, Low Loss, Scalable Throughput (L4S)
High Throughput without Respect to Latency	Bursty traffic flows and high data rate traffic flows	Software Updates, Cloud Backups, P2P File Sharing	Queue-Building (QB)

Table 1: Latency Sensitivity of Example Applications

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.

1.3. Application Performance Benefits

The benefits of low latency networking to end user applications, especially interactive ones, is significant. In the Comcast network in the United States, this technology is deployed to over ten million homes as of January 2026. That encompasses over 300 million end user devices, based on the typical number of devices on a user Local Area Network (LAN). Comcast has shared in many IETF meeting presentations the data showing latency and jitter reductions (such as in [IETF-122-Slides]). At a high level, 99th percentile loaded latency for downstream traffic declined from ~65 ms to ~33 ms from implementation of downstream Active Queue Management (AQM) on the Cable Modem Termination System (CMTS). The new low latency queue (for L4S and NQB) further lowered downstream loaded latency to ~18 ms and upstream loaded latency to ~20 ms. The DOCSIS link layer itself in the Comcast is responsible for a baseline idle latency of 13-17 ms, which means that loaded latencies for L4S and NQB traffic are very close to idle latencies. Jitter is also significantly lower and more consistent.

This means that video conferencing applications such as Apple FaceTime and cloud gaming such as NVIDIA GeForce NOW and Valve's Steam platform will perform much better on a network that supports low latency networking. It benefits any marked application that is interactive in nature, where a user is interacting with a screen or other device, via mouse, keyboard, voice command, gesture, or other interactivity mechanism, and the interaction involves remote entities. Apple recently shared statistics [IETF-123-Slides] that they collected that confirm this performance, using Responsiveness Under Working Conditions tests [I-D.ietf-ippm-responsiveness].

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. It was also thought that networks and applications had to choose between either high throughput or low delay - but dual queue low latency demonstrates that both can be achieved simultaneously.

2.1. Throughput is Bounded by Congestion Control Algorithms

The Mathis Equation [Mathis] explains how throughput is bounded in the Reno congestion control algorithm that was in use at that time. Reflections on this approach were published later and hint at the potential for new approaches like those explored in this document [Mathis-Reflections].

This equation shows that throughput scales at $1/\sqrt{p}$, where p is the packet loss probability. The square root is a problem, because achieving a ten-fold increase in throughput requires a one-hundred-fold reduction in packet loss probability, and thus a ten-fold reduction in packet loss rate on a per-second basis. This becomes unscalable as path capacities increase, as the required time between congestion signals becomes longer than the typical duration of a connection.

This nonlinear relationship described in the equation means that, to take a link from 1 Gbps to 10 Gbps for example, packet loss would need to be reduced one-hundred-fold. This makes it increasingly hard to hit ever-higher data rates in real world environments, especially with WiFi and 5G links, where those levels of loss are unachievable at the current time.

While the dominant congestion control algorithm today (Cubic) improves upon this, it still requires long intervals between congestion signals as rates increase [RFC9438]. Alternative congestion controllers like BBR [BBR-Paper] aim to infer the link rate without relying on packet loss as a signal but do so with very limited information from the path, and can lead to fairness issues, oscillations in throughput, starvation and other unwanted effects [Starvation] [BBR-Analysis].

L4S replaces packet loss with an explicit signal (CE mark) and it takes advantage of the fact that, in contrast to packet loss, the explicit congestion signal is not a degradation of communication and thus can be provided liberally by the network.

L4S also requires that the congestion control algorithm scales its throughput proportional to $1/p$, where p here is the congestion marking probability. What this means in practice is that if you increase throughput by 10x, you only need to reduce the congestion signal probability by 10x rather than 100x, and the rate of congestion signals on per-second basis remains constant. As a result, applications get very frequent and fine-grained feedback that they use to make adjustments to their sending rate, and congestion control is fully scalable, with no loss in the frequency of signals as the path capacity increases.

This table illustrates the implications of this for an example of a TCP connection with 50 ms RTT (example idle and working latencies can be seen in [FCC-MBA-2023]).

Data Rate	Reno: Interval Between Congestion Signals	Cubic: Interval Between Congestion Signals	L4S: Interval Between Congestion Signals
1 Mbps	285 ms	160 ms	12.5 ms
10 Mbps	2.85 sec	350 ms	12.5 ms
100 Mbps	28.5 sec	740 ms	12.5 ms
1 Gbps	4.75 min	1.6 sec	12.5 ms
10 Gbps	47.5 min	3.5 sec	12.5 ms
100 Gbps	7.9 hrs	7.5 sec	12.5 ms

Table 2: Comparison of Congestion Signal Frequency at 50ms RTT

2.2. Best Effort Priority

Low latency traffic is not prioritized over other (best effort priority) "classic" Internet traffic. 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].

Those approaches are grounded in an assumption that there is scarce bandwidth and so the network must grant higher priority to some flows, creating bandwidth winners and losers. That has raised the ire of network neutrality proponents, who have focused on negative effects this can have on flows that are not prioritized and the terms under which certain applications or destinations may be granted priority.

Another contrast with 5G network slicing is that IETF standards such as L4S do not require network operator-specific APIs, direct coordination, or the granting of permission by technical or legal means. Rather, marking packets with ECT(1) or CE as part of L4S will work on any network that has implemented L4S, with no need to coordinate with each network or seek permission. This exemplifies the loose coupling across layer (network and application) and permissionless innovation (at the edge by application developers) that are core tenets of the Internet's architecture.

The only exception for priority is within a user's in-home Wi-Fi network (in the case of a fixed network) due to the particulars of how the IEEE 802.11 wireless protocol [IEEE] functions at the current time - see [RFC8325]. 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].

2.3. Shared Capacity

Low latency networking flows do not get access to greater capacity than "classic" flows. Thus, a user's total provisioned or permitted capacity on an ISP access network link is shared between both classic and low latency queues. This removes an incentive to game the system through low latency packet marking.

Another way to think of L4S is that it is not really an instruction to a dual queue network link to ask for low latency treatment per se. Rather, applications marking for L4S are essentially making a promise to the network that the application will rapidly reduce throughput if needed (when detecting CE marks); it promises to help keep latency low by being a good network citizen.

2.4. 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], WiFi, Low Earth Orbit (LEO) satellite and many other types of networks.

2.5. End-to-End Diagram

This diagram shows an example of the end-to-end path for two different user devices; one phone connected to the home WiFi network and one PC connected via Ethernet to the home LAN. Along this full path, L4S or NQB marks must be able to pass between an application on a client device and an application server. While in theory dual queue low latency networking can be implemented on any of the routers on this path or the access point, the most common bottlenecks are in

the home network (WiFi), the CPE router, and the aggregation router. This alone will provide a tremendous benefit to interactive end user applications and dual queue need not in that case be deployed at any of the other hops.

(Artwork only available as SVG: see <https://www.ietf.org/archive/id/draft-livingood-low-latency-deployment-15.html>)

Figure 1: Network Architecture Diagram

3. Application Developer Implementation Suggestions

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. The application 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 (removing) the value set in 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

NQB uses the DSCP-45 code point in the DiffServ part of the packet header to ensure NQB works end-to-end. To achieve end-to-end marking, 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 (removing) a DSCP-45 mark. The server does not need to run a special responsive congestion controller. Since it is common for networks to bleach 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.

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 sensitive 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 necessary.

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.

3.5. Example: iOS Application for L4S

Implementations will vary depending upon variables such as server operating system, transport and application protocols, and client software and operating system. To use the example of an Apple device such as an iPhone, running iOS, L4S support is built into iOS by default [L4S-Testing]. If the developer uses the QUIC protocol, they can use "NSURLSession" [NSURLSession] which enables L4S.

4. ISP Implementation Suggestions

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

These suggestions focus on not bleaching (removing) ECN or DSCP-45 marks for L4S and NQB, respectively. That is because low latency networking will not work on any connected downstream network if these marks cannot travel end-to-end (e.g., a customer of the network, including fixed and mobile end users, enterprises, and other users).

There are also applications that will work on a peer-to-peer, intra-domain basis in a network, and those also would not be able to take advantage of L4S and NQB.

ECN generally works across all domain boundaries, so unless a network has gone out of their way to bleach the ECN header, this is unlikely to be a major concern.

The key issue for modification or bleaching (removal of marks) is really DSCP. NQB depends upon DSCP-45 passing across domain boundaries. Historically, there is little alignment between networks over common DSCP marking since historically DSCP marks have not been used across domain boundaries, and each network has implemented unique approaches to how DSCP is used. In practice this means most network have a router policy that will replace most or all DSCP values on ingress and mark DSCP-0 or 8 or whatever code point the receiving network uses for best effort or default traffic. That means a router policy change will be needed here. Any concern over this can be mitigated by ensuring that DSCP-45 traffic is treated as best effort priority like regular internet traffic.

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 L4S-aware and able to add a CE mark to signal that it is experiencing congestion at that hop. The end-to-end integrity of ECN marks is explored in Appendix C.1 of [RFC9331].

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). The end-to-end integrity of the DSCP-45 mark is explored in Section 6.4 of [I-D.ietf-tsvwg-nqb].

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 these 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 a 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. 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 than 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. 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 Section 7.3 of [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. 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.

4.6. Packet Loss Effects

As networks implement low latency networking, including L4S, NQB, and new forms of AQM, they may notice somewhat less packet loss. This is because applications now have new signals, such as a CE mark, to signal congestion. This signal is faster than the typical TCP packet loss mechanism.

Thus, this is expected, as noted in Section 5.1 of Section 5.1 of [RFC9330]:

| Latency is not the only concern of L4S. The 'Low Loss' part of
| the name denotes that L4S generally achieves zero congestion loss
| due to its use of ECN. Otherwise, loss would itself cause delay,
| particularly for short flows, due to retransmission delay
| [RFC2884].

| RFC9330

In addition see Section 3 of Section 3 of [RFC2884]:

| Since ECN marks packets before congestion actually occurs, this is
| useful for protocols like TCP that are sensitive to even a single
| packet loss. Upon receipt of a congestion marked packet, the TCP
| receiver informs the sender (in the subsequent ACK) about
| incipient congestion which will in turn trigger the congestion
| avoidance algorithm at the sender.

| RFC2884

4.7. Avoid 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 middlebox in the ISP network. In this case, these are network elements that have not been configured by an end user application but rather inserted in the path by a network operator as an element where active network management policies could be applied. An application-configured element, in contrast, could be something like an Oblivious HTTP Relay Resource (server) [RFC9458].

As explored below, avoiding a middlebox and leaving marking to application clients is the most efficient, least prone to misclassification, and is most acceptable from the standpoint of network neutrality. So, while a middlebox may exist on an end-to-end path, it is best not to have a middlebox actively classify traffic for a low latency queue; let applications make that decision to limit the potential for problems that middleboxes can cause.

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. IANA Considerations

RFC Editor: Please remove this section before publication.

This memo includes no requests to or actions for IANA.

6. Security Considerations

The key security consideration relates to access to capacity by flows using a low latency queue when those flows do not follow the guidelines or L4S or NQB behavior. If flows mark for low latency queue treatment but are not correctly rate-responsive or low bitrate, then there is a potential risk that they may disrupt the low latency goal of the low latency queue. [RFC9331] introduces a Queue Protection function to mitigate this risk. More can be learned about this in Section 5.2 of Section 5.2 of [I-D.ietf-tsvwg-nqb].

The necessity of Queue Protection (also referred to as Traffic Protection) remains in 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. 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 so much traffic to the low latency queue that the queue or both queues become unusable. This is quite similar to other "traditional" denial of service (DoS) attacks, so it does not necessarily seem unique to low latency networking. But due to the possibility of this occurring, and low latency networking being such a new approach, it seems prudent to implement Queue Protection.

7. Acknowledgements

Thanks to Bob Briscoe, Stuart Cheshire, Gorrry Fairhurst, Mat Ford, Vidhi Goel, Mirja Kuhlewind, Eliot Lear, Matt Mathis, Sebastian Moeller, Sebnem Ozer, Jim Rampley, Dan Rice, Jason Schnitzer, Greg Skinner, Joe Touch, Greg White, and Yiannis Yiakoumis for their review and feedback on this document and related presentations.

8. 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
- v10: Final changes from review
- v11: Suggestions and nits from Greg White, other nits
- v12: Sec 2.1 suggestions from Greg White, other nits
- v13: Nits from Jason Schnitzer, suggestions from Matt Mathis
- v14: Changes from Eliot Lear
- v15: fix minor SVG nit

9. Open Issues

RFC Editor: Please remove this section before publication.

10. Informative References

- [RFC2884] Hadi Salim, J. and U. Ahmed, "Performance Evaluation of Explicit Congestion Notification (ECN) in IP Networks", RFC 2884, DOI 10.17487/RFC2884, July 2000, <<https://www.rfc-editor.org/info/rfc2884>>.
- [RFC8325] Sziget, 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>>.
- [RFC9438] Xu, L., Ha, S., Rhee, I., Goel, V., and L. Eggert, Ed., "CUBIC for Fast and Long-Distance Networks", RFC 9438, DOI 10.17487/RFC9438, August 2023, <<https://www.rfc-editor.org/info/rfc9438>>.
- [RFC9458] Thomson, M. and C. A. Wood, "Oblivious HTTP", RFC 9458, DOI 10.17487/RFC9458, January 2024, <<https://www.rfc-editor.org/info/rfc9458>>.
- [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-09, 12 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-l4sops-09>>.
- [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-33, 16 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-nqb-33>>.

[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>>.

[I-D.ietf-ippm-responsiveness]

Paasch, C., Meyer, R., Cheshire, S., and W. Hawkins, "Responsiveness under Working Conditions", Work in Progress, Internet-Draft, draft-ietf-ippm-responsiveness-08, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-responsiveness-08>>.

[BITAG]

Broadband Internet Technical Advisory Group, "Latency Explained", 10 January 2022, <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5677705>.

[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>>.

[Mathis]

Mathis, M., Semke, J., Mahdavi, J., and T. Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm", DOI 10.1145/263932.264023, Association for Computing Machinery Sigcomm Computer Communications Review, 1 July 1997, <<https://dl.acm.org/doi/10.1145/263932.264023>>.

[Mathis-Reflections]

Mathis, M., "Reflections on the TCP Macroscopic Model", ACM SIGCOMM Computer Communication Review vol. 39, no. 1, pp. 47-49, DOI 10.1145/1496091.1496099, January 2009, <<https://doi.org/10.1145/1496091.1496099>>.

[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>>.
- [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>>.

[IETF-122-Slides]

Livingood, J., "Dual Queue Low Latency Networking Update", 17 March 2025, <<https://datatracker.ietf.org/meeting/122/materials/slides-122-tsvwg-41-jason-livingood-update-on-deployment-00>>.

[IETF-123-Slides]

Cheshire, S., "L4S with Apple Devices", 23 July 2025, <<https://datatracker.ietf.org/meeting/123/materials/slides-123-tsvwg-l4s-with-apple-devices-00>>.

[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>>.

[FCC-MBA-2023]

Federal Communications Commission, "Twelfth Measuring Broadband America Fixed Broadband Report", January 2023, <<https://www.fcc.gov/reports-research/reports/measuring-broadband-america>>.

[URLSession]

Apple, "Apple Developer Documentation, URLSession",
<<https://developer.apple.com/documentation/foundation/urlsession>>.

[L4S-Testing]

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

[BBR-Paper]

Cardwell, N., Cheng, Y., Gunn, C. S., Yeganeh, S. H., and
V. Jacobson, "BBR: Congestion-Based Congestion Control",
ACM Queue, Vol. 14, Issue 5, September 2016,
<<https://dl.acm.org/doi/10.1145/3022184.3012426>>.

[Starvation]

Arun, V., Alizadeh, M., and H. Balakrishnan, "Starvation
in End-to-End Congestion Control", Proceedings of the ACM
SIGCOMM 2022 Conference, August 2022,
<<https://doi.org/10.1145/3544216.3544223>>.

[BBR-Analysis]

Scholz, D., Jaeger, B., Schwaighofer, L., Raumer, D.,
Geyer, F., and G. Carle, "Towards a Deeper Understanding
of TCP BBR Congestion Control", 2018 IFIP Networking
Conference, May 2018,
<<https://doi.org/10.23919/IFIPNetworking.2018.8696830>>.

Author's Address

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