Lessons Learned from the World's First Deployment of Dual Queue Low Latency Networking

A technical paper prepared for presentation at SCTE TechExpo24

Takashi Hayakawa, Comcast

Jeff Howe, Comcast

Chad Humble, Comcast

Jason Livingood, Comcast

Eric Lorenzo, Comcast

Mike Rudolph, Comcast

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

In the post-gigabit world, networks have entered an *era of bandwidth abundance* where the capacity available to consumers far outstrips any available application use cases. Ironically, the only application that can use 1 Gbps symmetric is a speed test, and this bandwidth can comfortably support a great many users. Despite this bandwidth abundance, users still regularly encounter Quality of Experience (QoE) impairments, which are primarily attributable to poor (high) working latency of hundreds or even thousands of milliseconds[1].

The introduction of Active Queue Management (AQM) in the DOCSIS 3.1® specifications was a leap forward in reducing working latency in DOCSIS networks – achieving up to a 90% reduction in a single network queue. But parallel work at the IETF concluded that *two* network queues could enable even greater latency benefits, with a new queue that was designed specifically for flows that need low latency. These IETF standards have been incorporated into the Low Latency DOCSIS specifications.

In 2023, Comcast started the world's first field trials of these new dual queue low latency standards. This paper and presentation explain Comcast's experiences during this trial, which is useful to any network operator considering deployment of dual queue low latency networking – whether on a DOCSIS network or WiFi, 5G, or other type of network.

2. Before Low Latency DOCSIS: Observed Improvements with AQM

In DOCSIS, the initial approach to reducing latency was called buffer control[2]. This was implemented in a network queue where packets were serviced on First In, First Out (FIFO) basis. With this mechanism, as the queue fills, new packets are dropped if the queue is full (called "tail drop"). Correspondingly, enqueued packets that could be determined as stale are retained until they are serviced as they reach the front of the queue. Buffer control established a threshold for the depth of that queue, which varied depending upon the provisioned bandwidth to or from a given cable modem.

CableLabs continued to work on further latency reductions that were based on Active Queue Management (AQM)[3]. AQM is the policy of dropping packets inside a buffer before that buffer becomes full. The AQM goal is to reduce network congestion and improve latency.

Two different AQM algorithms are defined in DOCSIS. One for classic queuing (DOCSIS-PIE) that was introduced in DOCSIS 3.1, and the other is for the low latency queue (iAQM). When the AQM DOCSIS-

---

[1] To understand latency, especially the difference between working and idle latency, see https://bitag.org/documents/BITAG_latency_explained.pdf.

[2] "DOCSIS® Best Practices and Guidelines, Cable Modem Buffer Control, CM-GL-Buffer-V01-110915", see https://account.cablelabs.com/server/alfresco/a4c077e2-cc55-4a42-a557-32f92e637266.

[3] "Active Queue Management (AQM) Based on Proportional Integral Controller Enhanced (PIE) for Data-Over-Cable Service Interface Specifications (DOCSIS) Cable Modems", RFC 8034, see https://www.rfc-editor.org/rfc/rfc8034.html.

PIE algorithm deployment was completed by Comcast in the early part of the COVID-19 pandemic, the results were excellent[4].

3. Low Latency DOCSIS: Integrating L4S and NQB

Low Latency, Low Loss, Scalable Throughput (L4S)[5] flows are generated by clients or applications with the ability to monitor and act upon the congestion experienced (CE) bits within IP or IPv4 headers conveying this data. They are classified into the low latency queue by their population of the ECT1 identifier within the IP or IPv6 headers.

Non-Queue-Building (NQB)[6] applications provide low-rate, consistent-flow traffic also targeted for the low latency queue within the dual queue structure specified for DOCSIS flows. They are classified into the low latency queue with a DSCP marking (typically DSCP45). The LL queue AQM algorithm (iAQM) only marks those packets that have the ECT(1) bit populated, and then only if CE marking is applicable for a specific packet.

When an application sends packets at a rate fast enough to build the LL queue sufficiently, the CE marking is applied with probability related to the length of the remaining queue.

This can also be adjusted by the length of the classic queue, so that LL marking is sensitive to competing traffic.

Properly behaving L4S applications, when observing a CE mark in a response message will reduce the rates/numbers of transmitted packets to reduce the building queue size. An example of this might be video reducing the quality of its transmissions from 4K video to 720p – corresponding to several milliseconds that CE is being advertised.

Clients sending NQB packets, on the other hand, are not exposed to CE markings and therefore are not expected to reduce rates in a congestion scenario. However, these NQB clients are expected to send regular, low-rate traffic that is easily absorbed given the user's subscribed rate.

4. Identifying & Improving Common Bottlenecks

Dual queue low latency networking can be deployed on any network link (network interface), but will not have an impact on working latency unless it is a bottleneck link. While there is always a bottleneck (point of lowest bandwidth capacity) on the end-to-end path of any flow, these bottlenecks can change over the course of time – and even over the duration of an individual flow.

While it may seem unrealistic to try to deploy dual queue to all of these bottlenecks, the reality is that there are typically just *three* link domains – within the control of an Internet Service Provider (ISP) to varying degrees - where this is most common. Addressing these three areas will have a high return on

---

[4] "Improving Latency with Active Queue Management (AQM) During COVID-19", https://arxiv.org/abs/2107.13968.

[5] See RFC 9330 at https://www.rfc-editor.org/rfc/rfc9330 and RFC 9331 at https://www.rfc-editor.org/rfc/rfc9331.html.

[6] See https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/.

investment in improving end user QoE: (1) upstream from the home/device to the ISP access network, (2) downstream from the ISP access network to the home/device, (3) the user or ISP's WiFi network[7].

In DOCSIS, the Low Latency DOCSIS specification describes how to implement this downstream – in the Cable Modem Termination System (CMTS) – and upstream – in the cable modem. The upstream link is most simply improved when the layer 2 cable modem is integrated with a layer 3 gateway (router) function, such as is the case with devices like Comcast's XB8 gateway[8].

If the customer has either turned off the IP gateway function of such a device (called running in "bridge mode"[9]) or is using a customer owned and managed cable modem (COAM, also known as "retail" device), then to take full advantage of dual queue networking, the customer's router will need to be able to mark outbound traffic for low latency and be able to appropriately handle inbound low latency traffic.

5. Home WiFi Networks

Low Latency DOCSIS addresses the first two most common bottlenecks noted in the prior section. However, most customer traffic is transported wirelessly over the home network, using WiFi. This is important to recognize, because it has been observed that above several hundred megabits per second of access network bandwidth, the home WiFi network is the primary bottleneck[10].

In order to realize the benefits of dual queue in the home network, we needed to map dual queue packets in the low latency queue to a separate queue in the wireless home network. Unfortunately, there is not a precise mapping between WiFi's Wi-Fi Multimedia (WMM) in the 802.11 link layer and the IETF's IP layer packet marking for dual queue. WMM has four key queues: best effort (AC_BE), background (AC_BK), video (AC_VI), and voice (AC_VO)[11].

To provide a way for low latency IP flows to map into the 802.11 wireless domain, the IETF working group focused on dual queue has recommended that any L4S or NQB packets be handled in the WMM AC_VI wireless queue, providing that traffic with a lower latency path on the WiFi portion of the path. This was an integral part of field testing and it appears to be working as expected, with no negative effects for users.

However, that is not to say that testing did not encounter any issues. In fact, one significant issue that was discovered trial preparation had to do with the non-LLD (or classic) high-speed data (HSD) traffic incorrectly leaking an internal DSCP mark into customer home networks. The Comcast core network uses DSCP values to distinguish different types of network traffic by setting the DSCP value as traffic enters the Comcast network. The value used within the Comcast core network for normal HSD traffic is CS1. When this traffic then leaves the Comcast network, it was left as being marked as CS1. In discussion with some customers on this topic, it was pointed out that this could cause some problems in the Home

[7] Typically this is the user's in-home WiFi LAN, but in some cases can be an ISP-managed public WiFi network, such as Xfinity WiFi.

[8] See https://www.xfinity.com/learn/internet-service/supersonic.

[9] See https://www.xfinity.com/support/articles/wireless-gateway-enable-disable-bridge-mode.

[10] "Measuring the Prevalence of WiFi Bottlenecks in Home Access Networks", see https://arxiv.org/abs/2311.05499.

[11] See https://www.wi-fi.org/knowledge-center/faq/what-is-wmm.

Wireless Network, because DSCP CS1 traffic is typical mapped into WMM AC_BK traffic (background), which has the potential of being starved out by other AC_VO, AC_VI, and AC_BE traffic on the home network. While all traffic was handled in that queue, things were fine, but the prospect of using two queues with one being marked as "background" traffic may have led to user-perceptible Quality of Experience (QoE) impairments.

Since Comcast is setting the DSCP value on the way into the network, the solution to this issue was to modify the DOCSIS Type of Service (ToS) overwrite configuration for the HSD flows to set the DSCP value to CS0. This would cause the HSD traffic to be mapped into WMM AC_BE (best effort), allowing the downstream HSD traffic to compete equally with other best effort traffic on the wireless network.

As noted in the prior section, dual queue is most easily enabled in the home network and access network when an integrated cable modem gateway is used. In cases where a COAM cable modem is used, and thus the customer also runs their own WiFi network (either integrated into the cable modem, integrated into a home router, or in a standalone WiFi access point), they need to be sure that (1) the home router is able to mark & interpret L4S and NQB marks, and (2) their WiFi network will map L4S and NQB traffic into the WMM AC_VI queue.

6. More Permissive Outbound NQB Marking

During initial testing with some applications, we found DSCP 40, 46, and 56 upstream markings for flows intended to be classified into low latency queues. On Comcast's gateway devices, where the layer 3 router and layer 2 cable modem are integrated, we were able to make the entry criteria for the upstream low latency queue more permissive. Otherwise, and DSCP marks would typically have been bleached. We did so based on some feedback that until the IETF NQB draft was made an RFC and a IANA port assignment made, that developers preferred to maintain their existing marking – until they were certain that the marking would not change again in the future.

Thus, Comcast gatways – operating in gateway mode, not in bridge mode - now recognize these alternative DSCP marks and translate them to DSCP 45 as they transit the CMTS and treats these flows with low latency markings.

7. Coupled Classic Queue with DOCSIS-PIE AQM & L4S/NQB Low Latency Queue

As noted above, different AQM algorithms are defined in DOCSIS. One AQM is for classic queuing (DOCSIS-PIE) and a new queue has been defined separately for low latency (iAQM) – for L4S and NQB traffic. As a part of trial testing, we evaluated the performance of no AQM, DOCSIS-PIE AQM, and low latency, and the latency differences between these was striking.

We performed these tests using an internal tool, sequentially measuring both the idle and the loaded latency on each of the Classic and Low Latency queues. Figure 2 - Latency over time with AQM applied has three plots for each of seven users. Two from each plot, (blue – loaded latency in the LL queue, and red – idle latency) are interesting to this discussion. The blue represents loaded latency of measured in the low latency (LL) queue. The red represents idle latency. Our point is that the application of Low Latency DOCSIS reduced the loaded latency in the classic queue, and made it track with the idle latency more closely over time.

The AQM quickly responds to the classic queue build-up. Figure 1 is an example of how the DOCSIS upstream queue responds to a series of file uploads. The classic queue without AQM suffers from very high latency (300-400ms) during the entire duration of each upload whereas the latency goes down soon after the spike at the beginning of each upload. Meanwhile the low latency queue is not affected by the classic queue's congestion, maintaining the good latency (10~20ms).

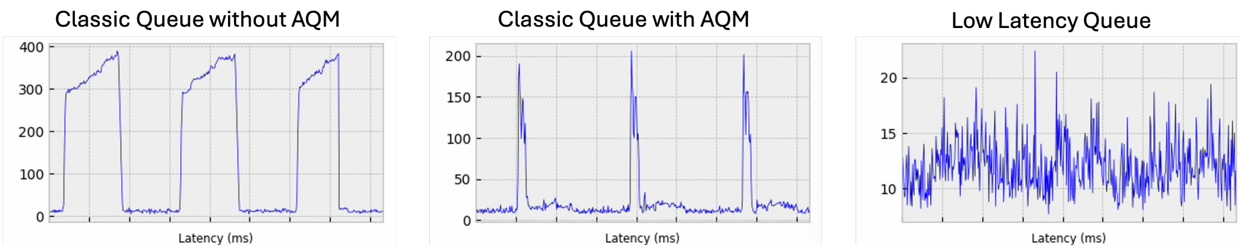

Figure 1 - How DOCSIS queue builds up during file upload via classic queue

(content continues on next page)

Similar to the above results, when LLD was applied to a DOCSIS Remote PHY Device (RPD) containing 57 Low Latency users, the aggregated latency reduction is illustrated in Figure 3 - Aggregated Latency over time with AQM applied. The time of the activation of the LL queue can be seen when the loaded latencies drop precipitously.

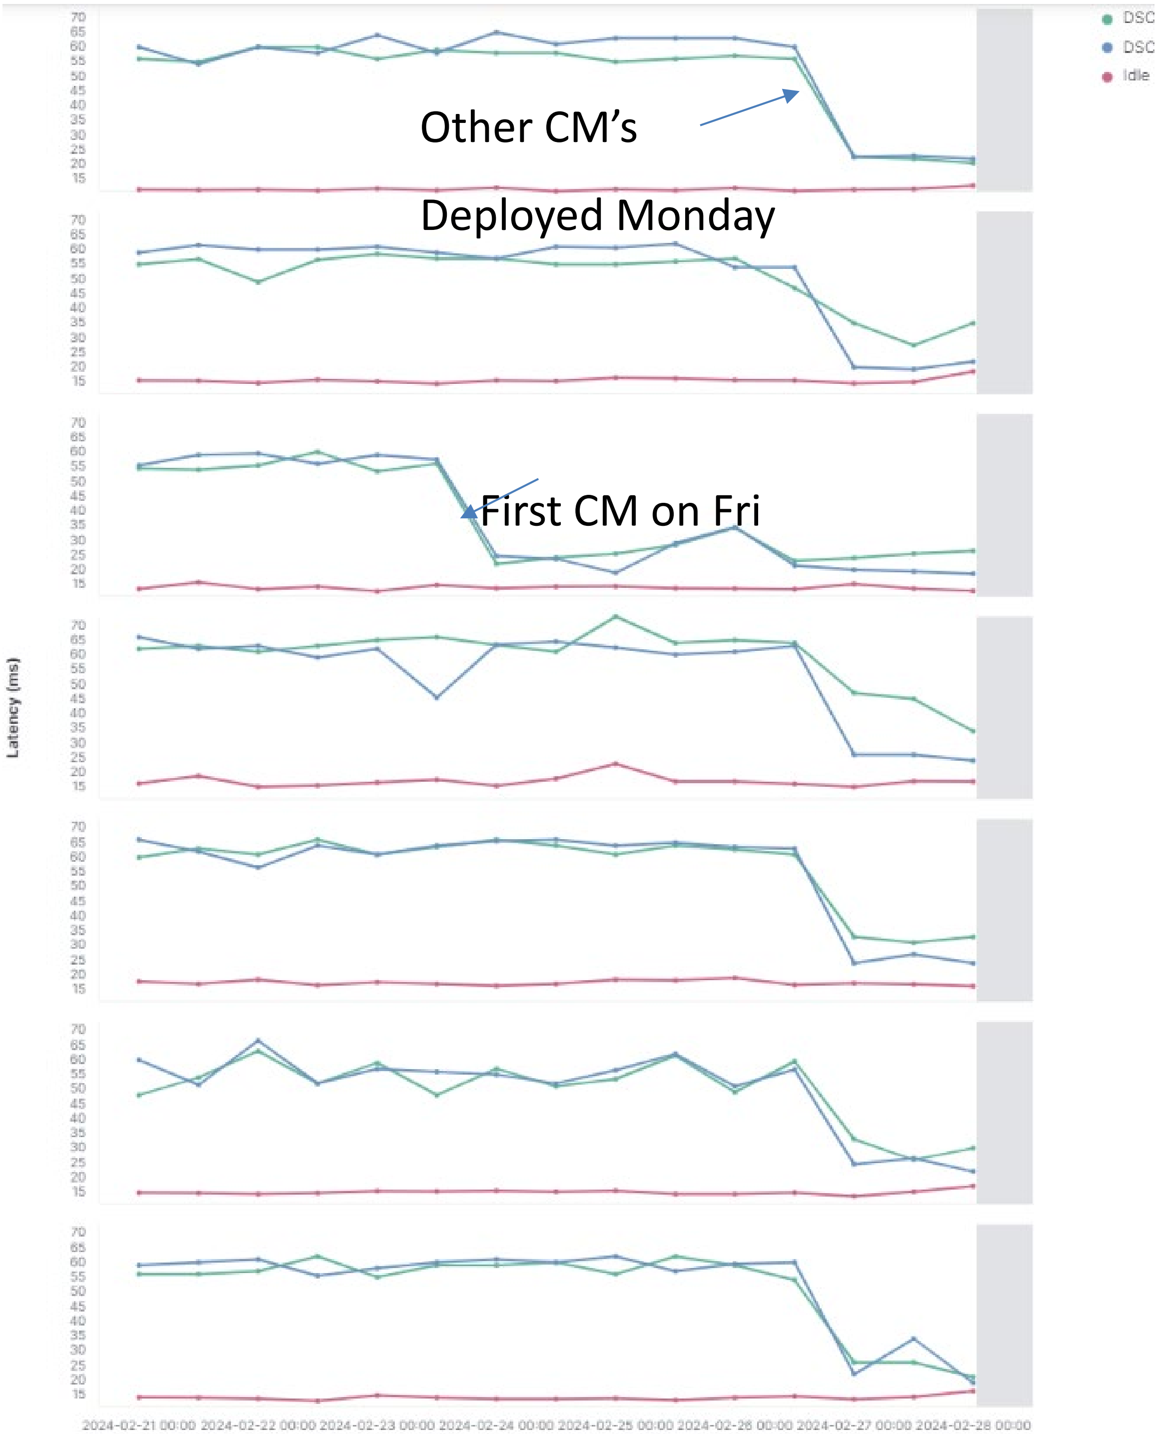

Figure 2 - Latency over time with AQM applied

In Figure 3 - Aggregated Latency over time with AQM applied, we can also see the drop in working latency when dual queue is enabled. This is aggregated by cable modem type, and we observe that different chipsets perform slightly differently. Divergent chipsets can likely be optimized over time to perform similarly.

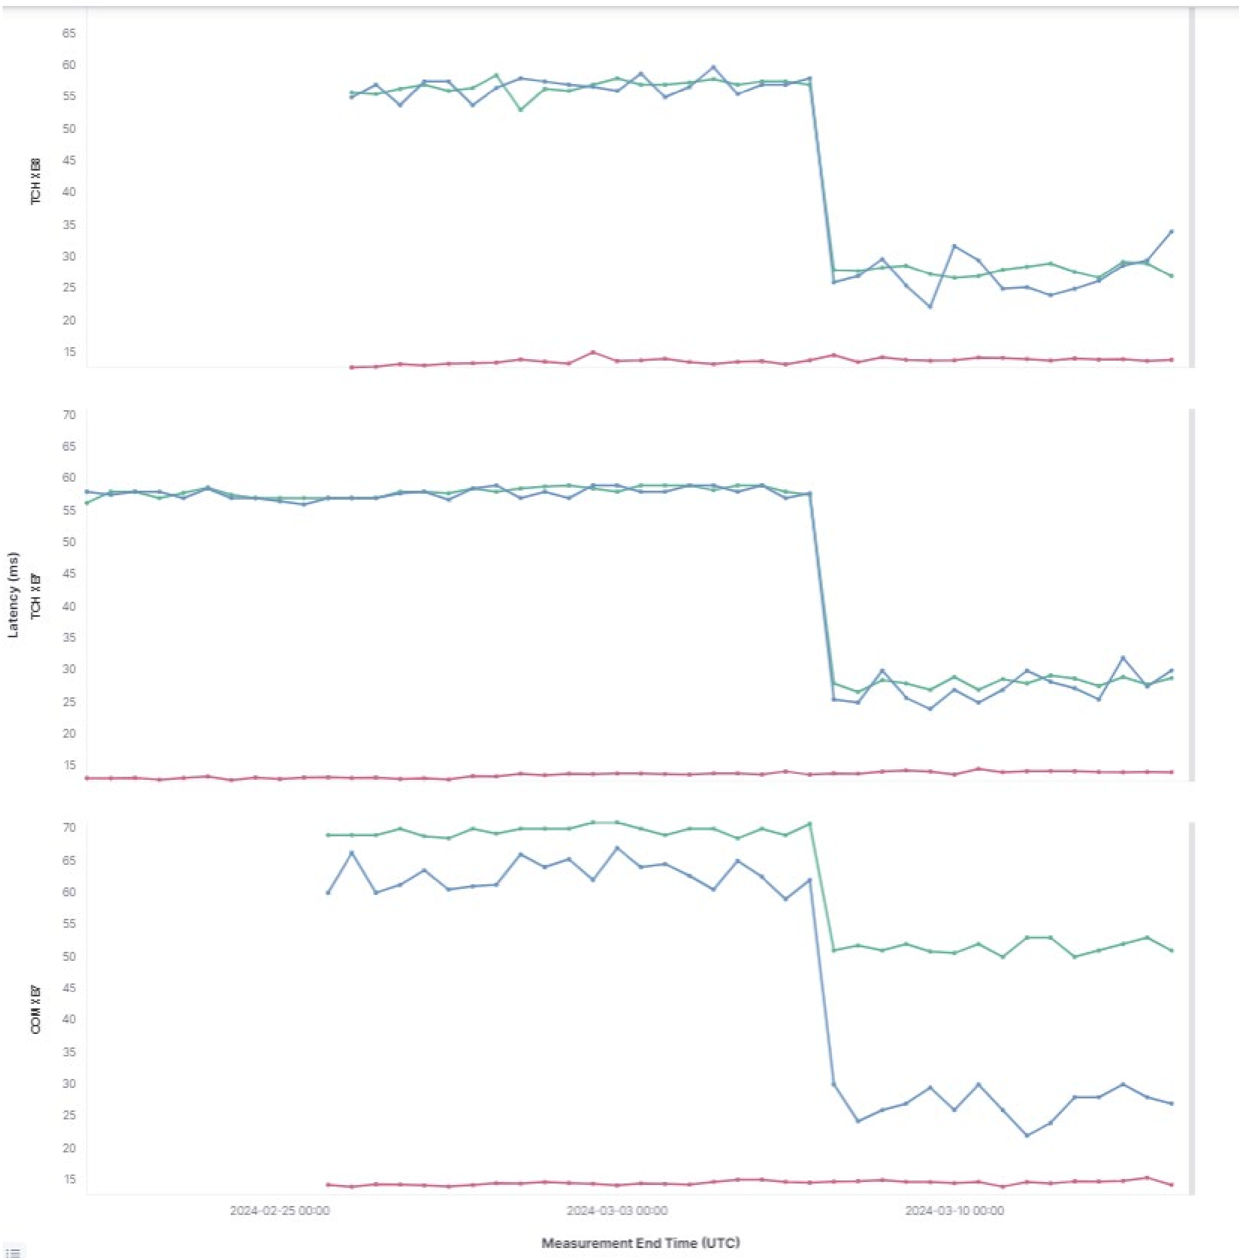

Figure 3 - Aggregated Latency over time with AQM applied.

8. High Level Trial Timeframe

Comcast announced the low latency field trial on June 16, 2023[12] in order to begin recruiting customer volunteers. We began provisioning employees on July 11, 2023, a typical first step in field tests. We provisioned customers beginning on August 13, 2023, with a first wave and on October 11, 2023 in a second wave, growing to roughly 200 total. All trial used needed to be on a virtual CMTS (vCMTS), a residential class of service, and be using one of four possible devices (two Comcast gateways and two COAM modems). In addition, we tested our first inter-domain NQB packet marking with Valve (DSCP-45) on October 5, 2023.

All of the testing in 2023 was focused on the two primary bottlenecks: the cable modem upstream and home WiFi network. In 2024, we expanded to test downstream from the CMTS once the vCMTS platform feature was ready for production on April 3, 2024 with employee testers and on April 19, 2024 with all trial users. The feature was defined as production ready in June-July 2024, with deployment plans being determined at that time.

9. No Bleaching: Enabling End-to-End Marking

A foundational first step was to enable L4S and NQB marking to work end-to-end, from an application on the home network to an off-network internet destination. For L4S that means not bleaching or modifying the ECN header, so that ECT(1) and CE marks pass in both directions. Since ECN bleaching isn't particularly common on the internet, once ECN marking is enabled in the access network (CM and CMTS) it will usually work fine end-to-end.

We initially believed that ECN marks worked, but later realized we were only testing in one direction (outbound to the internet). As a result, we had to update network configurations to stop bleaching ECN marks downstream. This was a easy and non-disruptive step to take, once we identified the issue.

For NQB, this was much more challenging, because most networks use DSCP marking only internally within their own network domain and there is not really any commonality between different network domains. It has long been the practice that a network will bleach DSCP marks on ingress (not egress), and convert packets to whatever DSCP marks the network has defined for best effort internet traffic (or other traffic, such as telephony).

In order to get NQB to work on an interdomain basis, a network first needs to establish a policy to allow DSCP-45 marks past the ingress router and all the way down to the user's CPE and home network. In cases where code point 45 is already used, internal hops may need to convert from 45 to some other mark, and then back to 45 when packets enter the CPE, and vice versa for outbound flows. Over the long term, networks should move other traffic off of code point 45 so that no remarking will be necessary.

This internal remarking approach is actually the one we needed to take. That is because DSCP-45 had been reserved for a potential type of traffic, which had actually never achieved meaningful scale. So on a temporary basis, there is internal remarking being done. But in the long-term, DSCP policies will be updated and DSCP-45 cleared for use by NQB.

---

[12] See https://corporate.comcast.com/stories/comcast-kicks-off-industrys-first-low-latency-docsis-field-trials

10. Upstream Enabled First, Downstream Later

LLD queuing algorithms are primarily implemented in the cable modem for upstream traffic and in the CMTS for downstream traffic. Both CM and CMTS play a role in either direction, but the data plane level implementations are primarily done separately. Downstream on the CMTS, upstream in the cable modem.

The control logic for low latency packet handling uses a DOCSIS Type-Length-Value (TLV). Implementation of support for TLV24 (upstream) and TLV25 (downstream) or similarly TLVs 70/71 was done in a common CMTS release. The Harmonic vCMTS used at Comcast was phased this way, with upstream first and downstream second. Initial support for the registration and control of upstream LLD was available in a software release prior to the vCMTS release containing downstream support. As a result, the upstream LLD capabilities preceded the downstream, and were trial tested in a cycle ahead of the downstream capabilities. The separation period spanned nearly half a year.

Low Latency configuration in the DOCSIS MAC and Upper Layer Protocols Interface (MULPI) specification is somewhat flexible. It allows parameter specification within the configuration file, or within the CMTS global data via TLV references. It also provides a precedence ruleset. This flexibility provided the opportunity to support earlier upstream LLD trials with explicit configurations supplied in the cm configuration file. Later, as we built tools to scale aggregate QoS profiles, we used the CMTS global data to contain much of our subscriber configurations. This is viewed as more nimble when considering operationally maintaining these parameters over time.

As a result, the upstream trials came first and focused on evaluating CM capabilities using CM-config file parameter adjustment and inclusion as necessary. Later we introduced the ability to use our per-CMTS structures – QoS Profiles (AQPs[13]) with new CMTS software to accommodate these, also introduced the downstream dataplane support.

11. New Parameters Required New Tools

LLD specifications require new parameters in both the .cm configuration file and (depending on deployment mode) in global data contained within the CMTS.

Rate and parameter based global data in a CMTS is managed by Comcast in a flexible tool/database system labeled G2. This tool aggregates changes of many types into configurations specific to each CMTS and uses conversions to translate that to syntaxes required by each CMTS vendor.

A similar suite of scripts is required to generate .cm configuration files as they iterate and evolve over rate changes, security evolutions, and functionality increases.

Parameters applicable to Low Latency can be included

1. in the .cm configuration file,
2. in the global data defined as a Service Class Name (SCN), or
3. within a table labeled Aggregate QoS profile. The Aggregate QoS Profile is global data with dimensionality aligned with a service class name, so it was reasonable to update the G2 system housing SCNs to also store, create, and manage the AQPs.

Therefore, the systems that generate these require upgrades to support these requirements.

[13] Technically AQP stands for Aggregate Service Flow (ASF) Quality of Service (QoS) Profile.

1) The scripts and tools employed to generate .cm parameters and TLVs was upgraded to include support for the TLVs listed below:

Downstream Service Flow Encodings

AQM Encodings

/* TLV 25.40.2 : */ Classic AQM Latency Target:

/* TLV 25.40.3 : */ AQM Algorithm:I

/* TLV 25.40.4 : */ Immediate AQM Max Threshold:

/* TLV 25.40.5 : */ Immediate AQM Range Exponent of Ramp Function:

/* TLV 71 */Downstream Aggregate Service Flow (ASF)

/* TLV 71.1 : */ Service Flow Reference:

/* TLV 71.7 : */ Traffic Priority:

/* TLV 71.8 : */ Downstream Maximum Sustained Traffic Rate:

/* TLV 71.9 : */ Maximum Traffic Burst:

/* TLV 71.27 : */ Downstream Peak Traffic Rate:

/* TLV 71.42 */ Low Latency Aggregate Service Flow Encodings

/* TLV 71.42.1 : */ Low Latency Service Flow Reference:

/* TLV 71.42.5 : */ AQM Coupling Factor:

/* TLV 71.42.6 : */ Scheduling Weight:

/* TLV 71.42.8 : */ QPLatencyThreshold:

/* TLV 71.42.9 : */ QPQueuingScoreThreshold:

/* TLV 71.42.10 : */ QPDrainRateExponent:

Corresponding parameters in the upstream TLV ranges (24.* and 70.* were also included in the new tools)

2) the G2 tool that is used to manage global data on the vCMTS supporting LLD deployments needed upgrades. First to support the Aggregate QoS profile defined for the dual queue system (mostly aligned with TLVs 70/71.40 changes and additionally to support the Service Class Name (SCN) parameter changes aligned with TLVs 24/25.40.

Internal tools have the potential to impact vast swatches of customers, so the introduction of both .cm toolsets and the G2 utility were introduced with crawl, walk, run approaches, wherein mechanisms to potentially impact minimal numbers of customers were designed into the first deployments.

12. Customer Measurements

Our user trials were conducted with written survey forms, to which volunteers were directed – as often as weekly. Our participants were not clustered in one geographic area, but spread across all of our operating divisions in order to get the broadest possible geographic distribution. Similarly, testers used a variety of both user-managed and Comcast-managed cable modems[14] in order to test across several chipset and hardware variations. Finally, testers also spanned many different service tiers (bandwidth levels). Though we initially targeted those at 1 Gbps and above, we eventually included our minimum postpaid service rates.

All testers were asked to perform varying tasks and were asked to report the results on a web-based survey form. Our tasks included performing speed tests using varying public sites and reporting the (non-low latency queue) corresponding reported latencies for upstream and downstream. They also included Apple's "networkQuality" tests[15], activated from a Mac command line (terminal). Additionally, the metrics available within Xbox, Valve (Counterstrike2/Dota2), and NVIDIA GeForce NOW were gathered after gameplay with LLD enabled and disabled.

Initial upstream tests were distributed weekly for a period, then as data collection and summarization slowed, the cadence slowed, and we knew testers were ready for the next test assignment.

After some initial analysis, we did discover some results that did not meet expectations, and were occasionally investigating each of the setup options to determine if contributed data was indeed valid. This was interesting, as some critics of L4S at the IETF expressed a concern that classic flows would be starved of bandwidth – when we had seen cases of the LL queue being starved of bandwidth instead. We determined this was a problem in how we had configured the cable modems and was corrected. New testing then confirmed that both queues functioned as expected.

However, we will observe that user testing can be labor-intensive for both the engineering team and customers themselves. The results can also vary, often due to each customer's level of technical sophistication. In addition, much as you may see with any other crowd-sourced testing, these tests are subject to self-selection bias and can be influenced by factors like cross-traffic[16].

13. Manual Testing of LL and Classic Queue Competition

Despite manual testing being labor-intensive, some such testing by the engineering team had benefits, especially when using packet generating tools such as ByteBlower[17] – which was enhanced to support low latency marking. In particular, one of the most important tests for us to perform in this way was to measure the effects on best effort flows that use the classic queue and with competing L4S flows. That is because L4S flows can be capacity-seeking and some L4S critics were worried that L4S flows could either starve classic flows of bandwidth and/or drive up their working latency.

---

[14] Comcast devices were CM4482 (XB7),CM3441 (XB7), CM4981 (XB8), and COAM devices included Arris S33v2, Arris S33 and Netgear CM1000.

[15] See https://datatracker.ietf.org/doc/draft-ietf-ippm-responsiveness/.

[16] For an exploration of the many factors that can affect such testing, see https://cacm.acm.org/research/measuring-internet-speed/.

[17] See https://www.excentis.com/products/byteblower/.

What we can actually see is the clear benefit of L4S's (coupled) congestion control with Congestion Experienced (CE) marking. L4S flow senders can manage the transmit rate and queue depth by responding to the CE marking feedback from the receiver, instead of packet drop or delay feedback. In our measurements, the L4S flows have avoided packet drops and maintained the latency very low by adjusting the transmit rate. At the same time, classic flows perform well.

14. Automated Testing

A better approach that user-initiated testing is the use of software embedded into cable modem gateways or ethernet-connected devices that run controlled tests on a scheduled basis, collecting significantly more data and higher quality data. To this end, we updated the Internet Measurement Platform (IMP)[18] test agent installed on Comcast gateways to support L4S marking. These tests had previously performed both speed (capacity) tests and working latency tests, among other measures. We added L4S marking to test working latency of the LL queue.

Over the course of thousands of tests, we observed a roughly 50% reduction in working latency via this IMP test agent.

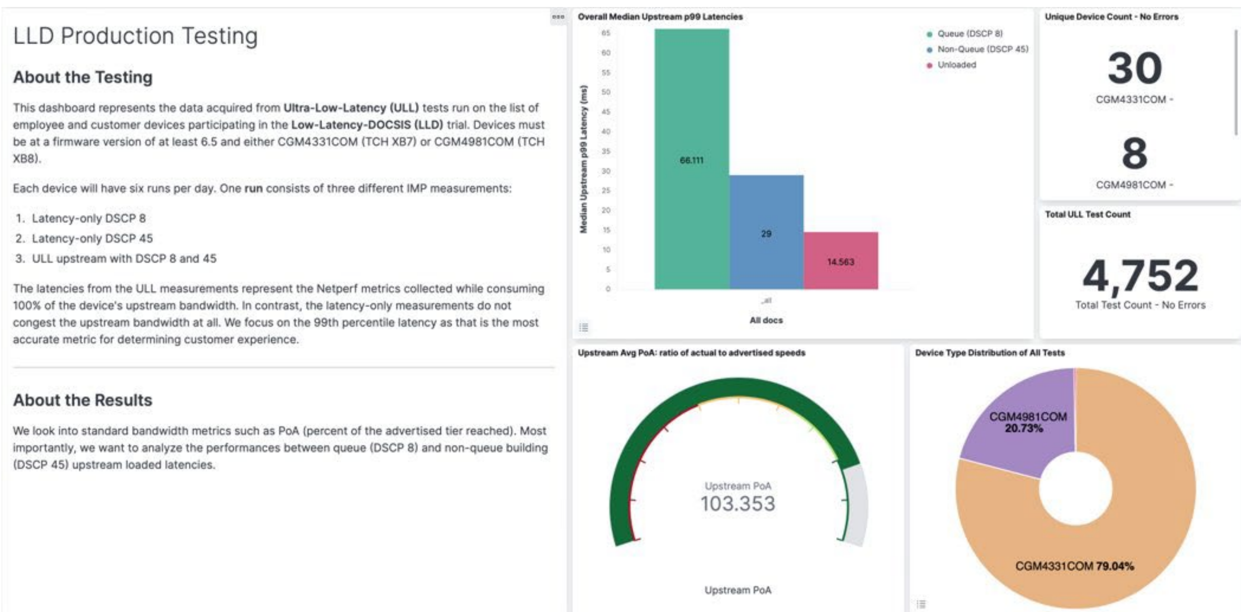

In addition, we also deployed test boxes from RIPE Atlas, SamKnows, and the University of Chicago (Netmicroscope). While none of these tests yet supported L4S or NQB marking, this enabled us to collect baseline data for the performance of the classic queue's working latency and idle latency, to confirm that it did not significantly vary from non-trial customers – an important validation. In addition, it also provided SamKnows and NetMicroscope with end users that had dual queue for research into the development of L4S or NQB tests.

While user-driven testing and their qualitative impressions have certainly been very valuable – indeed they pointed us to the starvation issue – these more controlled and voluminous tests became more useful as the trial progressed.

[18] See https://www.netforecast.com/wp-content/uploads/Comcast-Design-Audit-Report-NFR5133F.pdf.

We plan in the future, as SamKnows and NetMicroscope are ready, to test working latency of the classic queue as compared to the LL queue.

15. Measurement Via Other Network Components

In preparation for production readiness and eventual scale-up, we have also instrumented the network to collect volumetric statistics concerning L4S and NQB adoption. Given the small number of trial users and applications being tested, this volume is a small fraction of overall network volume. But as we deploy to millions of customers and major applications support low latency, this will be an interesting data point on existing traffic dashboards.

16. Measurement Via Application QoE

Since materially all traffic is encrypted, is is not possible for ISPs to accurately infer application QoE. Rather, this is better done on the application platforms themselves. Some applications, such as games, may display some QoE stats to users – such as working latency (lag). But generally all applications collect aggregated statistics on the QoE for users of their applications. These have proven valuable to the trial, both because there are aggregates of a good volume of tests (good sample size) but also because tests like those in IMP are ultimately artificial, and the only thing that really matters to customers is the QoE of their applications.

To this end, we have gotten such statistics from partners, such as NVIDIA and Valve. Both examples are provided below, and suggest even better improvements that our artificial IMP test suggested.

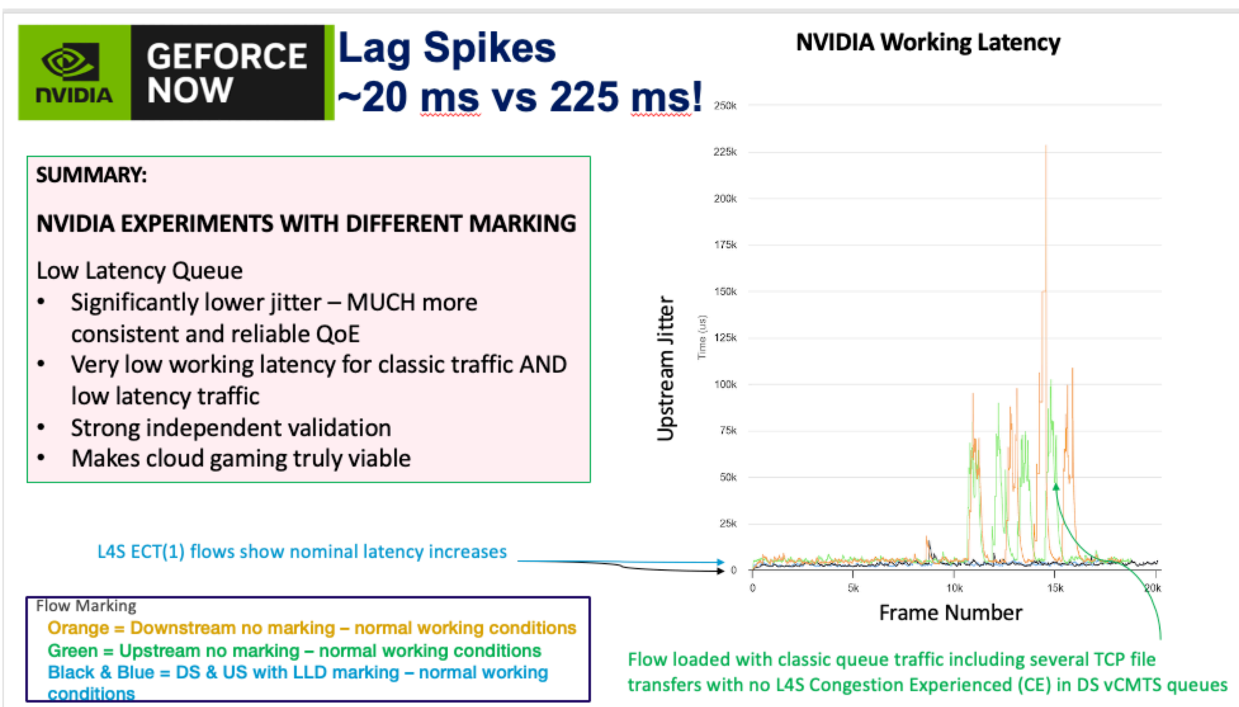

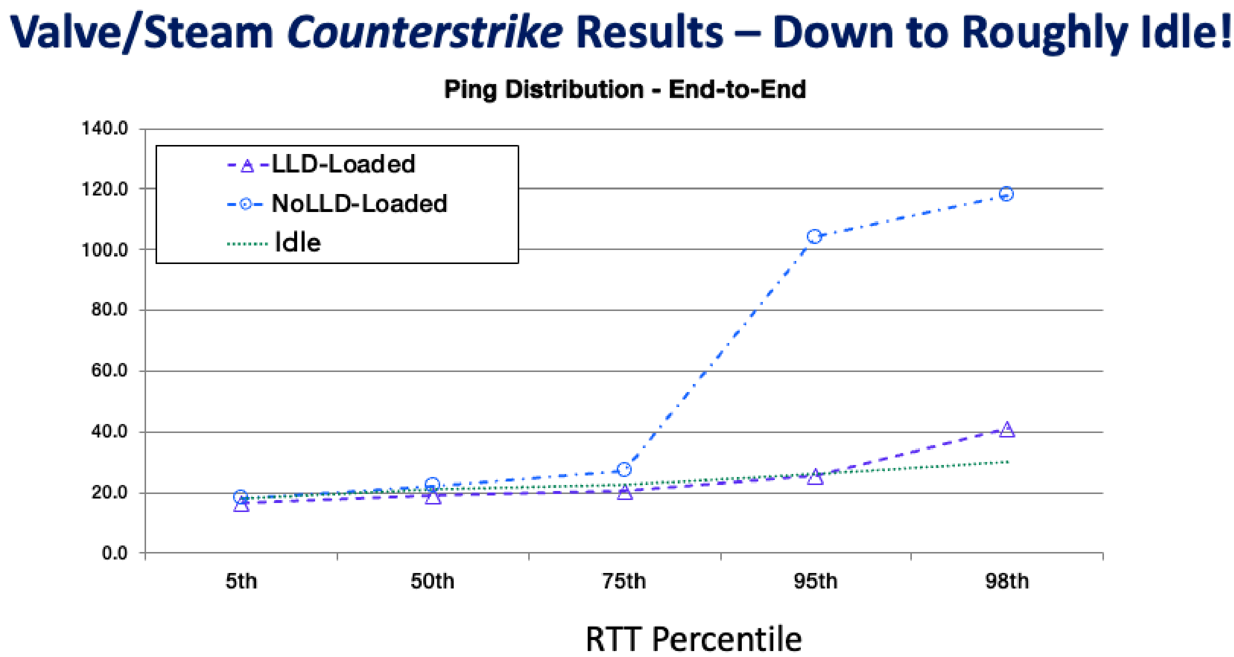

17. Net Neutrality Considerations

As we pondered deployment, we reflected on earlier net neutrality issues and current regulations and decided that the manner of deployment could bear significantly on whether or not low latency networking was considered compatible with net neutrality. These considerations are documented extensively in an IETF document[19], so they will only be summarized here:

- Only applications should mark traffic for L4S or NQB.
- All application providers should be welcome (without payment, legal agreement, etc.).
- L4S and NQB should work on both ISP-provided CPE and customer-owned CPE.

18. Conclusion

Based on an extended field trial of low latency networking, including close monitoring of classic traffic flows and collection of a range of measurement data, it is clear that (1) adding a low latency network queue does not harm classic traffic flows, (2) low latency flows perform better in the low latency queue – in some cases with working latency approaching close to idle latency. As a result, this technology is ready for broader production deployment, which we expect to do in 2024.

In addition, this network capability is only useful and valuable if many applications support L4S or NQB (which is to say it benefits from network effects). Thus, it is in an ISP's interest to make low latency network as widely available to customers as possible and to encourage more application providers to support it. Ultimately an ISP can only control their own deployment – but the scale of deployment will naturally attract interest and pull in more developers.

---

[19] See https://datatracker.ietf.org/doc/draft-livingood-low-latency-deployment/.

Abbreviations

ASF	Aggregate Service Flow
AQM	Active Queue ManagemenBE
BE	best effort
CE	congestion experienced
COAM	customer owned and managed cable modem
CMTS	Cable Modem Termination System
DOCSIS	Data-Over-Cable Service Interface Specifications
FIFO	first in, first out
HSD	high-speed data
IETF	
IMP	internet measurement platform
ISP	internet service provider
LL	low latency
MULPI	MAC and Upper Layer Protocols Interface
NQB	non-queue building
PIE	Proportional Integral Controller Enhanced
QoE	Quality of Experience
RPD	Remote PHY Device
SCN	Service Class Name
SCTE	Society of Cable Telecommunications Engineers
TLV	type-length-value
ToS	Type of Service
vCMTS	virtual CMTS
WMM	Wi-Fi Multimedia

Bibliography & References

1. To understand latency, especially the difference between working and idle latency, see https://bitag.org/documents/BITAG_latency_explained.pdf.
2. "DOCSIS® Best Practices and Guidelines, Cable Modem Buffer Control, CM-GL-Buffer-V01-110915", see https://account.cablelabs.com/server/alfresco/a4c077e2-cc55-4a42-a557-32f92e637266.
3. "Active Queue Management (AQM) Based on Proportional Integral Controller Enhanced (PIE) for Data-Over-Cable Service Interface Specifications (DOCSIS) Cable Modems", RFC 8034, see https://www.rfc-editor.org/rfc/rfc8034.html.
4. "Improving Latency with Active Queue Management (AQM) During COVID-19", https://arxiv.org/abs/2107.13968.
5. See RFC 9330 at https://www.rfc-editor.org/rfc/rfc9330 and RFC 9331 at https://www.rfc-editor.org/rfc/rfc9331.html.
6. See https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/.
7. Typically this is the user's in-home WiFi LAN, but in some cases can be an ISP-managed public WiFi network, such as Xfinity WiFi.
8. See https://www.xfinity.com/learn/internet-service/supersonic.

9. See https://www.xfinity.com/support/articles/wireless-gateway-enable-disable-bridge-mode.
10. "Measuring the Prevalence of WiFi Bottlenecks in Home Access Networks", see https://arxiv.org/abs/2311.05499.
11. See https://www.wi-fi.org/knowledge-center/faq/what-is-wmm.
12. See https://corporate.comcast.com/stories/comcast-kicks-off-industrys-first-low-latency-docsis-field-trials
13. Technically AQP stands for Aggregate Service Flow (ASF) Quality of Service (QoS) Profile.
14. Comcast devices were CM4482 (XB7),CM3441 (XB7), CM4981 (XB8), and COAM devices included Arris S33v2, Arris S33 and Netgear CM1000.
15. See https://datatracker.ietf.org/doc/draft-ietf-ippm-responsiveness/.
16. For an exploration of the many factors that can affect such testing, see https://cacm.acm.org/research/measuring-internet-speed/.
17. See https://www.excentis.com/products/byteblower/.
18. See https://www.netforecast.com/wp-content/uploads/Comcast-Design-Audit-Report-NFR5133F.pdf.
19. See https://datatracker.ietf.org/doc/draft-livingood-low-latency-deployment/.