Latency Outcomes Across Access Network Architectures

A technical paper prepared for presentation at SCTE TechExpo24

Eric Tijerina
Director of Network Architecture & Access Engineering
Midco
eric.tijerina@midco.com

Jeff Shields
Principal Engineer I - Architecture
Midco
jeffrey.shields@midco.com

Matt Zerfas
Network Intelligence Analyst Lead
Midco
matt.zerfas@midco.com

Karthik Sundaresan
Distinguished Technologist and Director of HFC Solutions
CableLabs
k.sundaresan@cablelabs

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

Passive Optical Networks (PON) and Hybrid Fiber Coax (HFC) networks are commonly deployed for delivering Internet data, voice, and video services to subscribers. Operators in the cable industry would like to understand the real-world latency numbers across access technologies to ultimately affect the customer experience. This paper will present latency measurements taken across production networks. It will analyze the impact of enabling LLD (Low-Latency DOCSIS®) architecture across different HFC architectures commonly used in different MSO environments. This includes legacy/traditional DOCSIS networks as well as DAA (specifically RPHY) nodes.

The data gathered will be primarily focused on measured latency benefits when enabling AQM (Active Queue Management) and ASF (Aggregate Service Flows) - features of LLD in DOCSIS 3.1. The impact of enabling PGS (Proactive Grant Service), another key component of LLD, will be estimated to the extent possible. The results of our testing on the production cable plant will be compared against real world PON deployments utilizing the same tools and metrics. The measurements will reasonably assess the impact of different IP DSCP and ECN values and how those packets are treated in the network to improve customer quality of experience (QoE). We will also share some results on the latency measurements on our core network. This paper will give the cable operator community a good understanding of network latencies as well as the benefits of enabling technologies like LLD.

1.1. Low Latency DOCSIS Technology Background

At a high level, the low-latency DOCSIS architecture consists of a dual-queue approach, a “low latency” queue for non-queue-building traffic and a “classic” queue for queue-building traffic, where both queues share a single pool of bandwidth. These queues are implemented using service flows that are simply optimized for two different types of traffic behavior [1]. This mechanism can provide consistent low latency for low-data-rate non-queue-building applications, and it also supports a new non-queue-building version of TCP congestion control, “L4S”, that allows them to send data at the full link rate while maintaining statistically low queuing delay.

To take advantage of LLD, developers of traditional low-data-rate non-queue-building applications (e.g. traditional online games and real-time communication apps) will need to mark those packets with a special “NQB” value in the DSCP field of the IP header, and operators will need to ensure that this value is carried across their networks. Developers of high-data-rate latency-sensitive applications (e.g. cloud games) will need to adopt L4S mechanisms, which enables applications to ramp up its sending rate to the full network capacity, yet quickly back off when the network signals impending congestion [1].

DOCSIS scheduling services are designed to customize the behavior of the request-grant process for particular traffic types. LLD introduces a new scheduling service called Proactive Grant Service (PGS), which can eliminate the request-grant loop entirely, [1]. In PGS, a CMTS proactively schedules a stream of grants to a Service Flow at a rate that is intended to match or exceed the instantaneous demand. In doing so, a majority of packets carried by the Service Flow can be transmitted without being delayed by the Request-Grant process.

1.2. Round-Trip Latency

Round-trip time (or latency) is the time taken for a packet of data to travel from the sender to the receiver, across one or multiple hops, plus the total length of time it takes for the receiver to send the packet back to the sender, through one or multiple hops. [4]. All the measurements in this paper are RTTs across the network segment of interest. All of our testing was done using IP packets marked with ToS values set to 0 and 1.

2. Design of Latency Testing

The testing was focused on the area of the network that is commonly referred to as the Access Network and the core network by service providers.

The Access Network consists of the network segment from the Cable Modem Termination System (CMTS) or Optical Line Terminal (OLT) to the customer's gateway, either a Cable Modem (CM) or Optical Networking Unit (ONU).

To facilitate measurement of latency on the access portion of the network, approximately 20 network probes or reflectors were installed in employee homes. These probes/reflectors were directly connected to the gateway via wired 1Gbps ethernet connections. The measurement agent or test server, placed nearest the CMTS, was the point to which round trip time/latency was measured from the reflectors.

To facilitate measurement of latency on the core portion of the network, network probes or reflectors were installed at locations close to the core routers. These probes/reflectors were also connected via wired 1Gbps ethernet connections. The same measurement agent or test server, placed near the CMTS, was used to measure latency from these reflectors through the core network.

2.1. Network Segments

The network segments under test were the Core and Access Networks. In-home latency can be highly variable, is dependent on several factors, and was not included in the testing described herein. Figure 1 (below) shows a high-level view of each network segment.

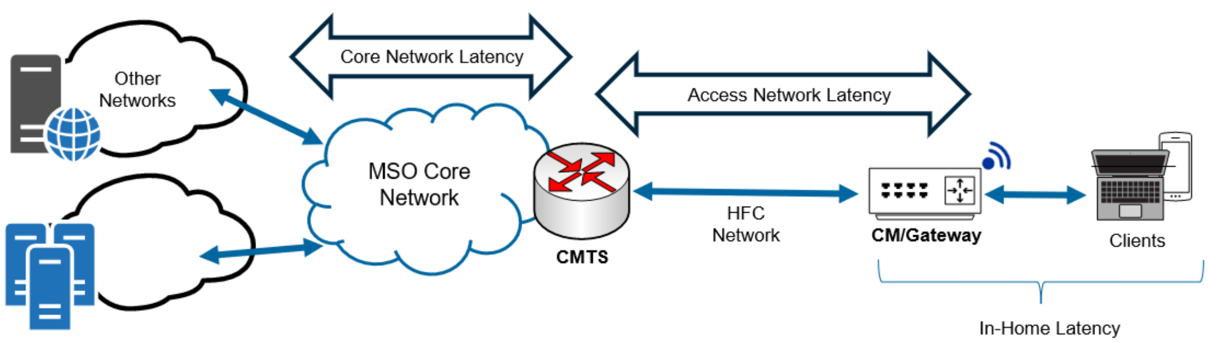

Figure 1 - Network Segments Under Test

Core Network: A core network or service provider backbone is inherently less latent largely due to the 100, 400, and 800GbE links connecting equipment. The core/backbone is unlikely to be a source of increased latency under normal operating conditions where network utilization is effectively managed through upgrades to meet growing throughput and peak demands.

Access Network: There are often varied technologies that may exist in service provider Access Networks that differentiate the network enough to result in different latency outcomes for subscribers. The production networks utilized for this testing included the following variations:

- **PON** (EPON) – a FTTH (fiber-to-the-home) delivery method was used primarily as a comparative measure to the latency outcomes of the HFC networks described in the next two

bullets. This also served as a baseline to ensure the data obtained during testing was valid and closely mirrored known latency measurements of PON delivery methods.
- **Legacy HFC** – the traditional HFC network in this environment consisted of legacy integrated CMTS hardware (not virtualized or server-based) and HFC plant that did not contain DAA components. This comprises the largest segment of access networks that many operators use to serve the majority of subscribers connected via the HFC network.
- **DAA** (RPHY) – the DAA network utilizes an RPD (Remote-PHY Device) located in the field, and connected via digital fiber originating at the CMTS. This is the Access Network architecture that many cable MSOs are actively migrating to. This may or may not be facilitated via a traditional "big iron" CMTS or by virtualized server-based CMTS platforms.

2.2. Latency Measurement Components

The components utilized to facilitate testing are listed:

Session Reflectors – ODROID N2+, CPE performing the RTT latency testing.
The Odroid is a small Unix based device, (similar to a Raspberry Pi), used to execute the latency tests for each customer (served by an ONU or CM). The latency tests are run at intervals that vary depending on desired configuration. Much of the latency data presented in this paper was collected from 10-minute test cycles running every ten minutes. This configuration allowed for full visibility of the latency on both the low-latency and classic flows at all times. Test packets sent to and reflected back from these devices were either tagged with a ToS value of 0, or a ToS value of 1. The test packets marked as ToS 1 will traverse the low-latency queue.

Measurement Agent – server test point for session reflectors.
This server is the test agent from which test traffic is sent downstream to the reflectors and reflected back on the upstream direction. This server was placed nearest to the CMTS and OLT to ensure that RTT latency being measured was primarily that of the Access Network and included as few additional network devices and hops as possible. This would typically be placed in a hub or headend environment and ideally uplinked to the same network device that provides WAN connectivity to the CMTS or OLT.

Controller/Collector – server for managing the test schedules and data collection/storage of test results. This server could be placed anywhere in the network such that it was reachable by the measurement agent server. The results of all tests, from all reflectors, were aggregated on this server for reporting purposes. Additionally, the control messages initializing the set of tests scheduled for the individual reflectors are created on the controller and sent to the measurement agent.

Software - The latency measurement software utilized on the components described in the previous section was written by CableLabs® and utilizes STAMP (Simple Two-Way Active Measurement Protocol, [RFC 8762]) and LMAP (Large-Scale Measurement of Broadband Performance, [RFC 7594]). The references at the end of this document link to more detailed documentation that cover the intended application and implementation.

The image below depicts where in the network these components are connected and the traffic flows between them.

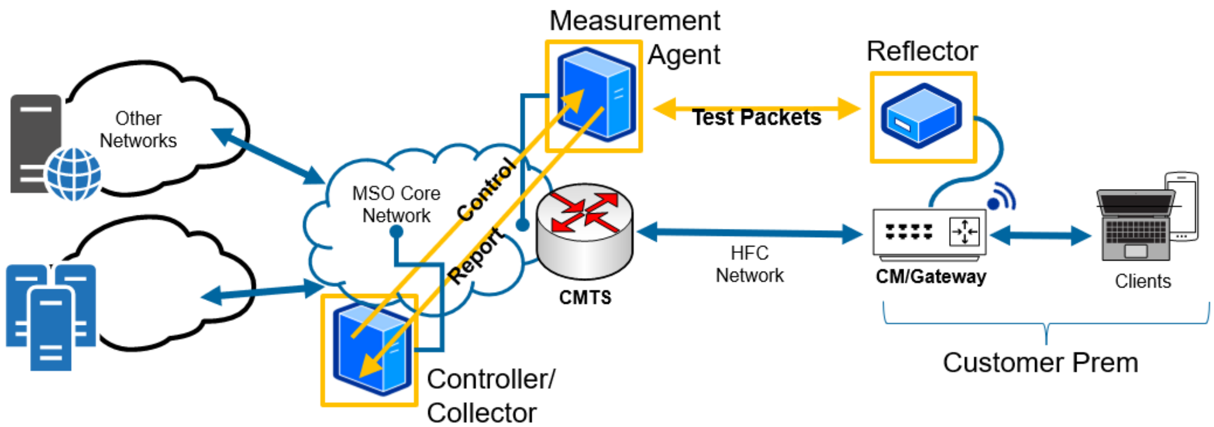

Figure 2 - Network Diagram & Test Components

2.3. Testing Overview

The structured test plan serves to outline the methods used for measurement and the information to be gleaned from the testing. The initial goal was to get a baseline latency measurement for the network segments under test. This allowed for quantitative analysis across the different network segments, technologies, and architectures. The Core Network and EPON Networks did not undergo any changes throughout the test period. As such, the measurements of both of these networks were consistent throughout this measurement effort. These networks were measured primarily for the purposes of drawing conclusions about other architectures and the impact of changes such as the implementation of LLD on HFC networks.

2.4. Service Flow Configurations

The DOCSIS service flow configuration illustrated below is a common configuration in MSO networks and serves as the starting point for the testing from which all baseline measurements were taken.

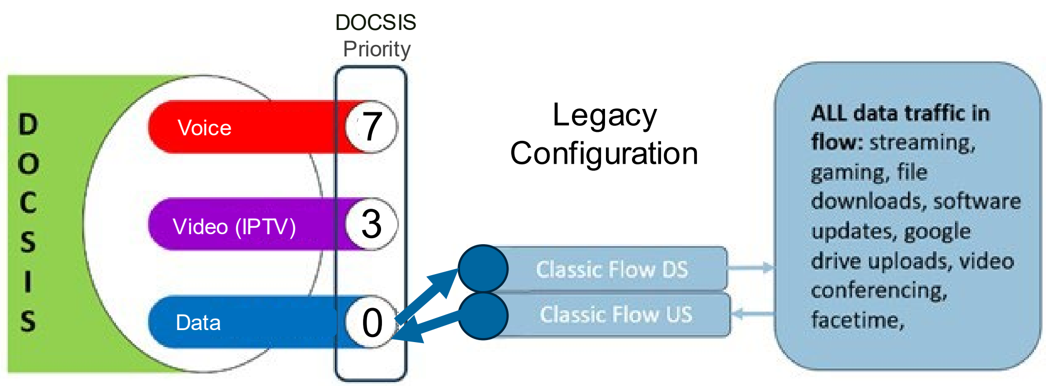

Figure 3 - DOCSIS Service Flow - Legacy Configuration

Once the baseline latencies were established, LLD configuration was enabled across several CMTS and cable modems (of varying makes/models) serving the employee homes equipped with the ODROIDs/session reflectors. This allowed an opportunity to measure the expected differences in latency between reflector traffic with a ToS value of 1, which would traverse the low-latency NQB flow, and reflector traffic with a ToS value of 0 , which would traverse the classic flow along with the rest of the queue-building traffic. The added benefit in this test scenario was the regular Internet traffic and usage patterns of typical single-family homes which can be difficult to simulate in lab environments.

Figure 4 illustrates the Aggregate Service-Flows (ASF) in each direction (upstream and downstream) and lists the traffic types that may utilize respective flows. Note that the application developers must mark the traffic appropriately to ensure it utilizes the low-latency flow only if appropriate. In addition, traffic that is marked to pass through the low-latency flow should behave within expected characteristics by sending traffic at a consistent or fairly even rate.

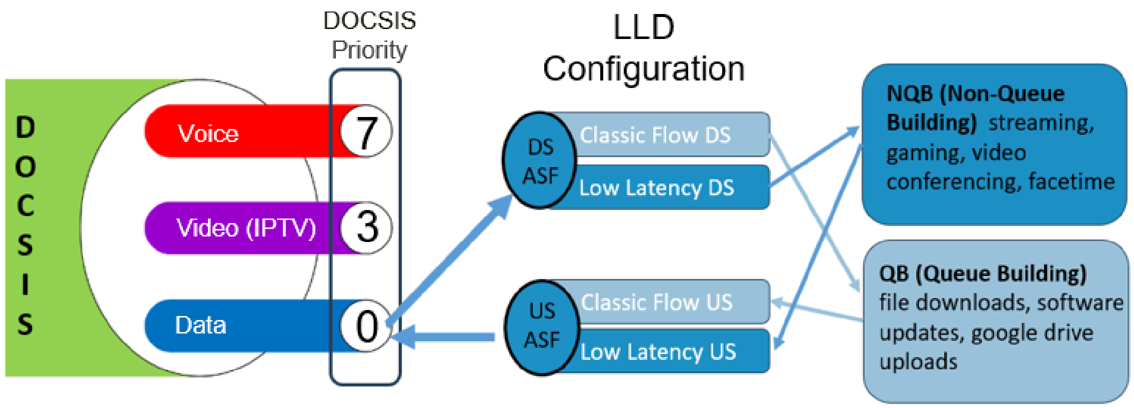

Figure 4 - DOCSIS Service Flow - LLD Configuration

2.5. PGS and DPS

Due to current CMTS software limitations, PGS (Proactive Grant Service) was not yet available for use in our testing. This feature of Low-Latency DOCSIS can eliminate traditional request and grant cycles.

An alternative (and proprietary) software feature in the CMTS is DOCSIS Predictive Scheduler(DPS). DPS utilizes unallocated MAPs to issue predictive grants to active service flows. PGS allows for scheduling a stream of grants that matches or exceeds demand for a particular service flow and thus traffic can avoid the request and grant cycles. PGS and DPS function differently, and both serve to reduce latency. DPS was explored as part of this testing and trial.

The ability for DPS to reduce latency is largely dependent on upstream utilization. With greater utilization, there are less available unallocated MAPs for which to issue grants. However, a lightly loaded service group (meaning a lesser number of modems) would experience greater benefit. DPS allows for a reduction of latency for all types of traffic, not just network traffic that meets the criteria to be classified as low-latency or NQB. Due to the fundamental differences in which DPS and PGS function, in theory it may be supported to enable/configure both features, but this was not able to be tested.

2.6. Software Platform for Measurement

The data presented herein was obtained utilizing a CableLabs® developed, STAMP based, tool for latency measurement. Measurements across multiple devices are presented with images from Tableau, though the measurements themselves were obtained via the CableLabs latency measurement suite. [4]

Additionally, some test data is included from SamKnows®. SamKnows agents (built into the router firmware) can run test schedules that are used for measuring latency in a similar architecture where the reflector is located at the customer premise and the measurement agent is nearest the CMTS.

Having two different tools to measure latency allowed for investigation of any anomalies and/or confirmation of results across measurement platforms.

2.7. Summary of Testing Sequence

The high level test plan was mainly built around collecting initial baseline latency measurements from the access network and the core network. The access network measurements includes the PON network as well as an HFC based network. On the HFC network, we had measurements against an integrated CMTS, which we name as HFC legacy in this paper. We also performed latency measurements against DAA architectures using RemotePHY technologies, i.e a CCAP-Core and an RPD out in the field. The goal was to perform a comparative analysis. The outline of the testing is as follows:

- Phase 1: Collect Latency Baselines
    - Core Network
    - Access network
        - PON (EPON)
        - HFC
            - Legacy
            - DAA (RPHY)
- Phase 2: Implement Low-Latency Changes
        - DPS (DOCSIS Predictive Scheduler)
        - LLD (AQM, ASF)
        - DPS & LLD

The individual latency tests themselves are typically 5 or 10 minute long tests in which the measurement agent sends 100 byte packets, at a rate of 20 packets per second to the session reflector. This test is repeated every 10 or 15 minutes (configurable) over many days/weeks. The session reflector timestamps these incoming packets and sends them back to the measurement agent which computes the round trip times. The packets are formatted as defined in the IETF RFC 8762 Simple Two-Way Active Measurement Protocol (STAMP) [6].

3. Latency Measurement - Baselines

Baseline measurements for the Core Network and variations of Access Network are presented in this section. Any pertinent detail regarding the measurements or conditions in which the measurements were taken are also included. Note, for the histogram figures, the start of the x-axis is not 0 but the lowest RTT measurement seen for that test interval.

3.1. Core Network

The testing used to collect the core network measurement is nearly identical to measurement of the Access Network, with the only exception being that the session-reflector (Odroid) was placed nearest a BFG (Border Facing Gateway) rather than at an employee home. Another way to think about the placement of this reflector is that it was placed at the opposite edge of the service provider network, closer to where a provider is exchanging EBGP (External Border Gateway Protocol) routes with other operators and/or networks. This allowed for the test traffic, and thus the measurement of latency, to travel the full network path from interconnects through the (operators) Core Network to the Measurement Agent. All traffic through the Core is treated equally in this environment such that the ToS values were irrelevant (i.e. the same treatment) for measurement on this network segment.

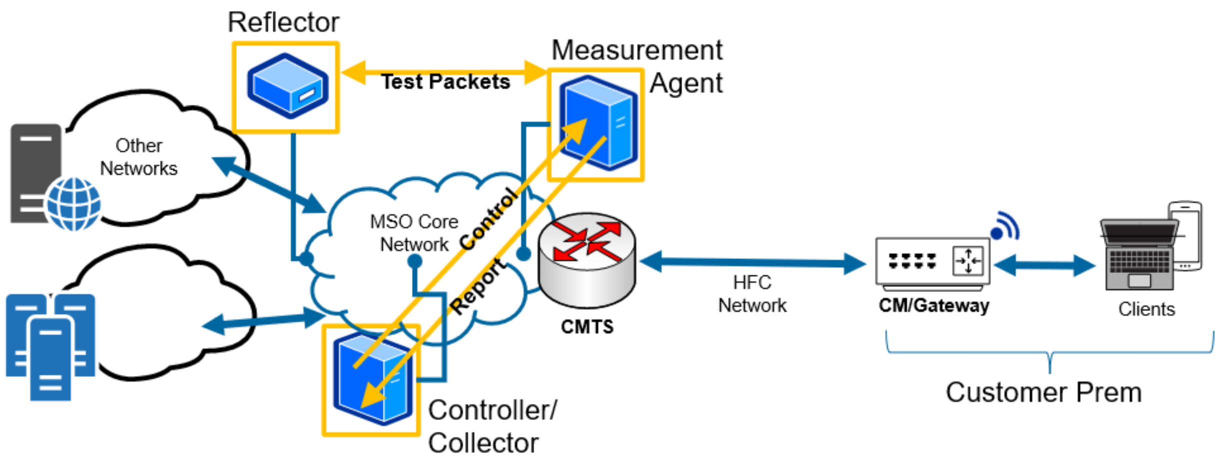

Figure 5 - Network Diagram with Core Latency Reflector

The histogram measuring Core Network Latency shows sub 2ms of latency for most traffic over a 5-minute period.

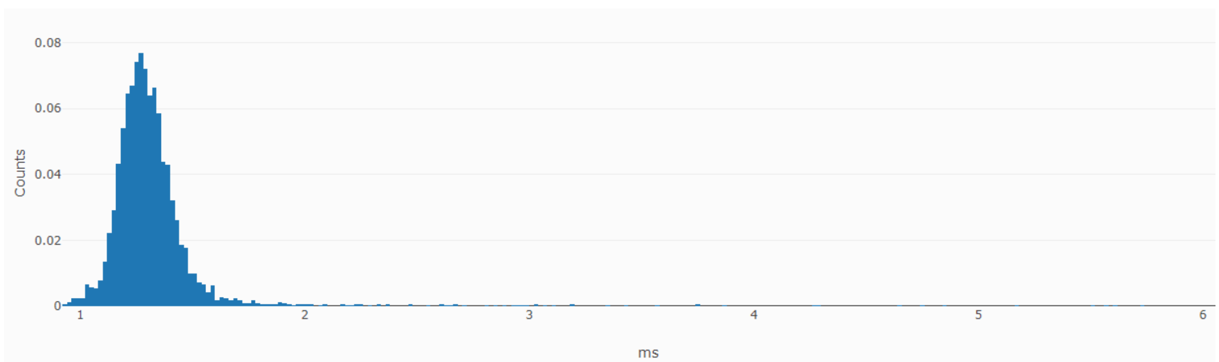

Figure 6 - Core Latency - Histogram

CCDF indicates that the 99th percentile for all traffic tested in this 5-minute interval was below 2.065697.

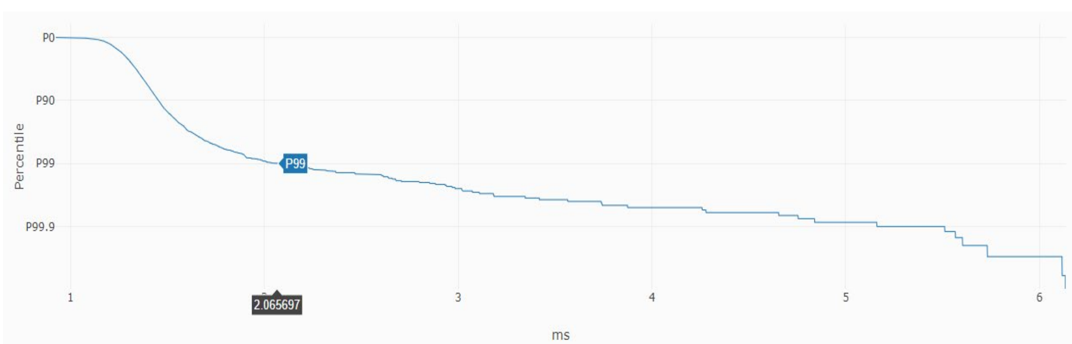

Figure 7 - Core Latency - CCDF

3.2. Access Network

The PON and HFC baseline latency measurements are presented below. Note that despite running two specific tests with different ToS values, only one is presented for each type of network. During baseline testing, without LLD enabled, all traffic takes the classic service flow, as such the results have no significant difference.

3.2.1. PON (EPON)

All of the testing on the PON network, was done on our fiber Internet deployments using 10G EPON technology. The session reflectors were deployed in employee homes which were served by fiber based services and was placed behind the optical network unit (ONU) in the home, connected via wired ethernet connections. The following graph shows the time series of latency over 10 minutes.

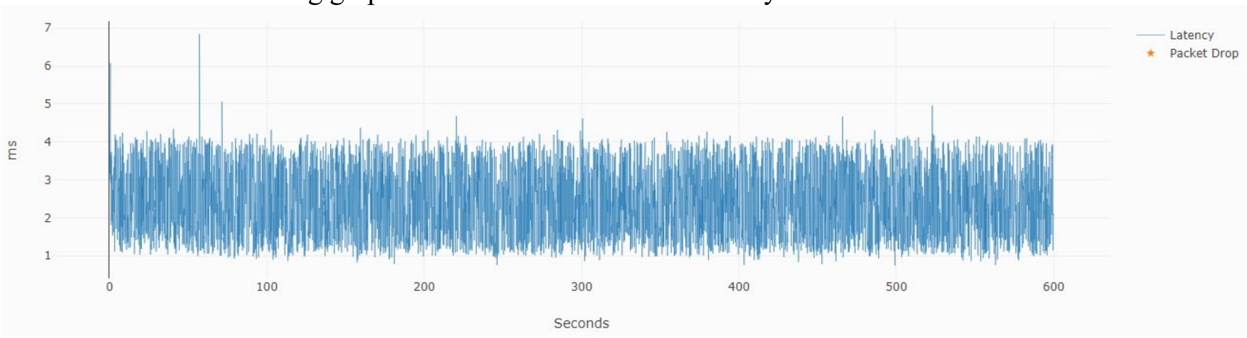

Figure 8 - PON Latency - Time Series

The following graph shows the RTT data above plotted as a histogram. (Note: the start of the x-axis is not 0 but the lowest RTT measurement seen at about 0.8 ms).

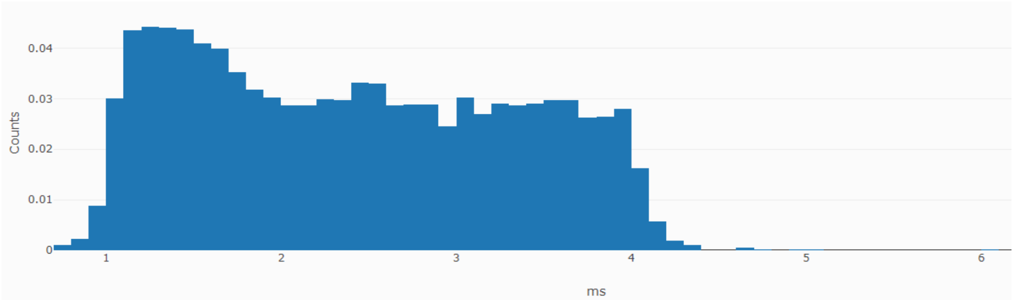

Figure 9 - PON Latency - Histogram

The CCDF indicates that the 99.9 percentile for all traffic tested in this 5-minute interval was below 4.674661ms, while the 100th percentile (maximum latency) was still under 7ms.

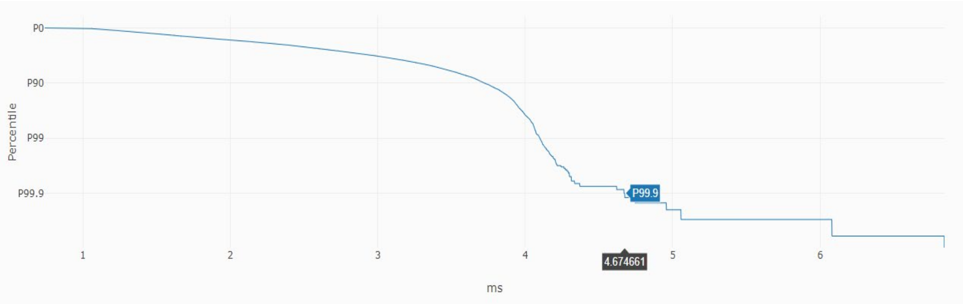

Figure 10 - PON Latency - CCDF

3.2.2. HFC Legacy

In this section we describe the measurements from the HFC network. These measurements were done on employee homes which had Internet services delivered over the DOCSIS 3.1 technology using integrated CMTSes. The session reflectors were deployed in employee homes which were served by HFC network and were placed behind the Cable Modem (CM) in the home, connected via wired ethernet connections.

The following graph shows the time series of latency over 10 minutes.

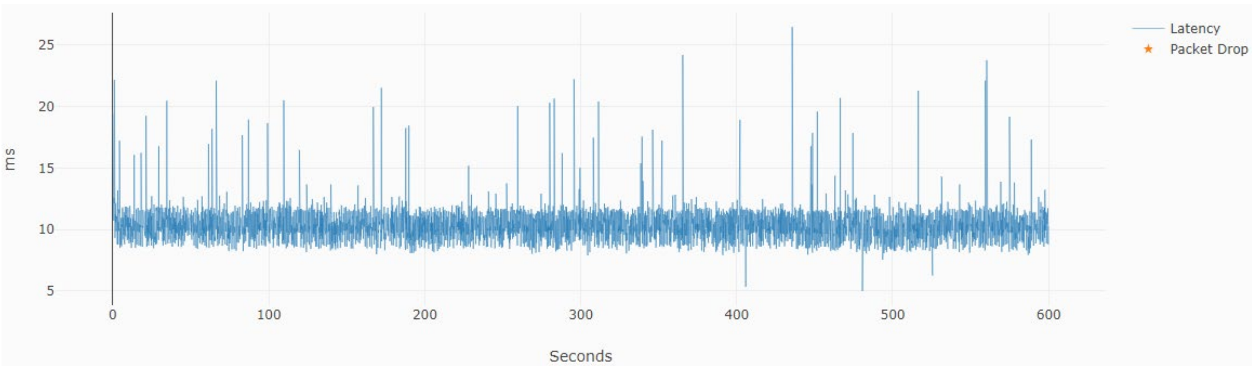

Figure 11 - HFC Legacy Latency - Time Series

The following graph shows the RTT data above plotted as a histogram. (Note: the start of the x-axis is not 0 but the lowest RTT measurement seen at about 5 ms).

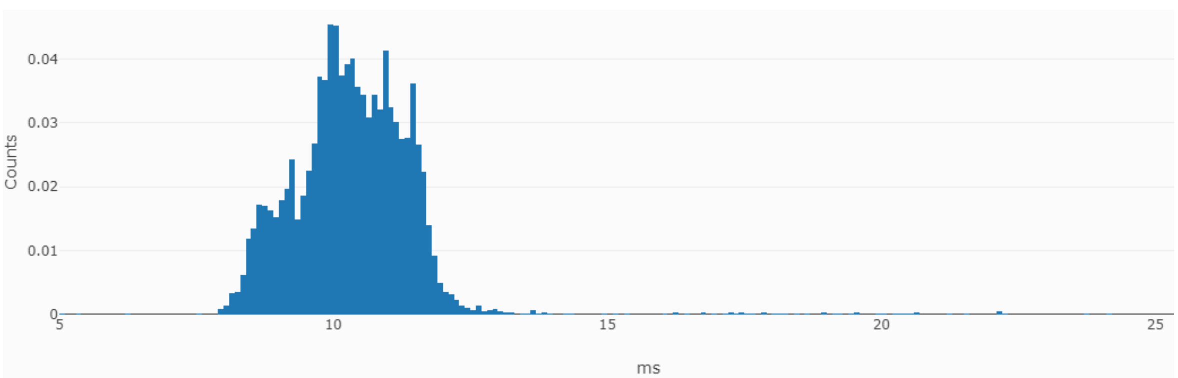

Figure 12 - HFC Legacy Latency – Histogram

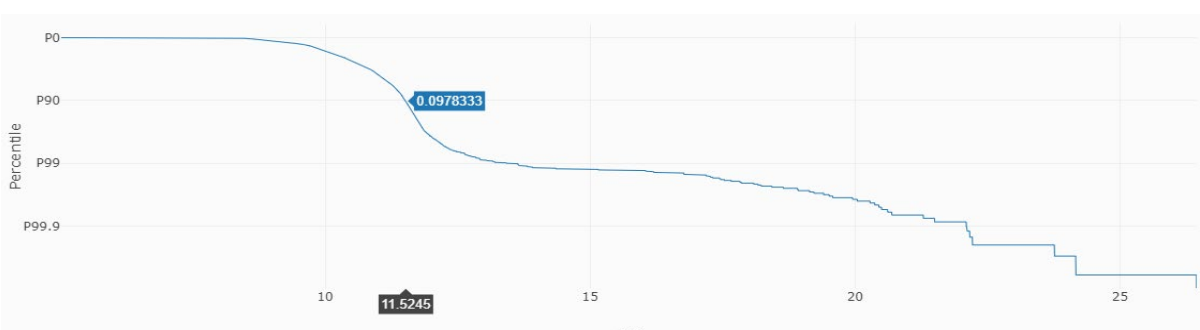

Figure 13 - HFC Legacy Latency - CCDF

The test period presented above is best case scenario (idle latency) for our deployed DOCSIS network not enabled for LLD. There appear to be no network events that drive latency up during this 10-minute test period.

When testing for latency under load, e.g. if a speed test or large file download were to occur during testing, the latency measurement spikes on all traffic for the duration of that activity. Figure 14 captures latency (of the test traffic) during a similar event, a console game update/download, starting at about 50 seconds.

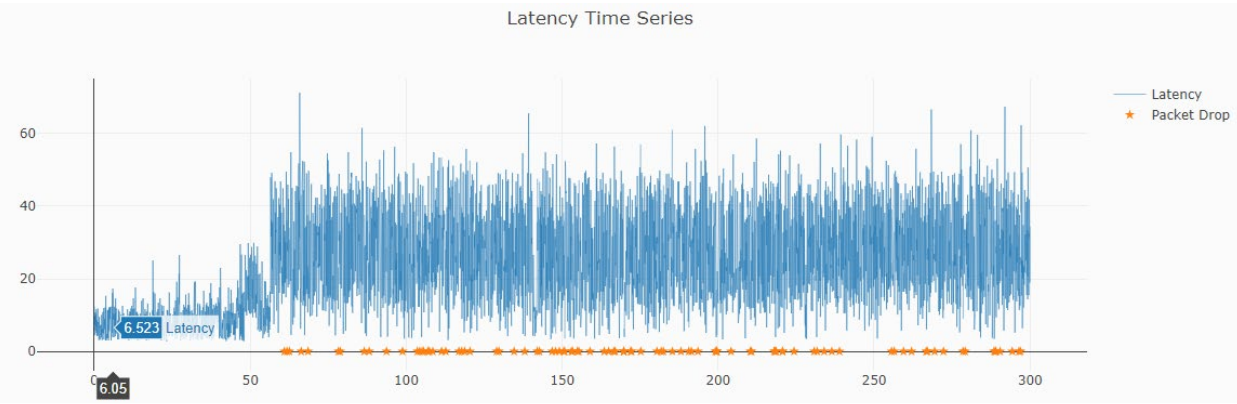

Figure 14 - HFC Legacy -– Latency During a Game Update

The packet drops indicated during the download period are likely attributed to the inherent and inefficient functions of legacy TCP design. Essentially, it is maxing out to the point of congestion, dropping packets, backing off, and repeating the process. This creates the sawtooth pattern that is well-known and well documented behavior of TCP. A transport protocol like TCP employs a “congestion control” algorithm (e.g. Reno, Cubic,) to adjust to the available bandwidth at the bottleneck link (i.e., the DOCSIS link). In these congestion control algorithms, the sender ramps up the sending rate until it’s sending data faster than the bottleneck link can support. Packets then start queuing in the buffer at the entrance to the link, i.e., the CM or CMTS. This queue of packets grows quickly until the queuing device has to discard some newly arriving packets, which triggers the sender to pause in order to allow the buffer to drain somewhat before it resumes sending. This process repeats over and over again causing increased latency and packet loss for all of the traffic that share the buffer.

3.2.3. HFC DAA (RPHY)

In this section we describe the measurements from the HFC network where the measurements were done on employee homes which had Internet services delivered over the DOCSIS 3.1 technology using distributed CMTS architectures. This typically includes a converged interconnect network, typically a layer 2 digital fiber link between the CCAP-core in the head end and the RPD at the node location.

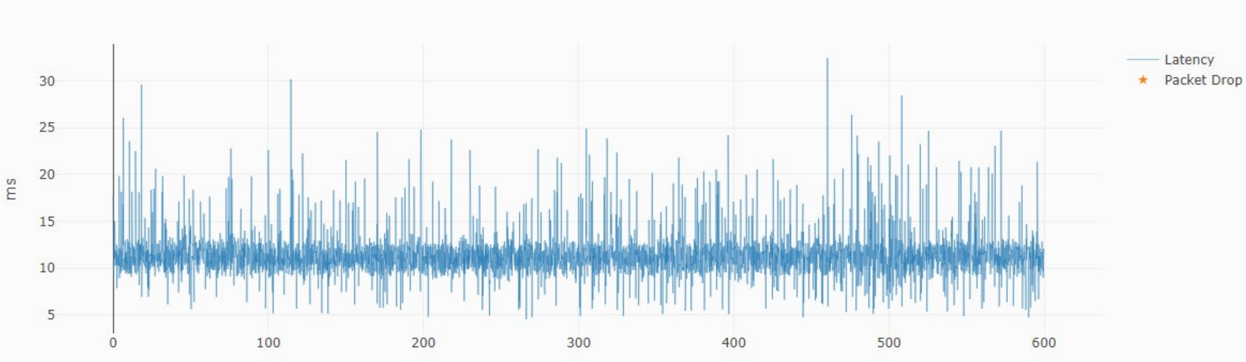

Figure 15 - HFC DAA (RPHY) - Latency Time Series

The graph above shows the time series of latency in the RPHY deployment over 10 minutes. The following graph shows the RTT data above plotted as a histogram. (Note: the start of the x-axis is not 0 but the lowest RTT measurement seen at about 5 ms).

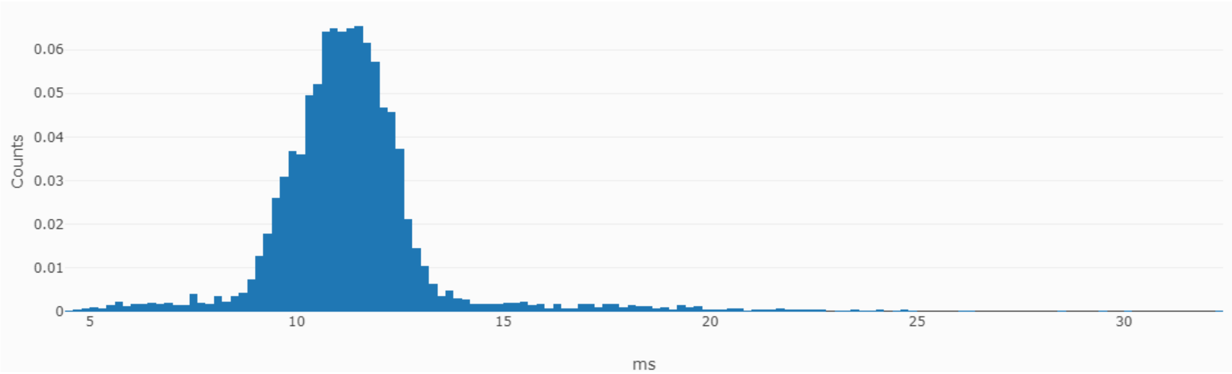

Figure 16 - HFC DAA (RPHY) - Latency Histogram

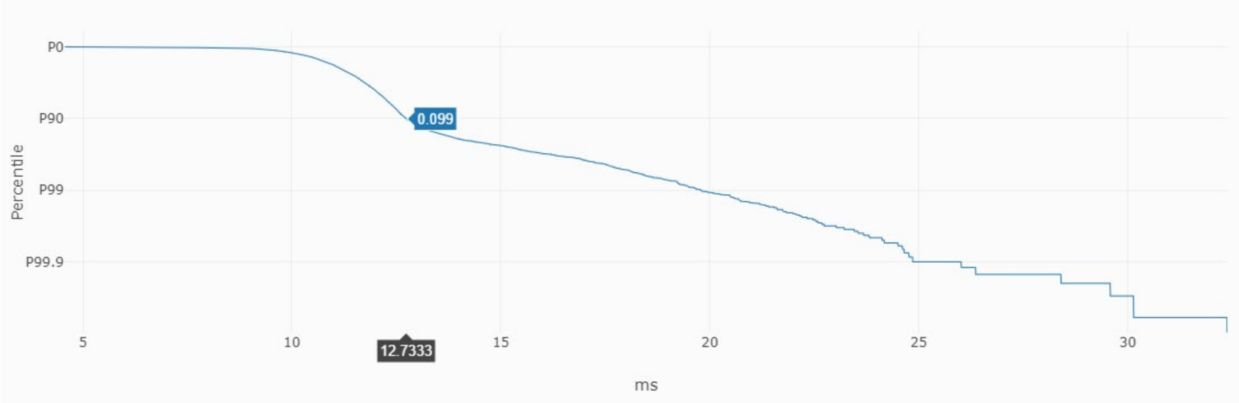

Figure 17 - HFC DAA (RPHY) - Latency CCDF

The baseline latency measurements across HFC Legacy and RPHY are similar, but there is a measurable increase of roughly 1-2ms of average latency. This increase was consistent when reviewing multiple test results, across multiple reflectors , for the DAA (RPHY) versus the legacy HFC. Though this was not investigated further, increased latency could be introduced by the CIN (Converged Interconnect Network) used in DAA and the additional hops packets must take through the leaf and spine architecture. Another theory was that OFDMA (versus SC-QAMs) upstreams may be causing a minor uptick in latency.

4. Impact of Enabling LLD on HFC Access networks

Once we completed the latency baseline measurements, we could now enable latency features on the HFC network and study their impact on the latency. The change actions (new configurations) to reduce latency were only taken on the HFC Access Network Architectures. PON was excluded from any effort to further reduce measured latencies, as was the Core Network. The main change actions that we chose to study were the following:

- Enabling DPS (DOCSIS Predictive Scheduler) (This is a proprietary feature)
- Enabling LLD (ASF, AQM Only)
- Enabling DPS & LLD

4.1. Enabling DPS

Given the lack of PGS support, we chose to enable DPS, which is a nonstandard way of trying to approximate the standardized DOCSIS PGS functionality. The impact of enabling DPS was immediate and positively impacted all traffic types as indicated by significant drops in average latency on ToS0 and ToS1 packets. The results shown in Figure 18 were taken from the measurement agent, like all individual unit tests shared previously, but graphed in Tableau for better presentation. The architecture and modems per SG are shown for each. The average decrease in latency across all testers was ~5.5ms and this was true across legacy (non-R-PHY) architectures as well.

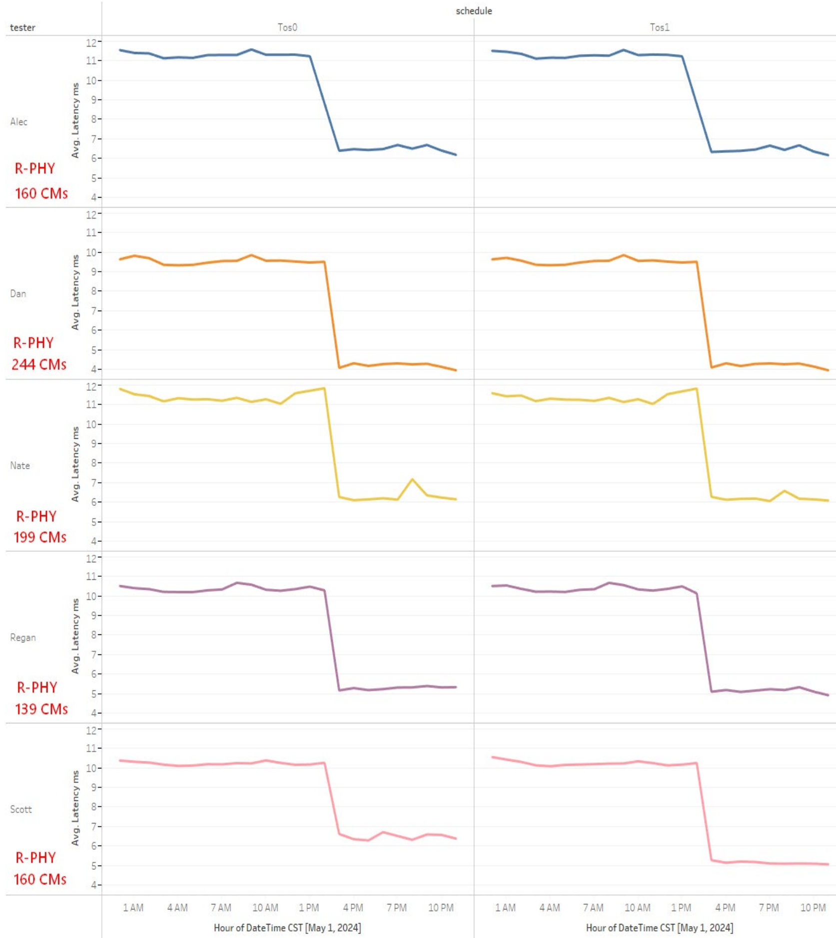

Figure 18 - Average Latency Decrease with DPS Enabled

4.2. Enabling LLD

For the next step in the testing process, we enabled Low Latency DOCSIS as supported on our CMTS platform.(Please see section 4.3.1 on our LLD service configuration parameters). Enabling LLD and then running the latency tests show us that the latency across both the LLD service flow and the classic service flow are largely the same when there are no network events that would cause increased latency. This is expected behavior, as there should not be differences in latency across the two different traffic types when there is no load. In the CCDF shown in the figure below, there are actually two lines shown, ToS0 and ToS1. However, because the results are so similar, they are nearly indiscernible. Even when zoomed in to show roughly 2ms of latency on the X-axis.

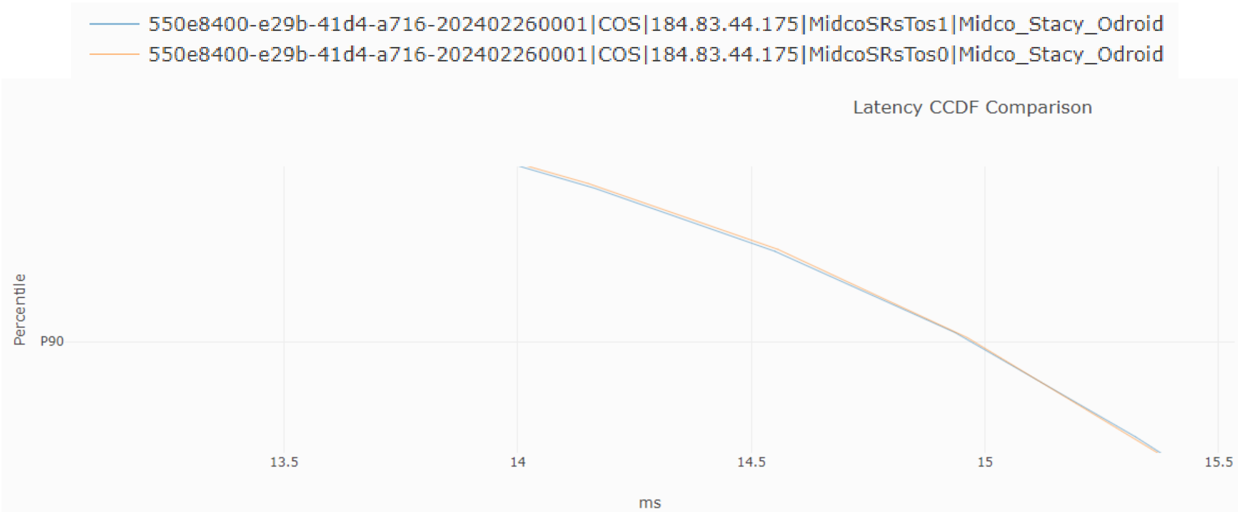

Figure 19 - Latency without Load ToS0 and ToS1

4.2.1. HFC Legacy

As shown in Figure 19, in the absence of queue building traffic the idle latency measurements of both the LLD and classic flows are very much the same. To best illustrate the outcomes of enabling LLD across these architectures, we ran “latency under load” tests. This was done by using speed tests that were run to drive up latency on the classic flow, while observing that the test traffic traversing low-latency flows held steady.

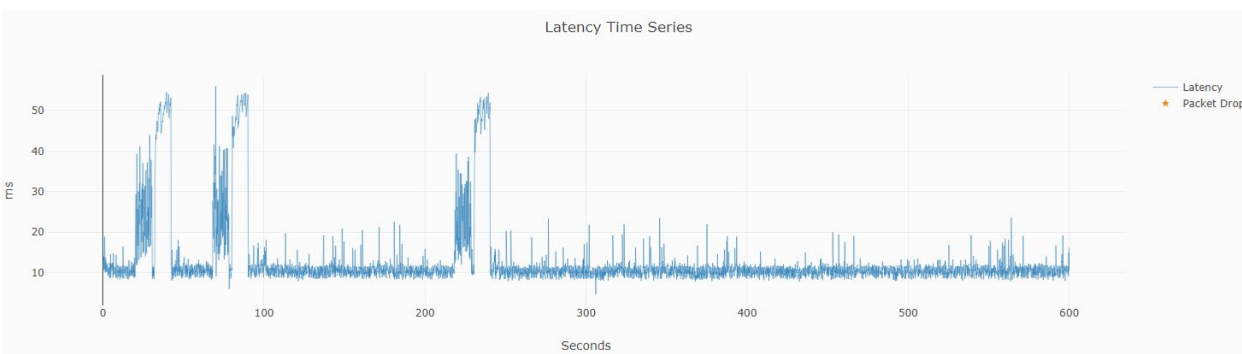

Figure 20 - HFC Legacy – LLD Enabled, Classic Flow during Speed Tests

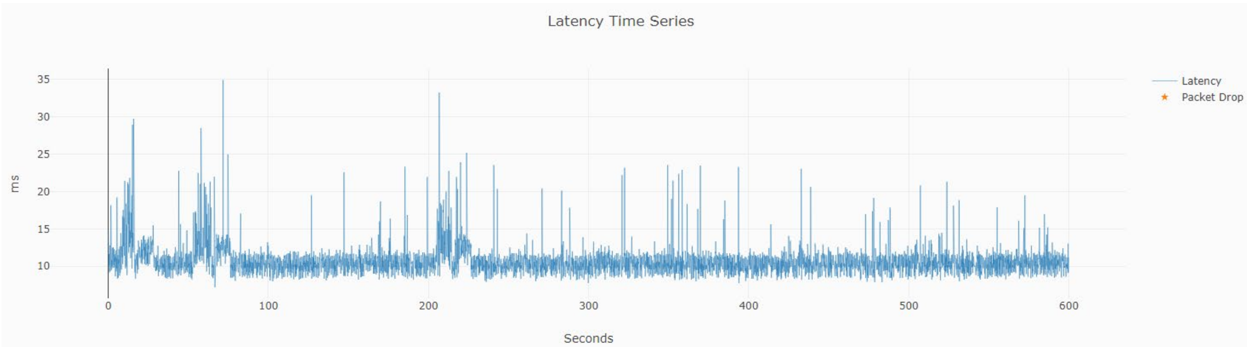

Figure 21 - HFC Legacy – LLD Enabled, LL Flow During Speed Tests

Figure 21 shows the Low-Latency measurements during the same 10-minute period (as figure 20) which maintains a much lower and consistent latency throughout. The brief spikes that correlate in the LL flow were of interest, but not fully understood at the time of testing. One theory was that modem resources were being taxed enough to drive latency higher on both the LL and classic flows (albeit briefly) but this was not explored fully or proven out.

The speed tests were initiated remotely from a SamKnows agent integrated on the Wi-Fi router. The ODROID was directly connected to the router, which was directly connected to a standard DOCSIS 3.1 modem. Aside from the ODROID and the SamKnows speed test traffic, there were no other devices competing for resources.

4.2.2. HCF DAA (RPHY)

The graphs below (Figures 21 and 22) show latency on ToS0 test traffic versus ToS1 with LLD enabled on an RPHY architecture. While the 90th percentile is highlighted in these images (16.4 ms vs 11.2 ms), also of note is the lack of any traffic extending beyond 30ms in the LL ToS1 graph (Figure 23).

This shows that the LLD technology not only lowers the latencies across all percentiles, but it also reduces the maximum latency experience by the data, (by almost half, in this setup).

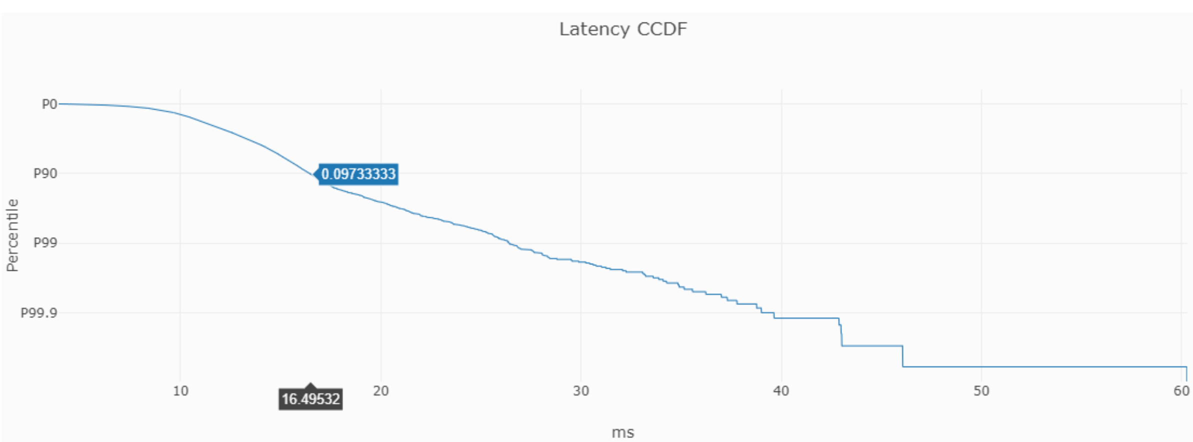

Figure 22 - LLD Enabled (RPHY) – ToS0

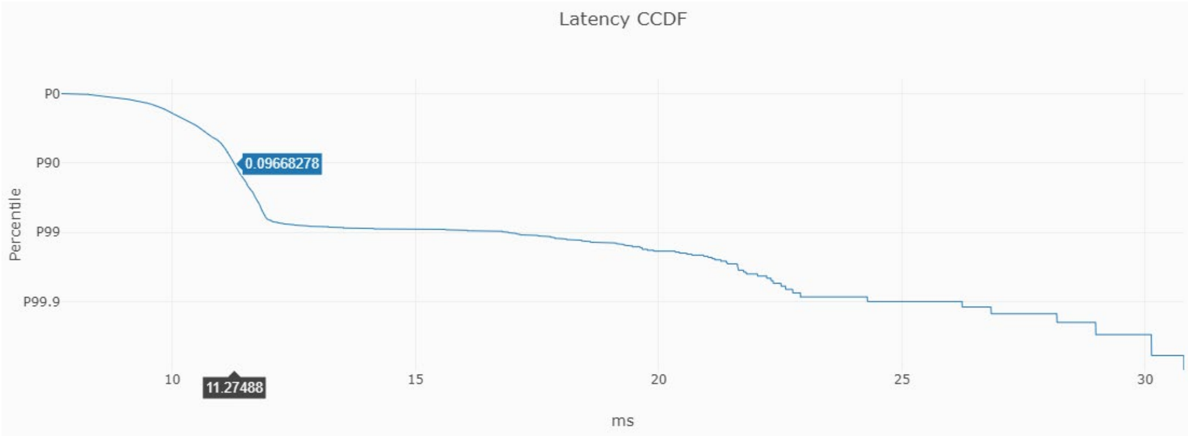

Figure 23 - LLD Enabled (RPHY) - ToS1

4.3. Enabling DPS & LLD

For the third set of experiments, we enabled low latency DOCSIS and are also using the DPS functionality (in lieu of the unavailable PGS functionality).

4.3.1. HFC Legacy

This section presents latency measurements from two separate perspectives, each with classic and LL performance metrics.

1. Idle Network – no significant network events, best case scenario.
2. Network Events (large file UL/DL) – These types of events drive latency up, but only on the classic/ToS0 measurements.

4.3.1.1. HFC Legacy - Idle Network

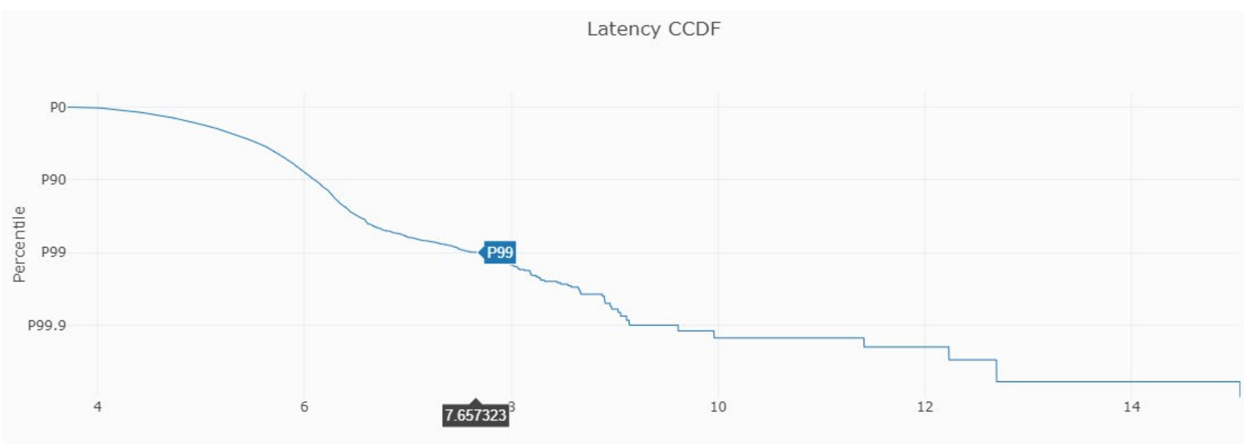

Figure 24 - Idle Network - LLD & DPS - ToS0

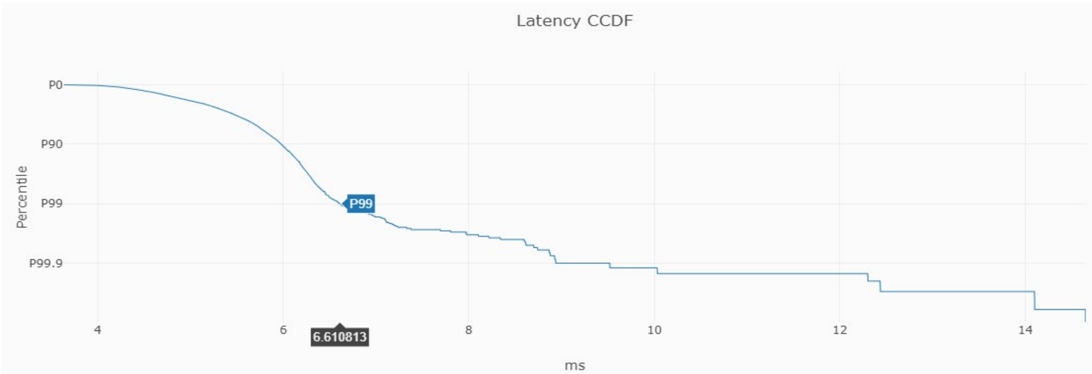

Figure 25 - Idle Network - LLD & DPS - ToS1

This is further evidence of very similar performance across the LL and classic flows when there is hardly any traffic to cause queueing delays at the modem. The lower latency across both is largely attributed to DPS simply reducing the latency across the board.

4.3.1.2. HFC Legacy - Network Events

When we introduce network events, like large file uploads, the benefits from LLDs are more meaningful and obvious. The image below shows both ToS0 and ToS1 latency measurements overlayed during a 1GB file upload. Throughout this test period a real-time multiplier gaming session was in progress, which saw a minimal latency increase of roughly 3ms.

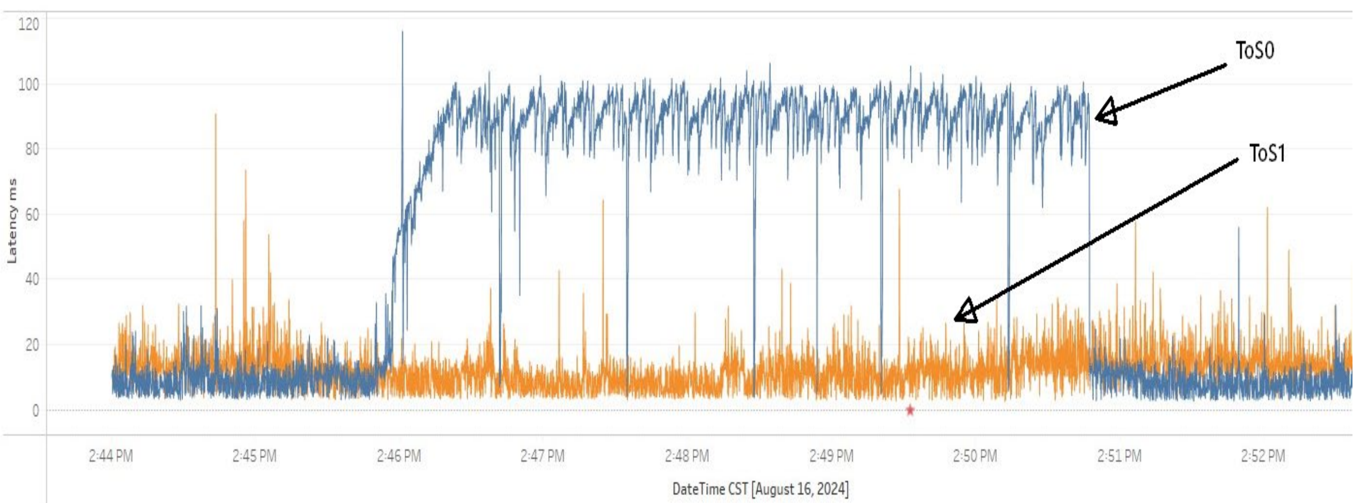

Figure 26 - LLD & DPS Enabled - ToS0 & ToS1 Overlay

4.3.2. HFC Legacy - LLD Configuration

The modem was provisioned for a 550Mbps DS and 33Mbps US on the classic flow, with 100Mbps DS and 10Mbps US on the LLD side. The QOS and ASF outputs from the CMTS are shared below for reference:

```
#show cable modem f834.5a52.e5ba qos

Sfid  Dir Curr  Sid   Sched  Prio MaxSusRate  MaxBrst     MinRsvRate  Throughput
          State       Type
18271 US  act   1825  BE     0    33000000    9800000     0           17035
27145 US  act   6870  BE     0    10000000    30000       0           13747
27146 US  pro   0     RSVD   0    0           0           0           0
18272 DS  act   N/A   N/A    0    550000000   3044        0           16971
27147 DS  act   N/A   N/A    0    100000000   3044        0           13064
27148 DS  pro   N/A   ASF    0    0           3044        0           0
```

Figure 27 - LLD Individual Service Flow Configuration for Test CM

```
#show cable modem f834.5a52.e5ba asf

DOWNSTREAM AGGREGATE SERVICE FLOW DETAIL:

            Aggregate SF
            ASFID:   27148
            ASF Ref: 00112
                /  \
               /    \
              /      \
             /        \
   Classic SF          Low Latency SF
   SFID:   18272       SFID:   27147
   SF Ref: 00002       SF Ref: 00102

ASFID                                : 27148
ASF Reference                        : 112
Priority                             : 0
Max Rate                             : 0
Max Burst                            : 3044
Min Rate                             : 0
Minimum Reserved Rate Packet Size    : 0
Peak Rate                            : 0
Low Latency SFID                     : 27147
Low Latency SF Reference             : 102
Low Latency SF SCN                   :
Classic SFID                         : 18272
Classic SF Reference                 : 2
Classic SF SCN                       :
AQM Coupling Factor                  : 20
Scheduling Weight                    : 230
Queue Protection Enable              : 1
QP Latency Threshold                 : 1500
QP Score Threshold                   : 32000
QP Drain Rate Exponent               : 21

UPSTREAM AGGREGATE SERVICE FLOW DETAIL:

            Aggregate SF
            ASFID:   27146
            ASF Ref: 00111
                /  \
               /    \
              /      \
             /        \
   Classic SF          Low Latency SF
   SFID:   18271       SFID:   27145
   SF Ref: 00001       SF Ref: 00101

ASFID                                : 27146
ASF Reference                        : 111
Priority                             : 0
Max Rate                             : 0
Max Burst                            : 3044
Min Rate                             : 0
Minimum Reserved Rate Packet Size    : 0
Peak Rate                            : 0
Low Latency SFID                     : 27145
Low Latency SF Reference             : 101
Low Latency SF SCN                   :
Classic SFID                         : 18271
Classic SF Reference                 : 1
Classic SF SCN                       :
AQM Coupling Factor                  : 20
Scheduling Weight                    : 230
Queue Protection Enable              : 1
QP Latency Threshold                 : 1500
QP Score Threshold                   : 32000
QP Drain Rate Exponent               : 21
```

Figure 28 - LLD ASF Configuration on CMTS

4.3.3. HCF DAA (RPHY)

The results with both DPS and LLD enabled put the 99th percentile for NQB traffic at just over 7ms. The original baseline for DAA/RPHY was roughly 19ms at the 99th percentile, see Figure 17. Not only is this a significant improvement, but it's also within 3ms of the baseline PON measurement at the 99th percentile!

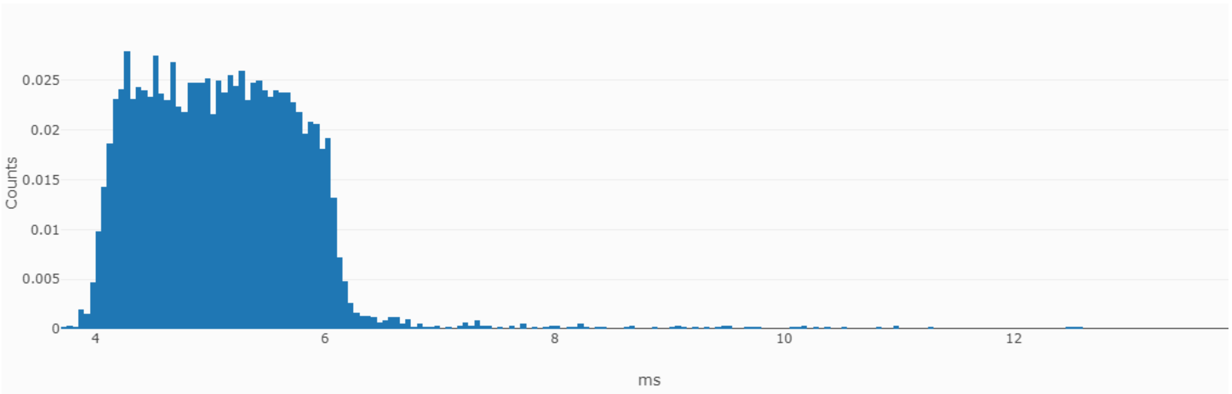

Figure 29 - HFC DAA (RPHY) with DPS & LLD Enabled – Latency Histogram

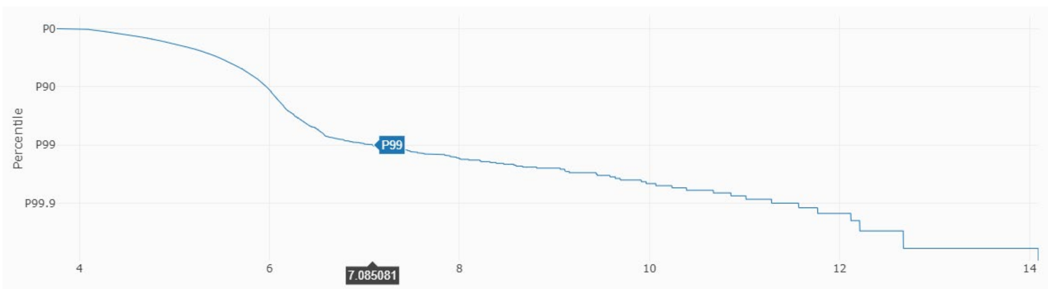

Figure 30 - HFC DAA (RPHY) with DPS & LLD Enabled – Latency CCDF

4.4. Latency Under Load

4.4.1. Game Update

The game download shared previously (Figure 14) is below once more, but this time with the Low-Latency measurements for the same 5-minute test period directly below.

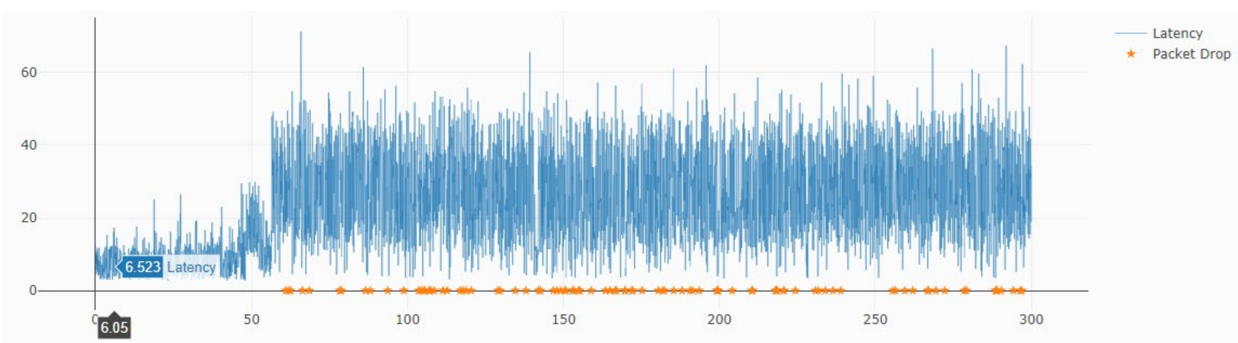

Figure 31 - Latency Under Load - Game Update -Classic SF

Below is the ToS1 test packet latency during the same period which shows almost no impact during the same 5-minute test period:

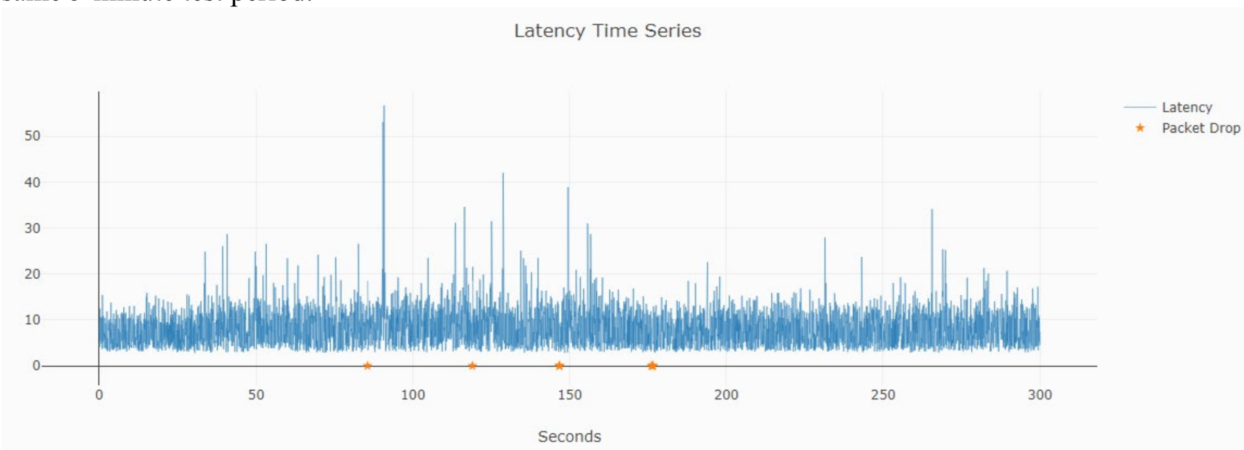

Figure 32 - Latency Under Load - Game Update - LL SF

4.4.2. Speed Tests

Speed tests were used throughout the testing, as a quick means to compare LLD test traffic while latency spikes were occurring on the classic SF. In addition to the comparisons presented in Section 4.2.1, below shows test packet latency during four consecutive speed tests ran with DPS and LLD enabled on Legacy HFC Architecture.

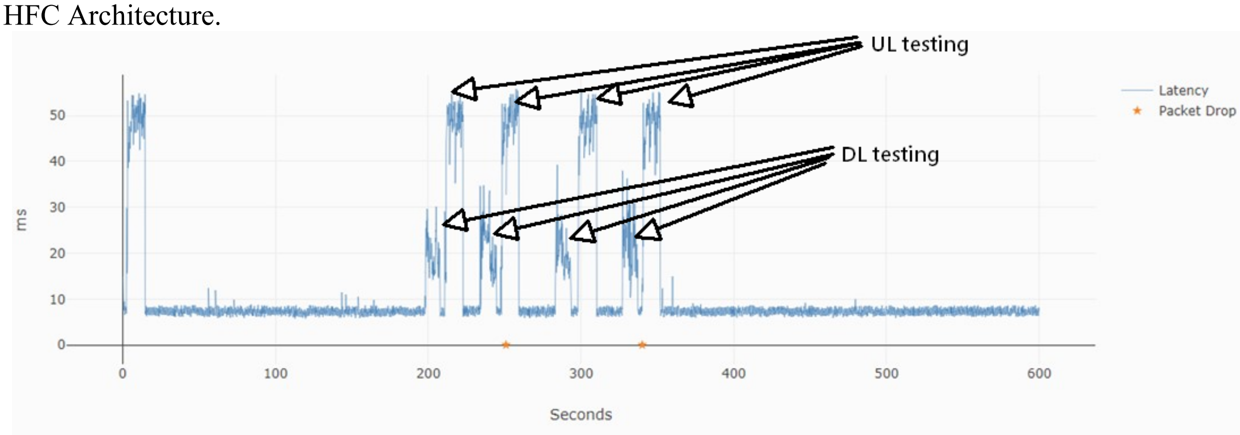

Figure 33 - Speed Test Time Series - LLD & DPS - Legacy HFC – ToS0

As seen in Figure 33, the Downstream speed test (or load) affects the latency to increase from 10 to about 30 ms, but the upstream speed tests affects the latency to increase to about ~50 ms. The impact of the upstream queueing causes more significant changes in latency.

For traffic on the low latency queue, **Figure 34** , the latency does not increase during DS speed tests, but we see a small increase in latencies (~2ms) during upstream speed tests.

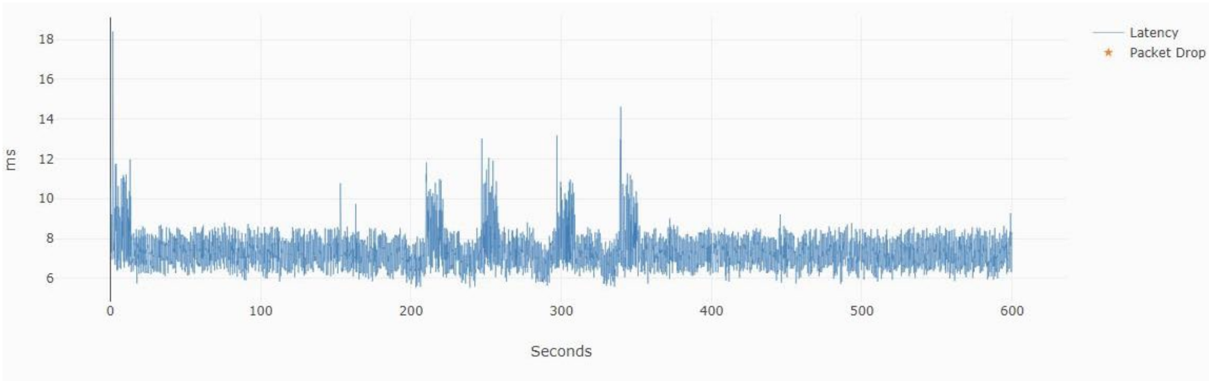

Figure 34 - Speed Test Time Series - LLD & DPS - Legacy HFC – ToS1

4.4.3. SamKnows Latency Under DS & US Load

The SamKnows® test suite was used as an alternate means to validate results obtained via the CableLabs® tool that served as the primary point of reference throughout this paper. It may be helpful for operators to have at least two methods to validate latency performance in the lab and the field.

The test results presented in this section were obtained via the SamKnows agent on the router performing a download test, while simultaneously sending test traffic marked with ECN bit set to 1. Prior to the significant drop in latency shown in both tests, the test traffic was unmarked (ECN 0) and subject to the same latency experienced by the download and upload test traffic. In other words, the drop in latency is a result of the test traffic being marked with ECN 1, to be begin utilizing the LL flows. The different colored lines represent different make and models of modems enabled as part of the field trial.

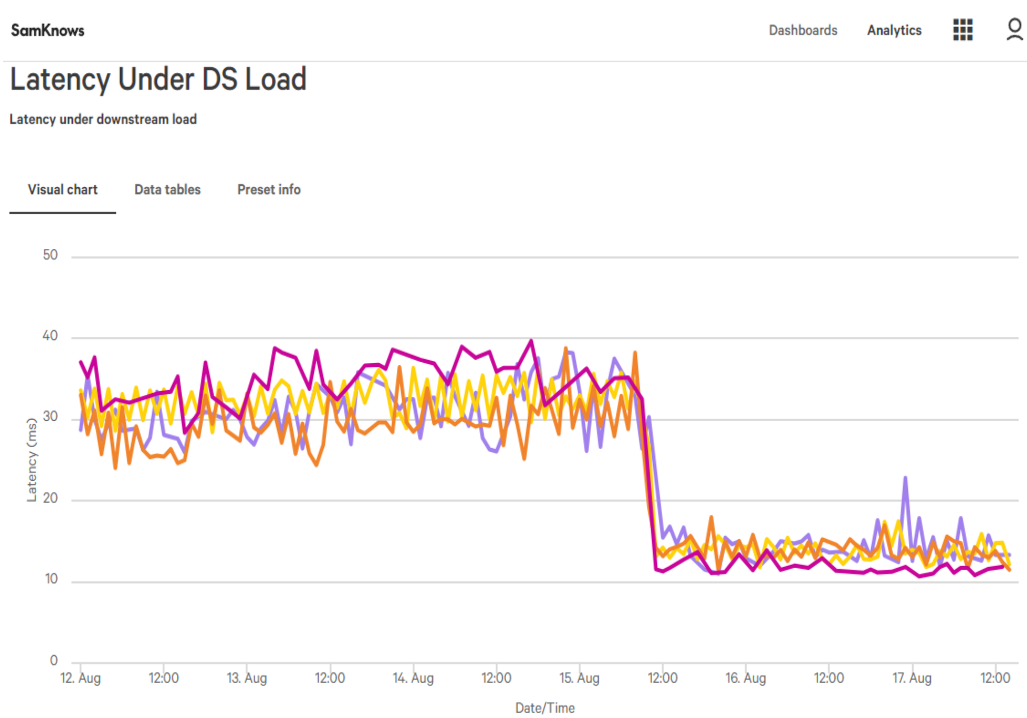

Figure 35 - SamKnows Latency Under DS Load w/ ECN 1

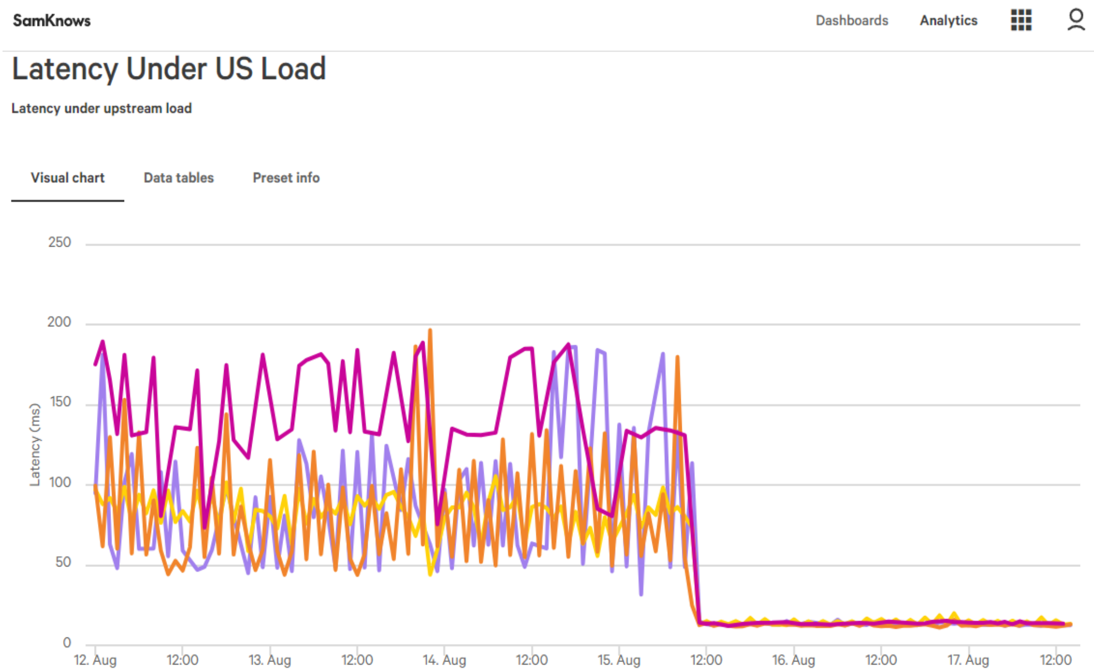

Figure 36 - SamKnows Latency Under US Load w/ ECN 1

5. Conclusions

The trial results and testing indicate that there are sufficient means to deliver high-speed data services on HFC networks with latencies that compete with or best PON performance for NQB traffic.

This assumes that enabling PGS would have further reduced latencies to levels ranging from 0.9ms to sub 5ms. Given the results of 7ms of latency for NQB traffic on RPHY architecture, it certainly seems both reasonable and attainable. This is caveated with the understanding that only applications utilizing L4S and traffic marked appropriately will ultimately utilize the low-latency flow. Scheduling features like DPS were shown to reduce latency on all traffic types by making use of unallocated resources, if desired.

The CableLabs developed latency measurement tools are useful for the purposes of measuring the impact of enabling LLD on MSO networks. This measurement suite can be used in production networks to provide quantitative data on latency baselines and outcomes when enabling features like LLD and DPS to improve latency for subscribers. There can be varied operator configurations to support different marked traffic types traversing the LLD flow and it's important to ensure the expected marked traffic types are utilizing the low-latency flows.

There were also some make/model of cable modems that did not benefit from low-latency implementation to the extent of others. Though not explored in detail herein, it's worth calling out that operators should validate deployed CM firmware and CMTS software for proper support of LLD features. There were performance differences noted across cable modem models, however any modems with perceived performance issues (or possible LLD implementation issue) were excluded from the test results. It is highly recommended operators have a means to easily disable LLD in the event it's needed for any reported network issues that need further investigation or troubleshooting.

Despite the LLD technology existing for several years, it's only now beginning to see more widespread deployment efforts by operators. This change creates many calls to action for several parties. MSOs should unite around traffic markings that are standardized whenever possible to ease the greater adoption of LLD. DSCP 45 (also recommended by the IETF [5]) is the logical choice for many, however if some networks are already using this for other purposes it may require rework. Additionally, operators should allow these markings to pass through their networks without remarking when possible. Application developers should begin marking traffic that is appropriate for the LL flow and fully understand the expected traffic behavior for compliance purposes.

The current generation of DOCSIS3.1 equipment deployed in the field experiences typical latency performance of ~8 to 12 ms on the access network link, and ~10-14 ms when using a DAA /RPHY architecture. Under heavy load, the link can experience delay spikes of ~50 ms or more. Enabling LLD features, AQM and ASF in our case, delivered a more consistent 4~7 ms delay on the DOCSIS network for non-queue building traffic. This makes the experience for real time applications more consistent with much smaller latency variation.

Abbreviations

AQM	Active Queue Management
ASF	Aggregate Service Flow
CCAP	Converged Cable Access Platform
CCDF	Complementary Cumulative Distribution Function
CE	Congestion Experienced
CM	Cable Modem
CMTS	Cable Modem Termination System
DAA	Distributed Access Architecture
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data-Over-Cable Service Interface Specification
DPS	DOCSIS Predictive Scheduler
DSCP	Diffserv Code Point
ECN	Explicit Congestion Notification
GbE	Gigabit Ethernet
HFC	Hybrid Fiber-Coaxial
ICMP	Internet Control Message Protocol
IP	Internet Protocol
L4S	Low-Latency Low-Loss Scalable Throughput
LMAP	Large-scale Measurement Platform
LL	Low Latency
LLD	Low Latency DOCSIS
MAP	Map
ms	millisecond
MSO	Multiple-System Operator
NQB	Non-Queue-Building
PGS	Proactive Grant Service
QB	Queue-Building
OFDMA	Orthogonal frequency-division multiple access
QoE	Quality of Experience
QP	Queue Protection
RFC	Request For Comments
RTT	Round-Trip Time
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SF	Service Flow
SNMP	Simple Network Management Protocol
STAMP	Simple Two-Way Active Measurement Protocol
TCP	Transport Control Protocol
TFTP	Trivial File Transfer Protocol
TLV	type-length-value Encoding
ToS	Type of Service
WAN	Wide Area Network

Bibliography & References

[1] CableLabs, Low Latency DOCSIS: Technology Overview, G. White, K. Sundaresan, B. Briscoe

[2] MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 04, 2020, Cable Television Laboratories, Inc.

[3] Low Latency DOCSIS Overview and Performance Characteristics, G. White, K. Sundaresan, B. Briscoe, SCTE Cable-Tec Expo 2019, https://www.nctatechnicalpapers.com/Paper/2019/2019-low-latency-docsis/download.

[4] CableLabs Network Performance, Latency Measurement Metrics and Architecture, G. White, K. Sundaresan - CL-TR-LM-Arch-V01-221123

[5] A Non-Queue-Building Per-Hop Behavior (NQB PHB) for Differentiated Services https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/

[6] RFC 8762 Simple Two-Way Active Measurement Protocol https://datatracker.ietf.org/doc/html/rfc8762