

## **Understanding Latency across PON systems**

**(With few comparisons to DOCSIS systems)**

A Technical Paper prepared for SCTE by

**Karthik Sundaresan**

Distinguished Technologist  
CableLabs  
858 Coal Creek Circle, Louisville  
3036613895  
k.sundaresan@cablelabs.com

**Evariste Some**

R&D Wired Intern, PhD Student,  
CableLabs, & University of Colorado, Boulder  
858 Coal Creek Circle, Louisville  
3036619100  
e.some@cablelabs.com

**Special Acknowledgments for Lab support to:**

**Doug Jones, Jay Zhu, Sheldon Webster, Aaron Quinto, CableLabs**

**James Lin, Radhouan Allani, Kyrio**

**Randy Wiggins, Tibit**

## Table of Contents

Title	Page Number
1. Introduction.....	4
2. PON Technologies Background.....	4
3. Background of EPON and GPON networks.....	6
3.1. EPON Latency background.....	7
3.2. GPON Latency Background.....	8
4. Results from Latency Measurements.....	9
4.1. Lab experimental Setup .....	9
4.2. Test Scenarios .....	10
4.3. PON Throughput Test Results .....	11
4.3.1. XGSPON Throughput .....	11
4.3.1. 10GEPON Throughput.....	12
4.4. Baseline Latency (RTT) Test Results .....	13
4.4.1. Baseline XGSPON Latency .....	14
4.4.2. Baseline 10GEPON Latency.....	14
4.4.3. Baseline DOCSIS Latency.....	15
4.4.4. Baseline DOCSIS (with LLD) Latency .....	15
4.5. UDP Load Test Results.....	16
4.5.1. GPON UDP 1Gbps Test Results .....	16
4.5.2. GPON UDP 9Gbps Test Results .....	17
4.5.3. EPON UDP 1Gbps Test Results.....	18
4.5.4. EPON UDP 9Gbps Test Results.....	19
4.5.5. DOCSIS UDP Test Results.....	20
4.6. TCP Load Test Results .....	22
4.6.1. XGSPON TCP Test results .....	22
4.6.2. 10GEPON TCP Test results .....	23
4.6.3. DOCSIS TCP Test results .....	24
4.7. GPON/ EPON /DOCSIS Test Results summary.....	25
4.8. Multiple LLIDs Test Results (EPON).....	27
4.9. Challenges and Further testing.....	29
5. Conclusions.....	29
Abbreviations .....	30
Bibliography & References.....	30

## List of Figures

Title	Page Number
Figure 1 – Evolution of PON Acces technology .....	5
Figure 2 – Simplified EPON Network.....	6
Figure 3 – Simplified Media access in EPON: Gate-Report-data .....	7
Figure 4 – GPON Upstream Transmission mechanisms .....	9
Figure 5 – PON Network Lab Test Setup .....	10
Figure 6 – XGSPON Throughput & one-way Latency(UDP) .....	12
Figure 7 – XGSPON Throughput & one-way Latency in Numbers (UDP).....	12
Figure 8 – 10GEPON Throughput & one-way Latency(UDP).....	13
Figure 9 – 10GEPON Throughput & one-way Latency in Numbers (UDP).....	13

Figure 10 – Baseline XGSPON latency (no load).....	14
Figure 11 – Baseline 10GEPON latency (no load) .....	14
Figure 12 – Baseline DOCSIS (with No LLD enabled) and no load .....	15
Figure 13 – Baseline DOCSIS with LLD latency and no load.....	16
Figure 14 – XGSPON latency – with 1 Gbps UDP background traffic, non-load ONU .....	16
Figure 15 – XGSPON latency – w 1 Gbps background traffic, Load ONU .....	17
Figure 16 – XGSPON latency – with 9 Gbps background traffic on non-load ONU .....	17
Figure 17 – XGSPON latency – with 9 Gbps background traffic on load ONU .....	18
Figure 18 – 10GEPON latency – no load.....	18
Figure 19 – 10GEPON latency – w 1 Gbps UDP background traffic- load ONU.....	19
Figure 20 – 10GEPON latency – w 9 Gbps background traffic - non-load ONU .....	19
Figure 21 – 10GEPON latency – w 9 Gbps background traffic - load ONU .....	20
Figure 22 – D3.1 (with AQM) latency – (4.7 and 1.7 Gbps UDP Load).....	21
Figure 23 – D3.1 with LLD, UDP test at 10% above service flow limit .....	21
Figure 24 – XGSPON latency – w TCP(1 Gbps, byteblower) - non-load ONU .....	22
Figure 25 – XGSPON latency – w TCP(9 Gbps,byteblower) - load ONU .....	22
Figure 26 – XGSPON latency – with Iperf TCP on non-load ONU .....	23
Figure 27 – XGSPON latency – with Iperf TCP on Load ONU .....	23
Figure 28 – 10GEPON latency – with Iperf TCP on non-load ONU .....	23
Figure 29 – 10GEPON latency – with Iperf TCP on load ONU.....	24
Figure 30 – DOCSIS no AQM config– with 1 Gbps TCP.....	24
Figure 31 – DOCSIS AQM Enabled Config, no LLD and Iperf TCP load .....	25
Figure 32 – DOCSIS LLD Enabled, Iperf TCP load .....	25
Figure 33 – GPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU .....	27
Figure 34 – EPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU .....	27
Figure 35 –10GEPON throughput (5 LLIDs 2Gbps per LLID) .....	28
Figure 36 –10GEPON Latency (5 LLIDs 2Gbps per LLID).....	28
Figure 37 –10GEPON visual of Throughput & Latency (for one US and DS LLID).....	28

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – PON Configuration settings .....	10
Table 2 – Latency Tests.....	11
Table 3 – GPON latency Test results , P0,P50,P99 .....	26
Table 4 – EPON latency Test results , P0,P50,P99 .....	26
Table 5 – DOCSIS latency Test results , P0,P50,P99.....	26

## 1. Introduction

HFC/DOCSIS networks are the most widely deployed technology for delivering Internet data services to the consumers. A passive optical network is another option for access technology deployed by operators using EPON, or GPON technology.

One of the performance questions which many in the operator community want to understand is the latency performance of each type of technology and how they compare. We setup a 10GEPON and XGSPON systems in our labs at CableLabs. We ran speed and latency benchmarks on these types of networks in a lab environment and compared the latencies that we get under a variety of conditions. DOCSIS has various options and configurations (D3.0 SC-QAM channels, D3.1 OFDM/OFDMA channels, AQM, LLD etc., while PON has various flavors: 10G-EPON, XGSPON etc. and configurations as well). This paper will give us a better understanding of the efficiencies and the latencies of each of these technology options and the tradeoffs this entails. Based on time constraints during lab testing, this paper is focused primarily on the PON latencies, with some limited testing on DOCSIS systems and a more complete testing and analysis will be part of future work. This paper will give a brief overview of each technology and its evolution and bring out a theoretical comparison of the latency characteristics of each type of technology and mainly reports on the lab testing results for each technology under different operating and load conditions.

## 2. PON Technologies Background

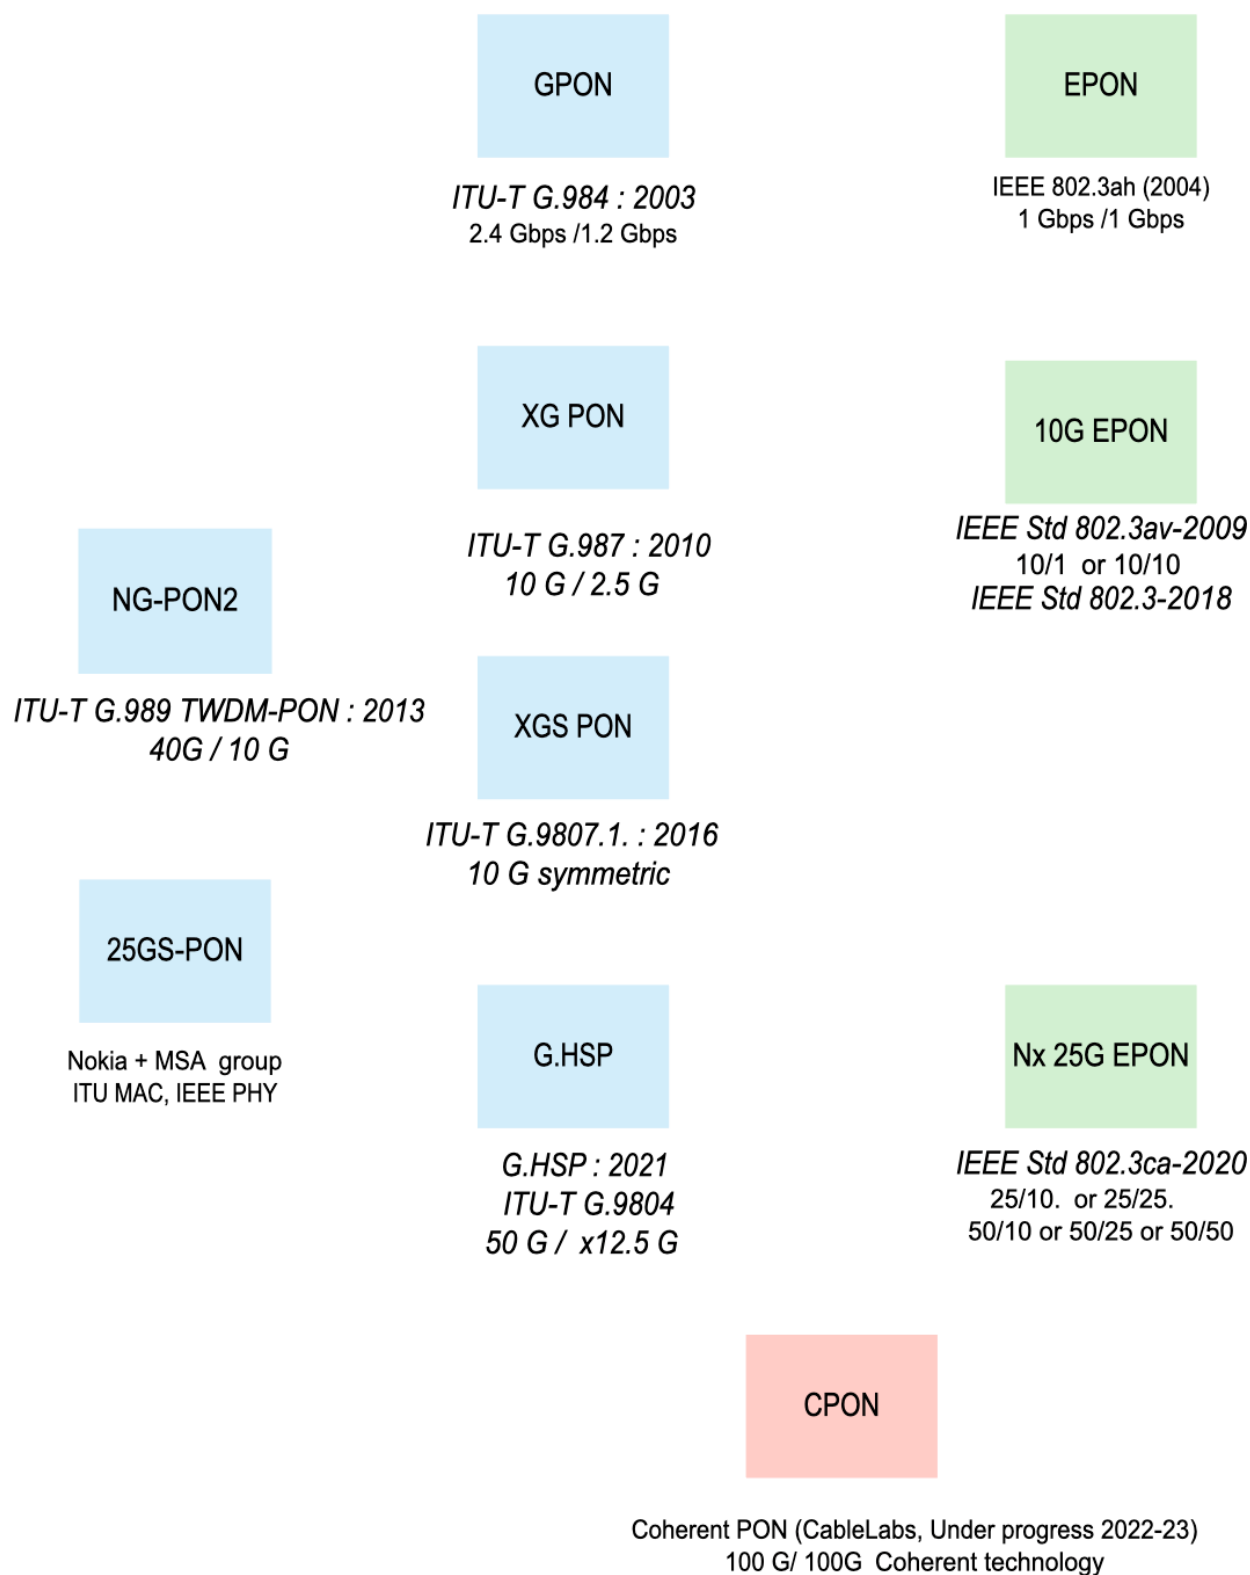
The continuous growth in bandwidth usage has put demands on the core, metro, and access networks. Gigabit access speeds are now the norm with a large share of US population having an option. Per NCTA, 88% of US Homes have access to gigabit Internet speeds available to them.

Passive optical networks (PON) have been used to serve the end users with high data rate, high split ratio, and high bandwidths. PON provides this to a large number of customers at a low cost. In PON, the Time Division Multiplexing (TDM) method raises the number of end users by using power splitters that split the power of the channel into multiple segments.

P2MP PON based architectures has proven to be the most popular among operators for FTTx deployments because it typically results in lower infrastructure costs. With P2MP, a shared fiber carries traffic to a passive splitting complex close to the end-customers, which can be single stage or cascaded. Consequently, there is no need to deploy a single dedicated fiber all the way to the customer premises from the central office/Headend.

Optical fiber access systems based on PON are currently being deployed on a mass market scale by numerous network operators worldwide. These systems typically exploit Gigabit-class PON systems such as G-PON or EPON as standardized by the ITU-T and IEEE respectively. Both these standards bodies have also defined 10 Gigabit-class PON technologies XGS-PON and 10G-EPON and are widely available today. So, we focus our energy on understanding the latency performance of the XGS-PON and 10G-EPON systems.

The figure below shows the evolution of PON technologies as defined by the IEEE and the ITU-T. both the EPON and GPON standards have been evolving over the last 20 years going from 1 Gbps to now 50 Gbps and beyond.



**Figure 1 – Evolution of PON Acces technology**

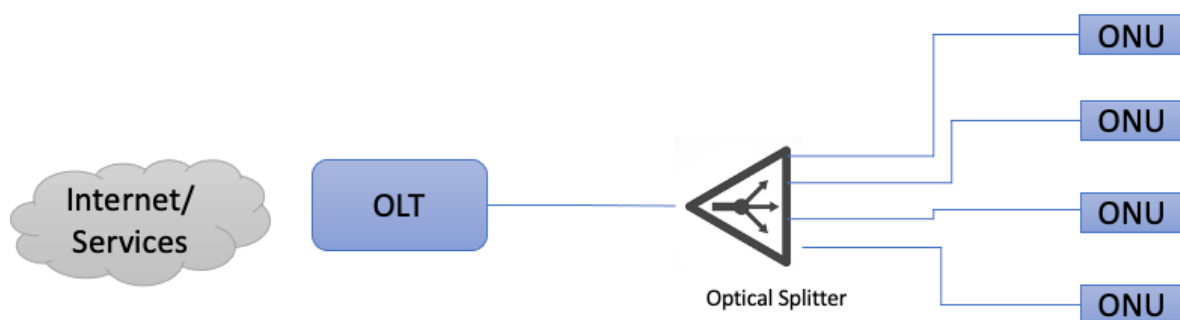
Over the last two decades, PON based technology evolved greatly and rapidly, and is one of the most attractive access network solutions for delivering high-speed data and video services. After initial 1 Gbps speeds, PON standards evolved into 10 Gbit/s Ethernet PON and 10-Gigabit-capable PON (XGS-PON). Over the last few decades, the access speed of PON over single wavelength has increased almost 100 times to support the ever-growing bandwidth demand from the emerging services. The 25/50G next-generation Ethernet PON (NG-EPON) specification has been approved by IEEE in 2020, based on 25 Gbps per single wavelength. Most recently, the ITU-T 50G PON standard, based on 50G single-wavelength is close to being a reality as consent has been achieved on multiple part of the series in 2021. Figure 1 above shows a summary of the existing PON standards and the supported downstream/upstream transmission rates.

### 3. Background of EPON and GPON networks

Optical Access Networks have typically been deployed using point-to-multipoint PON based architectures. In PON networks, Optical networking units (ONUs) communicate with the Optical Line Terminal (OLT) via a passive optical distribution network consisting of a feeder fiber (connecting the OLT and the passive splitter), a passive splitter, and numbers of distribution fibers (connecting each ONU to the passive splitter), as shown in Figure 2.

In the downstream direction (from OLT to an ONU), signals transmitted by the OLT pass through a 1: N passive splitter (or cascade of splitters) and reach each of the ONUs. In the upstream direction (from ONUs to OLT), the signal transmitted by an ONU will only reach the OLT, and not any other ONUs. To avoid data collisions and increase the efficiency of the subscriber access network, ONU's transmissions are arbitrated by the OLT.

This arbitration is achieved by allocating a transmission window (grant) to each ONU. An ONU defers transmission until its grant arrives. When the grant arrives, the ONU transmits frames at wire speed during its assigned time slot. A simplified P2MP topology example is shown in the figure below.



**Figure 2 – Simplified EPON Network**

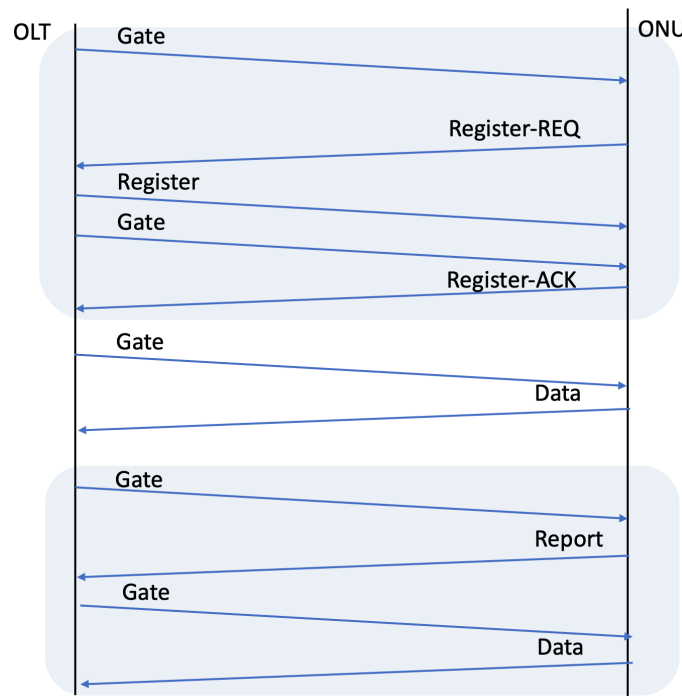
The packet delay in a PON network is the time elapsed between a frame's generation at the ONU and its arrival at the OLT. In the most general case, the delay components experienced by a frame include the Delay between frame generation and the transmission of the next REPORT message to the OLT, Propagation of the REPORT message from a given ONU to the OLT, (different for each ONU), processing delay at the OLT between the arrival of the REPORT and the transmission of the GATE message, Propagation of the GATE message to the ONU, delay between the arrival of the GATE message at the ONU and the beginning of the granted transmission window, propagation of the frame to the OLT and the duration of the frame transmission itself.

### 3.1. EPON Latency background

In EPON, the Multipoint MAC Control Protocol (MPCP) defines the messages and timers, to control access to the P2MP ODN topology in 10G-EPON. Every P2MP ODN topology consists of one Optical Line Terminal (OLT) plus one or more ONUs, as shown in the Figure.

The MPCP and underlying PHY allows an underlying P2MP network to appear as a collection of point-to-point links to the higher protocol layers (at and above the MAC Client). The MPCP achieves this by providing a Logical Link Identification (LLID) to each MAC which is dynamically assigned by the registration process.

The MPCP has the ability to arbitrate the transmitters of a number of ONUs. The OLT controls an ONU's transmission by assigning grants. The transmitting window of an ONU is indicated in gate message where start time and length are specified. An ONU will begin transmission when its local time counter matches start time value indicated in the gate message. An ONU will conclude its transmission with sufficient margin to ensure that the laser is turned off before the grant length interval has elapsed.



**Figure 3 – Simplified Media access in EPON: Gate-Report-data**

Multiple outstanding grants may be issued to each ONU. The OLT does not issue more than the maximal outstanding grants as advertised by the ONU during registration. In order to maintain a watchdog timer at the ONU, grants are periodically generated. For this purpose, empty gate messages may be issued periodically.

The purpose of gate message is to grant transmission windows to ONUs for both discovery messages and normal transmission. Up to four grants can be included in a single gate message. The length of the signaled grant, this is a 16-bit unsigned field. The length is counted in 16-bit time increments. The laser-on time, sync time, and laser-off time are included in the grant.



In the [IEEE 802.3] EPON technology, the MPCP protocol relies on strict timing based on distribution of timestamps. The implementations guarantee a constant delay through the MAC and PHY in order to maintain the correctness of the timestamping mechanism. The actual delay is implementation dependent; however, a complying implementation maintains a delay variation of no more than 16-bit times through the MAC stack.

The OLT does not grant less than 1024-time quanta into the future, in order to allow the ONU processing time when it receives a gate message. The ONU shall process all messages in less than this period. The OLT shall not issue more than one message every 1024-time quanta to a single ONU. The unit of time quantum is defined as 16 ns.

A report message has several functionalities. Time stamp in each report message is used for roundtrip time calculation. In the report messages ONUs indicate the upstream bandwidth needs they request per IEEE 802.1Q priority queue. Report messages are also used as keep-alives from ONU to OLT. ONUs issue report messages periodically in order to maintain link health at the OLT. In addition, the OLT can specifically request a report message. Status reports are used to signal bandwidth needs as well as for arming the OLT watchdog timer. Reports shall be generated periodically, even when no request for bandwidth is being made. This keeps a watchdog timer in the OLT from expiring and deregistering the ONU. For proper operation of this mechanism the OLT shall grant the ONU periodically.

EPON implementations typically adopt the traditional Interleaved Polling with Adaptive Cycle Time (IPACT) scheme (with refinements) as the method to allocate bandwidth across multiple ONUs and try to minimize this Gate-Report-data cycle.

### **3.2. GPON Latency Background**

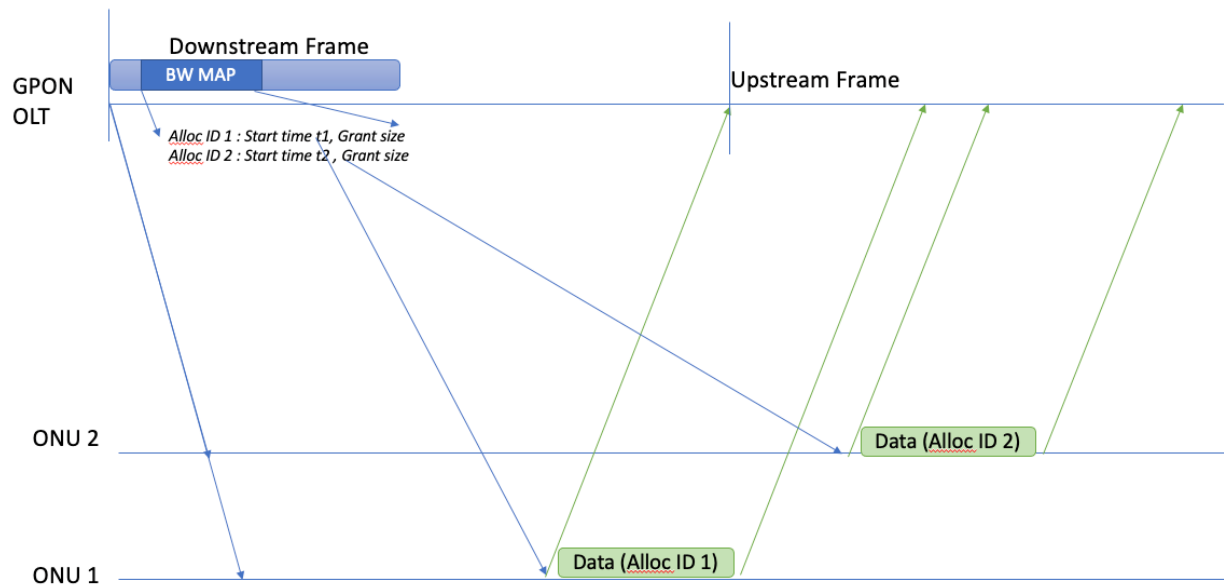
Recommendation ITU-T G.9807.1 describes a 10-Gigabit-capable symmetric passive optical network (XGS-PON) system in an optical access network for residential, business, mobile backhaul and other applications. This system operates over a point-to-multipoint optical access infrastructure at the nominal data rate of 10 Gbit/s both in the downstream and the upstream directions.

A GPON network will be used to carry different types of services, the delay requirements are designed and standardized according to the characteristics of these services.

TDM PON systems use the TDMA mechanism for communicating with multiple ONUs connected to the same OLT PON port. Therefore, Dynamic bandwidth Allocation (DBA) mechanisms and ONU activation mechanisms are needed in the convergence layer, which periodically introduce additional delay to the data transmission in the upstream direction. In real implementations, the data transmission delay of the PON link is considered to fluctuate within about 1.5 ms between service node interface (SNI) and user-network interface (UNI). When the OLT creates a quiet window as part of the ONU activation mechanism, the data transmission delay is significantly increased as a consequence

The XGS-PON OLT support DBA for the efficient sharing of upstream bandwidth among the connected ONUs and the traffic-bearing entities within the individual ONUs based on the dynamic indication of their activity. The dynamic activity indication can be based on the following two methods: (1) status reporting DBA employs the explicit buffer occupancy reports that are solicited by the OLT and submitted by the ONUs in response; (2) traffic monitoring DBA employs OLT's observation of the actual traffic amount in comparison with the allocated upstream transmission opportunities. To guarantee multi-vendor interoperability, the standard specifies the formats of the SR DBA status enquiries and buffer occupancy reports and the associated protocol. An ONU supports DBA status reporting, and transmit upstream reports as instructed by the OLT.





**Figure 4 – GPON Upstream Transmission mechanisms**


## 4. Results from Latency Measurements

### 4.1. Lab experimental Setup

We built a PON plant with a 10GEPON and XGSPON system at CableLabs, as shown in Figure 5. These systems consisted of SFP+ modules of 10GEPON and XGSPON OLTs. These SFP+ OLTs were connected to a 48-port 100Gps switch (Arista 7280R). The fiber cable connection coming out of OLTs were attenuated with a series of optical attenuators and fed into a 5-port optical splitter. An alternate setup was 10 Km of optical fiber to the ONUs. Coming off the splitter we had the two ONUs connected. We had two 10GEPON ONUs and two XGSPON ONUs connected to each of the respective OLTs.

The latency measurement system used here includes a Measurement agent (a STAMP Session sender) and a STAMP session reflector. Additionally, a Controller/Collector was used for visualization of results from the Measurement agent. See [STAMP SCTE21] for further details. These prototype software components are available at [C3 CableLabs] for operators to use.

A hardware traffic generator (ByteBlower with two 10 Gbps optical ports) was used to load the system with traffic. An iPerf server and client were also installed on dedicated servers with 10G NIC cards.



**Figure 5 – PON Network Lab Test Setup**

The goal of the lab testing was to understand the practical latency on the PON networks by sending different types of load traffic and measuring the latency through the network. Table below shows the default PON configuration on the OLTs

**Table 1 – PON Configuration settings**

XGSPON		10GE PON	
Discovery Period [ms]:	3000	Discovery Period [ms]:	3000
Encryption:	Bidirectional	Grant Spacing [16ns]:	16
Encryption Key Time [s]:	600	Encryption:	Downstream
Max Frame Size [bytes]:	9600	Encryption Key Time [s]:	900
Downstream FEC:	true	Sync Time [16ns]:	16
Upstream FEC:	true	Laser OFF [16ns]:	32
Guard Time [12.8ns]:	64	Laser ON [16ns]:	32
Preamble Length [12.8ns]:	64	Max Frame Size [bytes]:	9600
Error Det Max Ratio:	20%	FEC:	true
Error Det Min Sample [bursts]:	100	Fiber Reach:	Standard (0..20 km)
Fiber Reach:	Standard (0..20 km)		
Downstream & Upstream SLA Guaranteed: 512 Kbps, rest Best Effort		Downstream and Upstream SLA Guaranteed: 1.28 Mbps, rest is Best Effort	

## 4.2. Test Scenarios

The main test scenarios and traffic definitions for the tests were as follows.

- Latency Measurement Test Traffic (round trip time - RTT tests)
  - o Raw measurement: 10 pps, 256-byte packets for 2 mins = 20Kbps
  - o Percentile test: 30 pps, 256-byte packets, for 2 mins = 62Kbps
  - o Histogram test : 60 pps, 256-byte packets, for 2 mins = 122Kbps
- UDP Background traffic
  - o Send UDP traffic at different rates as load / background traffic and then measure the round-trip latency of the test traffic
  - o The idea is to send UDP traffic in steps of 1, 3, 5, 7, and up to 9 Gbps (just above the capacity of the two PON networks)
- TCP Background traffic
  - o Send TCP traffic and then measure the round-trip latency of the test traffic
  - o The idea was to have the TCP flows capped at certain rates (1, 3, 5, 7) and also unlimited (capped only by the network capacity ~9 Gbps)

**Table 2 – Latency Tests**

Scenario name	Load traffic Downstream Data rate	Load traffic Upstream Data Rate
UDP		
S0 Baseline	No load	No load
S1 UDP	1 Gbps UDP	1 Gbps UDP
S3 UDP	3 Gbps UDP	3 Gbps UDP
S5 UDP	5 Gbps UDP	5 Gbps UDP
S7 UDP	7 Gbps UDP	7 Gbps UDP
S9 UDP	9 Gbps UDP	9 Gbps UDP
TCP		
S1 TCP	1 Gbps TCP	1 Gbps TCP
S3 TCP	3 Gbps TCP	3 Gbps TCP
S5 TCP	5 Gbps TCP	5 Gbps TCP
S7 TCP	7 Gbps TCP	7 Gbps TCP
SU TCP	Unlimited TCP	Unlimited TCP

### 4.3. PON Throughput Test Results

#### 4.3.1. XGSPON Throughput

We ran a baseline throughput test, through the XGSPON system (OLT +ONU), just to understand what the maximum throughput is and what the OLT to ONU one-way latency is. A first port on the traffic generator device is directly connected via the ethernet switch to the OLT and a second port is connected to the ONU as a CPE device. Traffic is being sent in both upstream and downstream directions. The XGSPON system, shows a maximum throughput off 8.564 Gbps in the downstream direction and 8.416 Gbps in the upstream direction. Based on the latency numbers as calculated by the traffic generator device, the average one-way latency is 0.277 milliseconds on the downstream and 0.708 milliseconds in the upstream. The 10GEPON configuration was in the default state as detailed in Table-1.

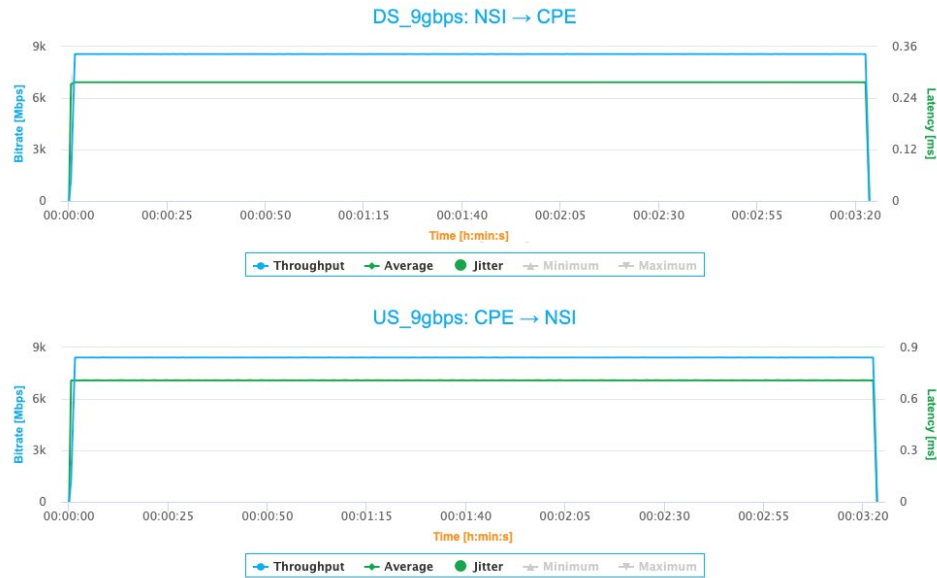


Figure 6 – XGSPON Throughput & one-way Latency(UDP)

Frame Blasting Flows : Throughput

Flow	Source	Destination	TX Frames	Rx Frames	Frame Loss	TX Bytes (+VLAN)	RX Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_9gbps	NSI	CPE	151,822,917	144,442,983	4.86%	227,734,375,500 (+607,291,668)	216,664,474,500	4.86%	3m, 22s 379ms 305µs 757ns	8,564.69
US_9gbps	CPE	NSI	151,823,300	141,561,578	6.76%	227,734,950,000	212,342,367,000 (+566,246,312)	6.76%	3m, 22s 378ms 708µs 613ns	8,416.25


Frame Blasting Flows : Latency

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_9gbps	NSI	CPE	0.024	0.277	0.281	0.001
US_9gbps	CPE	NSI	0.532	0.708	1.091	0.007

Figure 7 – XGSPON Throughput & one-way Latency in Numbers (UDP)

#### 4.3.1. 10GEPON Throughput

We ran a similar baseline throughput test, through the 10GEPON system (OLT +ONU), again to understand what the maximum throughput and the OLT to ONU one-way latency is. The traffic generator device has the same configuration as the above test. The 10GEPON system, shows a maximum throughput of 8.596 Gbps in the downstream direction and 8.436 Gbps in the upstream direction. Based on the latency numbers as calculated by the traffic generator device, the average one-way latency is 0.636 ms on the downstream and 1.739 ms in the upstream. The XGSPON configuration was in the default configuration as state in Table-1.



**Figure 8 – 10GEPON Throughput & one-way Latency(UDP)**

Frame Blasting Flows : Throughput

Flow	Source	Destination	Tx Frames	Rx Frames	Frame Loss	Tx Bytes (+VLAN)	Rx Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_9gbps	NSI	CPE	138,613,303	132,370,647	4.50%	207,919,954,500 (+554,453,212)	198,555,970,500	4.50%	3m, 4s 771ms 334µs 308ns	8,596.83
US_9gbps	CPE	NSI	138,613,615	129,561,132	6.53%	207,920,422,500	194,341,698,000 (+518,244,528)	6.53%	3m, 4s 769ms 526µs 608ns	8,436.89

Frame Blasting Flows : Latency

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_9gbps	NSI	CPE	0.024	0.636	0.641	0.003
US_9gbps	CPE	NSI	1.336	1.739	3.604	0.080

**Figure 9 – 10GEPON Throughput & one-way Latency in Numbers (UDP)**


#### 4.4. Baseline Latency (RTT) Test Results

The next set of tests that we ran is a set of standardized latency measurement tests as defined in section 5.1. Here the basic idea is to have a measurement agent (on the NSI side of the OLT), which sends STAMP (UDP) packets to a session-reflector device (located behind an ONU as the CPE). The session-reflector essentially sends the packet back to the measurement agent who then computes the roundtrip time. The sequence of 1200 packet latency measurements will be used to show a time series of the latency, a histogram of the latency and a latency CCDF as defined in [SCTE21 Latency Measurement]. The figures below also show any packet drops during the test. We also ran a couple of additional tests at the same time, which compute only the histogram of latency and the second one computes only the percentile values though we are not sharing the results from those tests, as for now the raw latency test

has all the information we need. These baseline tests are run without any load on the PON network, i.e., there is no background load traffic

#### 4.4.1. Baseline XGSPON Latency


As shown in the figure below, we see the latency of the XGSPON system to be between 1 to 2.1 ms, with a few outliers going all the way from 2.5 to 3.5 ms.



**Figure 10 – Baseline XGSPON latency (no load)**

#### 4.4.2. Baseline 10GEPON Latency

As seen in the figure below on a 10GEPON system, we see the latency values between 0.8 and 1.9 milliseconds




**Figure 11 – Baseline 10GEPON latency (no load)**



#### 4.4.3. Baseline DOCSIS Latency

We ran the same baseline latency test on a DOCSIS CMTS (no Low Latency DOCSIS (LLD) configured), to observe a baseline latency between 6.2-9 ms, with outliers in the 9-11 ms range.



**Figure 12 – Baseline DOCSIS (with No LLD enabled) and no load**

#### 4.4.4. Baseline DOCSIS (with LLD) Latency

We ran the same baseline latency test on a DOCSIS CMTS which is configured with the new Low Latency DOCSIS features (LLD) and observe a baseline latency between **2 and 3 ms**, with some outliers in the 3-4 ms range. This was just out-of-the-box configuration with LLD enabled on both the upstream and downstream. The latency measurement traffic in this test was marked with an IP DSCP Marking of 46 (expedited forwarding). With this we can see that enabling LLD features on a DOCSIS 3.1 CMTS definitely puts the latency in the same level as the EPON and GPON systems




Figure 13 – Baseline DOCSIS with LLD latency and no load

## 4.5. UDP Load Test Results

Now we have a set of UDP tests where we run different levels of UDP traffic in the background. The load traffic is sent through one ONU, while the latency RTT measurements are being done on both the reflectors that are behind the load ONU, as well as the second ONU which doesn't have any load traffic passing through it.

### 4.5.1. GPON UDP 1Gbps Test Results

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case.





Figure 14 – XGSPON latency – with 1 Gbps UDP background traffic, non-load ONU

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is carrying the load, the latency distribution does change and the histogram and the CCDF curves actually get better. From the


percentiles test we can tell that the latency for the P50 and P99 numbers actually improved. The 50<sup>th</sup> percentile latency improved from 1.557 to 1.278 ms and the 99<sup>th</sup> percentile latency from 2.78 ms to 1.534 ms. With the increased load on the PON network we can see that the polling-request-grant delay cycle is reduced, with a ONU buffer status report being a part of the upstream data frame header itself.



**Figure 15 – XGSPON latency – w 1 Gbps background traffic, Load ONU**


#### 4.5.2. GPON UDP 9Gbps Test Results

When we run background UDP load traffic of 9 Gbps (above the capacity of the network), we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case.



**Figure 16 – XGSPON latency – with 9 Gbps background traffic on non-load ONU**


When we run background UDP load traffic of 9 Gbps, we see that for the ONU that is carrying the load, the latency distribution does change much and the histogram and the CCDF curves get worse as expected. From the percentiles test we can tell that the latency for the 50<sup>th</sup> percentile latency (P50) increased from 1.557 to 2.423, and the 99<sup>th</sup> percentile latency changes from 2.78 ms to 3.482 ms. With the fully loaded PON network we can see that there is packet loss and the increased latency of the packets that did make through. This also points to a relatively short buffer on the ONUs and OLTs.




**Figure 17 – XGSPON latency – with 9 Gbps background traffic on load ONU**

#### 4.5.3. EPON UDP 1Gbps Test Results

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is carrying the load, the latency distribution shifts to the left (decreases) by 0.4 ms. For the ONU that is not carrying the load the latency remains the same as the baseline case.




**Figure 18 – 10GEPON latency – no load.**



**Figure 19 – 10GEPON latency – w 1 Gbps UDP background traffic- load ONU.**


#### **4.5.4. EPON UDP 9Gbps Test Results**

When we run background UDP load traffic of 9 Gbps (above the capacity of the network), we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case. This behavior for this 10GEPON system, is similar to what we saw with the XGSPON system.



**Figure 20 – 10GEPON latency – w 9 Gbps background traffic - non-load ONU**

When we run background UDP load traffic of 9 Gbps, we see that for the EPON ONU that is carrying the load, the latency distribution does change and the histogram and the CCDF curves get worse as expected. From the percentiles test we can tell that the latency for the 50<sup>th</sup> percentile latency (P50) increased from 2.132 to 2.423 ms as expected, though the 99<sup>th</sup> percentile latency improved from 2.769 ms to 2.548 ms. With the fully loaded PON network we can see that there is some packet loss and the increased latency of the packets that did make through. The packet loss on these test flows is much less than what we saw on the XGSPON network, this implies that the EPON implementations have a per-flow queuing mechanism. This also points to a relatively short buffer on the EPON ONUs and OLTs.




**Figure 21 – 10GEPON latency – w 9 Gbps background traffic - load ONU**

#### **4.5.5. DOCSIS UDP Test Results**

For this test we used a high-split CMTS, which has a downstream capacity of 4.8 Gbps downstream and 1.7 Gbps upstream. When we run background UDP load traffic of 4.7 Gbps downstream and 1.65 Gbps (just fully loading the DOCSIS link capacity configured for this setup), we see that for the CM, the latency distribution does change and the histogram and the CCDF curves shift to the right as expected. From the percentiles test we can tell that the latency for the 50<sup>th</sup> percentile latency (P50) increased from 7.2 to 15.5 ms, and the 99<sup>th</sup> percentile (P99) latency changes from 11.27 to 20.4ms. This test essentially shows the power of the AQM feature which all operators must ensure is turned on in their network.


As seen in the figure below, we can see that the baseline latency with fully loaded network is at the 7-8 ms mark. When we start the UDP load at about 16 seconds into the test we see that the latency spikes up to 30 milliseconds and then once the AQM algorithm kicks in, the latency stabilizes at about 15 ms.





**Figure 22 – D3.1 (with AQM) latency – (4.7 and 1.7 Gbps UDP Load)**

In the graph below we have an experiment with Low Latency DOCSIS (LLD) is turned on and the latency measurements are between 2 to 5 milliseconds even when we are sending traffic at 10% above the configured data rate. (In this configuration, service flows are limited to 200 Mbps down/20 Mbps up).



**Figure 23 – D3.1 with LLD, UDP test at 10% above service flow limit**

The test itself runs traffic at 10% above the 200 Mbps downstream rate limit and 20 Mbps upstream rate limit.

## 4.6. TCP Load Test Results

The next set of tests includes TCP traffic as load or background traffic either on the ONU under test or on a neighboring ONU on the same PON.

### 4.6.1. XGSPON TCP Test results

The next three figures show three different TCP Load test scenarios. In the first case we rate limit the TCP flows to be 1 Gbps, in the second we rate limit the TCP flows to 9 Gbps and in the third we run TCP flows using iPerf which are limited only by the network capacity. In all three cases we see that the XGSPON system handles bidirectional TCP flows without a significant change to the latency numbers.

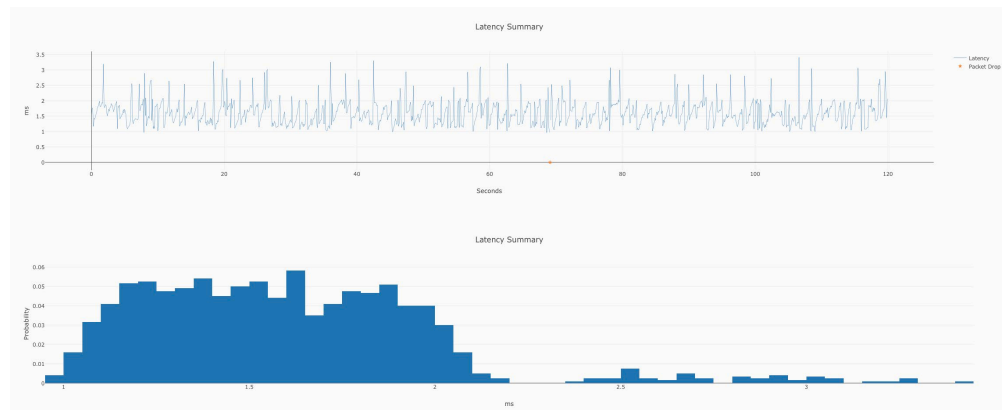


Figure 24 – XGSPON latency – w TCP(1 Gbps, byteblower) - non-load ONU






Figure 25 – XGSPON latency – w TCP(9 Gbps, byteblower) - load ONU

Again, as before when the latency measurement is being done on the ONU which is carrying the load traffic, we see that the latency numbers are actually lower, compared to the ONU which is not carrying any load traffic.



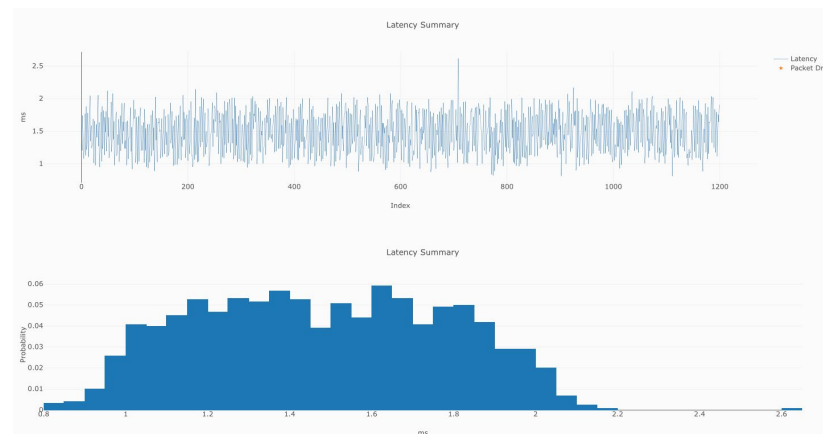
**Figure 26 – XGSPON latency – with Iperf TCP on non-load ONU**



**Figure 27 – XGSPON latency – with Iperf TCP on Load ONU**

#### 4.6.2. 10GEPON TCP Test results

Similarly, for the 10GEPON system, we see that the system handles bidirectional TCP flows without a significant change to the latency numbers.



**Figure 28 – 10GEPON latency – with Iperf TCP on non-load ONU**

When the latency measurement is being done on the EPON ONU which is carrying the load traffic we see that the latency numbers get worse, compared to the ONU which is not carrying any load traffic. This behavior is similar to the results we observed for the UDP tests on the EPON network

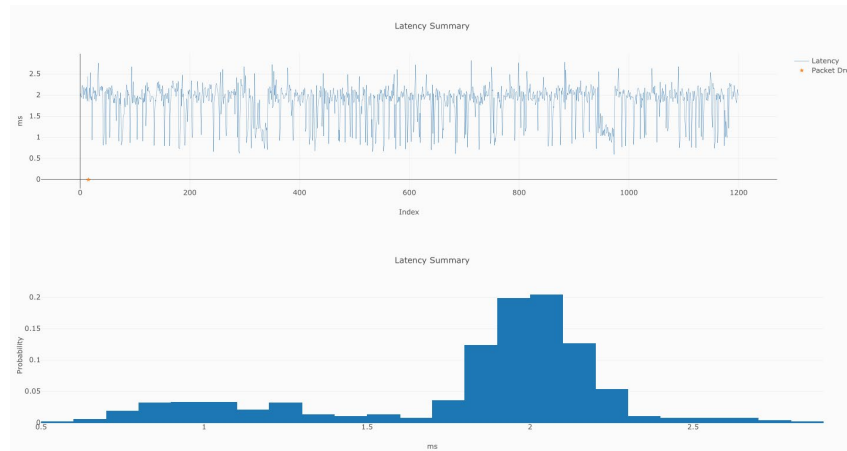


Figure 29 – 10GEPON latency – with Iperf TCP on load ONU

#### 4.6.3. DOCSIS TCP Test results

We ran a few of the TCP load set of test cases on a DOCSIS CMTS. Now due to various networking issues and corporate VLAN switching issues to connect to the DOCSIS CMTS in a different lab the number of test cases were limited and are left for a future testing effort.

In an older DOCSIS system configuration (without AQM configured or AQM Latency target configured to a high value) we see the latency with a TCP load test increases dramatically compared to the baseline DOCSIS latency seen in the previous sections. This is **no longer an acceptable configuration** and operators should ensure that AQM is enabled on all of their CMTS /CM devices.





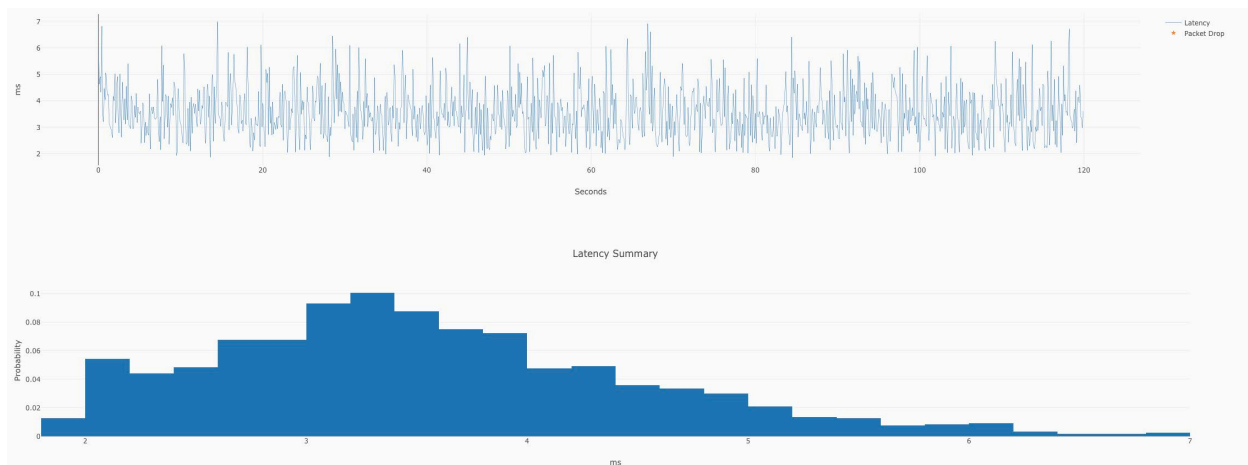
Figure 30 – DOCSIS no AQM config– with 1 Gbps TCP

In a D3.1 system correctly configured with AQM, (See figure below), we see the latency with a TCP load test does increase (the histogram shifts to the right by a millisecond and also extends a bit more) to 7-12 ms compared to the baseline DOCSIS latency (6-8 ms) seen in the section: Baseline DOCSIS Latency.



**Figure 31 – DOCSIS AQM Enabled Config, no LLD and Iperf TCP load**

In a D3.1 system configured with LLD, we see the latency with a TCP load test does increase, the histogram shifts to the right and is a bit more spread (1.9-5ms) compared to the baseline DOCSIS LLD latency (1.8-3 ms) seen in the previous section Baseline DOCSIS (with LLD) Latency.



**Figure 32 – DOCSIS LLD Enabled, Iperf TCP load**

#### 4.7. GPON/ EPON /DOCSIS Test Results summary

Looking at each of these above graphs and CDF's is very enlightening. The below table tries to summarize each of the tests that we ran by documenting three percentile values (0<sup>th</sup>, 50<sup>th</sup>, 99<sup>th</sup>), designated as P0 (minimum), P50 and P99. Please see the next three figures for a visual representation of the three tables. Please note that the DOCSIS testing was done on 3 different CMTS's/CMs, and as such these results from one DOCSIS CMTS cannot be compared with another due to change in the CMTS/CM hardware/Software versions and different plant configurations.

**Table 3 – GPON latency Test results , P0,P50,P99**

Scenario name	Reflector 4 (load ONU)	Reflector 3
		All units are ms
S0 Baseline	0.960, 1.557, 2.78	0.960, 1.550, 2.91
S1 UDP	1.019, 1.278, 1.534	0.948, 1.558, 2.93
S3 UDP	1.000, 1.177, 1.417	0.942, 1.550, 2.92
S5 UDP	1.000, 1.140, 1.340	0.961, 1.580, 2.99
S7 UDP	0.998, 1.150, 1.340	0.941, 1.580, 3.03
S9 UDP	2.840, 2.950, 3.482	1.670, 2.330, 4.023
S1 TCP	0.920, 1.280, 1.54	0.920, 1.569, 2.97
S3 TCP	0.961, 1.150, 1.36	0.945, 1.550, 2.90
S5 TCP	0.991, 1.140, 1.32	0.956, 1.580, 2.89
S7 TCP	1.005, 1.167, 1.35	0.960, 1.580, 2.98
S9 TCP	1.027, 1.227, 1.45	1.007, 1.640, 3.05
SU TCP (iPerf)	1.043, 1.667, 2.137	1.089, 1.900, 3.289

**Table 4 – EPON latency Test results , P0,P50,P99**


Scenario name	Reflector 2 (load ONU)	Reflector 1
		All units are ms
S0 Baseline	0.569, 2.132, 2.769	0.792, 1.384, 1.900
S1 UDP	0.547, 1.906, 2.125	0.673, 1.397, 1.932
S3 UDP	0.545, 1.925, 2.115	0.641, 1.419, 1.962
S5 UDP	0.522, 1.859, 2.041	0.656, 1.421, 1.957
S7 UDP	0.505, 1.824, 1.975	0.690, 1.400, 1.970
S9 UDP	1.118, 2.423, 2.548	0.990, 1.770, 2.341
		All units are ms
S1 TCP	0.549, 1.901, 2.126	0.633, 1.393, 1.927
S3 TCP	0.559, 1.899, 2.095	0.655, 1.389, 1.951
S5 TCP	0.421, 1.959, 2.724	0.625, 1.390, 1.947
S7 TCP	0.471, 1.950, 2.717	0.708, 1.393, 1.946
SU TCP (iPerf)	0.545, 2.122, 2.807	0.679, 1.403, 1.971

**Table 5 – DOCSIS latency Test results , P0,P50,P99**


Scenario	Reflector 5 (Different CMTS/CM ***)
S0 Baseline	2.80, 7.28, 11.27
S0 Baseline w LLD	1.76, 2.47, 3.834
UDP: No AQM **	89.3, 91.1, 91.97
UDP: with AQM	6.28, 15.5, 20.43
UDP with LLD	1.84, 3.42, 6.14
TCP: With AQM	3.11, 7.57, 10.64
TCP : With LLD*	1.94, 8.37, 12.625

*\* Prototype LLDCMTS software, not mature implementation. \*\* Recommend MSOs to always enable AQM. This test was just for historical purposes. \*\*\* 3 different CMTS/CM, so results cannot be compared.*





**Figure 33 – GPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU**



**Figure 34 – EPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU**

#### 4.8. Multiple LLIDs Test Results (EPON)

The next set of tests that we wanted to do on the 10GEPON system was the use of multiple LLIDs. We created some DPOE (DOCSIS Provisioning of EPON) configuration, which essentially classifies packets based on the incoming VLAN tags and forwards each of those packets on to a different LLID on the PON system. These LLIDs were configured on both the upstream and downstream direction.

The load test we ran was sending five different flows, 2Gbps each, in both the upstream and downstream directions which matched those VLANs. This essentially meant we were pushing the limits of the EPON system by sending 10 Gbps in each direction. Given that the capacity of the PON network was 8.5 Gigabit per second on the downstream and 8.4 Gigabit per second on the upstream we expect to see that the scheduler on the OLT equally shares the bandwidth between the five LLIDs, in both the upstream and downstream direction. In the test results we see that the average throughput on the downstream is about 1.712 Gbps for each of the 5 downstream flows. On the upstream we notice the first LLID almost has no drops with a rate of 1.968 Gbps and the remaining 4 flows share the remaining bandwidth with 1.585 Gbps each.


Flow	Source	Destination	Tx Frames	Rx Frames	Frame Loss	Tx Bytes (+VLAN)	Rx Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_DSCP_80_f7a0	NSI_4001	CPE_4001	20,000,000	17,398,331	13.01%	30,000,000,000 (+80,000,000)	26,097,496,500 (+69,593,324)	13.01%	2m, 2s, 243ms, 90µs, 737ns	1,712.46
DS_DSCP_112_f7a0	NSI_4002	CPE_4002	20,000,000	17,398,076	13.01%	30,000,000,000 (+80,000,000)	26,097,114,000 (+69,592,304)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.30
DS_0_f7a0	NSI_4000	CPE_4000	20,000,000	17,397,410	13.01%	30,000,000,000 (+80,000,000)	26,096,115,000 (+69,589,640)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.24
US_0_f7a0	CPE_4000	NSI_4000	20,000,000	19,996,993	0.02%	30,000,000,000 (+80,000,000)	29,995,489,500 (+79,987,972)	0.02%	2m, 2s, 251ms, 376µs, 496ns	1,968.11
UP_DSCP_80_f7a0	CPE_4001	NSI_4001	20,000,000	16,107,957	19.46%	30,000,000,000 (+80,000,000)	24,161,935,500 (+64,431,828)	19.46%	2m, 2s, 246ms, 146µs, 871ns	1,585.42
UP_DSCP_112_f7a0	CPE_4002	NSI_4002	20,000,000	16,106,642	19.47%	30,000,000,000 (+80,000,000)	24,159,963,000 (+64,426,568)	19.47%	2m, 2s, 248ms, 709µs, 251ns	1,585.25
DS_DSCP_160_f7a0	NSI_4003	CPE_4003	20,000,000	17,397,824	13.01%	30,000,000,000 (+80,000,000)	26,096,736,000 (+69,591,296)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.28
UP_DSCP_160_f7a0	CPE_4003	NSI_4003	20,000,000	16,106,557	19.47%	30,000,000,000 (+80,000,000)	24,159,835,500 (+64,426,228)	19.47%	2m, 2s, 244ms, 13µs, 4ns	1,585.31
DS_DSCP_104_f7a0	NSI_4004	CPE_4004	20,000,000	17,397,827	13.01%	30,000,000,000 (+80,000,000)	26,096,740,500 (+69,591,308)	13.01%	2m, 2s, 237ms, 765µs, 881ns	1,712.49
UP_DSCP_104_f7a0	CPE_4004	NSI_4004	20,000,000	16,105,553	19.47%	30,000,000,000 (+80,000,000)	24,158,329,500 (+64,422,212)	19.47%	2m, 2s, 250ms, 278µs, 316ns	1,585.13

**Figure 35 –10GEPON throughput (5 LLIDs 2Gbps per LLID)**

Using the measurements done by the traffic generator we observe that the average one-way latency on the downstream for each of the LLIDs is 0.291 ms, while in the upstream direction the one-way latency is 0.781 ms with the first Upstream LLID with a latency of 0.830ms.

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_DSCP_80_f7a0	NSI_4001	CPE_4001	0.026	0.291	0.307	0.005
DS_DSCP_112_f7a0	NSI_4002	CPE_4002	0.025	0.291	0.307	0.005
DS_0_f7a0	NSI_4000	CPE_4000	0.026	0.291	0.307	0.005
US_0_f7a0	CPE_4000	NSI_4000	0.328	0.830	1.814	0.153
UP_DSCP_80_f7a0	CPE_4001	NSI_4001	0.315	0.781	1.850	0.153
UP_DSCP_112_f7a0	CPE_4002	NSI_4002	0.320	0.781	1.789	0.153
DS_DSCP_160_f7a0	NSI_4003	CPE_4003	0.027	0.291	0.307	0.005
UP_DSCP_160_f7a0	CPE_4003	NSI_4003	0.307	0.781	1.837	0.153
DS_DSCP_104_f7a0	NSI_4004	CPE_4004	0.025	0.291	0.307	0.005
UP_DSCP_104_f7a0	CPE_4004	NSI_4004	0.326	0.781	1.793	0.153

**Figure 36 –10GEPON Latency (5 LLIDs 2Gbps per LLID)**



**Figure 37 –10GEPON visual of Throughput & Latency (for one US and DS LLID)**

#### 4.9. Challenges and Further testing

This was the first time we were setting up an 10GEPON and XGSPON in our labs. So as expected there was a very steep learning curve in all aspects of the setup configuration and operation of the PON network. Once we got past setting up the physical network using the needed attenuators (10 to 15 dB), the main challenge was related to VLAN configuration. These SFP+ OLTs were connected to a switch which needed careful configuration of VLANs, to make sure we were bridging the traffic to the devices it needed to go. We also needed to build up the lab capabilities with optical to copper 10 Gbps port converters, 10 Gbps NIC cards on servers and clients etc. Learning the provisioning of services on these GPON and EPON systems and getting the right configurations in was another step in the learning process. Also getting packets to flow through the corporate network and reach the DOCSIS CMTS in a different lab had its own set of challenges and limited the testing to some extent.

In terms of future testing, we would like to study in more detail the performance of TCP flows on these PON systems. For an XGSPON system, we would like to enable and test multiple flows (AllocIDs in XGSPON) in each ONU. We would like to do more detailed testing of the DOCSIS system with LLD configured and varying the different configuration parameters. We would also like to measure the performance of TCP flows and the latency, when the measurement agents are located much farther in the network (i.e., adding additional round trip time or delay in the network path). We would also like to enable testing of PON systems with a larger number of ONUs (more than the two that we had currently on each system). We would also like to mix and match different on ONUs with different OLTs to understand the performance. There are many scheduling parameters which can be tweaked on the OLT scheduler, this includes parameters like polling interval, minimum and maximum granting period etc., and we would like to understand the differences in performance by varying each of these parameters.

### 5. Conclusions

10GEPON and XGSPON networks show remarkably good performance with 8.6/8.5 Gbps symmetrical throughput. Both PON systems show consistent latency performance on an unloaded or loaded PON network. The majority of the latency numbers fall within 1 to 2 milliseconds on a lightly loaded PON network, while we observe latencies up to 2 to 3.5 milliseconds under heavily loaded conditions. (Note that this loaded condition is for the home network i.e. this ONU's load, whereas when the whole serving group is heavily loaded and each ONU is transmitting above the minimum guaranteed rate, the latencies will likely increase.) At least in the configurations we tested it looks like 10GEPON has slightly lower latencies than the XGSPON setup we tested, but given they are all in the 1-3 ms range they both are very similar latency ranges. (The 10GEPON did have more guaranteed bandwidth than the XGSPON setup). We tested the network using UDP load traffic and TCP load traffic and the latency numbers stayed relatively consistent and in the same range. We also compared the latency in a DOCSIS network and see that while without Low Latency DOCSIS(LLD) features, the latency in a DOCSIS network can vary from 6-7 ms (baseline) to ~15 to 30 ms (under load). Enabling LLD brings the latency numbers to be at comparable levels, 2-3 ms (baseline) to 2-6 ms (under load), to that of a PON network. As latencies drop to the sub-5ms levels, the connection to content servers (outside of access network) through the Internet can become bigger contributors to overall end-to-end latency.

*All testing was done with limited variations from default out-of-the-box configurations, varying the configuration parameters(for each of EPON, GPON, DOCSIS) may provide more variation in the results, leading to further analysis and understanding.*

## Abbreviations

10GEPON	IEEE 10 Gigabit EPON (Ethernet Passive Optical Network) technology
XGSPON	ITU-T 10 Gigabit Symmetric PON (Passive Optical Network) technology
RTT	round trip time
CDF	cumulative distribution function
DSCP	Diff Serv Code Point
ms	millisecond
LLD	Low Latency DOCSIS
STAMP	Simple Two-Way Active Measurement Protocol
Gbps	Gigabits per second
OLT	Optical Line Terminal
ONU	Optical Networking Unit
CMTS	Cable Modem Termination System
CM	Cable Modem
Pps	Packets per second
NSI	Network Side interface (i.e., Interface north of the CMTS or OLT)
CPE	Customer Premise Equipment ( i.e., device behind CM or ONU

## Bibliography & References

[ITU-T G.9807.1] 10-Gigabit-capable symmetric passive optical network (XGS-PON).

[IEEE 802.3] IEEE Std 802.3™-2018 IEEE Standard for Ethernet, Section 5 (10GEPON)

[LM SCTE 20] Latency Measurement: What is latency and how do we measure it? Karthik Sundaresan, Greg White, Steve Glennon, SCTE 2020

[STAMP SCTE21] A Latency Measurement System Using STAMP, Karthik Sundaresan, SCTE 2021

[C3 CableLabs] CableLabs Common code community, <https://community.cablelabs.com/wiki/display/C3>