

## **Turbocharging DOCSIS 3.1 Technology: An Incremental Step on the Way to DOCSIS 4.0**

A Technical Paper prepared for SCTE by

**Colin Howlett**

Chief Technology Officer  
Vecima  
[colin.howlett@vecima.com](mailto:colin.howlett@vecima.com)

**Jay Rolls**

Chief Technology Officer  
Broadband Success Partners  
[jrolls@broadbandsuccess.com](mailto:jrolls@broadbandsuccess.com)

## Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. Network Evolution .....	4
3. DOCSIS 3.1 Minimum Capability .....	5
4. DOCSIS 4.0 Outlook and Considerations .....	5
5. Turbocharging DOCSIS 3.1 Technology .....	7
5.1. Distributed Access Architecture (DAA) .....	8
5.2. Profile Management Application (PMA) .....	9
5.3. Low Latency DOCSIS (LLD) .....	10
5.4. DOCSIS 4.0 Cable Modems in 1.2 GHz Plant.....	13
6. Operational Considerations.....	14
6.1. Spectrum Reclamation – Sunset QAM Video .....	14
6.2. Overlapping OFDMA Channels.....	15
6.3. Leakage Detection .....	17
7. Upgrade Costs .....	18
8. Conclusion.....	19
Abbreviations .....	20
Bibliography & References.....	21

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - DOCSIS 4.0 Frequency Split Options .....	7
Figure 2 - Distributed Access Architecture Examples.....	9
Figure 3 - PMA Deployment Architecture Example from [Karthik1].....	10
Figure 4 - Bandwidth Usage for Capacity Seeking Applications with TCP .....	11
Figure 5 - Dual-Queue Approach for LLD from [White1] .....	12
Figure 6 – Round-Trip Latency with LLD for NQB-Marked Traffic from [White1] .....	13
Figure 7 - Capacity by Upstream Split and Plant Maximum DS Frequency with D4.0 CM .....	14
Figure 8 – Need for Overlapping OFDMA Channels .....	16
Figure 9 – Overlapping OFDMA Channel Arrangement .....	16
Figure 10 – Overlapping OFDMA Channel Solution – UCD and MAPs .....	17
Figure 11 - Leakage Detection Test Signals from [Coldren1].....	18

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – DOCSIS 3.1 HFC Outside Plant Minimum Capabilities .....	5
Table 2 – DOCSIS 4.0 Capabilities.....	7
Table 3 – DOCSIS Cable Modem Minimum Capabilities .....	13
Table 4 – Relative Upgrade Costs .....	19

## 1. Introduction

In North America, cable operators, delivering broadband over their hybrid fiber-coax (HFC) infrastructure, have become leading providers of broadband in the communities that they serve. It's not uncommon for a multi-system operator (MSO) to enjoy 65% market share. And subscriber additions accelerated during the first 18 months of the pandemic.

However, MSOs are now facing growing competition, in multiple forms across their service areas.

These threats mainly originate in three areas. First, fiber competition, both from large telco incumbents such as AT&T, Verizon, and Lumen, and Frontier and small upstart fiber operators. Second, the Mobile Network Operators (MNOs) have been pushing broadband over their mobile platforms very heavily over the past 24 months. (A sub-component also arises from Fixed Wireless Access providers but is a much smaller percentage of wireless broadband). And lastly, a not insignificant industry has materialized that specializes in just servicing multi-dwelling units (MDUs), and that is normally also served via fiber. MDUs can often represent 30-40% of an operator's customers.

Amidst this burgeoning competition, MSOs and their HFC architecture can point to an admirable record of remaining resilient and competitive. In fact, the limitations of the evolving DOCSIS standards have always been discussed, and those limitations have always been shattered.

Up to this point, none of the previous versions of DOCSIS<sup>®</sup> (1.0, 2.0, 3.0, and 3.1) specifications required any major changes to an operator's Outside Plant<sup>1</sup>. Implementations have mainly relied on changes to electronics on either end of the 'wire'; either in the headends and hub sites (Inside Plant), or in the home (the cable modem).

Outside Plant changes have been operator-specific and generally driven by overall increases in frequency capability operating within the limits of existing active device (amplifier) housings. These changes have included:

- DS frequency migrations from 550 MHz to 750 MHz, 860 MHz, 1 GHz or even 1.2 GHz by updating node and amplifier modules
- US frequency extension from 42 MHz to 85 MHz or in rare cases to date, 204 MHz, by updating amplifier modules with new frequency diplexers

However, that is about to change.

DOCSIS 4.0<sup>2</sup>, which will be introduced between 2023 and 2025, will require changes to the outside plant. Introducing this new version of DOCSIS, which will activate frequencies all the way up to 1.8GHz, is going to require an overhaul to today's HFC cable plant. This will be a far more invasive and expensive undertaking. Every device in the outside plant will need to be

---

<sup>1</sup> Perhaps this overlooks the initial 2-way activation of the cable plant in the 1990's, which was a precursor to offering broadband services.

<sup>2</sup> When discussing DOCSIS 4.0 specifications, this paper will focus on the part that many operators are planning to implement: ESD, or Extended Spectrum DOCSIS, which bumps the cable plant spectral capabilities up to 1.8GHz.

replaced (or in more rare occasions, upgraded). To put this in perspective, using an example of a rather large “fiber node” with a 350 to 400 home serving area, this equates to replacing between 10 and 20 amplifiers, between 10 and 25 splitters, and over 100 coaxial cable taps. The cost to undertake such a retrofit will start to be driven more by the labor involved than the equipment itself.

In this paper, we reference “turbocharged” DOCSIS 3.1 technology as a collection of specific improvements that have been added to DOCSIS 3.1 and DOCSIS 4.0 specifications since the original DOCSIS 3.1 specification was published:

- New deployment architectures (Distributed Access Architecture or DAA)
- Operation with varying impairments at each customer (Profile Management Application or PMA)
- Application latency improvements (Low Latency DOCSIS or LLD)
- Use of enhanced DOCSIS bonding group channel counts in DOCSIS 4.0 cable modems along with higher speed Ethernet interfaces

The “turbocharged” collection of improvements is not a separate specification, but a useful way to categorize an incremental step possible for operators before they deploy networks capable of realizing all parts of the DOCSIS 4.0 specifications. This step is a valuable way to achieve the next step in DOCSIS performance at lower cost than both full DOCSIS 4.0 network deployment and far lower cost than full fiber to the home rollouts.

## **2. Network Evolution**

Broadly speaking, there are a handful of potential technology paths for evolving an HFC network. In between and amongst those options, many smaller choices will need to be made.

The two most significant factors to consider in meeting growing demands on a broadband network are capacity (scale) and speeds. Capacity is the ability to meet the growth in the number of subscribers on the network, as well as the growing consumption needs of each subscriber. Speeds relate to speed tier offerings, and in particular the top speed offered on the network.

To meet the demands outlined above, operators have these options to consider, as they evolve their network:

- Create capacity by segmenting the shared nature of an HFC network into smaller parts. Most often this is facilitated by splitting neighborhoods into smaller service groups, via node segmentations and node splitting.
- Expand the capacity of the network by increasing the breadth of frequencies supported. Choices here include expanding today’s 1GHz plant to 1.2GHz and/or to 1.8GHz.
- Leveraging technology to get more out of the network at hand. This includes running more of the signaling with advanced modulation (orthogonal frequency division multiplexing or OFDM), running higher modulation 2k and 4k QAM, and leveraging systems that can operate on more channels.
- Lastly, changing to a different network altogether; in this case fiber to the premises (FTTP).

### 3. DOCSIS 3.1 Minimum Capability

DOCSIS 3.1, introduced to field deployment in ~2015 (see [MULPIv3.1] and [PHYv3.1]), has been able to keep pace with advancements needed in capacity and speed to serve both increasing customer usage (widespread uptake in streaming video platforms, proliferation of consumer devices to name a couple) and competition from operators widely deploying FTTP in the same service area.

The current HFC outside plant is faced with several constraints. Today's 1GHz and below plant will eventually create a capacity constraint. And the current low split (~42MHz in North America) configuration, leads to high bandwidth asymmetry, with far more capacity in the downstream than in the upstream. Loosely this is along the lines of a 10:1 ratio.

Upstream speed for subscribers is limited based on frequency split, the condition of the plant (achievable US modulation) and need to support legacy DOCSIS 3.0 and earlier Advanced Time Division Multiple Access (ATDMA) channels and is summarized in Table 1. Operators looking to move to support gigabit upstream (~1 Gbps) speeds are now actively testing, trialing, or deploying High Split networks to move the upstream upper band edge frequency to 204MHz and the downstream lower band edge frequency to 258 MHz.

In the downstream direction, DOCSIS 3.1 cable modems have a minimum bonding group of 32 channels of single carrier QAM (SC-QAM) and two 192 MHz orthogonal frequency division multiplexing (OFDM) channels. With 256QAM modulation for SC-QAM and 4096QAM modulation for DS, the maximum throughput of this bonding group is 32 x ~35 Mbps + 2 x ~1.7 Gbps or 4.5-4.6 Gbps. The availability of spectrum for this total capacity depends on maximum DS frequency limit and the amount of spectrum used to delivery QAM video service, both of which are operator- and location-specific.

Commercially available DOCSIS 3.1 cable modems today are limited to supporting the minimum bonding group noted above and are also limited to maximum Ethernet interfaces using 2.5 Gigabit Ethernet (2.5GE). The combination of these commercially available minimum implementations results in service offerings with the values summarized in Table 1:

**Table 1 – DOCSIS 3.1 HFC Outside Plant Minimum Capabilities**

Upstream Frequency Upper Band Edge	Upstream Capacity	Maximum Upstream Service Tier	Downstream Bonding Group Capacity	Maximum Downstream Service Tier
Low Split (42MHz)	100-250 Mbps	50-100 Mbps	4.5-4.6 Gbps	3 Gbps
Mid Split (85MHz)	~400-550 Mbps	250 Mbps	4.5-4.6 Gbps	3 Gbps
High Split (204MHz)	~1.3-1.6 Gbps	1 Gbps	4.5-4.6 Gbps	3 Gbps

### 4. DOCSIS 4.0 Outlook and Considerations

The next version of DOCSIS, 4.0, is due to see equipment released between 2023 and 2025.

The primary objectives in DOCSIS 4.0 as the next version include:

- Support for customer upstream service tiers beyond 1 Gbps
- Support for customer downstream service tiers beyond 2 Gbps

The DOCSIS 4.0 standard (see [MULPIv4.0] and [PHYv4.0]) presents two different implementation options depending on an operator's objectives. The options are:

1) FDX - 1.2 GHz Full Duplex (FDX)

FDX is generally intended for Node+0 HFC with no active amplifiers but industry initiatives are underway to build FDX repeater amplifiers to stretch operation into networks including amplifiers. FDX targets maximum DS frequency with the same 1218 MHz limit as DOCSIS 3.1.

2) FDD/ESD - 1.8 GHz Frequency Division Duplex (FDD)

FDD/ESD is generally intended as an incremental upgrade to a traditional HFC network of actives and passives with an arbitrary number of amplifiers. This second option is commonly referred to as ESD, or Extended Spectrum DOCSIS. Industry demonstrations in 2022 have shown ESD operation to be compatible with traditional amplifier cascades and plant designs. This paper specifically focuses on ESD as an incremental path from traditional DOCSIS 3.1 to DOCSIS 4.0 technology.

Upgrading an HFC cable plant to 1.8GHz will be challenging. The signal attenuation on a cable plant to signals above 1.2GHz are significant. The passive components currently deployed are designed for 1GHz or 1.2GHz and may extend slightly higher, but many of these components will need to be changed out to support operation to the 1.8GHz maximum capability of the DOCSIS 4.0 FDD specification:

- Taps required to direct a portion of the signal energy to/from specific subscribers
- Power inserters for feeding AC into the plant
- Directional couplers and splitters used within the trunk portions of the coaxial


Likely some design changes will be required to other parts of the HFC cable plant. Some areas may require smaller booster amplifiers, estimated to be in the range of 10-20% in early design studies from multiple North American operators. In some cases (less than 5%), older deployed coax cable itself might not support 1.8GHz, requiring new coax to be deployed. These are areas of study that will be necessary for an operator to perform as each company assesses their readiness for DOCSIS 4.0, and their ability to meet the challenges that will arise.

Since the outside plant is changing and frequency is being pushed higher than what DOCSIS has done before in the HFC network, the implementation of DOCSIS 4.0 will require more extensive lab and field testing than previous versions. Many operators are likely to take a "wait and see" approach to DOCSIS 4.0 to see how early field tests and deployments work before making decisions on the future of the HFC plant.

Even when deciding to move forward, with DOCSIS 4.0 deployment, expectations are that operators will progress through multiple incremental phases of DOCSIS 4.0 ultra-high frequency



splits (UHS) as capacity and speed requirements increase. These frequency split options are shown in Figure 1:



**Figure 1 - DOCSIS 4.0 Frequency Split Options**

Note that the DOCSIS 4.0 standard also includes other splits but these are not included based on operator feedback on probable deployments:

- UHS-300 (300/372MHz) offers only a small improvement above 204 MHz so operators are much more likely to switch to UHS-396 or UHS-492 as the next increment
- UHS-684 (684/834MHz) has no additional downstream bandwidth while requiring more spectrum and providing very high peak upstream rates that are unlikely to be used by service offerings

DOCSIS 4.0 capabilities for these tiers are summarized in Table 2 based on up to 900 Mbps per 96 MHz OFDMA (700 Mbps in low band) and 1.8 Gbps per 192 MHz OFDM channel:

**Table 2 – DOCSIS 4.0 Capabilities**

Upstream Frequency Upper Band Edge	Upstream Capacity	Maximum Upstream Service Tier	Downstream Bonding Group Capacity	Maximum Downstream Service Tier
High Split (204MHz)	1.6 Gbps	1 Gbps	14.4 Gbps	10 Gbps
UHS-300 (300MHz)	2.6 Gbps	2 Gbps in great plant	13.3 Gbps	10 Gbps
UHS-396 (396MHz)	3.5 Gbps	2 Gbps	12.2 Gbps	10 Gbps
UHS-492 (492MHz)	4.4 Gbps	3 Gbps	11.1 Gbps	8 Gbps
UHS-684 (684MHz)	6.2 Gbps	5 Gbps in great plant	9.0 Gbps	6 Gbps

## 5. Turbocharging DOCSIS 3.1 Technology

Since the first field deployment of DOCSIS 3.1 networks in ~2015, cable operators and vendors have continued to innovate to improve the performance and operation of DOCSIS-based systems. Bundled together, these enhancements on top of the original DOCSIS 3.1 turbocharge DOCSIS 3.1 technology as a concept of a noticeably improved DOCSIS implementation suited to long term deployment, and as an incremental step towards full DOCSIS 4.0 deployment.

These options available in the DOCSIS 3.1 specifications, if leveraged efficiently, can instantly turbocharge any DOCSIS 3.1 deployment:

- High Split frequency split – full support of the upstream capability specified in DOCSIS 3.1 specifications to enable gigabit upstream service tiers

- Distributed Access Architecture (DAA) – extend all-digital Ethernet fiber deep into the network for reduction in hub space and power needed for an increased number of service group, along with improvements in RF performance
- Profile Management Application (PMA) – optimized subcarrier modulation in downstream and upstream to improve channel capacity for impaired modems or when overall RF plant conditions are non-optimal
- Low Latency DOCSIS (LLD) – reducing overall end-end packet delay through optimized queuing and standardized packet marking, especially for applications such as gaming
- DOCSIS 4.0 cable modems in 1.2 GHz plant – taking advantage of extra OFDM and OFDMA channels likely to come in DOCSIS 4.0 CMs to increase overall bonding group capacity in both DS and US while also supporting service tiers above 2 Gbps in the DS

### **5.1. Distributed Access Architecture (DAA)**

The cable industry has long relied on a centralized architecture in configuring traditional HFC networks. Signals that were broadcast on the cable plant originated in centralized headend and hub site facilities. This practice continued with the introduction of broadband services delivered via the DOCSIS standard. However, the industry is in the middle of a major transition away from this approach, moving electronics from centralized to more distributed locations, with key elements (electronics) being placed in the outside plant fiber nodes. The broad acronym in use to describe this new approach is DAA – Distributed Access Architecture.

A primary initial use case for DAA was to reduce hub space and power by moving the modulation and demodulation of physical layer (PHY) signals (QAM and OFDM) from the hub to the node location. This reduction is necessary to avoid hub expansions as the number of subscribers per service group continues to shrink to provide more capacity and speed to individual subscribers so more DOCSIS and RF equipment is needed. DAA allows the removal of RF equipment including RF combiners, analog/digital optical RF systems, and the extensive wiring needed to connect everything together. Instead, all hub DOCSIS-related equipment is now digital with simplified Ethernet/IP interconnects or removed altogether in the case of Remote-MACPHY architectures.

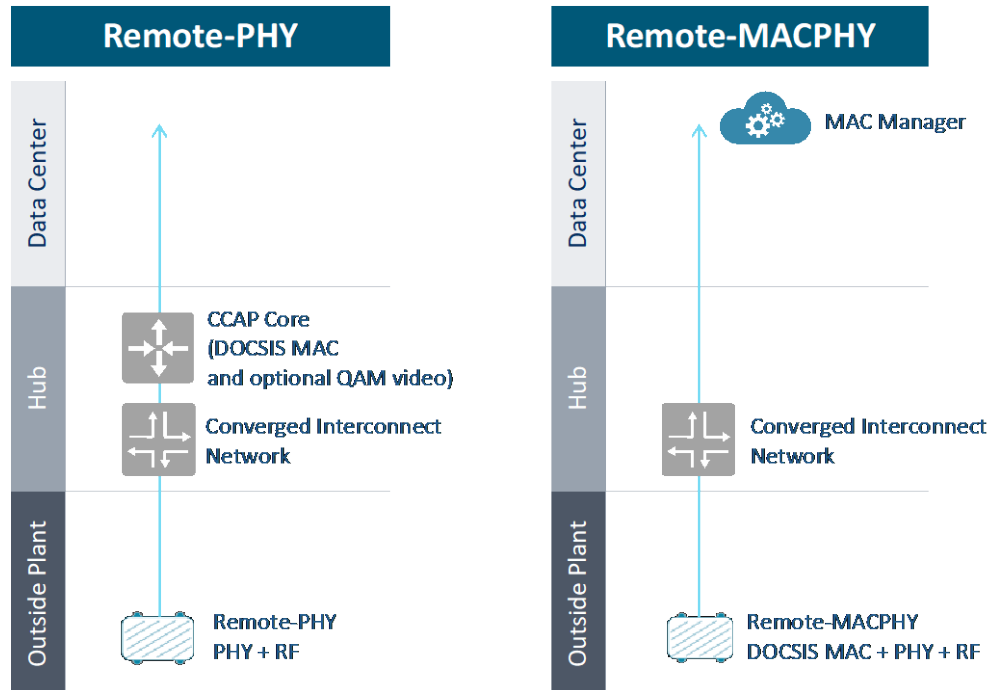
Along with the reduction in hub space and power, one of the most important advantages of this new approach is that it allows operators to move from “analog optics” to “digital optics” – thus simplifying how signals are transmitted on the fiber portion of an HFC network, and greatly increasing the fidelity of the signals that are transmitted. Field experience from a number of DAA deployments has shown 5 to 7 dB improvement in C/N (carrier to noise ratio), which in essence represents a doubling in RF performance. Such an improvement allows for higher modulation (more bits per hertz), thus increasing overall system capacity.

The advent of DAA also holds promise in addressing some operational problems. Some areas that might be operating on the margin today (ie. Signal quality), could see improved performance from a DAA implementation and reduced trouble calls.

DAA implementations come in two primary configurations specified by CableLabs under the umbrella of Distributed CCAP Architecture (DCA) (see [DCA]). These configurations are



known as Remote-PHY, and Remote-MACPHY. Both have the advantage of moving to digital optics. Both can serve to allow for better RF performance out of the cable plant. Entire technical papers have been devoted to the merits of one over the other, with details being beyond the scope of this analysis. Figure 2 shows a high level view of the DAA architectures.




**Figure 2 – Distributed Access Architecture Examples**

One extremely important consideration, accepted by the industry while developing the specifications, is that DAA architecture will be required to implement DOCSIS 4.0. DAA will be paramount in addressing challenges with supporting spectrum as high as 1.8GHz. As an incremental step towards DOCSIS 4.0, turbocharging DOCSIS 3.1 technology assumes the use of DAA as well.

## 5.2. Profile Management Application (PMA)

OFDM DS and OFDMA US channels as specified in [MULPIv3.1] and [PHYv3.1] support the use of multiple “profiles” which allow for different data modulations (256QAM, 1024QAM, 4096QAM, etc.) to be configured for each modulated subcarrier. The use of these profiles allow the HFC system performance, in both robustness and overall capacity, to be optimized to the current conditions in the network.

The Profile Management Application (PMA) (see [PMA-TR]) as shown in Figure 3 is an external software solution which uses data on receiver MER and codeword errors from cable modems and CMTS, along with sophisticated algorithms and significant server processing power to optimize the set of available profiles to match current conditions. Without the dynamic changes in available profiles through PMA, modems that are temporarily impaired or operating outside normal design targets would generally use a default modulation of 256QAM.



**Figure 3 - PMA Deployment Architecture Example from [Karthik1]**

Since the cable plant channel model is quasi-static with slow variations over time due to temperature effects or specific impairment events, PMA solutions update profiles on a periodic interval measured in minutes.

The types of impairments or plant conditions that benefit from PMA include:

- Ingress from wireless communications within the cable band
- Operation in the rolloff region above the current plant frequency. This is particularly interesting as part of turbocharging DOCSIS 3.1 technology since operators can use PMA to extend into a rolloff region (say 1.1 GHz for 1 GHz tap and gain several hundred Mbps of additional capacity without the labor costs of swapping taps in the field for 1.2 GHz.
- Standing waves from impedance mismatches, grounding issues, or other non-ideal operation

Gains achievable, as seen in [Karthik1] are based on a significant set of operator data across several geographies. Capacity gains in excess of 30% improvement over a default 256QAM profile are achievable using PMA techniques.

### **5.3. Low Latency DOCSIS (LLD)**

Many applications such as gaming, interactive videoconferencing, and web browsing do not require significant capacity, but instead require timely responses from remote servers to provide the best user experience. Low Latency DOCSIS technology (LLD), as part of [MULPIv3.1], has been developed by CableLabs, cable operators, and vendors to improve overall application latency, especially for those applications that are not using significant bandwidth.

The primary problem to solve from an end-to-end perspective relates to applications and transport protocols (e.g. TCP) which are capacity seeking. High bandwidth applications like file transfers and streaming IP video run over TCP, using up available network bandwidth until they

experience congestion. Congestion is recognized by way of dropped packets, at which time the application backs off until congestion ceases and then ramps back up until congestion is experienced again, at which time they once again back off. This cycle is repeated, creating a sawtooth pattern of bandwidth usage as shown in Figure 4.



**Figure 4 - Bandwidth Usage for Capacity Seeking Applications with TCP**

The problem this creates for applications that need low latency is that their packets get queued up during periods of congestion and experience latency and jitter patterns that follow the sawtooth bandwidth usage pattern of the capacity seeking application protocols. The industry has settled on the terms "queue building", to describe capacity seeking applications and transports, and "non-queue building" (NQB), to describe the transport that is needed by low latency applications.

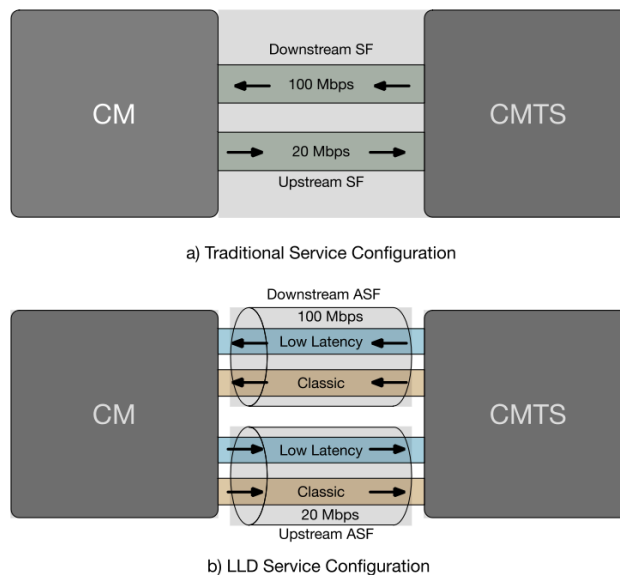
Early methods for dealing with queue building applications included creating larger queues in network devices to absorb bursts without loss of packets. However, these methods were easily defeated by applications which just continued to ramp up bandwidth demands until packet loss was experienced. This made delay and jitter even worse and was referred to unkindly as "buffer bloat".

The solutions included in the overall LLD ecosystem include the following. Many of these solutions are just now progressing to readiness in the broader consumer ecosystem including gaming platforms and major device operating systems:

- A new NQB packet marking that applications can use to indicate that they are non-queue building so that their traffic can be treated differently in the network (e.g. classified into a Low Latency Service Flow), see [NQB1]
- Recommended use of Explicit Congestion Notification (ECN) and ECN Capable Transport (ECT) markings in the Traffic Class field of the IP header by low latency applications and networking equipment (e.g. routers), see [ECN1]
- Support for a new congestion control scheme called Low Latency, Low Loss, Scalable Throughput (L4S) (see [L4S1] and [L4S2]), which leverages Active Queue Management (AQM) techniques and which is intended to be applied in next generation transport

protocols supporting low latency applications. L4S congestion controls are leveraged in low latency applications and in network equipment as follows:


- Low latency applications will mark ECT in their packets
  - Network equipment experiencing congestion involving that ECT traffic will mark ECN Congestion Experienced (CE) in packets before forwarding rather than dropping packets
  - Low latency applications supporting ECT and receiving ECN CE will respond by marking ECN CE in traffic toward the originating application
  - The originating application will respond by reducing its transmission rate
- DOCSIS queuing improvements address queueing delay by allowing applications to avoid waiting behind the delays caused by the current TCP or its variants. At a high level, the low-latency architecture consists of a dual-path approach that treats both queues as a single pool of bandwidth as shown in Figure 5:



**Figure 5 - Dual-Queue Approach for LLD from [White1]**

- DOCSIS scheduler improvements address media acquisition delays by:
  - Lowering the request-grant delay with a shorter MAP Interval broadcast by the CMTS and associated MAP Processing Time in the cable modem
  - Adding a new scheduler service known as Proactive Grant Service (PGS) to proactively grant to a service flow without incurring the request-grant delay

Simulations published in [White1] show the dramatic improvements in round-trip latency and especially consistency of round-trip latency possible for NQB-Marked Traffic. These results show the 99<sup>th</sup> percentile of traffic having round-trip latency (DOCSIS part of network) well below 10ms without PGS and below 1ms with PGS, results that are 2-3 orders of magnitude better than standard DOCSIS 3.1.



**Figure 6 – Round-Trip Latency with LLD for NQB-Marked Traffic from [White1]**

#### 5.4. DOCSIS 4.0 Cable Modems in 1.2 GHz Plant

One key change in the DOCSIS 4.0 specifications is an increase in the minimum number of OFDM DS and OFDMA US channels that can be bonded by the cable modem as shown in Table 3.


**Table 3 – DOCSIS Cable Modem Minimum Capabilities**

Modem Generation	Upstream OFDMA Channels	Downstream Channels
DOCSIS 3.1	2 x 96 MHz	32 x SC-QAM + 2 x 192 MHz OFDM
DOCSIS 4.0	7 x 96 MHz	32 x SC-QAM + 5 x 192 MHz

In addition, commercial DOCSIS 4.0 cable modems are expected to shift from 2.5GE Ethernet interfaces to 10GE Ethernet interfaces to take advantage of the bonding group capacity to deliver very high single cable modem speeds and enable new service tiers not possible today.

As shown in Figure 7, leveraging this capability without incurring the extra costs of full 1.8 GHz outside plant changes allows an operator to provide a premium service tier of 5 Gbps or higher instead of the 3 Gbps limits resulting from maximum bonding in DOCSIS 3.1.

The UHS-396 with 1.2 GHz plant and UHS-492 with 1.2 GHz plant options in Figure 7 are unlikely to be interesting for deployment unless upstream usage and customer need changes significantly from current projections. The overall DS capacity is limited in 1.2 GHz plant. In addition, amplifiers capable of supporting UHS-396 and UHS-492 diplexers are likely 1.8 GHz capable.



**Figure 7 - Capacity by Upstream Split and Plant Maximum DS Frequency with D4.0 CM**

## 6. Operational Considerations

Several operational considerations need careful attention from operators to ensure the largest benefit can be achieved from turbocharging DOCSIS 3.1 technology. These include:

- Spectrum Reclamation – transitioning DS spectrum from QAM video to DOCSIS data use
- Overlapping OFDM Channels – ensuring Mid Split and High Split modems can utilize their full available bandwidth when mixed in the same service group
- Leakage Detection – ensuring regulatory compliance when sensitive aeronautical bands in the 130 MHz range move from being DS signals to US signals as part of High Split transitions

### 6.1. Spectrum Reclamation – Sunset QAM Video

A key service still provided by most operator networks today is QAM digital video delivery to traditional set-top boxes. The number of channels used is dependent on market and QAM video technology (i.e. broadcast vs. switched digital video (SDV)) but can occupy 30-60 or even 80 channels.



Concerns that arise from keeping QAM video available in the network:

- Insufficient spectrum to transition to High Split

In this case, the transition from low or mid split to high split requires freeing up 150 MHz (108-258 MHz for mid to high split) or 204 MHz (54-258 MHz) of downstream spectrum. If downstream spectrum is relatively full already due to multiple OFDM channels or outside plant DS frequency limits, this spectrum reclamation may not be possible.

- Insufficient DOCSIS spectrum for service tier goals

This case is a concern when pushing service tiers to 2 Gbps and above, especially if using DOCSIS 4.0 cable modems to increase DS bonding group size. QAM video spectrum needs to be reclaimed to achieve maximum capacity.

- Legacy set top box carriers – SCTE 55-1 or SCTE 55-2


Support for operation of the legacy set top box DS carriers above 130 MHz for SCTE 55-1 or 55-2 is a possibility with some set-top boxes but this varies by the specific set top box. Adaptation solutions do exist to deal with these issues: downconverters at the STB from ~250 MHz to 70-130 MHz for seamless operation or use of DOCSIS Set Top Gateway (DSG) to out-of-band (OOB) carrier generators can help.

The primary solution for Spectrum Reclamation is transition to IP video service while sunsetting QAM video except possibly for niche use cases. This transition, already done at many operators, with the full digital QAM video lineup available in IP format to stream as data over DOCSIS channels, also helps operators deliver video service to FTTP customers without the complications associated with RF over the fiber.

## **6.2. Overlapping OFDMA Channels**

The need for Overlapping OFDMA Channels (OOC) comes from DOCSIS 3.1 requirements and common implementations which limit the number of available OFDMA channels in an Upstream Service Group to two.

Deployment of High Split ideally has one of the two OFDMA channels placed at 108 to 204 MHz, and the other OFDMA channel placed immediately below that (e.g., 12 to 108 MHz). The problem is that many operators have fielded Mid-Split DOCSIS 3.1 CMs, which have a diplexer at 85 MHz and which are unable to make use of any OFDMA channel that spans 85 to 108 MHz (or which goes above 85 MHz, to be precise). See Figure 8 for further details on the concern.




**Figure 8 – Need for Overlapping OFDMA Channels**

As a result, in moving to High-Split operators may be forced to downgrade service levels for subscribers with Mid-Split CMs as these CMs would no longer be able to use any OFDMA. Conversely, the operator would be forced to place one of the OFDMA channels completely below 85 MHz (e.g. 12 to 85 MHz), which leaves spectrum from 85 to 108 MHz vacant.


To make matters worse, most operators are reluctant to deploy time and frequency division multiplexing (TaFDM) which overlaps ATDMA and OFDMA channels and ensures that there are no overlaps in the scheduler. If avoiding using TaFDM, these operators would carve out explicit spectrum for ATDMA channels below 85 MHz and explicit spectrum for the OFDMA channel below 85 MHz, making the OFDMA channel that exists there very small.

DOCSIS 4.0 CMs are also not required to be able to make use of 85 to 108 MHz. This leads to an analogous problem for those CMs when Ultra-High Split is in play (i.e., they cannot make use of an OFDMA channel that spans 85 to 108 MHz).

The solution to this issue is shown in Figure 9 and Figure 10 and now included in [MULPIv3.1]. Each CM (Mid-Split or High-Split) sees the channel that it is capable of transmitting, while the CMTS generates different long and short Upstream Channel Descriptors (UCDs) and MAPs. No schedule conflicts or channel descriptor conflicts exist since Mid-Split just sees a shortened version of the same messages.



**Figure 9 – Overlapping OFDMA Channel Arrangement**




**Figure 10 – Overlapping OFDMA Channel Solution – UCD and MAPs**

Implementation of OOC for a plant transitioning from Mid-Split to High-Split can maintain an additional ~200 Mbps of capacity while allowing both Mid-Split and High-Split CMs to be simultaneously deployed.

### 6.3. Leakage Detection

System leakage monitoring and detection is required to ensure regulatory compliance and avoid interfering with the sensitive aeronautical band in the 108 to 137 MHz range. In traditional Low-Split and Mid-Split deployments, legacy methods for accomplishing this using downstream spectrum and portable field meters detecting specialized CW carriers have been in place for many years and are well-understood.

With High-Split and DOCSIS 4.0 Ultra-High Split frequency plans, the aeronautical band is no longer in the downstream spectrum and instead falls within the upstream spectrum. With these splits, leakage test signals must instead be generated by the CM just above the aeronautical band as shown in Figure 11. The CM uses OFDM Upstream Data Profile (OUDP) test probes, normally assisting with upstream profile validation, to create specific pilot patterns which are detected by the field meters. [Coldren1] provides a summary of the approach.



**Figure 11 - Leakage Detection Test Signals from [Coldren1]**

Since the standard updates were published in 2021, significant progress has been made towards industry interoperability between operators, field meter vendors, and CMTS/node vendors. Field results has proven the ability to generate the signals in the CM, detect the signals with the field meter, and effectively troubleshoot leakage situations in the outside plant.

## 7. Upgrade Costs

Absolute costs for upgrading a particular deployed network are location-specific and depend on many factors such as existing installed equipment and cables, operator service offerings, regional regulatory requirements, labor rates, and operational practices.

In assessing turbocharged DOCSIS 3.1 technology as a step in HFC deployment, four upgrade points are considered for relative cost per home passed:

- 1) Turbocharged Baseline: This is an upgrade of the HFC actives (amplifiers and nodes) to high split with 1.2 GHz maximum DS. “Turbocharged Baseline” Assumptions:
  - Passives are assumed to support >1 GHz operation and are not deliberately a part of the upgrade
  - No significant changes to plant design or spacing
  - Amplifiers are upgraded in place to 1.2 GHz – only obsolete or unsuitable amplifiers are replaced
  - Nodes are transitioned from analog and/or digital return nodes to DOCSIS 3.1 DAA nodes
  - Cable modems are DOCSIS 3.1
- 2) Turbocharged Baseline with D4.0 Cable Modem
  - Above and beyond the Baseline scenario, customers are provided with DOCSIS 4.0 cable modems to enable higher DS tiers above 2 Gbps
- 3) Full DOCSIS 4.0 Upgrade
  - All passives are proactively replaced in the network to support 1.8 GHz or greater operation
  - Amplifiers are upgraded in place or swapped for 1.8 GHz

- Booster amplifiers are added in 10-20% of locations
- Some cable replacement (1-3%) is required
- Labor is a very significant portion of the overall cost in this scenario due to the amount of plant touch

4) FTTP Upgrade

- Service group is swapped over to fully fiber to the premise
- Average/median costs across a wide range of geographies and types are used for this analysis

Based on industry and operator feedback for each of these items, we estimate the relative upgrade costs as shown in Table 4:

**Table 4 – Relative Upgrade Costs**

Upgrade Scenario	Relative Upgrade Cost
Turbocharged Baseline	1.0
Turbocharged Baseline + D4.0 CMs	1.2-1.5
Full D4.0 Upgrade	1.8-2.4
FTTP	10

## 8. Conclusion

Operators have many tools available to increase capacity and performance of their existing HFC networks to address evolving customer experience targets and competition. DOCSIS 4.0 technology is clearly the long term future of HFC to provide multi-gigabit downstream and upstream capability, but may require significant plant investment to roll out 1.8 GHz capability across the plant to take full advantage of all that is available in the specification.

A more incremental approach with turbocharged DOCSIS 3.1 technology is possible – take advantage of the changes that have been added in the DOCSIS 3.1 specifications with recent industry performance enhancements and include the higher multi-channel capability of DOCSIS 4.0 modems if needed. This turbocharged approach allows an operator to gradually roll out changes to the outside plant while enabling DS tiers to 5 Gbps and significant improvements in latency.

HFC will continue to deliver the service needed by customers for many years – and turbocharging DOCSIS 3.1 technology provides yet another way to get there incrementally.

## Abbreviations

AC	alternating current
AQM	Active Queue Management
ASF	Aggregate Service Flow
ATDMA	Advanced Time Division Multiple Access
bps	bits per second
CCAP	Converged Cable Access Platform
CMTS	Cable Modem Termination System
DAA	Distributed Access Architecture
dB	Decibel
DCA	Distributed CCAP Architecture
DOCSIS	Data-over-Cable Service Interface Specifications
DS	Downstream
ECN	Explicit Congestion Notification
ECN CE	ECN Congestion Experienced
ECT	ECN Capable Transport
ESD	Extended Spectrum DOCSIS
FDD	Frequency Division Duplexing
FDX	Full Duplex
FTTP	fiber-to-the-premises
HFC	hybrid fiber-coax
Hz	Hertz
IP	Internet Protocol
L4S	Low Latency, Low Loss, Scalable Throughput (L4S)
LLD	Low Latency DOCSIS
MAC	media access control layer
MDU	multi-dwelling unit
MNO	mobile network operator
MSO	multi-system operator
NQB	non-queue building
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PGS	Proactive Grant Service
PHY	physical layer
PMA	Profile Management Application
QAM	quadrature amplitude modulation
RF	radio frequency
SCTE	Society of Cable Telecommunications Engineers
SDV	switched digital video
TCP	Transmission Control Protocol
UHS	ultra-high frequency split



## Bibliography & References

[Coldren1] – *Leakage Detection in a High Split World: Industry Progress Toward a Viable Solution*, Rex Coldren et al., SCTE Expo 2021

[DCA] – *Distributed CCAP Architectures Overview Technical Report*, CM-TR-DCA-V01-150908, CableLabs

[ECN1] - *Explicit Congestion Notification (ECN) Protocol for Very Low Queuing Delay (L4S)*, draft-ietf-tsvwg-ecn-l4s-id-28, K. De Schepper et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-ecn-l4s-id-28>

[Karthik1] – *Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)*, Karthik Sundaresan et al., SCTE Expo 2019

[L4S1] - *Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service Architecture*, draft-ietf-tsvwg-l4s-arch-19, Bob Briscoe et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-l4s-arch/>

[L4S2] - *DualQ Coupled AQMs for Low Latency, Low Loss and Scalable Throughput (L4S)*, draft-ietf-tsvwg-aqm-dualq-coupled-24, K. De Schepper et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-aqm-dualq-coupled-24>

[MULPIv3.1] – *DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1-I23-220328, CableLabs

[MULPIv4.0] – *DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv4.0-I05-220328, CableLabs

[NQB1] - *A Non-Queue-Building Per-Hop Behavior (NQB PHB) for Differentiated Services*, draft-ietf-tsvwg-nqb-10, Greg White et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-nqb-10>

[PHYv3.1] – *DOCSIS 3.1 Physical Layer Specification*, CM-SP-PHYv3.1-I19-211110, CableLabs

[PHYv4.0] – *DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-I05-220328, CableLabs

[PMA-TR] – *DOCSIS 3.1 Profile Management Application Technical Report*, CM-TR-PMA-V01-180530, CableLabs

[White1] – *Low Latency DOCSIS, Overview and Performance Characteristics*, Greg White et al., SCTE Expo 2019