The First Anniversary of @Real10G: What Have We Learned?

A technical paper prepared for presentation at SCTE TechExpo24

Dr. Robert Howald
Fellow
Comcast
robert_howald@cable.comcast.com

Errol D'Souza
VP SW Dev and Engineering
Comcast
errol_d'souza@comcast.com

Jeanne Ciampa
Director, Strategic Program Delivery
Comcast
Jeanne_ciampa@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

That stirring that reverberated throughout the Colorado Convention Center at last year's SCTE Expo, and subsequently across the industry, can be ascribed to the timely launch of the World's First Data over Cable Service Interface Specification 4.0 (DOCSIS® 4.0) services. This proud industry moment, years in the making, happened that same week! New "X-Class" speeds were made available to customers in Colorado Springs, just down Interstate 25 from the event. The ground-breaking phenomenon was followed soon after by launches in the Atlanta and Philadelphia areas. Network upgrades and Full Duplex DOCSIS (FDX) activations have since continued. Over 10,000 FDX nodes are already deployed, and the pace increases throughout 2024.

In this paper, we will provide a backwards, current, and forward-looking perspective around this monumental broadband milestone. The operationalizing of new technology, in particular with the nature of innovation inherent in D4.0 FDX, builds on previous transformational change introduced with distributed access architecture (DAA) and the virtual cable modem termination system (vCMTS). Unique development, integration, and field challenges drove creative plans and tactics across teams to meet the 2023 launch target and prepare for scalability in 2024. The path ahead requires a close look at the current strategy and the foundation it provides, and subsequently the relationship among performance, product speeds, capacity, and spectrum management to plan effectively for the multi-year migration ahead.

In this paper, we will detail that journey, including:

- Capacity, Speed, Spectrum, and Network Strategy
- FDX Traffic and Utilization
- Hardware (HW) and software (SW) Integration, Test, and Trial Methodology
- Key Observations, Field Measurements, and Lessons Learned
- Future Feature and Technology Roadmap

The next era of DOCSIS – DOCSIS 4.0 – has officially begun! It represents a significant leap forward for broadband, and it turns the 10G vision first announced at the Consumer Electronics Show in 2019 into reality. Attendees will get an up-close view of that reality and the path ahead, shining a light on this remarkable achievement that once again demonstrates the industry's commitment to broadband leadership.

2. Network Migration Strategy

2.1. Distributed Access Architecture (DAA) Acceleration

At the 2022 SCTE Expo, Comcast announced its plans to upgrade 80% of its footprint by the end of 2025. The press release announcing this objective is here:

The initiative incorporates upgrades to multiple platforms. It includes a DAA upgrade – vCMTS platforms replacing legacy integrated cable modem termination system (I-CMTS) platforms, digital nodes (RPDs) based on the Remote Physical Layer (R-PHY) standard from multiple original equipment manufacturer (OEM) partners, and radio frequency (RF) actives in the plant upgraded to the Mid-Split (5-85 MHz) upstream (US) and 1218 MHz downstream (DS) bandwidth.

The upgrade, shown from a platform and spectrum level in Figure 1and Figure 2, provides all the well-documented benefits of DAA and virtualization, and delivers operational efficiency, transparency, simplicity, and automation that can only come through virtualizing coupled with cloud compute applications. It increases the capacity of the US by up to four times when deploying DOCSIS 3.1 (D3.1) orthogonal frequency-division multiple access (OFDMA), and more once DOCSIS 3.0 (D3.0) single carrier quadrature amplitude modulation (SC-QAM) and increases the DS bandwidth by up to 60%. With these upgrades, Comcast can increase broadband speed tiers to 2 Gbps DS and 300 Mbps US, while also delivering speed upgrades to existing customers across other lower speed tier broadband packages. The 80% coverage amounts to approximately 50M homes and business passed, and more than half of that has been completed on schedule and is on track to meet the stated homes passed objectives.

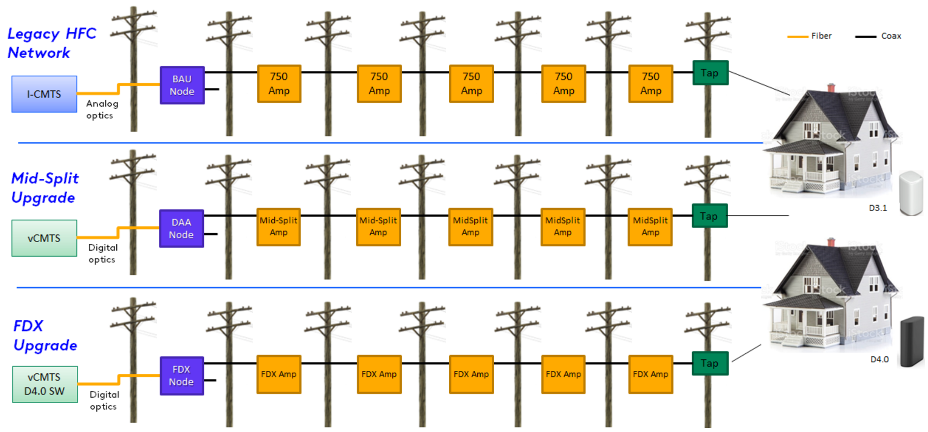

Figure 1 - DAA Network Upgrade - Platforms

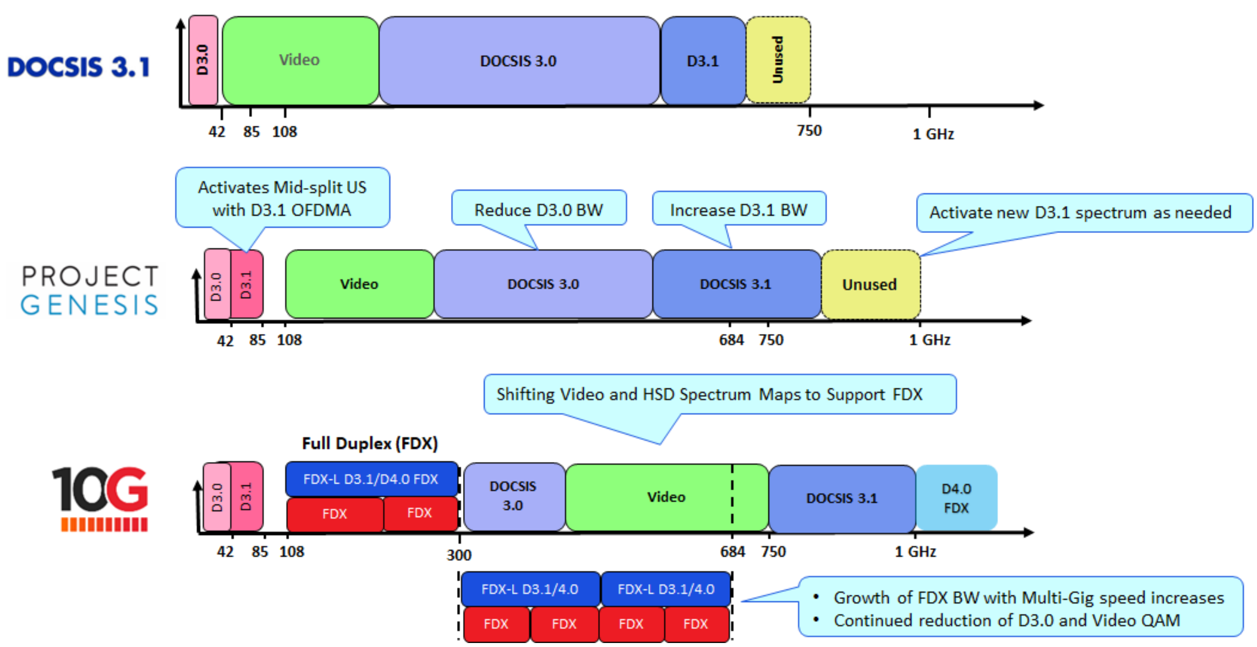

Figure 2 - Comcast DAA Network Upgrade -Spectrum

DAA is, of course, a necessary foundation for D4.0. After many years of dueling DAA standards –R-PHY and Remote MAC-PHY – the clear winner in the industry is R-PHY, which enables a lighter weight node and a more centralized core hardware and software platform for packet processing, which is most amenable to virtualization. With DAA being accelerated across the footprint, the fast-following technology of D4.0 FDX can then be deployed. The D4.0 upgrade in a virtualized, DAA systems is very efficient. It requires a software upgrade of the vCMTS core code to support D4.0 FDX features. The DAA RPHY nodes are upgraded to FDX-capable platforms with module swaps. And, of course, a customer who wants multi-gigabit DS and US speeds, branded X-Class at Comcast, will receive a new D4.0 FDX capable cable modem (CM).

Note that the DAA switching platforms must also be capable of supporting the higher speed, X-Class, services in aggregate. However, it is the nature of the Ethernet switching infrastructure, and among the core value propositions of DAA, that it benefits from commercial-off-the-shelf (COTS) technology riding the continuous Moore's Law-based throughput gains. This contrasts sharply with the upgradeability of multi-year, generational line-card upgrades that burden I-CMTS platforms, and even of proprietary I-CMTS cores driving digital nodes. As such, the necessary DAA infrastructure platform upgrades that support FDX speeds were rolling out into the field well before FDX nodes were even available.

2.2. On To DOCSIS 4.0 Full Duplex (FDX)

The FDX roll-out builds on the DAA migration. The logic, of course, is that the prioritization for where and when to upgrade the network, including managing a given labor pool size, considers where capacity is most in demand across the network as measured by the highest utilizations and the regional competitive environment. It follows logically that areas prioritized for the DAA/Mid-Split/vCMTS upgrade (dubbed "Genesis" for simplicity internally) would also be prioritized for the FDX upgrade.

Figure 1 showed the platform migration to FDX. Figure 3 shows two of the key hybrid-fiber coax (HFC) components that bring FDX to life – the FDX Node and the FDX Amplifier. Notable, one of the main advantages of D 4.0 FDX over D4.0 frequency domain duplex (FDD) is that no Tap upgrades are required. Tens of millions of Taps are deployed in the network, and while new Taps themselves would have a modest premium, the time and cost to recover them is a large hurdle to successfully implementing FDD. There is an additional uncertainty and consideration for how much cable may need to be replaced for these extended frequencies to be used effectively, given the specifications for today's cable does not extend to performance at 1.8 GHz. Rather than deploy new plumbing in the plant in 1990's fashion, FDX takes advantage of new technology to use the *existing* cable bandwidth more efficiently – for both DS and US. It is worth noting that the two D4.0 approaches are complementary, as we will discuss later in this paper.

Figure 4 shows an example spectrum migration approach for FDX which we will describe in more detail in the context of capacity and speed offerings

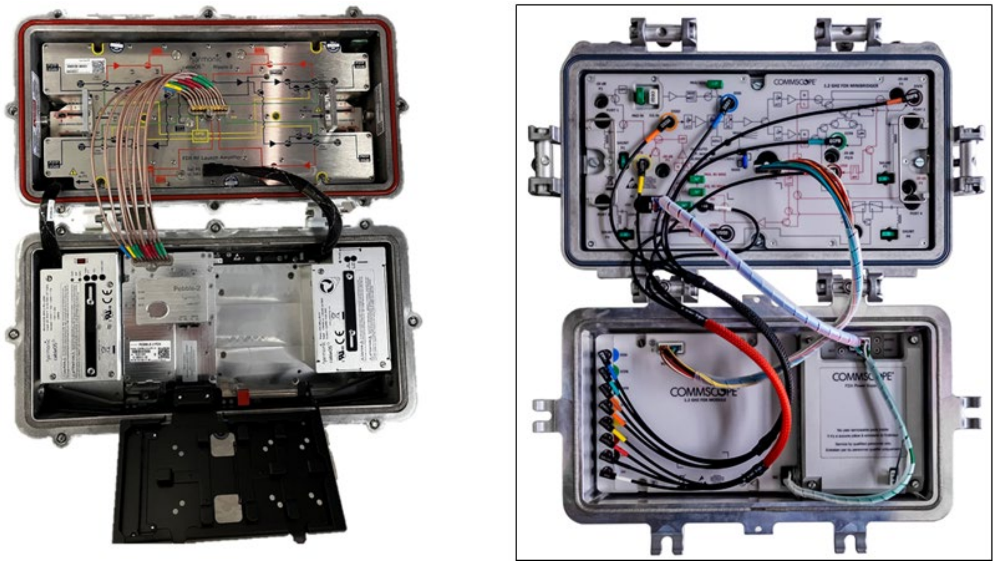

Figure 3 - HFC Node (L) and Amplifier (R) Components for FDX

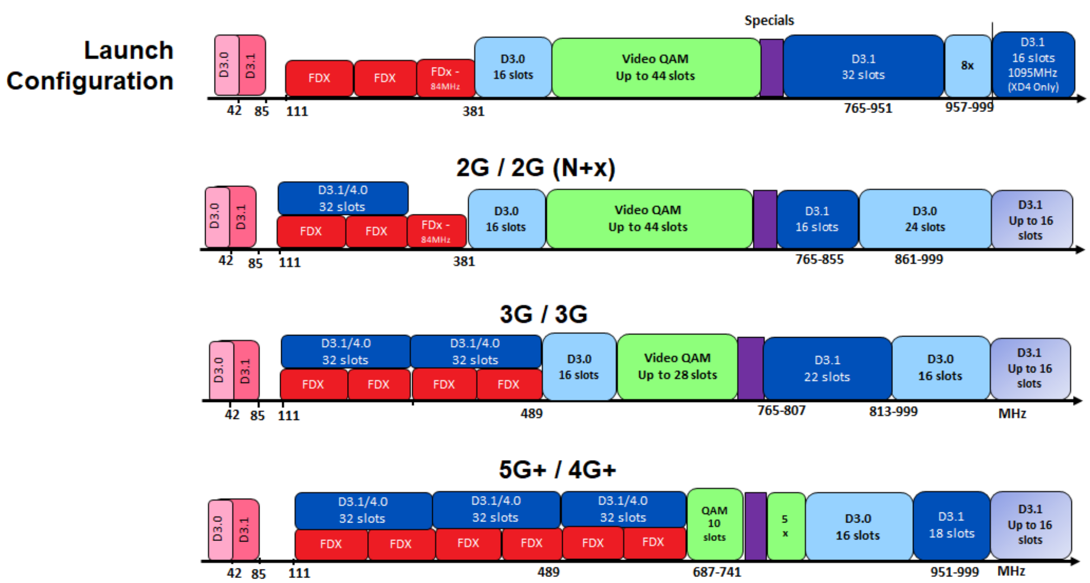

Figure 4 - Example FDX Spectrum Migration (note axis uses center frequencies)

2.3. Capacity and X-Class Speeds

It is clear from Figure 2 that the priority for D4.0 FDX is to massively expand the US capacity. The DS capacity essentially remains unchanged, although there is the potential for D4.0 devices to operate in spectrum above 1 GHz, which not every Comcast D3.1 consumer premises equipment (CPE) platform is able to do, due to Multimedia of Coaxial Alliance (MoCA) considerations. Because the DS bandwidth and capacity is essentially unchanged, FDX *itself* plays an insignificant role in DS-related capacity dynamics. However, because of FDX and its core objective to deliver multi-gigabit symmetric speeds, it does drive video quadrature amplitude modulation (QAM) reclamation in favor of D3.1 channels, indirectly adding capacity to the network. Furthermore, FDX peak speeds exceed average utilization significantly more than the historical norm, so the result is, in fact, excess capacity for future growth. Because of this aspect of FDX and X-Class, the primary factor to manage with respect to DS capacity is ensuring there is sufficient *D3.0* capacity kept in place to manage the utilization and speed tiers on that platform, while those devices continue to leave the network.

Historically speaking, the US had been both a bottleneck to upload speeds and the driver of new node splits. However, that history was based on a Low Split US of 42 MHz or less (in North America), D3.0 (or lower) high-speed data (HSD) platforms, analog fiber plant, and in many cases, years ago, also serving very large fiber nodes in terms of subscriber count. Over the last 5+ years at Comcast, the capability, and effectively the clock on the HFC network lifecycle, has been reset with four significant initiatives:

1. Upgrading the network from Low Split to Mid-Split, with an 85 MHz US
2. Activating the Mid-Split with D3.1 OFDMA
3. Developing the Profile Management Application (PMA) to optimize D3.1 capacity
4. Migrating to DAA, providing large fidelity benefits both DS and US
5. Continuing to pull fiber deeper into the network, and segment or split nodes as required

The result of this is smaller nodes, robust and predictable DAA fidelity, PMA-optimized bandwidth-efficient spectrum use, and ultimately more than 5x the US capacity available. These initiatives significantly extend US-driven network lifespan estimates. Peak-hour US utilization is now able to be satisfied for many years.

With the addition of FDX bandwidth, as D4.0 FDX penetration increases over time, there is essentially unlimited bandwidth with respect to capacity lifespan, as FDX adds approximately 8x more US capacity on top of Mid-Split. The FDX upgrade is about delivering faster upload speeds, of course. The Mid-Split US tops out today with 300 Mbps product speeds (D3.0 SC-QAMs are still a significant portion of the US spectrum) – which is enough speed for many users and applications. However, it is important to always stay one step ahead of the demand curve, and the network speed capability is being ratcheted up to multi-gigabit symmetric to do that. As such, sufficient capacity to hit those speeds must be put in place.

In summary, FDX is about delivering new and much higher upload speeds, and less about expanding capacity and network lifespan. While capacity does increase as part of the upgrade's spectrum migration plans, the "Genesis" initiative was built around a migration strategy that would satisfy the network capabilities from a capacity vs. utilization perspective, while introducing a powerful, future looking DAA/virtualization technology reset into the HFC network. FDX builds on this foundation to enable multi-gig symmetric speeds.

We will discuss speed-based network migration dynamics in a subsequent section of this paper. For now we note that the launch of X-Class speeds added the following speed tiers to the Comcast product portfolio:

X-300: 300 Mbps Symmetric

X-500: 500 Mbps Symmetric

X-1000: 1 Gbps Symmetric

X-2000: 2 Gbps Symmetric

As a reminder, note that the top asymmetrical product speed today is 2 Gbps / 300 Mbps. Also, note that X-300, while a new product, is accomplished without requiring D4.0 FDX. A D3.1 device deployed in a Mid-split configuration can comfortably deliver 300 Mbps US speeds, while accommodating peak-busy-hour background traffic, which at the start of 2024 was averaging approximately 500 kbps per user. As shown in Figure 4, the Mid-Split configuration used in production today allocates four D3.0 SC-QAMs, minimum, from 5-40 MHz, while above 40 MHz is all D3.1 OFDMA. The net capacity with these allocations is 450-500 Mbps, which is easily able to accommodate a 300 Mbps peak with speed with sufficient capacity to spare. Over time, the D3.0 carriers will be reduced as these modems come off the network in favor of D3.1 devices, the OFDMA will expand, and the available Mid-Split capacity will increase accordingly.

2.4. Spectrum Roadmap

Step 1 (Row 1): Launch up to 2G

As observed in Figure 4, the initial launch of the technology resembles a D4.0 FDD approach. That is, the US band is expanded only, while the DS vacates this band entirely. One of the major drawbacks of FDD is the inefficient allocation of spectrum, forcing the DS bandwidth also to be expanded beyond 1218 MHz, and thereby forcing upgrades of all Taps and passives. As mentioned previously, it also increases the risk of having quality spectrum across the extended DS over existing coax trunks, feeders, drops. However, for a first FDX step, by limiting the speed to 2 Gbps and limiting the deployments to node+0 actives (N+0) footprint, with accompanying small node sizes, the "FDD" start for FDX provides several benefits. It offers full visibility into the new US to understand RF dynamics of this never-before-used band for US. It allows for an opportunity to exercise and observe the node's echo canceller (EC). The essential new FDX digital signal processing (DSP) technology, with more flexibility than if it was carrying payload traffic. Lastly, it defers the more complex features of scheduling Resource Block Assignments (RBAs) and FDX-Limited (FDX-L) on legacy D3.1 devices until a solid RF characterization foundation is well-understood.

The FDX bandwidth allocation per the D4.0 standard options are shown in Figure 5.

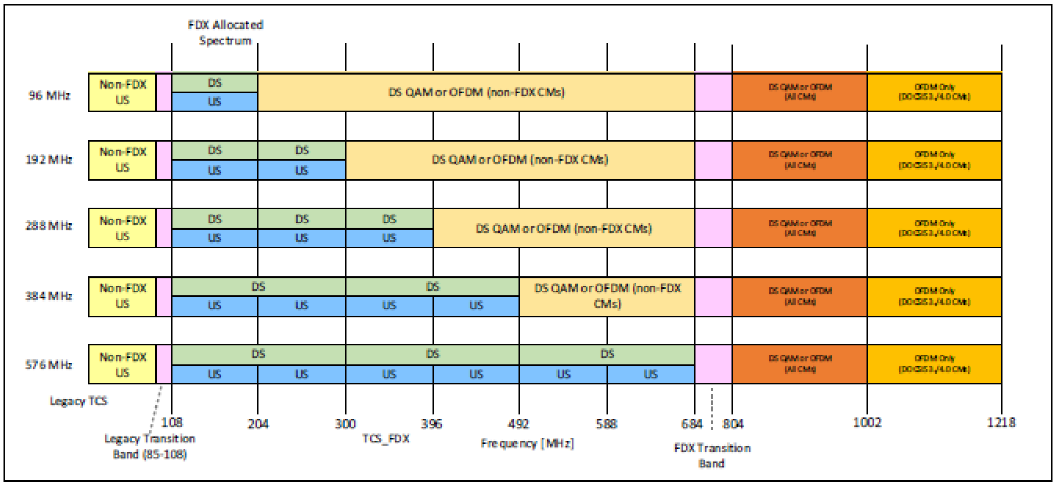

Figure 5 - FDX Spectrum Allocation Options

As shown in Figure 5, it was anticipated that operators would need the flexibility to incrementally add FDX bandwidth over time. Indeed, referring to Figure 4 once again, that is the expectation at Comcast as well. The launch configuration used to deliver 2 Gbps is very close to row 3 – 288 MHz of defined FDX spectrum. The exact configuration deployed excludes the last 12 MHz to make room for enough video and D3.0 SC-QAM slots to support the DS aggregate demand and speeds for D3.0 devices and allocate sufficient video QAM to meet existing video product suites without additional channels being moved to the internet protocol (IP) tier. The 44-slot video QAM allocation of Figure 4 is sufficient in many areas, while in some markets it may require Headend re-multiplexing to reduce a 48-52 QAM video line-up.

The FDX US allocation is therefore 276 MHz total, for a total of 356 MHz of US bandwidth. Conservatively, the available US is then:

- Mid-Split Band: 450-500 Mbps
- FDX Band: 1900-2200 Mbps [10-11 bps/Hz x 70% gross-to-net throughput all in overhead]
- Total: 2350-2700 Mbps

This is more than sufficient to deliver 2 Gbps speeds for the X-2000 customers, and of course easily supports X-1000 and X-500 customers:

In the DS, for DOCSIS 3.1 users, we can similarly calculate:

- D3.0: 24 x 37.5/6 MHz = 900 Mbps
- D3.1: 32 x 51.9/6 MHz = 1661 Mbps (production number with active PMA optimization
- Total : 2561 Mbps

Downstream peak hour usage was roughly 5 Mbps at the start of 2024. Like US, the small nodes associated with N+0 means there is ample capacity to steady state user traffic as well as peak speeds for D3.1 users of 2 Gbps DS.

For D4.0 FDX users, we have

- D3.0: 8 x 37.5/6 MHz = 300 Mbps
- D3.1: 48 x 51.9/6 MHz = 2491 Mbps (production number with active PMA optimization
- Total : 2791 Mbps

Lastly, note that the vCMTS, which has been supporting D4.0 since mid-2023, incorporates scheduler behavior that, for the D4.0 FDX devices, will prioritize use of the FDX band over the legacy band. Thus, non-FDX devices have access to more capacity than they otherwise might because the D4.0 CM traffic is biased to "offload" into the FDX band.

Step 2 (Row 2): 2G with node + x actives (N+x)

In the second row of Figure 4, we see that the "magic" of FDX begins – overlapping of the DS and US spectrum. However, this step does not increase the speeds. What is happening is that, with the introduction of FDX Amplifiers, we are no longer able to benefit from the small node sizes.

By counting on N+0 and small node sizes (60-75 hhp/node typically), we were able to take advantage of two favorable options to manage spectrum

1. Reduce the D3.0 spectrum allocated for the population of D3.0 cable modems, going from 36-44 slots (depending on specifics of traffic utilization) down to 24 slots
2. Use of spectrum > 1 GHz. The N+0 footprint was built out to 1218 MHz. This added spectrum provides sufficient spectrum for the FDX CMs to deliver 2 Gbps DS. Unlike the D3.1 devices, the D4.0 FDX devices cannot "see" D3.0 SC-QAMs in the FDX band, so the >1GHz band becomes important, accessible capacity. The Comcast family of D3.1 devices, when operated as Mid-Split devices and provisioned for 2 Gbps DS, can see all of the DOCSIS channels *except for* frequencies > 1 GHz, since that is their maximum DS. These devices, when operated as Mid-Split devices, all support use of the MoCA home networking technology, which relies on frequencies between 1 GHz and 1.218 GHz to support in-home video distribution to lightweight, subtending STBs.

Both benefits are eliminated for N+x networks. The nodes may no longer be small. Average node sizes are on the order of 250 households passed (hhp) in non-N+0 footprint. And, in N+x systems, while all the amplifiers that are upgraded with the "Genesis" project are 1218 MHz capable, the Taps in these systems are not being replaced or upgraded, so are design limited to 1 GHz. They will, of course, support some capacity above that in the roll-off, similar to as is done today with 750 MHz systems limited by 750 MHz amplifier platforms. The PMA will optimize use of this bandwidth by selecting the appropriate modulation profiles.

For N+x therefore, with broadcast video QAM spectrum independent of N+0 vs N+x, the FDX band must now support the DS *and* the US to enable a more robust D3.0 allocation, and to provide a high, guaranteed capacity that the bandwidth above 1 GHz cannot assure. This is what is shown in Row 2 of Figure 4. The figure shows the movement of the 32 DOCSIS 3.1 channels from 765-951 MHz (note these are center frequencies) to the low end of the FDX band. There is an additional 84 MHz of available DS capacity that may also be used in implementation.

While amplifiers themselves introduce noise and distortion as any active device must, the ratio of DS transmit level to US receive level at the Node is much more favorable for N+x than for N+0. This means that in N+0, the EC must work much harder than it does in N+x systems. It was recognized early in N+0 prototype testing [**1,2,3,6,7**] that the D4.0 US receiver was much more sensitive than the specification minimum, and essentially performs as a D3.1 receiver when the EC is on. These findings are what led Comcast to conclude that aggregating ECs through amplifiers would still result in high fidelity. In short, we expect the same US spectrum allocation to deliver sufficient capacity for N+x, because it operates in a substantially better echo-to-signal environment – up to 12 dB better – than it does for N+0. Since EC noise floors add on a 10*log X basis, long cascades of Amplifiers that aggregate EC noise floors behave similar or better than a single EC with a very challenging (N+0) echo-to-signal level relationship.

In the DS, the node leg is a single interference group (IG), and as such there will be no DS modulation error ratio (MER) loss associated with a neighbor FDX US transmission due to simultaneous usage of complementary IGs. The FDX CM will be checked for risk of self-interference and poor in-home environments as part of the operationalizing of FDX and enablement and verification of self-install kit capability (SIK). These tools will check receive levels and echo levels at the CM DS receiver, and if necessary, actionize a professional install to eliminate these effects.

Step 3 (Row 3): 3G/3G

Row 3 of Figure 4 begins the incremental step of increasing speeds until the entire FDX band (108-684 MHz) is fully utilized bi-directionally. With the continued removal of DOCSIS 3.0, the spectrum payload is reduced from 40 SC-QAM to 32 SC-QAMs. The FDX band is extended by 96 MHz to Row 4 of Figure 5 from the DOCSIS 4.0 standard – 384 MHz of FDX bandwidth, and a total of 464 MHz of total US bandwidth. To make room for this necessary capacity, the next phase of video QAM reclamation takes place, moving from 44 slots to 28 slots.

The DS in the FDX band is now two full orthogonal frequency-division multiplexing (OFDM) blocks, or 64 slots of 6 MHz occupied bandwidth, and there are eight more outside the FDX band for a total of 72.

Revisiting the previous calculations, first for US:

- Mid-Split Band: 450-500 Mbps
- FDX Band: 2688-2957 Mbps [10-11 bps/Hz x 70% gross-to-net throughput all in overhead]
- Total: 3138-3457 Mbps

This is sufficient to deliver 3 Gbps US speeds for what will be X-3000 customers, albeit with less margin than the 2 Gbps case at the low end. Of course, it easily supports speed tiers below that. For more capacity margin, should it be necessary, the transition to 3 Gbps at a future time can be aligned to increasing the OFDMA in the legacy Mid-Split band, rather than continue to limit OFDMA to above 40 MHz only, while carrying four D3.0 SC-QAMs below 40 MHz.

Also, multiple options are being developed to increase the OFDMA efficiency of the FDX band:

- Optimization of the cyclic prefix (CP)
- Optimization of pilot pattern alternatives that carry less overhead
- PMA 2.0: New development for PMA that changes the way US paths are operated in the FDX band. Of particular significance will be a bias towards operating the D4.0 CPE at maximum US transmit and allow the Node receiver to accept a wide range of US receive levels. This is very different from today, where an US receive level set point (RLSP) is fixed and has a small operating window. With an intelligent, virtualized, and machine-learning based network, these constraints can give way to active management QAM profiles, as is done with PMA today, but with greater access to knobs to adjust transmit and receive levels to effect MER and forward error correction (FEC) metrics.

For the DS, we will use only one set of calculations for both D3.1 and D4.0 users. We will not make any assumptions about capacity available above 1 GHz – that is all extra capacity available for a D4.0 device:

- D3.0: 32 x 37.5/6 MHz = 1200 Mbps
- D3.1 (legacy band): 22 x 51.9 = 1142 Mbps
- D3.1 (FDX Band): 64 x 46.7/6 MHz = 2988 Mbps
- Total : 5330 Mbps

Downstream peak hour usage was roughly 5 Mbps at the start of 2024, with persistent but modest growth (15%) anticipated over the next several years. Growing to 7.5 Mbps, as an example, in three years means that a large node would have on the order of 1 Gbps of DS traffic at peak hour – again leaving substantial room for the 3 Gbps speed tier. Managing the mix of D3.0 and D3.1 devices to ensure that the D3.0 population can achieve peak speeds allowed on these platforms, and that there is sufficient capacity for the aggregate, is not shown but implicit to this migration strategy.

Step 4 (Row 4): 5G+/4G+

The final incremental step of increasing speeds is to activate the entire FDX band (108-684 MHz). With the continued removal of D3.0, the spectrum payload is reduced from 40 SC-QAM to 32 SC-QAMs. The FDX band is extended by 96 MHz to Row 4 of Figure 5 from the D4.0 standard – 384 MHz of FDX bandwidth, and a total of 464 MHz of total US bandwidth.

The DS in the FDX band is two full OFDM blocks, or 64 slots of 6 MHz occupied bandwidth each, and there are eight more outside the FDX band for a total of 72.

Revisiting the previous US calculations:
- Mid-Split Band: 450-500 Mbps
- FDX Band: 4032-4435 Mbps [10-11 bps/Hz x 70% gross-to-net throughput all in overhead]
- Total: 4482-4935 Mbps

This is sufficient to deliver 4 Gbps US speeds for what will be X-4000 customers, with substantial capacity margin, even without extracting more capacity from the legacy Mid-Split band, and of course easily supports speed tiers below that.

The total achieving 4935 Mbps is within striking distance of 5 Gbps, and as such becomes a potential future product based on the efficiencies previously noted coming to Mid-Split utilization, OFDMA efficiencies still ahead in the FDX band as the technology matures, and performance gains that become available with optimizations anticipated with PMA 2.0.

For the DS, we again not make any assumptions about capacity available above 1 GHz – that is all extra capacity available for a D4.0 device:
- D3.0: 16 x 37.5/6 MHz = 600 Mbps
- D3.1 (legacy band): 18 x 51.9 = 934 Mbps
- D3.1 (FDX Band): 96 x 46.7/6 MHz = 4483 Mbps
- Total : 6017 Mbps

As previously noted, multiple years into the future, DS user traffic at peak hour, under assumed growth rates on a large node could accumulate to 1 Gbps. This is within the capacity available as shown for a 5 Gbps DS speed, but with not much margin. At the point of these speed increases, options such as additional video QAM reclamation, D3.0 reclamation, and node splits come under consideration, depending on how the multi-year compounded growth rate plays out – if it is persistent or flattening. These are the same trade-offs that operators have been making for years to balance service groups size, spectrum allocations, capacity, and product offerings. The spectrum above 1 GHz has been ignored so far in our calculations, and this of course provides capacity opportunity for the D4.0 CMs.

FDX-L

Part of the D4.0 FDX standard, FDX-L, defines how, with a software upgrade to a D3.1 CM, it becomes compatible with a D4.0 FDX system. The actual PHY definition of the DS OFDM and US OFDMA is exactly that defined for D3.1. What FDX provides is a means to use them on the wire bi-directionally, backed by new DSP (Echo Cancellation) and logical architecture (IG / Transmission Group (TG)) functionality, which have been described in previous publications [**4**].

Of course, D3.1 CMs cannot themselves implement FDX, since there is new D4.0 FDX silicon required. A Mid-Split D3.1 CM cannot transmit at all in the FDX US band, while a High Split CM can transmit between 108-204 MHz US. Either D3.1 diplex type, Mid- or High Split, can receive in the FDX DS band – in the case of Mid-Split, 108-684 MHz, and in the case of High Split 258-684 MHz. For D4.0 FDX to be compatible with widely deployed D3.0 and D3.1 systems, the specification was written to re-use the OFDM and OFDMA of D3.1. It would not be reasonable to have spectrum allocation for D3.0, another for D3.1, and another for D4.0. What FDX-L does is make the D3.1 CM aware that it is part of a D4.0 FDX system. Namely, it can listen and receive OFDM channels in the DS of the FDX band. Furthermore, other FDX CMs on the same coaxial leg are aware of the FDX-L enabled D3.1 devices. With respect to managing interference groups, an FDX-L CM will have DS packets sent to it over OFDM when there is no risk of an FDX US transmitter in the same IG from causing interference that prevents that reception. The packets become scheduled as part of the FDX system with this new awareness that FDX-L provides.

Figure 6 shows the use of FDX-L for D3.1 devices, in comparison to D4.0 devices on that same node leg.

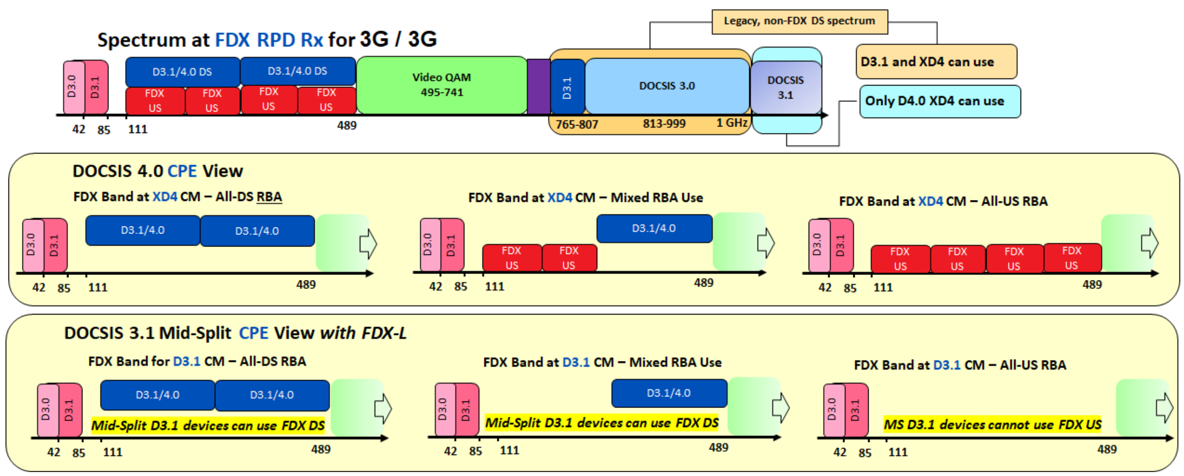

Figure 6 - FDX and FDX-L (Mid-Split) Enabled Devices

Figure 6 also shows an FDX bidirectional allocation in the top row from a node perspective. RBA examples of various FDX spectrum allocations are shown in the next two rows. Two full FDX channels are shown (2x 192-MHz OFDM, 4 x 96 MHz of OFDMA) in this example with various RBA options. Focusing on the FDX band only in the bottom two rows, we can see how a D4.0 device will utilize the FDX bandwidth under three different Resource Block Assignment (RBA) settings – all DS, all US, and mixed-use, DS and US. By comparison, the bottom row highlights how the Mid-Split based D3.1 CM will utilize the RBAs.

For a High-Split D3.1 CM, the capabilities extend further. It, too, can listen to the OFDM DS, and have packets destined for it come through unscathed because the FDX scheduler knows how to manage its traffic in the presence of IGs. Because it can transmit in 108-204 MHz, which is the first FDX US band (and not a coincidence that this is so!), it can also schedule US traffic in this band when FDX-L enabled.

In all the spectrum allocation and capacity calculations above, it is assumed that the DOCSIS 3.1 devices are all supported by FDX-L upgrades. FDX-L capable devices are undergoing interoperability testing, with trials planned in the network in 2H24.

2.5. FDX Utilization and Network Migration

As has been previously mentioned, FDX itself is more about building on top of the "Genesis" upgrade, whose roots are in long-term capacity planning and powerful technology reset. Part of the FDX technology suite is the concept of IG's, whereby the vCMTS adds within its scheduler capability the awareness of which devices can be transmitting and receiving at the same time, using the same RBA. Put another way, it keeps track of devices that should not do that, or they could interfere with one another, and it then ensures this does not occur.

A natural question comes up in this qualitative description of the potential interference risk – how much does this affect efficiency? How often does this happen? What does this mean for managing service group (SG) sizes? How big can an IG be? The introduction of the FDX amplifier brought more attention to this topic, because of its tendency to elongate IGs to all users behind the amplifier.

So, we set out to answer all of these questions, and the analysis was published [**4**]. The abbreviated results and summary below are from [**4**].

Speed-Based Traffic Analysis with Overlapping Upstream + Downstream FDX Channels

The FDX Amp use case is leveraging existing plant and may contain many 100's of subscribers. When FDX was conceived, this was not envisioned as a use case. The reference architecture is an N+0 network. It was also thought that there would be significant overlap to manage between the US and DS in the FDX band as shown in Figure 7.

The more that the DS and US peak period average consumption overlaps, then the more important it becomes to have multiple FDX TG.

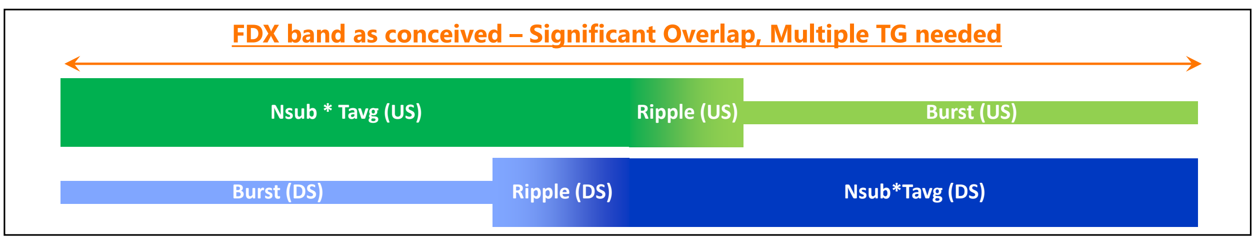

Figure 7 - FDX Band Overlap as Originally Conceived

What has subsequently come to light is that the multi-gigabit burst region dominates the traffic engineering formula. The peak burst region is much larger than the background traffic utilization component set by the number of subscribers and their average utilization (i.e. Nsub*Tavg). Figure 8, drawn to scale of what occurs in practice, shows how it is actually the DS and US *burst* regions that overlap in the FDX band. The majority of US traffic fits below 85 MHz while most of the DS traffic can be accommodated by the dedicated DS spectrum above 684 MHz.

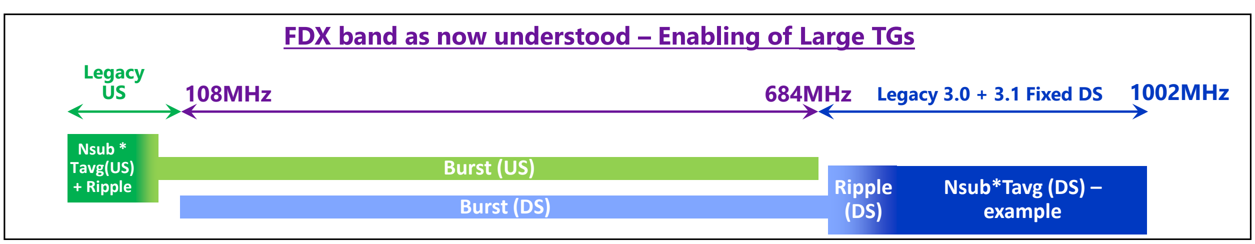

Figure 8 - FDX Band Overlap as Now Understood

In one practical example based on the 5G/4G scenario and discussed in [**4**], the US stays below 85 MHz for 99.8% of the time, and only needs to access the FDX band for 0.2% of the time! Meanwhile, the DS stays above 684 MHz for 99.9% of the time and only needs the FDX band for 0.1% of the time. The FDX band is basically providing the burst bandwidth (i.e. 5G DS, 4G US in this case) and the probabilities of US + DS bursts overlapping is microscopic.

The "Cloonan Capacity" equation – named for industry pioneer Dr. Tom Cloonan – is the basis for this analysis:

C ≥ (Nsub*Tavg) + (K-1)*Tmax_max + Tmax_max

There are three main components to this traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e. Nsub*Tavg)
2. Peak Busy Period "Ripple" (i.e. (K-1)*Tmax_max) – effectively capturing the deviation around the mean utilization of (1) due to burstiness
3. Headroom for maximum Service Tier Burst (i.e. 1 * Tmax_max)

The variable "K" is an empirical scaling factor derived from observation and can also be viewed as a way to assign the risk/margin desired for speed test success.

With this equation in mind, there are certain SG capacity requirements that must be met in any system, whether FDX or not.

These include:

1. DS SG Capacity ≥ Nsub(SG) * DS Tavg + DS Tmax_max + DS Ripple
2. US SG Capacity ≥ Nsub(SG) * US Tavg + US Tmax_max + US Ripple

These two traffic engineering (TE) conditions are shown pictorially in Figure 9**.**

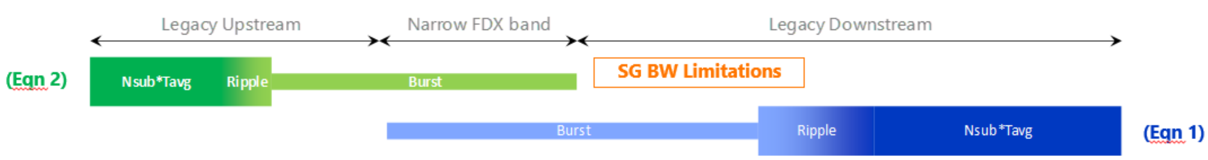

Figure 9 - FDX Band Size - Service Group Limit Visualized

In an FDX world, that means that the DS capacity above 684 MHz, plus the DS capacity in the FDX band, must be sufficient to meet TE condition #1. This must be met, independent of US requirements. Also, the US capacity below 85 MHz plus the US capacity in the FDX band must be sufficient to meet TE condition #2. This must be met, independent of DS requirements. In this respect, an FDX network's capacity requirements is no different than any other network.

FDX introduces the new twist with a TG sharing US + DS spectrum inside the FDX band. *Our traffic engineering analysis makes a foundational assumption that it is acceptable for US + DS burst regions to overlap within a TG, but it is NOT acceptable for one burst region to overlap with the other's consumption region* (i.e. Nsub*Tavg + Ripple). What this means effectively is that the usage in the consumption region is too frequent to effectively share FDX bandwidth (BW).

For FDX TG, then, there are now two additional TE conditions that must be met:

1. DS TG Capacity ≥ Nsub(TG)*DS Tavg + DS Tmax_max + DS Ripple – excess US consumption
2. US TG Capacity ≥ Nsub(TG)*US Tavg + US Tmax_max + US Ripple – excess DS consumption

These are shown in Figure 10:

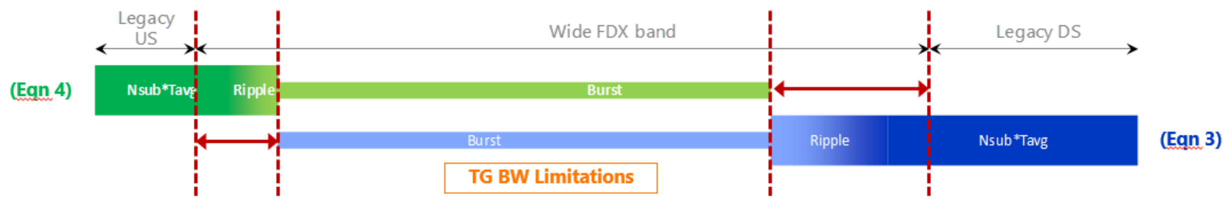

Figure 10 - FDX Band Size - Transmission Group Limit Visualized

Normally, the consumption BW is kept within the fixed legacy spectrum. This is below 85 MHz for the US and above the top of the FDX band (e.g. 684 MHz) for the DS. If these conditions are met, then neither of the TG conditions kicks in.

The DS TG Capacity, TE condition #3, only comes into play when the US consumption (i.e. US Nsub*Tavg+Ripple) exceeds the 85 MHz capacity. As this excess US consumption fills into the FDX band, it effectively takes away from FDX BW that the DS can use within this TG.

Similarly, the US TG Capacity, TE condition #4, only comes into play when the DS consumption (i.e. DS Nsub*Tavg+Ripple) exceeds the dedicated DS spectrum capacity above the FDX band (e.g. 684 MHz). As this excess DS consumption fills into the FDX band, it effectively takes away from FDX BW that the US can use within this TG.

Probability of DS + US Overlapping Tails

We will now take a quantifiable look at the probabilities of a DS tail overlapping with an US tail in the shared FDX band. The assumption is a single TG, so the US and DS timeshare any FDX channels they need to use. Network parameters are in the title of the charts below and it is also assumed 1 GHz maximum, has no QAM video, serves 5G/4G, and has 28 SC-QAMs set aside for DOCSIS 3.0 services. Utilizations shown for DS and US represent projections for 2025.

Figure 11 is a complex chart showing the probabilities of spectrum usage for both US + DS cumulative distribution function (CDF) with 200 subscribers. Our model assumes that the US capacity fills from the lowest channel in the spectrum and increases to the right until it reaches the top of the FDX band. The DS capacity is the opposite where it fills from the top channel in the spectrum down to the left. These are for visualization purposes only, a real scheduler may choose different specific frequencies to allocate traffic, but this does not affect the mathematical basis. The fixed DS capacity is completely used before the DS uses the FDX band.

The yellow curve on the left is the US ‘CDF’. Technically, it is the (1 – CDF) function. It tells us the probability that the US will need MORE than that amount of spectrum. The curve crosses the red line (85MHz) at 0.2%. This means that 99.8% of the time, the US capacity remains completely within 85 MHz, and only 0.2% of the time does it even need to request FDX BW. That equates to only 20 seconds every evening!

The blue curve on the right is the DS cumulative distribution function, 'CDF'. It tells us the probability that the DS will need spectrum below that point. The curve crosses the top of the FDX band (684MHz) at 0.1%. This means that 99.9% of the time, the DS capacity remains completely above 684 MHz, and only 0.1% of the time does it even need to request FDX BW. That equates to only 10 seconds every evening!

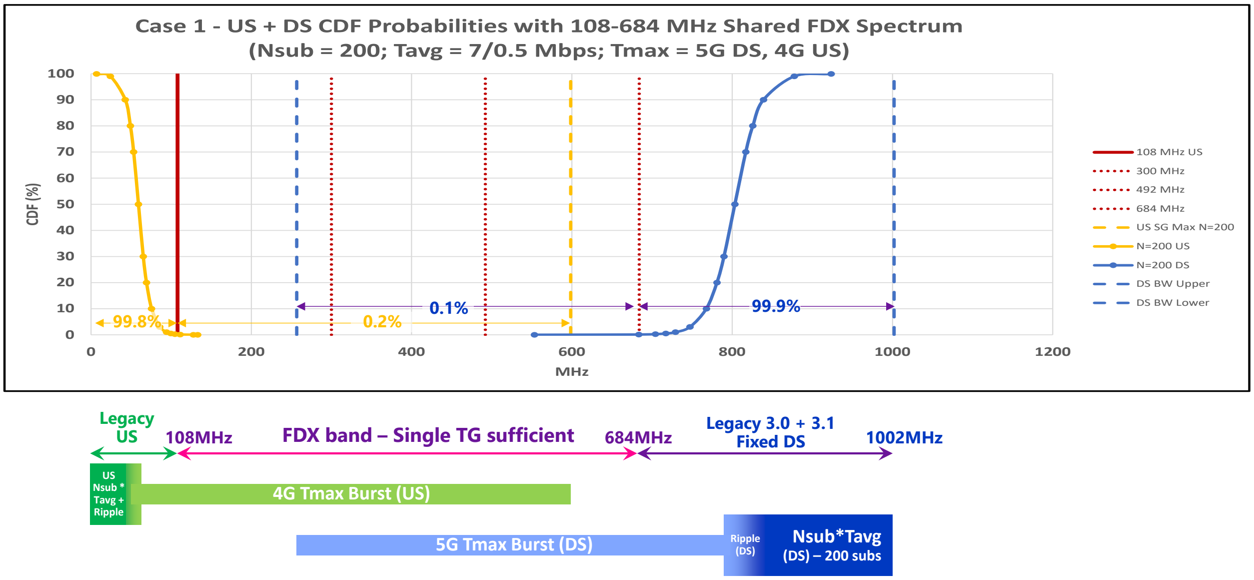

Figure 11 - Probability of DS and US Overlapping Tails (No QAM video, 200 subs)

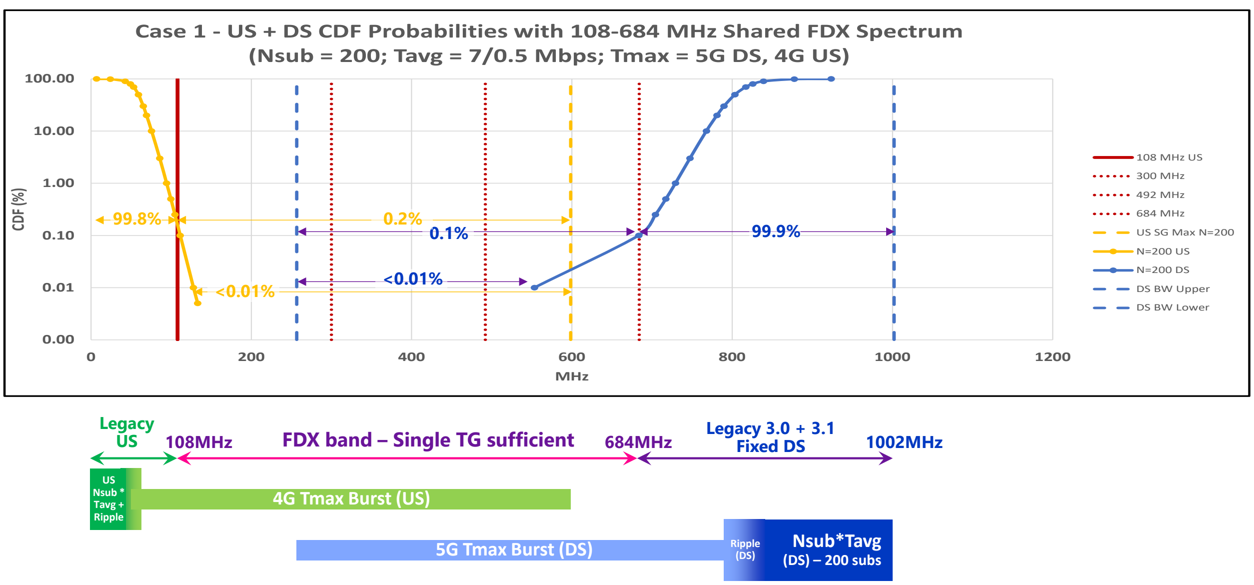

Figure 12 - Probability of DS and US Overlapping Tails (same @ Log Scale)

To take a closer look at the overlapping tails, Figure 12 shows the same data on a log scale. The US datapoint for 0.01% is at 128 MHz. What this means is that the US needs more than 20 MHz worth of capacity from the FDX band 0.01% of the time! For 20 MHz worth of capacity, the scheduler could assign this as 25% of a single 96 MHz FDX channel or as 4% of the entire FDX band. So, the US can still burst to the top of the FDX band. However, it is very infrequent and in short bursts that do not congest the system.

Looking at the DS CDF, the 0.01% datapoint is at 533 MHz. This means that 99.99% of the time, the DS BW needs are above this point and only 0.01% of the time does the DS need to go below this point. 533 MHz represents ~80% use of a single 192 MHz FDX DS channel or ~26% of the entire FDX band. The blue dashed line at 257 MHz represents the lowest spectrum BW point that the DS needs based on the traffic engineering formula (e.g. a speed test on top of normal traffic).

During the normal course of events, the probability of the US + DS tails overlapping is on the order of 1 second every 100+ years! Please refer to [**4**] for a comprehensive analysis including additional scenarios and longer-term projections.

Modeling Results Summary: Traffic Engineering Applied to Service Group / Transmission Group Size

The foundational principle enabling large transmission group sizes is that, because of the statistical characteristics of real traffic, networks can be oversubscribed – meaning that many customers can effectively share capacity that is a number less than the sum of their individual peak usage needs. *This has been the case for decades, before FDX existed.* As the peak speeds to average utilization skew to larger ratios in the Gigabit speed era, we see that the higher the peak burst, the increasingly infrequently it occurs. Years of evidence from broadband service delivery have been used to generate probability distributions that describe the traffic behavior, and which are leveraged and extrapolated to predict the expected performance in these models, and in particular how that applies to the dual-use FDX band.

The consolidated results that translate this analysis into speed-driven/TG-"limited" service groups size thresholds are shown in Table 1. Results show that large TGs can be supported. Furthermore, the TG sizes align with what are current node and service groups sizes, and that result from network augments that are already part of the Comcast HFC upgrade plan. No extraordinary network augmentations or segmentations are required. It took many years after the definition of IGs and TGs were formed to understand in practice their significance, or lack thereof, for practical traffic patterns. The focus during the development of the standard has what has turned out to be a narrow RF-centric focus and missed this larger picture because the simple question "how large can an IG be" was never asked.

Table 1 - TG Size Summary: Spectrum Scenarios and Speeds

TG Size Summary*

	Scenario	2G / 2G	3G / 3G	4G / 4G
Near Term	1	>350	306 / X	157 / X
Near Term	2	>400	250	157 / X
Near Term	3	> 400	>400	>400
Longer Term	4	310	225	150
Longer Term	5	350	290	230

*4k-QAM DS / 1k-QAM US

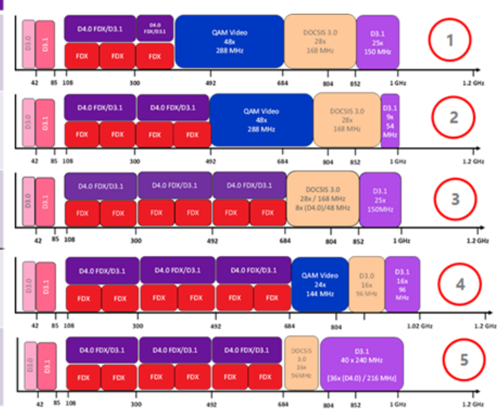

Scenario 5:
- 5 Gbps DS @ TG Size = 218
- 6 Gbps DS @ TG Size ~ 125
- 7 Gbps DS @ TG Size ~ 60

3. Traffic Engineering Conclusions: FDX Maintains Tractable, Predictable Capacity Management Tools

Traffic engineering of the HFC network has been at the root of our success deploying HSD services cost-effectively without over-engineering the network. The industry has decades of successful deployment experience. Over time, operators have developed mature empirically based statistical models of DS and US. This has translated into robust processes for operating and augmenting the network in the face of DS and US CAGR and continually increasing speeds tiers. The essential relationships of capacity supply vs peak speed and user demand have led to predictable thresholds for configuring service groups. These same learnings can be applied to traffic engineering the FDX band.

From a traffic engineering perspective, an FDX US population sharing with DS “looks like” a small number of less aggressive “users” looking to infrequently access frequency and time resources in a specific band that is managed jointly by the vCMTS. Similar to independent DS and US criteria, thresholds have been quantified for TG size versus peak speeds, for a given set of assumptions of various scenarios. Infrequent (statistically) peak bursts and DS Average BW >> US Average BW means the FDX Band’s RBAs will be deployed as DS blocks the vast majority of the time.

By contrast, if a large chunk of the coaxial spectrum such as 5-204 MHz, 5-300 MHz, 5-396 MHz, etc., is dedicated to US traffic as in DOCSIS 4.0 FDD, it will be ***idle the vast majority of the time***. Furthermore, for the increasingly high “ultra-high split” options, while this high-quality spectrum mostly idles as an outlet for a very occasional US burst, the approach will force more DS into the least predictable, never before activated part of the coaxial spectrum above 1 GHz, and only *after* all of the taps and passives have been upgraded.

Traffic analysis of the FDX bandwidth has shown that

- The initial intuitive instinct to minimize the size of an Interference Group turns out, in practice, to be evidentially unfounded
- Large TGs can be supported, similar to how oversubscription models have worked for operators for decades of broadband services
- The TG sizes determined for the multi-gigabit symmetric speed tiers of interest align well with current node sizing and the expected network augments in the Comcast upgrade plan in the years ahead
- Capacity, speed, penetration, and TG relationships can be used, as they are similarly used today, to provide guidance to network operators' network augmentation and business (speed) planning

With these findings, and with the innovative development of FDX amplifiers already in the works and showing promise, the industry can now feel confident that multi-gigabit symmetrical speeds can quickly be enabled in their existing N+x HFC deployments, founded on a deep understanding of the practical realities of burst traffic engineering and how it applies in the FDX band.

4. Test, Integration, and Trial Methodology

Comcast has multiple labs serving different purposes for the overall Quality Assurance (QA) validation of new products. There are development team labs, full end-to-end labs which support all CPE and traffic testing, and Physical and Environmental (P&E) labs for stressing equipment. Through the initial QA cycles leading up to First Customer Ship (FCS), from engineering prototypes to final node designs meant managing through seven combinations of evolving RPD and RF Tray hardware modules, in parallel to the software development lifecycle. Maintaining minimum quantities of hardware with continually evolving components meant strategically allocating equipment to labs such as to keep all QA workstreams on track with the latest versions of HW and SW required for their needs; eventually leading to all labs being equally outfitted with the Generally Available (GA), production, build of hardware and software. The Program Management team tracked the moving parts against software builds to ensure proper characterization of bugs to environment. A focused lead for lab resource management was put in place to ensure efficiency of lab utilization across the FDX program while competing with other programs.

To set goals for the team on the short- and long-term business deliverables, five key milestones were defined with a specific level of functionality to be proven at each stage leading to FCS, as shown in Figure 13:

Milestone 1: FDX Capable Node running in Genesis (Mid-Split) Mode and product speeds (up to 2G/300Mbps)

Milestone 2: FDX Capable Node running in Genesis Mode and Spectrum Alignment Complete (preparation step in FDX enablement)

Milestone 3: FDX Capable Node with FDX Configuration supporting Genesis speeds

Milestone 4: FDX Capable Node with FDX Configuration supporting FDX & Genesis speeds (Employee & Existing Genesis Customers)

Milestone 5: FCS – Ready for deployment – and ultimately customer sales

Clear entrance and exit criteria were documented for each of these milestones.

Keeping vendor partners engaged and aligned to the same priorities required a combination of Scrum Ceremonies, coupled with deep dive sessions as deemed necessary to debug issues, with close coordination monitored in multiple weekly program sessions supported by engineering leads. Comcast and vendor partners agreed on common terminology in the form of JIRA labels, and dashboards to keep SW Milestone Release Candidates aligned to ensure development teams working off the same priority features and burn down lists. All software issues were prioritized relative to the deployment milestone so the Comcast and partner development teams could focus on the issues relevant to each milestone, incrementally delivering on milestones that built overall confidence in the final deliverable – customer launch in October 2023.

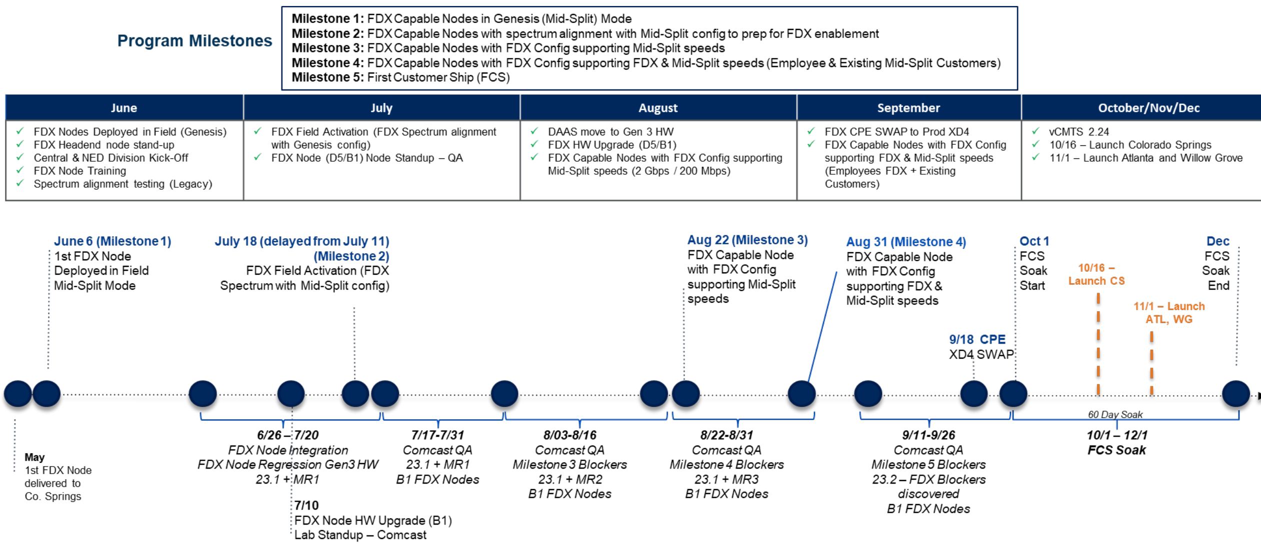

Figure 13 - Milestone-Based Program Plan to FDX Production

Readouts were and continue to be provided multiple times a week to different stakeholders ranging from the executive level to the working team to keep teams apprised of the current set of achievements and challenges.

The scope of what it means to bring a node online related to deployment tools changed over time. Agreements were established between our deployment and support teams as to the tolerance and scope for manual procedures to be used in advance of automated capabilities, that naturally lag platform HW and SW development. Considerations for manual scalability included both the potential impact to quality and risk of errors during stand up and configuration, as well as the sheer necessary scale of resources who would be enabling the technology.

The need to keep the business team aware of the health of the newly installed nodes and cadence of Go-No-Go (GNG) to launch the next set of nodes and enable sales is still in place today as we continue to scale our deployment of nodes and enablement of FDX.

5. Field Performance, Key Observations, and Lessons Learned

5.1. Speed

Close monitoring of the performance of FDX nodes and D4.0 devices continues with the expansion of FDX deployments. From a technical perspective, we have found that all symmetrical tiers of service consistently achieve their DS performance targets. On the US, similar consistency has occurred for 300 Mbps, 500 Mbps, and 1 Gbps service tiers. These are seamless installations, benefitting from these speeds being able to be hit even if the network and drop/home network RF characteristics are below average. Since all installations are via professional install today, drop/home issues are cleaned up regardless of speed tier. However, it points to the success of future SIK being very good for this most popular range of speed tiers.

At the 2 Gbps service tier, early challenges occurred periodically, where speeds would top out around 1.7 Gbps. The speed test platforms and test capability of the technicians' meters were part of this discovery. On system or SW issues, focused teams capturing logs, reviewing metrics, and working collaboratively end-to-end to understand observations has led to continuous code improvement on the RPD side and CPE side (both CM and Router).

Furthermore, configurations related to system settings carried over from Mid-Split OFDMA and needed to be modified for FDX band operation, including:

- US Rx RLSP and dynamic range window (DRW)
- PMA thresholds for MER and FEC
- Ranging algorithm
- Partial service declaration

The extended US characteristics for loss and frequency response are, of course, very different than the much more limited Mid-Split band, and optimization variables with more flexibility and range are called for to truly optimize use.

Learnings have contributed to steady improvement of speeds that deliver 2G more in line with expectations, although the consistency still is not as solid as the DS.

Less consistency in the US was to be expected, as this is what is truly "new" in X-Class launches. Comcast had been delivering 2 Gbps DS for approximately a year before FDX was launched, and it, too, went through speed consistency challenges early in deployment. Another challenge with X-2000 is that the penetration of this tier is relatively low – most customers choose the X-500 and X-1000 service, so there is a significantly smaller sample size of 2G customers. The use of employee testers where possible is used to provide better visibility, as these employees are provisioned with the X-2000 service.

As a single illustration, Figure 14 shows a speed test sample over time of one of our 2 Gbps tiered customers and the improvements over time observed (especially on the US speed).

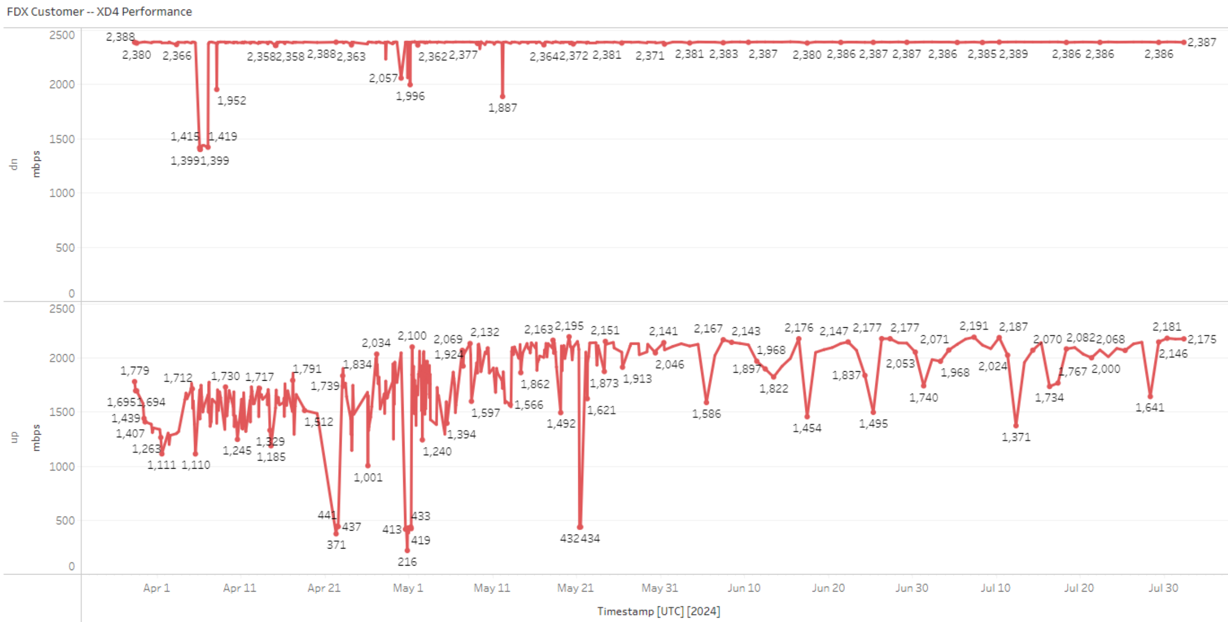

Figure 14 – User Speed Variation Observed on Recent X-2000 Installation

5.2. Echoes

The EC function, of course, is at the heart FDX operation. The performance of the EC had been well documented in prior publications [**1,2,3,6,7**], and in fact led to the conclusion that FDX operation in N+x environments – using FDX amplifiers – would also provide a high fidelity US in the FDX band

Two interesting and important observations were noted during test and integration activity that has led to HW and SW updates:

Very Low Echo – When the external Echo is extremely low, it was observed that the US MER could exhibit some unexpected variation around the mean. The EC preferred to have a robust echo to operate most effectively. Some algorithm updates were made to sense the low echo case and eliminate variation associated with the DSP hardware response, which is designed to track echo interference, and in some scenarios, very little echo interference exists.

Effective Internal Echo – As part of full range of EC characterization that led to the low-echo finding above, it was also important to characterize the internal node contributions to the aggregate echo interference. While zero internal reflection and internal isolation would of course be ideal, it turns out that simply knowing these parameters and their frequency dependent behavior well is all that is needed to minimize their impact on performance. It also helps in managing the aggregate total echo power and stabilize system performance expected, versus what would otherwise be a wholly external network dependency.

5.3. Installations

All FDX installations to date have been done via technicians onsite. The use of SIK will come with the availability of the single-box DOCSIS 4.0 gateway, which will support both FDX and Unified mode (to be discussed in subsequent section). The onsite installation and activation, of course, provides the opportunity to deploy the FDX CM per the reference architecture – a point-ot-entry installation with no in line splitters, and a clean, robust drop network. Because homes are not always wired to Comcast standard, drop clean-up and rewire happens on many of these installs. Since most DOCSIS 3.1 devices are installed and activated via SIK, this has led to the perception that a large percentage of installations will require re-wire for FDX to work. Updated training was required to clarify the "white glove" launch approach and educate the field on expected variations between FDX OFDMA US channels and the Mid-Split OFDMA US, in particular with respect to new ranging behavior and CM US transmit power thresholds of acceptability.

Additionally, part of FDX development has been to create a tool to predict viability of an X-Class speed being able to be delivered via the SIK process for a given broadband subscriber based on DS and US telemetry available on the current DOCSIS device. This tool – the FDX Home Assessment Test (fHAT) [**5**], looks at telemetry, does some calculations to predict the speed that will be possible, compares this to the speed tier asked for by the customer, and ultimately draws a conclusion about eligibility.

While this tool was envisioned as a tool for the business teams to utilize for sales enablement – which it still is – through learnings in the field and in the fHAT proof-of-concept, use for the RF measurements and calculations from fHAT can also be used to provide insight to the technician going onsite to make the installation most efficient when the determination is made that a technician install is required. Comcast expects to achieve a very high rate of self-install with FDX comparable to D3.1 products as the technology matures with the new integrated gateway.

5.4. Traffic Characteristics

Referring to the traffic management and analysis in the network migration section, an initiative was put in place to empirically determine what the analysis projected with respect to the infrequency of peak bursts. Comcast has a tool known as the Virtual Service Gateway (vSG) which granularly monitors traffic, including views of short-term peaks and long-term aggregate averages. Figure 15 shows data observations for both DS and US for peak burst extremes, and the results aligns well with the analysis results described and as predicted in [**4**].

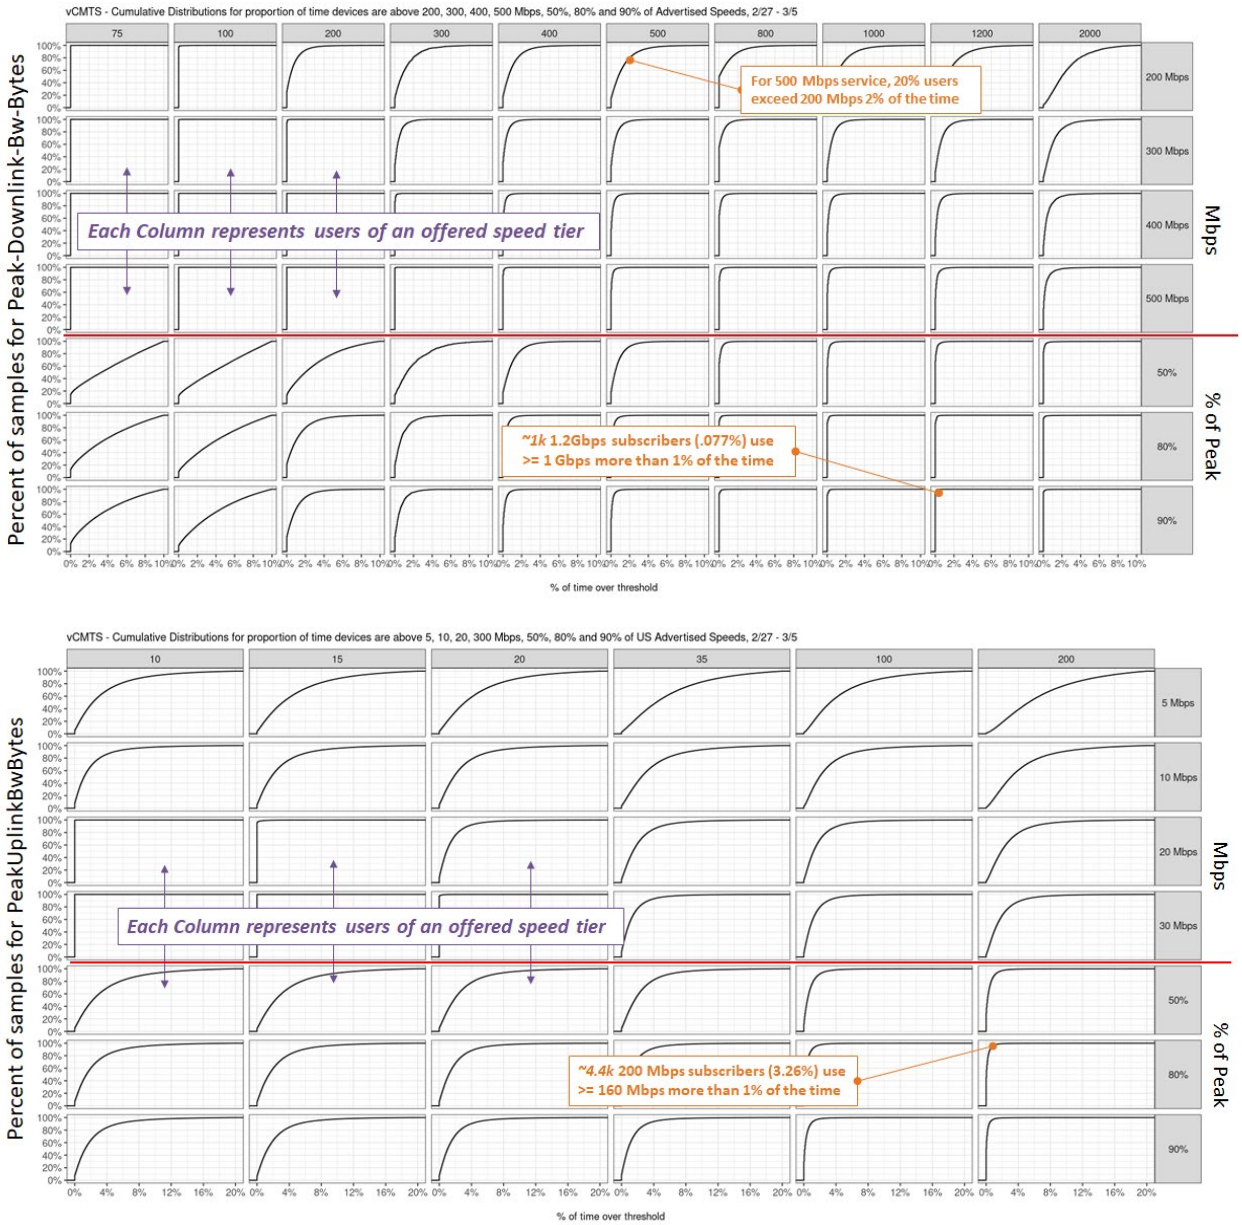

Figure 15 - Measured Data Confirms the Rarity of Max Speed Bursts on Peak Tiers

A wide range of speed tiers (columns) is shown in Figure 15. For each tier (column), what is shown is the percentage of the devices' proportional time spent above a given threshold of absolute Mbps (rows 1-4), or as a % of their speed tier (rows 5-7). The latter is done because the raw Mbps range can exceed the speed tiers in some cases. Therefore, those users can never exceed the Mbps, and the CDF is squared off (0 users exceed the max speed) accordingly. The % rows at the bottom remove this artifact.

Some stateable conclusions can be drawn from this analysis that clearly tell the story:

For a sample size of >1M subscribers with 1.2 Gbps DS speed

- 23 (**.002%**) use 1+ Gbps more than 10% of the time
- 1000 (**.077%**) use 1+ Gbps or more 1% of the time

With respect to US speeds, based on the 200 Mbps peak tier (135k subscribers at time of measurements)

- 88 (**.065%**) use 160+ Mbps more than 10% of the time
- 4400 (**3.26%**) use 160+ Mbps more than 1% of the time

And of course, whether DS or US, the probability of *two users* accessing peak speeds is exceedingly rare, and the impact, should that occur, at the application level will almost always be non-customer impacting.

These results empirically validate the traffic modeling and the conclusions drawn. Of course, since we are projecting multiple years into the future, and the model is based on historical traffic patterns and existing applications, characteristics could change over time. In particular, these extreme multi-gigabit speeds are likely to lead to new applications, and perhaps less human-experience centric and biased more towards more machine-to-machine, Internet of Things (IoT), and industrial automation types of applications. These could change traffic dynamic, and trends will be watched. One thing we know is that we are not very good at predicting the next killer app!

5.5. Project and Process

On this FDX journey, from first HW in lab to production deployment, there has been a broad range of observations and learnings:

- *Priority Alignment:* At the start of 2023, FDX was competing with multiple high priority initiatives that would put at risk the resources and focus required to deliver on the mid-October FCS (First Customer Ship) target. Creating leadership awareness and working cross-functionally with dependent teams enabled a better balancing of priorities to ensure sufficient resources were available for FDX
- *Initiative Leads:* Lead Engineer and Program Managers were assigned early on to drive the cross-functional technical engagement for FDX. Leads were supported by executive leadership to ensure they were effective in their cross-organizational roles via agreed-upon dotted-line engagements. Clear leads are especially critical when a) multiple technology initiatives need to be integrated and b) when field triage engagement is needed, and it is unclear what the root cause of an issue may be.
- *Limited Hardware Quantity:* FDX required a new RPD node and a new D4.0 CPE device. Hardware was being built as the program took hold. One of the key tactical learnings was the hard decisions on limited hardware quantities and deciding how best to utilize them in the development life cycle. At times, this resulted in moving FDX nodes between labs and the field to keep making progress.
- *Reach final hardware before deployment scaling kicks in:* During the trial and early deployment phase, multiple hardware revisions of the FDX node existed in various labs and field locations as we worked to mature the technology. Approaching deploying at scale, having the final hardware in place was important so as not have to deal with different issues across multiple node SKUs well into deployment, which would negatively impact our support teams and inefficiently use our technical resources. Because of that, non-final hardware was limited in deployment scope and to specific locales so it would be easier to swap out when the new hardware was ready and in production.

Hardware permutations in the labs prior to full-scale factory operations, and module swapping during revision updates and debug, became a challenge to platform calibration integrity. Calibration files that come with HW platforms. These files describe internal RF characterizations to compensate for the manufacturing variations, and are commonly used in DAA nodes and created in the factory. Without full automation of factory test and file generation, and multiple HW variants to contend with through development, calibration integrity is prone to error and required a dedicated team member to track and manage iterations of HW and appropriate calibration files.

- *Identify and mitigate blockers quickly:* Clearly defined milestone definitions allowed the focus on identifying blocker issues to specific functionality as quickly as possible. Effort then went into working mitigations and quick-fix workarounds, or reverting and identifying changes that might have triggered the blocker, while working on a solution. In this manner, we continued to make progress on other fronts and prevented stalling of the project due to any single blocking functionality.
- *Triage Process:* Starting with a small base of customers allowed key technical leads to have exposure at the customer experience for debug and triage, which illuminated areas for improved tools in the scaling phase. These deep findings during early launch and exposure will help build a robust FDX solution.
- *Health Metrics and Operational Telemetry/Logging:* Establishing node and device health in deployment was a key to allow for decisions to be made to increase deployment scale or to pause and assess what needed to be improved. When issues occurred where there was limited little visibility, changes were implemented to improve the visibility, and driving reproduction to quickly understand and fix the issue.
- *Health Monitoring:* In initial trials, when the product was less mature and automated monitoring was at a minimum, there were challenges relying on eyes-on-glass monitoring to make sure the system was operating correctly. This is a key role to have an engineer, or multiple engineers, assigned to so there is immediate awareness of issues. With automation, the need for this role goes away but at the start of a trial with minimal automation, this manual engagement is important.
- *Issue Prioritization:* Early on in the program, test and trial teams were generating a lot of defects that required appropriate prioritization and handling. Beyond indicating the priority of the issue based on the impact of that issue on the program, it was also important to identify a relative timeline of when the issue needed to be addressed before it hindered other progress. This helped to provide the right level of guidance to our development and test teams to focus on the issues that were important to current and future milestones. Awareness of exit criteria for each milestone was critical to be able to add this additional layer of guidance.
- *Issue Reproduction and Testing Improvements:* As we moved into trial and deployments with new technology, issues were being encountered which had never been seen before, as is usually the case when stepping out live in the field. In some cases (like the earlier mentioned 2 Gbps US speed challenges) these issues only exhibited themselves at a customer location. The primary focus in these situations was to mitigate the customer's negative experience as the top priority, but also to understand enough about the issue to drive a lab reproduction and avoid having to further triage the problem with the customer. For example, grabbing log information and capturing key metrics from the device before re-booting to recover the customer, so there are breadcrumbs to follow in the lab. Issue reproduction was also used to drive new test case definition, as learnings around how to reproduce problems could be spun back into relevant test cases and executed in future automated test cycles.
- *Two-box vs one-box CPE Devices:* The implementation of a two-box CPE solution, uncommon for DOCSIS platforms at Comcast, drove a lot of unique operational use cases to understand and overcome. They generally fall into two categories – technical and end-to-end integration.

On the technical side, the use of Media Access Control Security (MACSec) to secure the Ethernet link between the CM and router had stability issues of which the trigger for these events – leading to no connectivity – was difficult to track down.

In addition, if speed issues were observed during monitoring, it was difficult to identify the breakdown as the CM itself was unable to run a speed test initially. A Speed test client on the router or a Mac-minicomputer at an employee home served as the means to do this. The extra box and cabling for the 2-box system also provided opportunity for the technician or customer to mis-wire.

On the integration side, it was found initially after scanning of the device during installation. It would not be found in the back-office inventory, requiring intervention to re-establish that for provisioning and activation.

The single box solution (XB10) arrives later in 2024.

- *Ensure trial market selection aligns with deployment goals:* It is critically important to pick trial markets that align deployment target areas
    - A good example: our earliest FDX node trial market was a relatively clean market (minimal RF issues). It progressed from there to noisier markets as the solution took on more robustness in a smooth transition toward robustness.
    - An example with room to improve: The earliest D4.0 device trial market focused on a sub-split spectrum configuration to drive maturity of the device via large scale availability, quickly, of systems that aligned with the CPE test lab configurations that had been fully validated. However, once the device was installed in a mid-split configuration (a precursor to FDX enablement), many new device issues that were not experienced during the sub-split trials were found. Initiating device trials with a mid-split market may have in retrospect saved time.

6. What's Next

Comcast is currently engaged with technology partners to bring "Unified" DOCSIS 4.0 to life, including updating the DOCSIS 4.0 standard to include this Unified mode of operation. This initiative was announced at the 2023 SCTE Expo and is located here including demonstrations of reference designs. One year later, at the 2024 SCTE Expo, production ready equipment will be exhibited with the unified technology.

Unified DOCSIS 4.0 blends the core elements of DOCSIS 4.0 FDX and DOCSIS 4.0 FDD simultaneously. It supports the use of FDX channels in the FDX Band, but complements this by adding DS spectrum up to 1794 MHz. It has been the vision of DOCSIS 4.0 for many years [1,3] that FDX and FDD were inherently complementary to one another, and how an operator wanted to begin their journey might vary, but neither approach precluded the other in the long term. That vision becomes a reality in a Unified DOCSIS 4.0 system, as shown in Figure 16.

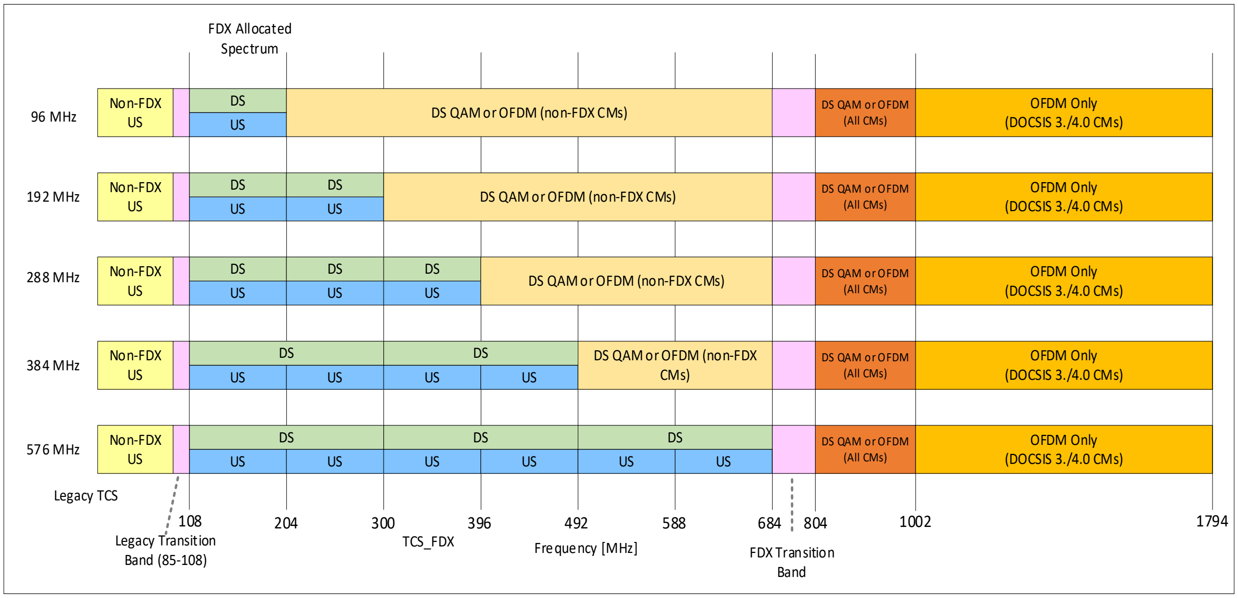

Figure 16 - Unified DOCSIS 4.0 Spectrum Allocation

Work is underway within the D4.0 standard to add a Unified mode to the standard. With respect to the Unified CM, Unified RPDs, and Unified vCMTS capability:

- The Unified D4.0 CM can operate in D3.1 Mode, FDX Mode, or FDD Mode
- Unified D4.0 CMs, Nodes, and RPDs support DS frequencies up to 1794 MHz
- A Unified D4.0 CM in FDX Mode adheres to all requirements for an FDX CM, and DS operation up to 1794 MHz
- A Unified D4.0 CMTS, Node, or RPD operating in FDX Mode adheres to all the requirements for an FDX CMTS, Node, or RPD respectively, and supports of DS operation up to 1794 MHz
- A Unified D4.0 CM in FDD Mode adheres to all requirements for an FDD CM
- A Unified D4.0 CMTS, Node, or RPD in FDD adheres to all requirements for an FDD CMTS, Node, or RPD respectively

One of the most powerful benefits of the Unified solution is that it is based on what is now mature D4.0 FDX and FDD solutions. No new development is required, only the proper integration of these to working technologies.

Strong optionality benefits accrue with the Unified solution. Operators are continually evaluating the network migration strategy, balancing the pros and cons of HFC upgrades that touch coaxial plant components versus expanding the fiber-to-the-home (FTTH) builds. The latter step has become much better aligned for operators with the DAA architecture and a vCMTS in place. Remote Optical Line Terminals (OLTs) are now modular platforms to install in DAA nodes, utilizing the same Ethernet DAA infrastructure, and landing on the same virtual HW core, where packets are processed through either the DOCSIS or Passive Optical Network PON pipeline managed by software before egress/ingress from the platform's wide area network WAN ports. The addition of the Unified solution extends the DS capacity, providing operators with another tool to weigh in the balancing of HFC upgrade versus FTTH decision criteria.

Similarly, but more tactically, the access last mile optionality is especially powerful in multi-dwelling units (MDU) scenarios. In these environments, high density residences combine with small business opportunities that are often negotiated with ownership groups and home-owners associations (HOAs) driving future-looking service requirements for their residents.

Comcast anticipates that Unified platforms will become available for use from multiple OEM partners in 2025.

7. Conclusion

Comcast is proud to have become the first Cable operator to introduce D4.0, and to pioneer a powerful new technology refresh for the industry with the launch of D4.0 FDX. This paper provided a backwards, current, and forward-looking view centered around that launch milestone in October 2023. As the X-Class footprint expands rapidly, and new FDX features and higher X-Class speeds are introduced, we will continue to keep the industry abreast of the progress, and to share both development and operations lessons learned that benefit all operators with D4.0 deployment plans – whether FDX or FDD.

While much has been accomplished, much remains to be done on several fronts – advanced feature development of FDX to maximize its use and optimize efficiency, automation to drive improved scalability to the level of today's non-FDX RPDs, and completion of fully-featured tools, new processes, and DOCSIS 4.0 training for line and fulfillment technicians.

Nonetheless, we are very excited to report that the era of D4.0 has officially begun! It is a significant step forward for the industry and delivers that 10G vision presented way back in 2019, in less than four years. A couple of those years had extreme challenges that took many eyes off the ball, for all the right reasons, emphasizing further what a remarkable achievement this the launch of D4.0 is for the industry.

Abbreviations

CDF	cumulative distribution function
CM	cable modem
COTS	commercial-off-the-shelf
CP	cyclic prefix
CPE	consumer premises equipment
D3.0	data over cable service interface specification 3.0
D3.1	data over cable service interface specification 3.1
D4.0	data over cable service interface specification 4.0
CMTS	cable modem termination system
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
DS	downstream
DSP	digital signal processing
EC	echo cancellation
FCS	first customer ship
FEC	forward error correction
FDD	frequency domain duplex
FDX	full duplex docsis
FDX-L	fdx-limited
FTTH	fiber-to-the-home
GA	generally available
GNG	go-no go
HFC	hybrid fiber-coax
HHP	households passed
HOA	home-owners association
HSD	high speed data
HW	hardware
IG	interference group
IoT	internet of things
IP	internet protocol
MACSec	media access control security
Mbps	megabit per second
MDU	multi-dwelling unit
MER	modulation error ratio
MHz	megahertz
MoCA	multimedia over coax alliance
MSO	multiple system operator
N+0	node+0 actives
N+x	node + x actives (amplifiers)
OEM	original equipment manufacturer
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OLTs	Optical line termination
P&E	physical and environmental
PMA	profile management application

PHY	physical layer
PON	Passive optical network
QA	quality assurance
QAM	quadrature amplitude modulation
RBA	resource block assignment
RF	radio frequency
RPD	remote phy device
RPHY	remote phy
R-MACPHY	remote mac-phy
SC-QAM	single carrier QAM
SIK	self-install kit
SG	service group
SW	software
TE	traffic engineering
TG	transmission group
US	upstream
vCMTS	virtual cable modem termination system
WAN	wide area network

Bibliography & References

[1] Howald, Dr. Robert and Jon Cave (Comcast), Olakunle Ekundare (Comcast), John Williams (Charter), Matt Petersen (Charter), Developing the DOCSIS 4.0 Playbook for the Season of 10G, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[2] Howald, Dr. Robert, Roaring into the 20's with 10G, 2020 SCTE Expo, Oct 13-16 (Virtual Event).

[3] Howald, Dr. Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, What is 10G – The Technology Foundation, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

[4] Howald, Dr. Robert L, and John Ulm, Saif Rahman, and Dr. Zoran Maricevic, Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway, 2022 SCTE Expo, Sept 19-22, Philadelphia, PA.

[5] Howald, Dr. Robert, L, John Chrostowski, Chris Marinangeli, and Dr. Richard Primerano, Hattery Will Get You Everywhere, 2024 SCTE Expo, Sept 23-26, Atlanta, GA.

[6] Prodan, Dr. Richard, 10G Full Duplex DOCSIS Implementation Exceeds Expectations, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[7] Prodan, Dr. Richard, Optimizing the 10G Transition to Full Duplex DOCSIS 4.0, SCTE Expo, Oct 13-16, 2020 (Virtual Event).

Press Releases

[7] *Comcast Demonstrates Fastest-Yet Speeds Over a Complete 10G Connection on a Live Network*, Comcast Corporate Press Release April 28, 2022.

[8] *Comcast Announces World-First Test of 10G Modem Technology Capable of Delivering Multigigabit Speeds to Homes*, Comcast Corporate Press Release Jan 13, 2022.

[9] *Announcing Another 10G Milestone Amidst a Flurry of Innovation*, Comcast Corporate Press Release Oct 14, 2021.

[0] *Comcast Conducts Groundbreaking Test of 'Full Duplex' Chip that Will Support Future Multigigabit Upload and Download Speeds*, Comcast Corporate Press Release Apr 22, 2021.

[11] *Comcast Reaches 10G Technical Milestone Delivering 1.25 Gig Symmetrical Speeds Over a Live, All-Digital HFC Network*, Comcast Corporate Press Release Oct 8, 2020.

Other

[12] *DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-D01-190628, CableLabs 2019

