Execute The Upstream Makeover Without Leaving Scars

A Technical Paper prepared for SCTE by

Dr. Robert Howald
Fellow
Comcast
1800 Arch Street, Philadelphi PA
robert_howald@comcast.com

Larry Wolcott
Fellow
Comcast
1401 Wynkoop Street, Denver CO
larry_wolcott@comcast.com

Leslie Ellis
President
EllisEdits
1401 Wynkoop Street, Denver CO
leslie@ellisedits.com

Table of Contents

Title **Page Number**

List of Figures

List of Tables

1. Introduction

For decades, and within the halls of this very event (Cable-Tec Expo), technologists, being a practical bunch, described the substantial task that is widening the 5-42/54 MHz upstream signal path, as the kind of monumental event that would happen but once in a lifetime. This was usually expressed as "not in my lifetime," or variants.

The title of this paper, and the breadth of technical literature happening concurrent with this paper, is, first, an acknowledgement that widening the reverse path is very much going to happen in our lifetimes. It's also an assurance that while going to a Mid-Split (85 MHz) or High-Split (204 MHz) upper spectral boundary for the upstream, home-outwards signal path is a network makeover, it is not a network rebuild. There are ways to accomplish a larger upstream signal path that are precise, reasonably swift, and forgiving – all vital elements to a "makeover without scars."

Informed by substantial lab, field and live/production environment experiences, this paper aims to illuminate why a roomier upstream path is happening now. It will describe the major things that matter, when preparing for and enacting a systemic widening of that narrow sliver of upstream spectrum at the low end of the frequency band, between 5-42 MHz. A spectral area renowned for its many signal-squelching quirks, like impulse noise. The intent is to share what works and what doesn't, when it comes to accomplishing an upper spectral boundary of 85 MHz or 200 MHz.

Because, unquestionably, the upstream path is intrinsic to all two-way applications: It is one of the two ways.

Mid- and High-split upstream configurations coincide with increasingly powerful DOCSIS 3.1 options, even as DOCSIS 4.0 is emerging. The optimal near-term expansion and long term DOCSIS 4.0 migration varies by operator. Each operator must necessarily consider its network starting point, given the interdependency of bandwidth initiatives such as node splits, Distributed Access Architectures (DAA), upper spectral boundaries, and fiber-deeper topologies. It's also worth noting that DOCSIS Annex-A (conversationally known as "Euro DOCSIS") reflects the fact that our colleagues "on the other side of the pond" send signals upstream in the spectrum between 6-65 MHz and do so very successfully.

The mathematics of Compounded Annual Growth Rate (CAGR) provide a straightforward way to quantify capacity growth and network lifespan. "Billboard speeds" must also inform the upgrade roadmap, but generally the "What" of traffic engineering and lifespan management is tractable analyses.

The "How-to" of spectrum migration is where it gets complicated. Operators understand investments in node and amplifier upgrades from previous cycles. However, these cycles didn't address upstream spectrum, largely because usage patterns didn't warrant it. The 42/54 MHz split has been in place for decades, and devices that adhere solely to it, particularly set-top boxes (STBs), are in many millions of homes. Production-scale tools, techniques, and processes must be developed to ensure that a new, wider upstream path can be efficiently operationalized, while being transparent to customers.

This paper will describe the analysis, tools, techniques, and processes to enable this upstream bandwidth transformation, focusing on production operationalization of an 85 MHz Mid-Split including:

- How homes may be impacted by a mix of device spectrum capabilities
- Mid-Split activation using SC-QAM and OFDMA
- Existing metrics and tools to assess home health

- New automation techniques to enable a seamless transition for customers with the activation of new upstream spectrum
- Cross-functional tools and processes for Tech Ops, Care, and Serviceability
- Identify and discuss some of the differences between Mid-Split and High Split (204 MHz), and of DOCSIS 4.0

Widening the upstream to stay ahead of heavy bidirectional consumption is a multi-dimensional topic. Readers will learn about new tools and operational practices that can smooth this transformation.

2. A Brief History of Cable's Upstream Path

The term "upstream path" is synonymous with the "reverse path" and the "return path" because it came second, after the "forward" signal path, from Headends to homes. For the first few decades of cable television's evolution, from the late 1950s to the late 1970s, the upstream signal path wasn't necessary. Television signals were broadcast downstream, through the plant, to homes; subscribers turned on their TVs, and watched. Nothing was "clicked," and none of those clicks moved upstream, from homes to Headends, because nothing was clickable.

In the late 1970s, some operators experimented with televisions and rudimentary data services that encouraged consumers to interact. Coincident with that, attention started to focus on building a two-way path to augment the existing one-way, downstream plant. That involved installing modules into existing amplifiers that fed a signal upstream, to the headend, then balancing that two-way signal path. From the late 1970s until the mid-90s, in fact, operators expressed their two-way-readiness in terms of what percentage of amplifiers were "two-way-capable." This meant that the amplifier housing had an empty slot for the reverse module.

Spectrally, the 5-42/54 MHz reverse path is an inhospitable zone, highly susceptible to signal ingress and impulse noise. What makes it worse is that most noise – upwards of 70%, by some estimates – originates inside homes. Because the upstream signal path is a multipoint-to-point architecture (the exact opposite of the downstream signal path), any noise generated in a home is funneled upstream, through taps to nodes, getting amplified as it moves to the headend. This effect is called noise funneling. Noise funneling is bad.

The harsh conditions of the upstream path required a sturdy modulation type, relative to the QAM-styled modulation used to carry signals downstream, towards homes; QPSK was an early workhorse. Using a lower-order modulation, like QPSK, is not unlike slowing down when driving on a road with deep potholes: It's the only way to get to the destination, without gaining any unplanned "adventure badges" on your vehicle.

Over time, as fiber reached deeper into neighborhoods, which shortened amplifier cascades, it became possible to move to higher and higher orders of modulation in the 5-42 MHz upstream: 16-QAM and 32-QAM and 64-QAM via DOCSIS 3.0 SC-QAM. Today using DOCSIS 3.1 OFDMA, up to 1024-QAM will be viable, especially in DAA systems. Use of 2048-QAM may also be achieved, and 4096-QAM is within the standard. These increasingly bandwidth efficient formats allow ever-increasing amounts of data to be carried from homes outwards, to the Internet or cloud.

3. Differences Between the Upstream and Downstream Signal Paths

There are a few notable differences between the forward/downstream signal path, and the reverse/upstream signal path. They are briefly noted here.

One is channel widths. Because the upstream path was never envisioned (or designed) to carry video, its channel widths aren't a static 6 MHz, as they are in the downstream (home-facing) path. Upstream channel widths typically use one of three sizes: 1.6 MHz, 3.2 MHz, and 6.4 MHz.

The upstream signal path was envisioned as a way to move small amounts of information, such as a click to order a movie, or, later, a click of a mouse to request a web page. When voice-over-IP entered the service mix, audio signals began moving upstream. All are negligible, relative to the "carrying size" of broadcast video. So, until recently (hello, webcams!), traffic type was also a differentiator between what moved upstream vs. downstream.

Modulation is a third difference between the downstream and upstream signal paths. The width differences are to accommodate multiple modulation rates for sending traffic: QPSK, 16-QAM, and 64-QAM.

A fourth difference – and an omnipresent conversation – is the matter of noise and ingress funneling in the upstream direction, which makes upstream more susceptible to over-the-air (OTA) signals. As the upstream bandwidth grows, some OTAs flip from downstream phenomenon to upstream, and in doing so become more troublesome. In the lower spectral regions, it used to be the off-air analog channels, which vacated the band coincident with digital. There's the FM band, which sits between 88-108 MHz. Potential issue: Interference. There's also the Aeronautical Mobile and Radio Navigation, between 108-137 MHz. Potential issue: Signal leakage. And let us not forget legacy out-of-band signaling, used by some set-tops and modems to move things like guide data, and command-and-control information.

Some readers may remember the big-growth days of high-definition TV, and the concerns about having enough downstream capacity to carry them all. Suddenly, we needed to add capacity, adjust channel lineups, advance another leap in video compression (at the time, to MPEG-4), and roll out things like Digital Terminal Adaptors, or, for some operators, Switched Digital Video.

These days, downstream capacity is reasonably under control (even as 8K TVs started rolling into retail this summer). It's keeping ahead of the growth in upstream demand that drives a larger part of our plant augmentations. As it turns out, after DOCSIS 3.1, the only viable tool in the non-fiber-deep playbook, besides constantly splitting nodes (which is increasingly inefficient) is to add spectrum.

Ironically or not, while this paper was being written (summer 2021), the author was participating in nightly overnight maintenance window sessions aimed at bringing to production the new "scar detector" tool on live nodes to activate on the Mid-split band. It was about as good as it gets for upstream geeks!

4. HFC Spectrum Relationships: Mid-Split, High Split and DOCSIS 4.0

The HFC network in North America has been limited to 42 MHz or less in the upstream direction since there has been HFC plant. The launch of HSD services increased the focus on the upstream because of the central role it now plays in providing a quality Internet experience for the ever-increasing range of demanding real-time applications. Fortunately, the growth of Internet traffic per year has been generally quite predictable, although slightly less so in the upstream than the downstream. In the upstream, year-on-year fluctuations historically experienced periods where traffic has been more dynamic, and periods where it has been flat. It all nets out to an average usage per year that is predictable enough to let capacity planners do their jobs effectively.

There has been over 20 years of growth managed almost exclusively by a fixed amount of upstream spectrum between 5-42 MHz. The amount and type of traffic moving upstream largely populates the quality spectrum available and managing new growth has transitioned from new QAM carriers and node splitting to node splitting and more node splitting. The COVID-19 pandemic accelerated this activity [2]. DOCSIS 3.1 can be used below 42 MHz, but better QAM bandwidth efficiency is no match for spectrum when it comes to adding capacity – according to that Shannon guy (http://www.inf.fu-berlin.de/lehre/WS01/19548-U/shannon.html). For high SNR (Signal-to-Noise Ratio) cable networks, Capacity ~ [BW/3] * SNR [dB]. The key thing to note is "dB." Capacity increases directly proportionally to bandwidth, and only logarithmically proportional to SNR.

As nodes get split smaller and smaller, it tends to become less efficient to continue to split. It is less likely to yield a 50/50 split, so the full benefit of the split is not realized. Whereas a 50/50 split buys ~3 years of growth at a CAGR of 25%, if the node is split 60/40 or 70/30, it is less. It is not unusual for one port of a node to be naturally more heavily loaded with traffic than another, since these ports feed different neighborhoods, and one, for example, may include a student housing complex or have a high density of business customers, while another may service less online-active customers.

Figure 1 shows a set of upstream expansion options available for MSOs. Many are active or imminent. They can be viewed as sequential in time, with some overlap and market-based criteria informing the timing and path to 10G. The architectures are described further below.

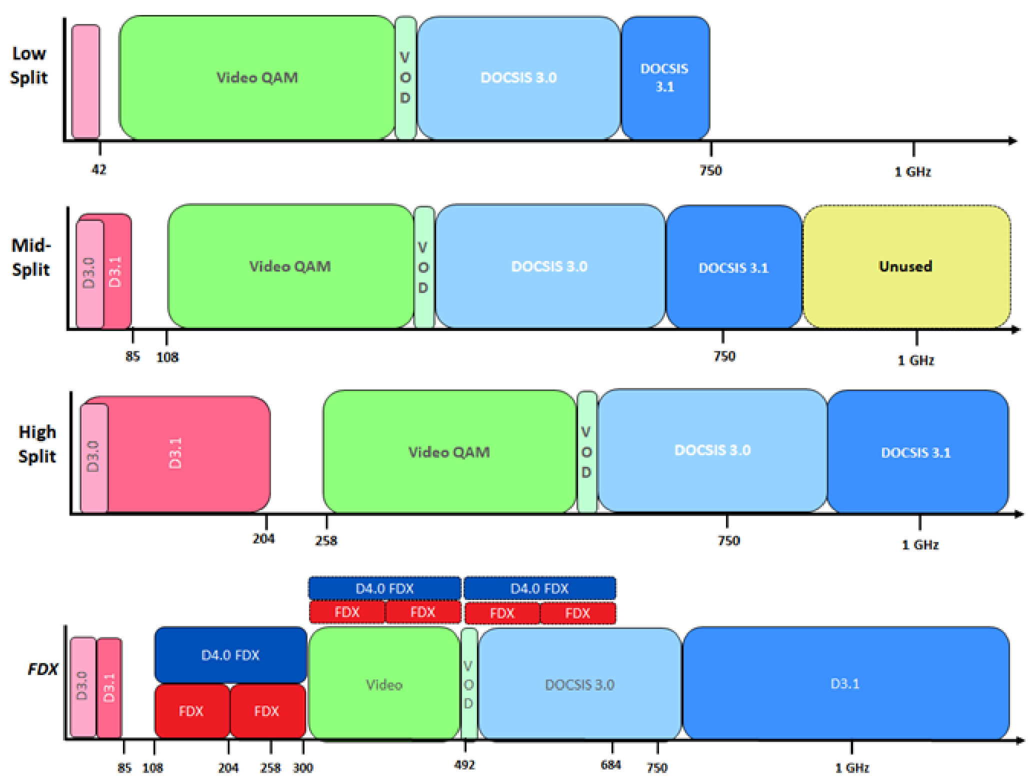

Figure 2 - Spectrum Migration Options through DOCSIS 4.0 (FDX only)

4.1. The Mid-Split: 5-85 MHz (defined initially in DOCSIS 3.0)

The Mid-Split is viewed as a practical steppingstone and with a relatively light touch because it resolves the upstream capacity challenge as we mathematically know it today. When combined with a node split, it defers additional augments to address congestion for at least 5 years, typically more (depending on D3.1 vs D3.0 assumptions). Furthermore, with an all-OFDMA channel, it can support around 600 Mbps. With a 4xSC-QAM DOCSIS 3.0 payload in the 5-42 MHz portion, about 450 Mbps is expected. Speeds up to 300 Mbps are expected in scale, under some traffic engineering guidelines tied to new utilization patterns.

An 85 MHz payload consistent with most MSO DOCSIS 3.0 usage today is 4x64-QAM DOCSIS 3.0 carriers, and a single OFDMA block from approximately 40 MHz to 85 MHz. This configuration is shown in **Figure 2.**

At this time (summer 2021), the Time and Frequency Division Multiplexing (TaFDM) feature, which allows spectrum to be DOCSIS 3.0 in some time slots and DOCSIS 3.1 in other time slots, has not been enabled. This remains an option, depending on the penetration mix of DOCSIS 3.0 and DOCSIS 3.1 modems and the net efficiency provided.

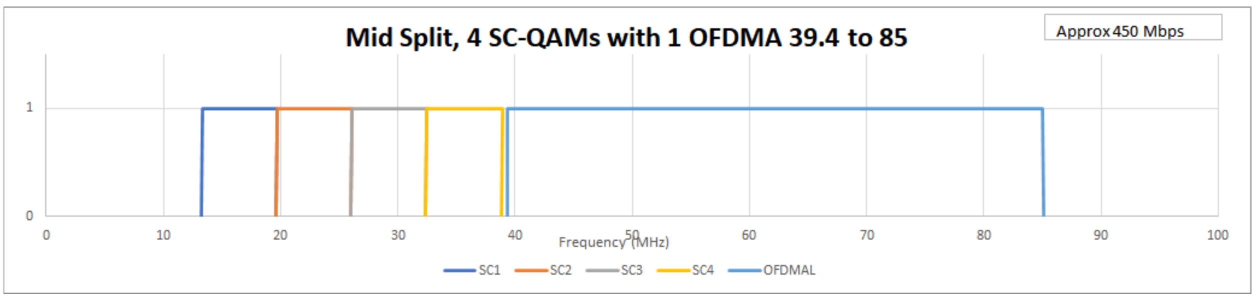

Figure 2 – Mid-Split DOCSIS 3.0 + DOCSIS 3.1 Loading Configuration

4.2. The High-Split: 5-204 MHz (efined initially in DOCSIS 3.1)

The High-Split is another popular option, as it stretches the speeds possible in the upstream to 1 Gbps or slightly more, as was demonstrated in the fall of 2020 [3]. Since the Mid-Split is such a powerful solution itself for capacity, going to a High-Split is a potentially very long-term solution with respect to capacity. However, practical capacity benefits are driven ultimately by the number of High Split-capable devices that can access that spectrum. There are many more Low-Split and Mid-Split modems deployed today compared to High Split, although this could change over time, and in particular for those that deploy with High Split.

4.3. DOCSIS 4.0

Like the High-Split, the primary value of DOCSIS 4.0, FDX or FDD, is upstream speeds. DOCSIS 4.0 fully attacks the historical asymmetry of downstream and upstream capacity, bringing multi-Gigabit symmetric capability to HFC. As it is defined today, the upstream will achieve 5-6 Gbps when fully activated. A path to 10 Gbps upstream is available by extending the bandwidth in FDX or FDD above today's 684 MHz limit with two more OFDMA blocks, to 1068 MHz. While this is easy to draw on a diagram, it creates challenges like upstream transmit power from a cable modem to overcome high coaxial losses.

Figure 3 summarizes the speeds associated with the options shown in **Figure 1**.

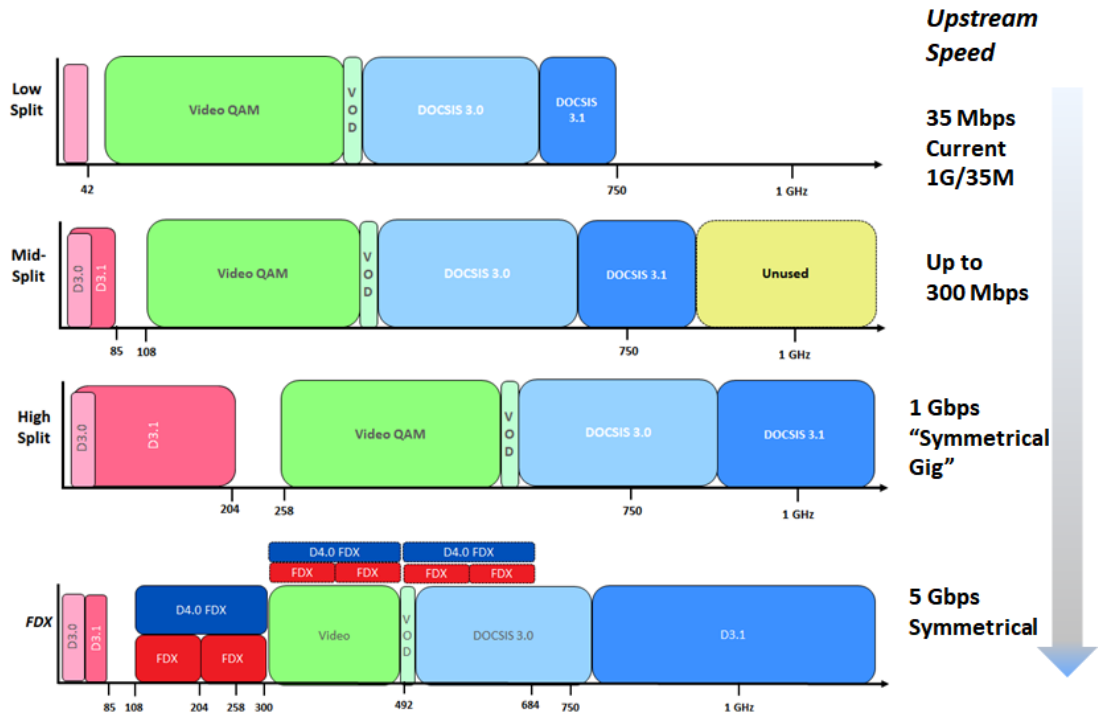

Figure 3 – Spectrum Migration and Implications to HSD Speed Tiers

We will get into the nitty-gritty details of the paper from this point on by mostly examining the Mid-Split scenario. Many of the same concepts are applicable to High Split, although there are some important differences that we will call out in the next section. There are deeper details, software and tool development, and mature processes that can be explained more readily using the Mid-Split case study due to its longevity, so we will lean on that for the bulk of the deep dives.

4.4. The Math

Mid-Split expansion takes the available upstream bandwidth from 37 MHz to a limit of 80 MHz. It was defined in DOCSIS 3.0, with the upper limit selected in part to fall just below the FM radio band in the US, while preserving the important downstream video out-of-band (OOB) signals widely used by legacy QAM set-top boxes (STBs). Per the earlier discussion, it is typically the upstream that drives network upgrade activity.

Because of the average per-user peak-busy-hour (pbh) upstream is still in the hundreds of kbps range, the upstream payload generally grows more slowly than downstream. Plus, because the new upstream spectrum is much cleaner, the Mid-Split impact on network lifespan is extremely powerful.

Figure 4 shows the time runway generated by three options – node split, node split plus upgrade to Mid-Split, and finally N+0 with Mid-Split. While N+0, with smaller service group size, offers the longest runway of the three, an N+x migration tied to a node split is also a very effective way to extend HFC lifespan to nearly 7 years in this analysis.

A key benefit of N+x with spectrum migration is its ability to add capacity quickly when compared to N+0. With the Covid-19 spike eliminating months of CAGR lifespan, N+x upgrades bring more US bandwidth to the network quickly to reset the lifespan timeline. The naturally slower pace of deeper fiber construction will leave too many areas without an augment for too long of a period of time. With the capacity growth "time" erased due to the pandemic, alternatives such as drop-in HFC upgrades that are both fast and effective make a sensible augmentation step. Having a diverse strategy, not one-size-fits-all, adds important flexibility to deal effectively with adjustments for situations like Covid-19.

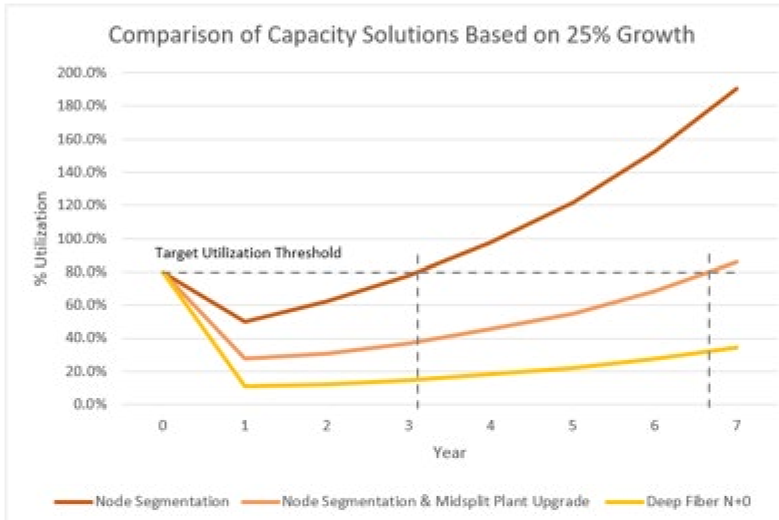

Figure 4 – Upstream Lifespan Expansion Options [2]

Lastly, looking ahead to future capacity and speed demand, and coupled with the objective to push fiber deeper into the network whenever possible, adjustments are being made to the architecture where it makes sense. For example, adding fiber in an underground network without the benefit of conduit is an inherently slow process. However, by providing the flexibility to allow a strategically placed amplifier (e.g., to allow an N+1 network) or two, there will be less construction, increased node size, and decreasing cost per household passed (HHP.) Combined, all speed the pace of the network upgrade and deliver the added bandwidth to more HHPs/year.

4.5. How About Some Real Visuals, Larry?

Comcast and other MSOs, such as Shaw Communications, have been building and activating Mid-Split spectrum for about 5 years. As much as can be gleaned from the crisp PowerPoint visualizations that led up to this moment is not nearly as exciting as displaying the real thing! **Figure 5** shows an activated Mid-Split upstream of 4xSC-QAM and the rest OFDMA. Upstream traffic is bursty, so this spectrum tool, the Yeti upstream spectrum analyzer application, has been placed in Max Hold mode to show the full 85 MHz band over time being utilized by the CM traffic. Most of the traffic is still in the DOCSIS 3.0 SC-QAM – the "Heatmap" view would show this – because most of the CMs in the system are DOCSIS 3.0. This balance is changing rapidly but DOCSIS 3.0 CMs are still the majority.

Figure 5 – Activated Mid-Split of 4xSC-QAM + OFDMA

5. New Spectrum, New Challenges

It is a sizeable project to upgrade the access network to support a new spectrum split. While some of the equipment is hosted in Hubs and Headends, where it is more easily accessible and centralized, upgrading outside plant (OSP) is more challenging. It requires going into the field, to every active device which has a diplexer – which is to say, every active device – and upgrading it to the new split. In some cases, this is something that can be done by changing a plug-in filter inside of the housing (not a live housing, removed from network) but in most cases it is not this simple. Many amplifiers in the field are decades old, made by vendors who no long support the product line, requiring swapping of devices altogether. Regardless, many operators have made the decision that the time is now for a frequency split upgrade and are committed to executing it.

Most of the above applies directly to an upgrade to a High-Split when it comes to upgrading actives in the field. However, the nature of the upgrade to Mid-Split is a lighter touch, with respect to other important variables. These are important to understand as part of a decision criteria on the upstream plan.

The seemingly intuitive “if more spectrum for Mid-Split is good, then even more spectrum for High-Split must be better” runs up against some significant new complications, outlined below.

5.1. Legacy QAM STB OOB Carrier

QAM video STBs that do not have a DOCSIS STB Gateway (DSG) control channel utilize the SCTE 55-1 or SCTE 55-2 protocol to get the necessary information to the STB. Signals returned from the STB over the upstream path in the 8-12 MHz range, typically. For fellow upstream geeks – yes, that part of the spectrum is terrible, but the protocol calls for very simple, robust modulation that is inefficient, but the traffic requirements of this modem are very low by today's HSD standards.

In the downstream, however, the "out-of-band" (OOB) carrier must be in the band 70-130 MHz according to the standard. When the upstream stops at 85 MHz, there is plenty of spectrum to place the OOB signal. When the spectrum extends to 204 MHz, legacy QAM STBs are stranded unless they can receive this OOB channel some other way. There are creative ways to do that, however it is some version of one-off solution.

Some of the thinking at the time of the DOCSIS 3.1 standard was that legacy QAM STBs were on the decline and would largely be out of the network by the time the High-Split was deployed. In addition, there was a move towards all-IP video delivery, which is still true today - although with somewhat less urgency based on changing business dynamics of the 10 years since the specification was being developed. One such class of box (General Instrument / Motorola DCT2000) is a model old enough that its tuner is at a fixed frequency and will not tune up or down from this frequency, which is about 72 MHz. Because it cannot tune up, it is incompatible even with Mid-Split. As a result, where Mid-Split is deployed, a pre-requisite is to swap these STBs out of the network. Due to the age of these STBs, the number of these STBs are very small, and the burden thus relatively low.

This is one of the very important aspects of High-Split compared to Mid-Split. Eliminating the OOB can make it a more invasive procedure for customers, as the best operator option is to extract these non-DSG-capable STBs, and this is more likely to "leave scars." The other import aspect attributable to High-Split is the Neighbor Interference (NI) phenomenon, which we will discussed later in this paper.

5.2. Aeronautical Leakage Band

The Aeronautical band, 108-137 MHz, is one in which there are requirements on operators to ensure there is not egress above a certain amount that could interfere with over-the-air (OTA) users in that band. Operators have mature processes and equipment to monitor this, placing "leakage carriers" in and around this band (and others) to measure leakage systemically to ensure compliance. It acts inherently as plant hygiene, so there is substantial benefit to operators - because where there is egress, there is the possibility of ingress. Ingress, of course, has been haunting the upstream for many years, and especially as the DOCSIS HSD upstream has grown more critical.

In practice, tones are placed close to the 108-137 MHz band in the downstream spectrum line-up and measured by specialized equipment. Of course, this works fine for Mid-Split. Mid-Split ends at 85MHz and the leakage band begins above that. However, this band, for High-Split is now in the upstream. To the letter of the FCC requirements, the upstream transmit power of a specification-compliant DOCSIS 3.1 modem cannot exceed the FCC limit, even at maximum transmit power. However, rather than do away with this aspect of plant maintenance altogether and lose the value it brings to operators and regulators alike, leakage measurements in this band will continue. Techniques which can measure leakage coming upstream from the modem are required, however, which is much different than today. Since these transmissions are bursty and short, the probability of catching them using today's auto-pilot drive-by method is not sufficient. Instead, techniques that use probe signals sent from CMs that are scheduled, and such that burst detecting equipment – also new for these meters – can capture the burst and assess the leakage performance is needed. It is a more complex and intricate solution than is needed on the downstream, but early proof-of-concepts have shown it viable.

5.3. Cable Modem Maximum Upstream Transmit Power

The maximum Total Composite Power (TCP) of a DOCSIS 3.1 cable modem is 65 dBmV. An "average" upstream transmitter in our footprint launches at about 43 dBmV/6.4 MHz. Extrapolating to 4x SC-QAM carriers, this becomes a TCP of 49 dBmV, still plenty of headroom to 65 dBmV. Extrapolating a uniform Power Spectral Density (PSD) over Mid-Split, this becomes roughly 53 dBmV using the Mid-Split configuration of Figure 2 – again, still plenty of headroom, but keeping in mind that 43 dBmV was the average.

The 90% point for upstream TCP is about 51 dBmV/6.4 MHz, meaning 90% of cable modems transmit 51 dBmV/6.4 MHz or lower. Extrapolated to Mid-Split, this is a TCP of about 61 dBmV. Our headroom is disappearing! Indeed, the 55 dBmV/6.4 MHz case, which would be the limit of the TCP that the modem can transmit, when extrapolated to Mid-Split, is about a 99% point on the cable modem upstream, Tx power cumulative distribution function (CDF). Of course, 1% represents a very small relative likelihood of running out of gas, but it is certainly not negligible for a large DOCSIS footprint when the population of DOCSIS 3.1 devices (all Comcast Mid-Split capable devices are DOCSIS 3.1) in the field is growing.

Now consider these extrapolations for the High-Split case:

Average US Tx: 43 dBmV → TCP (High-Split) = 58 dBmV

90% Point US Tx: 51 dBmV → TCP (High-Split) = 66 dBmV

This suggests that there could be a significant increase in the number of CMs that will be transmitting at their maximum and more, reaching the amplifier or node port at lower than designed levels, all else the same. This may impact network performance (lower MER) and throughput, in addition to creating challenges for operations and maintenance in aligning the network. It is unlikely this would be noticeable to a customer during normal use of Internet applications. But it could increase slightly the probability of a speed test failure for a 1 Gbps upstream service, which does not have a lot of capacity headroom above 1 Gbps.

5.4. FM Band

One of the benefits of the Mid-Split spectrum stopping at 85 MHz is that the FM radio band begins at 88 MHz. Indeed, by the time discussions about an expanded DOCSIS 3.0 spectrum were happening, the role of plant ingress and impact on the upstream was beginning to be felt and understood. FM radios broadcast from 88-108 MHz and can be very powerful signals when nearby transmitting antennas on major stations with the most powerful signals. It is expected that this band will suffer in terms of guaranteed MER across part of the network.

There is no better option for working effectively through an FM band with residual radio noise ingressing onto the cable than using DOCSIS 3.1 OFDMA. However, if FM radio signals create high interference, as reflected by a low MER, there is only so much that can be done by OFDMA. The 88-108 MHz span is a modest chunk of spectrum – about 12% of the newly added capacity for High-Split, so the impact also is expected to be modest.

As with the Total Composite Power (TCP) case, this scenario could also make it more difficult for the High-Split solution to achieve the 1 Gbps or greater target, given the small amount of headroom that exists.

6. Tool Time !

A suite of existing Comcast tools is essential to Mid-Split activation, which will become apparent as we describe the new tool development and tech ops processes supporting the initiative. This section is an introduction to the essential tools that we will reference along the way.

6.1. Premise Health Test (PHT)

Premise Health Test is a tool used to assist technicians in the diagnosis of any customer premise. PHT has evolved over nearly a decade, beginning as Home Integrity Check (HIC), starting with DOCSIS-only measurement values. It has been expanded to be more comprehensive, including many Proactive Network Maintenance (PNM), Wi-Fi, MoCA, EPON and other measurements. The test is usually invoked before and after installs and repairs to provide outlet-level readings from the installed equipment, where available. In addition to pass-or-fail, PHT also provides details about the service to facilitate the troubleshooting process. **Table 1** has a complete list of pass-or-fail criteria.

Table 1 – PHT Pass-or-Fail Criteria

Metric	Upper Fail	Lower Fail	Single Threshold Fail
Actual US TX	> 54 dBmV	< 25 dBmV	
Partial Bonding			Registration state <>4 Downstream state <>1
FLUX ICFR			>= 3 dB ICFR (ICFR Indicators)
DS RX	>13 dBmV	<-13 dBmV	
DS SNR			< 33 dB
SpectraCM Impairments*	* Full Spectrum Devices Only		Any Individual Impairment <>ACP
MoCA PHY Rate * MoCA Network *	* MoCA Capable Devices Only * MoCA Segmented Network Devices * MoCA Unexpected/Foreign Devices		< 200 Mbps (XG to XI, XB to Xi, XG to XB devices only; Xi to Xi or RNG150 to RNG150 not in scope)
FM Ingress	* All DOCSIS Devices		Severe Ingress Condition
EPON	>-8.0 dBm >-4.0 dBm	<-28.5 dBm Downstream <-28.0 dBm Upstream	Please Refer to Market RTM Process for any Failing Light Level Conditions
All Out			All Devices are Unresponsive
Wi-Fi – RSSI	* Re-Launch 10-29-19; None Pass/Fail		< -70 dBm RSSI Range at Xi5/Xi6

6.2. Yeti

The Yeti tool comes from the PNM suite, providing real-time upstream spectrum capture information to users. Historically, operators have relied on hardware-based spectrum analyzers to perform this function. Since the advent of PNM, spectrum capture capabilities are now available in the cable modem, from both downstream and upstream burst receivers. Upstream spectrum capture implemented within the burst receiver provides several distinct advantages over external, hardware-based solutions. First, it eliminates the need for additional hardware, which typically occupies valuable headend space and requires facilities power and cooling. It also allows operators to take advantage of the powerful burst receiver demodulators and CMTS core scheduling information. For example, in-channel demodulator performance metrics can be displayed along with the spectrum capture traces, with thresholds and colorization to aid in human interpretation. Another powerful feature is the CMTS scheduler's "quiet time" mode, which captures spectrum traces when no modems are transmitting. This provides users with a "noise-only" view of the upstream spectrum, simplifying the troubleshooting process by removing the cable modem bursts. **Figure 6** shows an example of the Yeti display, including SC-QAM and OFDMA bursts, in-channel demodulator statistics and threshold-based colorization.

Figure 6 – Yeti Upstream Spectrum Capture Display with SC-QAM and OFDMA

6.3. SpectraCM

This downstream spectrum capture tool provides a cable modem-oriented view of the RF spectrum. In the PNM suite, it's referred to as Full Band Capture (FBC) and provides the modems with downstream receiver spectrum capture. It's especially useful when upgrading from Low-Split to Mid-Split operation because of the switchable diplex filters in the cable modems. With the diplex filter operating in Low-Split mode, the cable modem is very effective at capturing the noise environment from within the home, which, as mentioned earlier, is often the source of RF ingress. Then when upgrading to Mid-Split operation, the frequency spectrum up to 85 MHz changes direction, creating a notorious funnel effect, which complicates troubleshooting. Having the ability to switch between diplexer modes of operation allows operators to automate this traditionally difficult troubleshooting process. **Figure 7** illustrates an example of SpectraCM being used with Yeti to match and locate ingress of VHF television signal ingress.

Figure 7 – SpectraCM (bottom) Locates Ingress with Yeti (top)

7. Activate Spectrum So No One Will Notice

It is an axiom of good network strategy and upgrade practices to think of success like a baseball home plate umpire does: You've had a good game when nobody notices that you were there. This is absolutely the case for migrating spectrum, with the caveat that they will notice in a good way that you were there, as the products enabled by these network upgrades – in particular upstream HSD speeds – become available. Until then, though, the internal satisfaction that is flat trouble call metrics and meeting schedule and budget targets will have to do.

We describe below the potential issues to manage and share practices that help to achieve a seamless customer experience.

7.1. Activate Spectrum So No One Will Notice

With the decades of spectrum split in North America being set at 42 MHz/54 MHz, all QAM video STBs deployed are configured this way. They were built to receive video channels beginning at 54 MHz. When the network is NOT configured this way, and instead is upgraded to enable the diplex split to expand the upstream and activate new spectrum above 42 MHz, the QAM STB's point of view for Mid-Split or High-Split changes. This is shown in **Figure 8**.

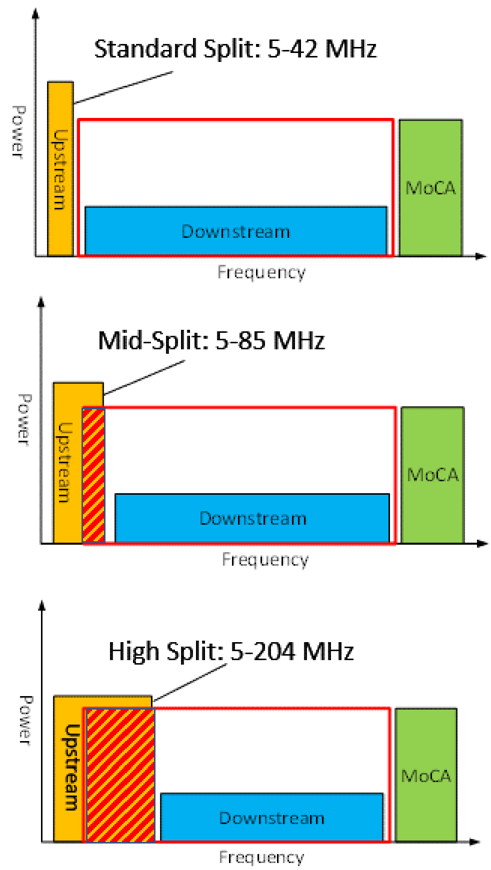

Figure 8 – New Spectrum Splits vs. Standard Deployed Equipment

The red cross-hatched areas in **Figure 8** represent the spectral overlap imposed on a QAM STB by a Mid-Split capable cable modem when utilizing that band. Any signal energy that appears above 54 MHz can be seen by the STB downstream receiver, because it is built expecting to operate on downstream signals that begin at 54 MHz. Unfortunately for the STB, in a home that also contains a Mid-Split capable cable modem (CM), the CM sees that band as "eligible" for placing carriers when the CMTS is configured to allow CMs to use it. In a CMTS that is properly load balancing, with much of the existing traffic volume generated by devices with a 42 MHz limit, and OFDMA turned on above 42 MHz, the Mid-Split capable devices would expect to be utilize the 40-85 MHz spectrum for transmissions.

Note that the RF processing front-end of the STB is not acting on any specific signal type – it is simply adapting its front-end gain to deliver the ideal level to the A/D converter using its Automatic Gain Control (AGC) function. AGC measures the total energy in the downstream band and doesn't care about its origin. It adds or subtracts gain to deliver the signal intact and at the optimized level – a balance between Signal-to-Quantization-Noise Ratio (SQNR) and clipping distortion – to the A/D converter, on its way to the digital processing. Thus, if new Mid-Split upstream energy on the STB receiver is very high, the STB receiver will add attenuation to keep the right operating point on the A/D converter. When this happens, the *desired* video channels will inadvertently be pushed into the noise floor through a phenomenon called "Adjacent Channel Interference" or ACI. If it attenuates too much, then the QAM video signals can be pushed down enough to cause low SNR in these channels, and video pixelization could ensue. Potential makeover scarring alert!

Note the above description is a "static" or time-fixed snapshot view of upstream energy and spectral overlap with signals moving downstream to a STB. Actual upstream traffic is called "bursty" because it bursts on and off. This is important because the AGC function has dynamic characteristics, but they tend to be slow (levels don't change very much, typically). As a result, the duty cycle (off/on ratio) and burst duration of upstream signals is a factor that can impact the AGC implementation of different STBs.

Note that, as represented in **Figure 8**, the nature of the levels is not favorable – the downstream receive level is low (DS Rx), while the upstream transmit level (US Tx) is high. Until now, there was a diplex filter to separate them, but now, between 54-85 MHz, this is no longer the case. As discussed, the US Tx average level is about 43 dBmV/6.4 MHz, with the vast majority at 51 dBmV/6.4 MHz or less. By contrast, downstream receive levels typically target 0 dBmV/6 MHz, but range widely by design from a minimum of about -12 dBmV/6 MHz to a maximum of +10 dBmV/6 MHz.

Fortunately, between a CM and a QAM STB there will be an RF splitter, the design of which will inherently isolate the port of a CM from the port of a STB by some amount. This scenario is illustrated in **Figure 9**, showing just one isolation path between modem (MTA) and a STB's RF inputs.

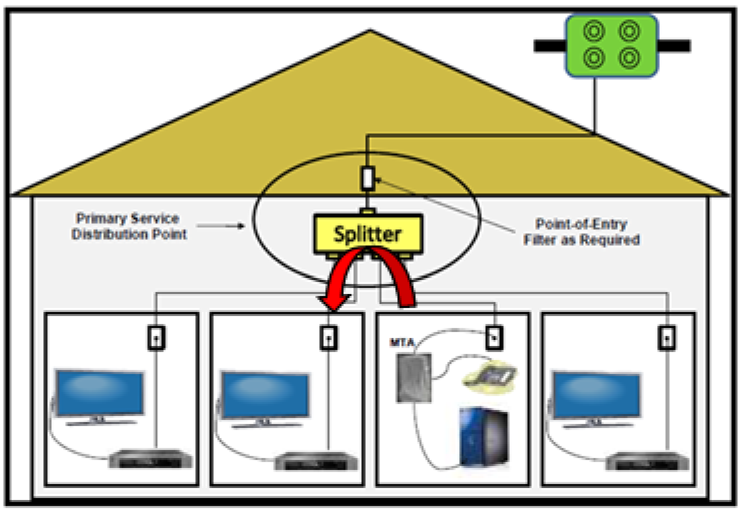

Figure 9 – Mid-Split Band Energy Isolation Path Across and RF Splitter

How much energy leaks through to the STB? That question depends directly on the splitter and home wiring shown in **Figure 9**. Home wiring – RG6 cable – has a very predictable dB/loss per foot and is easily modeled, from which some basic assumptions can be made about the range of run lengths in non-celebrity, which is to say "reasonably sized" homes. The most important factor with respect to the ACI phenomenon is the splitter(s) used to distribute RF to devices for video and data services. **Table 2** shows the specification for the isolation parameter (paragraph 13.0) for approved Comcast splitters.

The parameter used to evaluate the relative risk of video interference is called Carrier-to-Adjacent-Channel-Interference Ratio (CACIR). The threshold at which video degradation can be observed varies by STB model, and these have been individually characterized for every model still in use. Automated tools can discover the model type directly or through information in the billing systems and can apply a CACIR threshold according to the model type. For purposes of this paper, we will use the worst case empirically observed lab test value of -22 dB CACIR as the threshold value for deriving statistics. It seems like a 20 dB higher signal should completely blow up an RF front-end, right? However, note that the -22 dB is a single carrier-to-single carrier comparison. There is, of course, more total bandwidth in the downstream than the upstream, so the total difference in power of the interfering energy to total downstream energy is less than this.

Testing was done as both "always on" and over a range of fixed duty cycles meant to emulate the "on/off" bursts of real upstream traffic. It is only when the upstream is bursting on, of course, that there are signals that can interfere with a downstream STB.

Table 2 – Port-to-Port Isolation of Comcast Approved RF Splitter

COMCAST

Field Equipment Specification: Mid-split Passive RF Drop Splitters

<table>
<tr><th>ID</th><th>Specification</th><th colspan="4">Requirements</th><th>Test Method / Other References</th></tr>
<tr><td rowspan="3">12.0</td><td rowspan="3">Return Loss</td><td colspan="4">Minimum RF Port Return Loss, All Other Ports Terminated (MHz)</td><td rowspan="3">ANSI/SCTE 144 2007</td></tr>
<tr><td>5-1002</td><td colspan="2">1003-1218</td><td>1219-1675</td></tr>
<tr><td>≥ 18 dB</td><td colspan="2">≥ 15 dB</td><td>≥ 5 dB</td></tr>
<tr><td rowspan="3">13.0</td><td rowspan="3">Isolation</td><td colspan="4">Minimum Output Port to Output Port Isolation (MHz)</td><td rowspan="3">ANSI/SCTE 144 2007</td></tr>
<tr><td>5-10</td><td>11-85</td><td>86-1002</td><td>1003-1125</td></tr>
<tr><td>≥ 25 dB</td><td>≥ 35 dB</td><td>≥ 25 dB</td><td>≥ 20 dB</td></tr>
<tr><td rowspan="4">14.0</td><td rowspan="4">Maximum Output Port to Output Port Isolation 1,125 MHz-1,675 MHz</td><td colspan="2">Splitter Device</td><td colspan="2">1126-1,675</td><td rowspan="4">ANSI/SCTE 144 2007</td></tr>
<tr><td colspan="2">2-way</td><td colspan="2">≤ 28 dB</td></tr>
<tr><td colspan="2">3-way (unbalanced)</td><td colspan="2">≤ 30 dB</td></tr>
<tr><td colspan="2">4-way</td><td colspan="2">≤ 30 dB</td></tr>
</table>

Comcast specifies a minimum isolation through the Mid-Split band of 35 dB. That says that a lot of the upstream transmit power is going to be attenuated on the way to the STB. (Yay!) Is it enough? Usually. And, fortunately, when it is not, it is discoverable and easily remediated. For a scar-free upstream makeover, the key statement is "discoverable."

Because of the importance of quantifying the potential for ACI, an existing model of Comcast-approved splitters was measured for actual performance. Practical performance of these splitters is shown in **Figure 10**. Note that the band between 54-85 MHz is, in the case of this model, the "sweet spot" of RF isolation, with 45 dB being more characteristic of splitter performance. This is an extremely valuable 10 dB with respect to quantifying the potential to impact video.

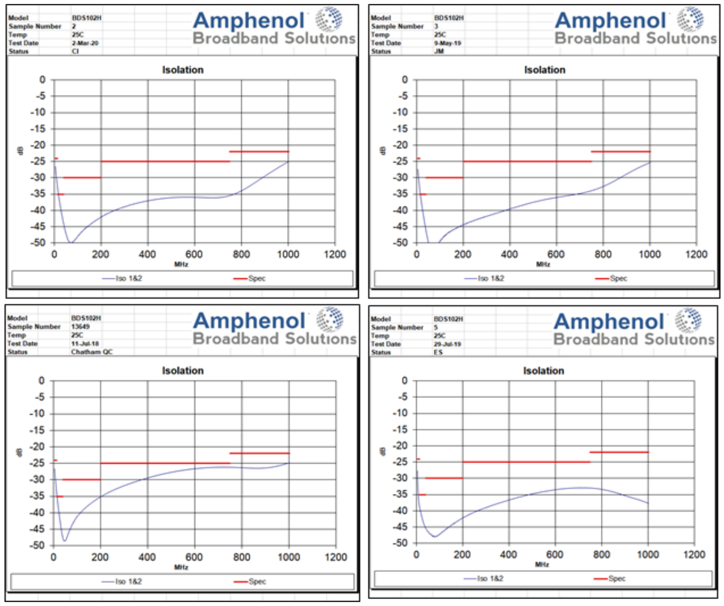

Figure 10 – Measured Port-to-Port Isolation of Comcast Approved RF Splitter

For every "sweet spot," there is of course a counter example. In the ACI scenario, the most concerning counter example, with respect to video service impact, is the use of a low-cost, off-the-shelf splitter that might be found in the cable TV accessories section of a home improvement store, such as Lowe's or Home Depot. (For old timers, this is where we used to say, "Radio Shack," but you may be hard pressed to find anything related to radio on the shelves there anymore ... if you can even find a storefront.)

Figure 11 shows the isolation performance of this type of splitter.

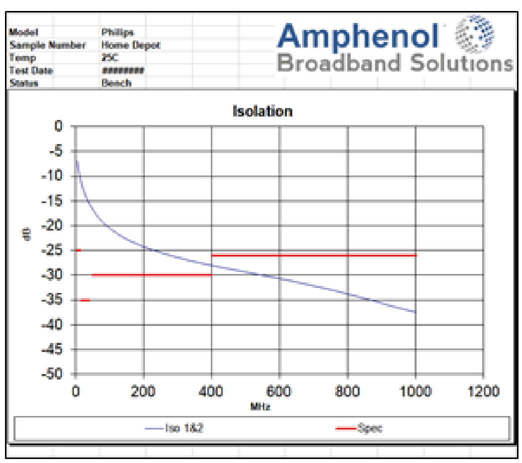

Figure 11 – Measured Port-to-Port Isolation of an Off-the-Shelf RF Splitter

We believe these retail splitter scenarios create the most likely risk to STB video degradation, as they perform with 15-20 dB worse isolation than the Comcast requirements. This in-home scenario is easy to envision occurring in practice and is generally out of the operator's control. Through the automated diagnosis of homes slated for a wider upstream, we can find such scenarios and take proactive measures prior to activating the Mid-Split spectrum.

Eliminating the possibility of video interference attributable to the activation of Mid-Split spectrum is a primary criterion for smooth spectral transition. Using empirical data of measured Comcast DS Rx and US Tx from production CMs, and typical home network assumptions for splitter and coaxial LAN runs, **Figure 12** shows the probability of ACI reaching the threshold for video interference to be 1.29% for the worst-case sensitivity among all STBs tested. This number looks small, and it is. However, when measured against the total number of broadband subscribers with video service, it is not negligible, and needs to be managed with proper tools and processes.

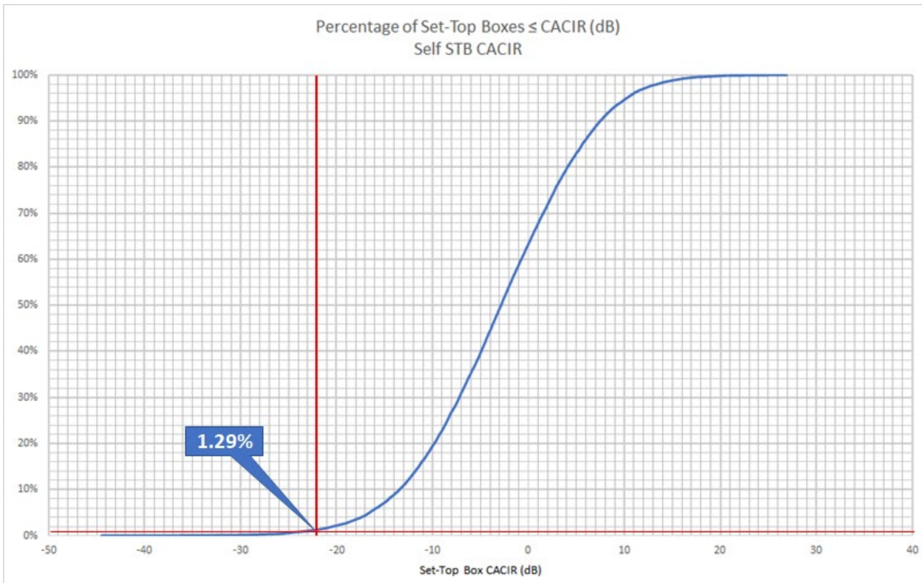

Figure 12 – Probability of Interference ≥ ACI Threshold of STB

One important note is that **Figure 12** shows the correlation of measured dB relationships with a lab-observed video impairment, under a set of fixed upstream transmission patterns. In real life, the upstream duty cycle is low, and the transmissions are relatively random in both size and duration. It is difficult to precisely correlate RF impairments measured in dB, to customer-impacting video degradation from real traffic, and further, to degrade it enough to generate a trouble call (versus the so-called silent sufferer – a worse scenario). While it is straightforward to create an impaired condition in the lab, how this translates to field exposure will be something that will be continually learned over the course of trials and the scaling-up of Mid-Split activation. That is the genuine way to test the hypothesis and lab measurements when encountering real traffic.

7.2. In-Home Amplifiers

In addition to OSP and traditional CPE devices that provide residential services that only know the 5-42 MHz Low-Split, many homes also use drop amplifiers to overcome losses across the in-home coaxial network. As you might expect, these are also built with a Low-Split diplexer.

These devices must come out…eventually…. but because they may or may not be customer-impacting, they do not *necessarily* have to be tackled coincident with the activation of Mid-Split spectrum. From a capacity perspective, every drop amplifier that can be removed so that it doesn't block Mid-Split energy from exiting the home is good for capacity. The operations perspective depends on the percentage of homes that include a drop amplifier – estimated at 15-20% but can cluster depending on geography and practices.

Methodically removing in-home drop amps over a period of time may make more sense than dealing with amplifiers transactionally, meaning only as part of a service call or product upgrade. A proactive plan to address drop amps will eliminate the perpetual limbo state that is mixed-mode devices working in mixed-mode spectrum.

With capacity and product in mind, we can itemize home amplifier management into two buckets:

1) Capacity-driven – Referring again to **Figure 4**, the coaxial lifespan when doing digital node splits and upgrading the network to Mid-Split is shown to be almost 7 years. This includes new capacity made available by Mid-Split (450 Mbps used), which depends on the DOCSIS 3.0 QAM bandwidth consumed and no considerations for TaFDM. Again, this only works if modems can access the Mid-Split bandwidth, which only happens if the CM is capable (new enough) of doing so. All DOCSIS 3.1 CMs we use are Mid-Split-capable, and our DOCSIS 3.0-only CMs are not. CMs are migrating to DOCSIS 3.1 status steadily, so over the 7-year period it is safe to assume the vast majority will be installed and capable of Mid-Split upstream connectivity. However, even if a Mid-Split-capable CM is present, any home that cannot allow the spectrum to pass out of the home is one that cancels the capacity gains. The "real" penetration of DOCSIS 3.1 OFDMA is decreased accordingly, and the 7-year lifespan is compromised. Thus, as mentioned, over time, these amplifiers must be removed, so that the capacity plan can deliver on its lifespan promise. How quickly this must be done is a mathematical analysis of utilization vs the "real" penetration trajectory.

2) Product-driven – One of the key benefits of the Mid-Split is the ability to deliver HSD speeds in the upstream such as 100 Mbps, 200 Mbps, 300 Mbps, even higher, as OFDMA begins to replace DOCSIS 3.0 QAMs in the upstream. Traffic engineering rules developed for these speeds account for utilization and total capacity and are considered reasonable expectations for potential product offerings. Once such products are made available, customers with home amplifiers will be (self)-blocked from receiving them. Interest in speeds that require Mid-Split would trigger immediate action, to remove the blocking amplifier. The challenge is how to manage this efficiently and, more importantly, in a way that is impact-free to the customer. The good news is these blocking devices (amplifiers or any filtering within the band that may have been installed inline) can be discovered remotely and in real-time. While the customer cannot get the new upstream speed immediately, a rapid and transaction-based process can serve to notify the customer that additional steps are required to support the speed upgrade. That we have detected the need for additional steps is a communications decision that is out-of-scope for the purposes of this paper. Either way, appointment scheduling can commence to eliminate the problem and get the new product speed to the customer.

   The product case is a different kind of operational task than the "capacity driven" case and is difficult to pre-plan because it is impossible to predict exactly which customers seeking the new upstream speed will also have in-home amplifiers. This adds to the long and growing list of reasons to proactively replace drop amplifiers during, for example, scheduled truck rolls to homes that also have been discovered as having a drop amp. Or, as mentioned above, build a program to recover in-home amplifiers proactively, over time, in a way that spreads out the cost of the effort, controls it better, and unifies the network for all homes -- rather than during product requests, which introduce a reactionary mode.

In a nutshell, the problem statement for Mid-Split activation caused by the 54-85 MHz spectrum overlap identified in **Figure 8**, is to develop a way to unobtrusively discover the state of a home with respect to these two criteria:

1) Potential for video interference
2) Ability to support DOCSIS upstream pass-through in the Mid-Split band

To enable this home-by-home assessment in scale, we developed an automated in-Home Assessment Test – aka iHAT – to enable a seamless migration of capable CMs to utilize Mid-Split when conditions 1 and 2 above are satisfied, as shown in **Figure 13**.

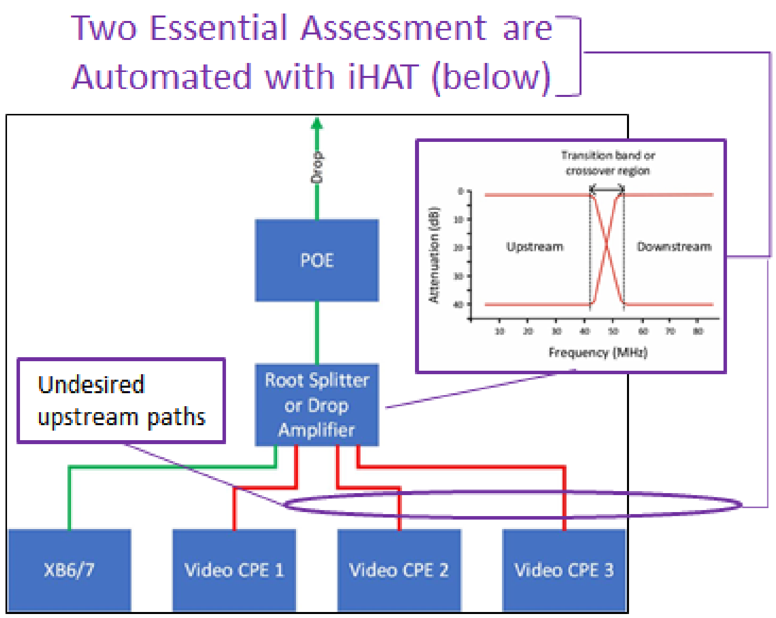

Figure 13 – The Two Basic RF Assessments Evaluated by iHAT

7.3. What's Under the HAT?

The "Black Box" view of iHAT that includes its functional core is shown in **Figure 14**.

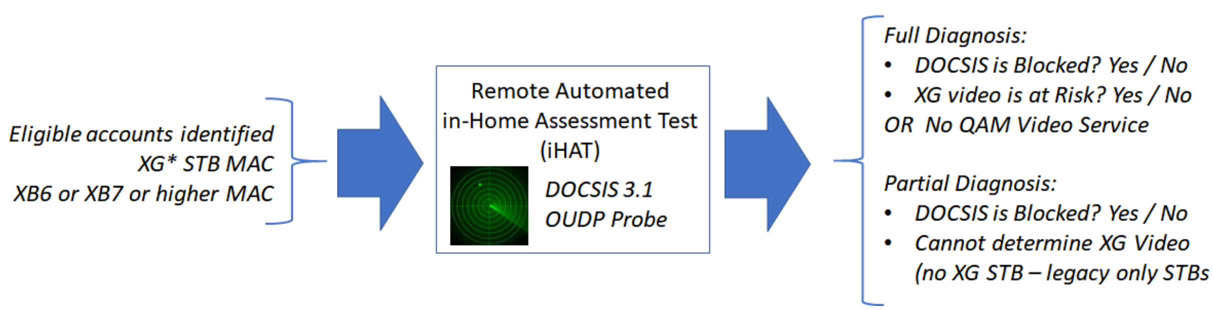

Figure 14 – The iHAT "Black Box" – Method and I/O

7.3.1. Incoming !

As shown in **Figure 14**, iHAT retrieves a list of devices, by account, on a particular Mid-Split-enabled RemotePHY Device (RPD) node, after the node is cutover, activated, and services restored that meet a pre-cutover state of performance (we are activating Mid-Split with OFDMA-only on DAA platforms). When an account is identified as having a Mid-Split-capable CM – for us, this includes the DOCSIS3.1 gateway family of XB6, XB7, and XB8 – it is deemed eligible for an iHAT test. With one of those devices present, it will be possible to place the CM in Mid-Split mode to determine whether its upstream transmissions in the Mid-Split band are able to be seen and received by a Mid-Split enabled vCMTS and DAA node, or if they are blocked.

When an account also includes the "XG" class of QAM STB, the iHAT evaluation will look both for DOCSIS Mid-Split pass-through and the potential for video interference. This XG family, the majority of QAM STBs in the Comcast network, supports the proactive network maintenance (PNM) and SpectraCM functionality needed to capture RF measurements that are the basis for iHAT scoring of video interference potential. Older QAM STBs do not support this capability. In a home that includes an XG class STB, that measurement taken is a reasonable proxy for the expectation for other non-XG STBs with respect to their isolation from Mid-Split spectrum energy.

If there is no XG-class STB present at all, but "legacy" QAM STBs are present, then no iHAT assessment can be made with respect to the potential for video degradation. At the outset, these homes will default to Mid-Split activation as data is accumulated. After some scale of statistical significance is built up, the policy will be revisited to determine if a course correction is needed. Also, as we observed in **Figure 12**, the risk of video impairment is very small. By NOT defaulting to activating in these homes, the alternative being committed to is to roll a truck to each home that only has QAM STB and take "iHAT" style isolation measurements manually when only a very small fraction may be impacted.

7.3.2. Start Your Engines

The method iHAT uses to make its determination is based on the DOCSIS 3.1 OFDMA Upstream Data Profile (OUDP) feature, which allows a pre-defined "probe" signal to be scheduled by a CMTS and generated as a test signal. The probe can be defined by center frequency, bandwidth, and duration. When iHAT runs, it schedules this probe signal, home by home, to be burst into a portion of the Mid-Split spectrum.

Figure 15 shows the probe signal centered at about 80 MHz. It is 1.6 MHz wide (a common reference bandwidth for OFDMA bandwidth used in the DOCSIS 3.1 requirements), has a PSD at the ranged OFDMA power, and lasts 3-5 seconds. These are empirically-derived values through trial-and-error testing and optimization in the lab.

Figure 15 – Probe Signal Used in iHAT via DOCSIS 3.1 OUDP Feature

When the probe is fired, the time stamp is used to instruct the XG STB when to execute a Full Band Capture (FBC), and with that capture, samples are returned to that include levels of the OUDP probe and the first few downstream QAM channels. By determining the relative levels of these components and comparing them to an interference threshold value, making offset adjustments that account for the test probe not occupying the Mid-Split band completely, the home can be classified as to whether it needs remediation.

The OUDP method provides three major advantages:

1) It is part of the DOCSIS 3.1 specification, so a required featured to be compliant to the specification (when asked for!)
2) It can be a scheduled event within a system's normal operation, and therefore is very non-intrusive, happening without a customer's awareness or service interruption
3) As a scaled down (in total power) representation of an actual upstream signal, it does not actually create enough interference to impact video. Instead, it emulates what a small portion of the filled spectrum would look like and extrapolates mathematically to draw the proper pass/fail conclusion.

Iterative optimization of the parameters yielded a repeatable, reliable result that correlates well as a mathematical extrapolation with the video threshold testing that forms the foundation of ACI analysis.

The probe signal can also be used to evaluate blocking of the Mid-Split upstream by a drop amplifier, because if this is so, the CMTS will not be able to observe the probe. However, as part of the iHAT test, Mid-Split becomes active on a modem prior to an OUDP probe being launched, once the CMTS has a configuration that supports it. Ranging information of the OFDMA band (DOCSIS 3.1 ranging) is available to determine if the upstream was successfully sounded. If not, this is typically sufficient cause to identify a home with a drop amp issue, at which point the CM can remain in partial service or reverted to Low Split mode. In either case, the home state is logged as "remediation required."

In future iterations of iHAT, a time-out on partial services re-tries will be used to force an auto-revert of the modem into the Low Split band. However, it is anticipated that, rather than do this with filter switching in and out, the vCMTS will support multiple bonding group (BG) operations that include both a 4-channel BG of all-DOCSIS 3.0 QAMs, and a 5-channel BG that is the former plus one OFDMA block of 40-85 MHz. This will simplify iHAT testing and shorten time consumed by avoiding modem reboots that force a diplex filter switch to Mid-Split, in order to execute the test. Instead, CMs will arrive on the scene in Mid-Split mode, by default, and if conditions such as drop amps block Mid-Split signal passage, then the lower 4-Channel BG will be deployed on that modem. In the case of a product need for that device (100Mbps upstream, for example), there will, of course, still need to be a rapidly-executed action scheduled for a good customer experience, to eliminate the blocking amplifier or filter, and to provide the service speed requested.

Figure 16 shows two sample outputs from iHAT. In these screen captures, the test was launched locally. To support deployment in volume production, these tests will run from the cloud.

In the top screen capture in **Figure 16**, we see the measurements being made by the iHAT tool. Notice the 2nd to the last value, which is what is compared to the 22 dB threshold "US OUDP Power minus DS Rx Power." This is well below the threshold and thus this measurement is a "pass." These values are stored for trend analysis and optimization. Another one of the values of particular importance for this is the last row, "Isolation." With iHAT, we now have the game-changing tool of being able to see the RF isolation between a gateway and XG STB in every home. Note also that 35.84 dB is very close to the Comcast isolation spec observed in Table 4.

In the bottom capture of **Figure 16**, we see the explicit result: "video interference was above threshold for at least one set top box," and also the isolation value, in this case called out by its variable name in the actual code as "dbCDelta" of 24.16 dB – above threshold.

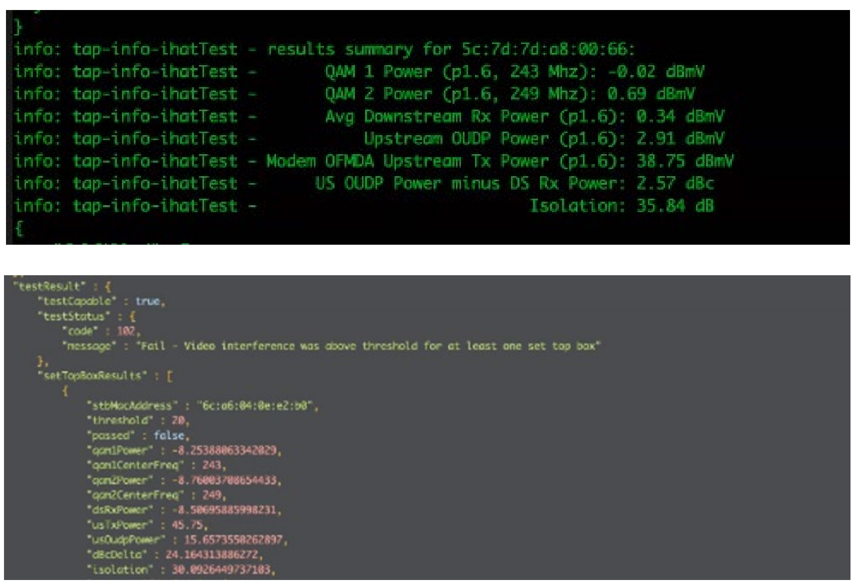

Figure 16 – iHAT Screen Captures of Video Assessment: Pass (Top), Fail (Bottom)

7.3.3. The Answer is…..

As shown in **Figure 14**, the output of iHAT is straightforward:

- DOCSIS Mid-Split pass-through (pass/fail)
- Potential for video degradation without intervention (pass/fail)
- Partial diagnosis – DOCSIS Mid-Split compatibility only; this is the case of a customer that has video services but no XG STB to support the telemetry needed to run the video assessment of iHAT

For a home to be declared ready to be activated, both DOCSIS and video tests must pass. If either does not pass, the home is left in Low-Split mode, and the home is dispositioned for remediation with the associated reason code (DOCSIS or video). How to optimally process the remediation queue itself is a discussion among many stakeholders. In addition to the pass/fail "answer" at the heart of iHAT, the RF measurements taken in the home, such as the isolation measured from the CM to the STB, are recorded and stored for purposes of trend analysis and iHAT optimization.

A fourth "state" that iHAT technically discovers is that the home is simply "ineligible" for Mid-Split because it has a CM which is only capable of Low-Split. In this case, the RF test engine does not run at all. This discovery occurs on the front-end, during the filtering of accounts connected to the RPD to only those that include a Mid-Split-capable CM.

Note that for a fully automated solution, iHAT receives input account/device data and develops output conclusions and an accompanying set of numerical parameters associated with the result for use elsewhere. In this sense, iHAT is the test function, with appropriate interfaces into other key back-office tools and subsystems, to operationalize the completely automated solution into the end-to-end ecosystem. In this context, iHAT is the "engine" of the overarching *Mid-Split Spectrum Upstream Launch* (MUSL) method, which we shall discuss next.

8. The Muscular Frame Supporting the iHAT Engine

8.1. MUSL-Up: End-to-End Device Activation Overview

As described above, the innovative iHAT tool provides a relatively non-intrusive view into a customer's home, and, on a home-by-home basis, will make a Go-No Go declaration with respect to readiness for activating spectrum in the Mid-Split band. iHAT scores a home's DOCSIS readiness for passing spectrum to 85 MHz, and its likelihood of creating video interference.

Referring to **Figure 17**, we can show how iHAT fits within the broader operational perspective, going from the trigger of a DAA Mid-Split node cutover on the far left, to the completion of activation on the far right. Note that "Tier 1" and "Tier 2" refer to different categories of markets which inform the upgrade strategy used at Comcast, however they have no real bearing on the flow otherwise.

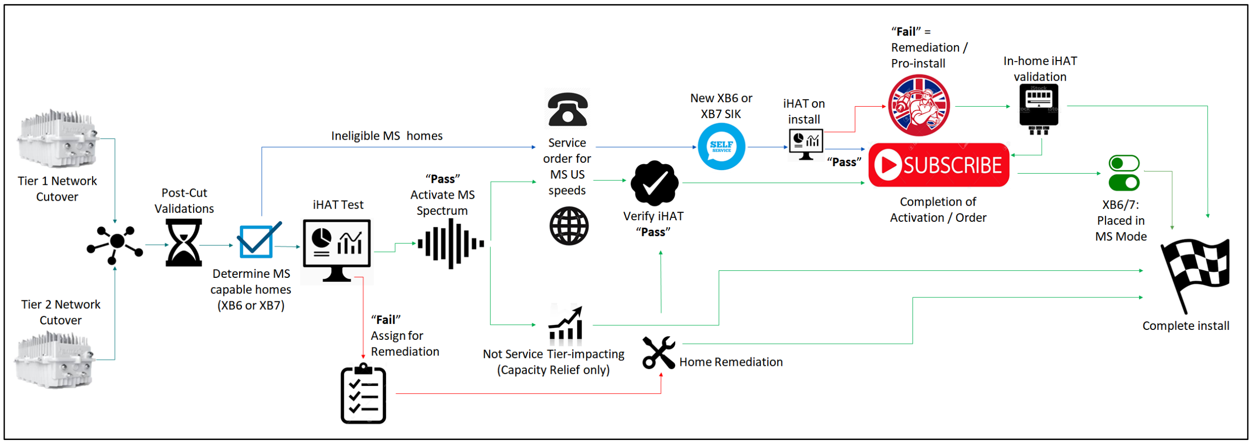

Figure 17 – Simplified Mid-Split Activation Flow: Cutover Through Activation

Beginning on the left, when a Mid-Split network upgrade occurs, internal tools will notify systems when construction is complete and officially closed out. This triggers the spectrum activation process, notifying other tools that the network is now be able to take this step. Two things must happen prior to letting iHAT sweep across the node and validate homes where Mid-Split can be turned on. They are:

1) *Post-Cut validation* – Ensures that the network has resumed to BAU metrics after the cutover. It is not uncommon to have a short period of elevated network activity shortly after a cutover, and it is desirable to resolve any residual cutover issues prior to moving to Mid-Split. This can be time-based, or it can be directly associated with, for example, observation of trouble TC metrics, pre-cut vs post-cut.

2) *Determine which homes are eligible for activation* – This boils down to whether the DOCSIS CPE is capable of Mid-Split. At Comcast, all DOCSIS 3.1 Gateways are Mid-Split-capable.

On item 2) above, if a home is ineligible, then iHAT does not run. Following this arrow to the top path in Figure 2, there is no immediate required step to get that customer a Mid-Split-capable modem. There is an effective loss of capacity for every CM that cannot access the DOCSIS 3.1 spectrum, because it forces utilization in the Low-Split band, rather than accessing the faster and wider OFDMA spectrum.

There is guidance in the in the field on what triggers a DOCSIS 3.1 upgrade for a customer – a particular speed tier for example. Over time, DOCSIS 3.0 CMs will organically disappear from the field, and it is likely at some point there will need to be a proactive effort to remove the older CMs in the network to maximize the DOCSIS 3.1 capacity.

Now, as shown in **Figure 17**, when a customer decides to upgrade their speed tier to one that requires Mid-Split, then getting them a gateway capable of that becomes a priority. Also, this customer's home needs to be evaluated for its ability to be placed in Mid-Split mode. So, as a new Mid-Split-capable gateway is brought onboard, one of the first things it needs to do is call upon iHAT to determine the state of the home for Mid-Split. If the iHAT "pass" is recorded, then the activation process continues, and iHAT sets the device into Mid-Split mode. It then becomes capable of using the OFDMA spectrum between 40-85 MHz. If iHAT records a "fail," then the customer is notified that a technician must come to the home to complete their install, and that their new speed tier will not be available until this "Pro Install" step happens. When the remediation is complete, the technician will validate onsite with iHAT, in this case triggered locally from the Performance Health Test (PHT) application.

If the eligibility conditions are in place – Mid-Split capable CM, and a STB model with the necessary telemetry capability – we move to the right of the blue checkmark of **Figure 17**: "iHAT Test." Let's now follow the lower path under the iHAT test icon – "iHAT fail."

8.1.1. The Remediation Queue

As noted, unless there is a speed upgrade required by a customer, there is not necessarily an immediate need to provide them with a Mid-Split capable gateway. However, it is still important that the iHAT score be logged. The fact that the home needs to be remediated is documented and populated into tools used by agents and technicians. Homes in this category are placed into a "Remediation Queue." iHAT will identify the specific failure mode, so that technicians know what needs to be done when they arrive. In general, remediation tasks are well-understood and known to technicians, and include changing out home amplifiers for alternative devices, checking home splitter configuration, models, and wiring, to bring the home to Comcast compliance standards. After remediation is performed, the iHAT test is run to validate readiness for Mid-Split spectrum, and the activation then completed.

When to schedule a home for remediation, assuming there is no speed tier motivation, is a business decision. They can remain in Low Split mode until that time, with some impacts on the network side. There are multiple variables to consider that have to do with capacity, efficiency, and proactive expense. Ultimately, however, all homes in the remediation queue will need to get serviced to extract the full DOCSIS 3.1 capacity and maximize the upstream runway these architectures are made to deliver.

Also note that a customer's iHAT "score" is not necessarily static. This is a very important point – the home has never been static, but now there is quantifiable information that can be used and leveraged to account for this, to improve the customer experience. Changes to the coaxial network in the home made by the customer, or new CPE brought into the home, can both affect the iHAT score. These events are "on demand triggers" that will call on iHAT to run off-cycle even after the initial iHAT sweep of the node at cutover.

8.1.2. iHAT Pass

The most straightforward flow in **Figure 17** is down the center, left to right. Both branches are logical and easily understood. An iHAT "Pass" means that the DOCSIS signals up to 85 MHz are able to be received by the vCMTS receiver, indicating that there is no home amplifier or filter blocking this transmission. AND it means that the home has been checked for RF isolation between the CM and the STB and determined not to be of concern.

Going to the lower green flow down the center of **Figure 17**, this is the case where there is no speed upgrade involved. The spectrum is being turned on to maximize efficient use of upstream capacity. The Tier 1 and Tier 2 plans are counting on use of this capacity to defer any future network augmentation by at least 5 years. So, while it may not be noticeably service-impacting to a customer, it is network- impacting. It may be indirectly service-impacting by lowering the congestion on that node overall (a good thing) and providing lower utilization spectrum for that customer to take advantage of for their current services.

The upper green flow is the case when a speed tier upgrade request is made, and there is already a Mid-Split-capable device present. Because a home's iHAT score is not static, a new iHAT score is taken prior to upgrading the customer. The customer expectations for the new service will be higher, and the awareness will be acute to service-impacting issues, so it is prudent to be certain that the home is still in a "ready" condition. In addition, because the customer now has, for example, a speed tier of 200 Mbps, they will have bursts of energy more likely to utilize a wide chunk of the Mid-split band at once, a condition that more aggressively exposes the STB to energy that can cause video degradation. If this "updated" iHAT result is still "pass," then activation is completed. If not (this is not shown), this home reverts to a

Remediation state, and because of the desire for a new service tier, it is a Remediation Queue with a higher priority.

8.2. MUSL-Building Logic

As powerful as the iHAT engine is, to use it in a production flow such as **Figure 17**, and, as importantly, seamlessly in production scale, it cannot be done on a home-by-home basis via human interaction. The information iHAT needs to run and the information needed by other systems to act on the iHAT outcome must be automated, and the interfaces to these other functions built for production scale. A logical flow diagram for the overarching MUSL ecosystem is shown in **Figure 18**. As shown, within the MUSL framework, like its role in **Figure 17**, iHAT is the engine. **Figure 18** speaks more to the software logic and definition of the adjacent subsystem interfaces that are implied by the flow of **Figure 17**.

The interfaces for iHAT for use in production are highlighted in the red box at the bottom of **Figure 18** and briefly described below. These represent interfaces for MUSL to distribute this important information to stakeholders, for the end-to-end operational success of Mid-Split activation.

Customer Accounts – Serviceability: When there are new upstream speeds that only a Mid-Split upstream can provide, it is important that the systems to upgrade a customer, whether online or through a service call, recognize the home's readiness state, as identified by iHAT. Alternatively, these tools can trigger an instant iHAT test for an updated result.

Biller – new CPE: When a customer changes CPE, possible iHAT variables that are affected are the device DOCSIS capabilities, the sensitivity to interference of a new video CPE, and the possibility of a wiring change in the home. It is prudent, given these potential changes to the iHAT state, to test (or re-test) the home.

XOC – Job Scheduler: When a home "fails" iHAT, it goes into a remediation queue, with a flag for what needs to be remediated (video or DOCSIS). For a Tech Ops plan based on proactive remediations, occurring routinely and not waiting for a house call to take care of iHAT-known issues, iHAT can report its findings per account to the local XOC tools that queue, prioritize, and schedule jobs.

Sales – Serviceability: Similar to Customer accounts, sales representatives should be able to quickly assess whether a customer is eligible for Mid-Split speeds by accessing iHAT status in existing sales tools.

Care – iHAT status, ITG Updates: When a care agent takes a customer call, after some amount of Interactive Troubleshooting Guide (ITG)-led questioning, the possibility of the issue being Mid-Split-related will be considered. A check on the iHAT status of that home, or an instantaneous iHAT test, can help the triage process.

Tech Tools – Tech360: Similar to care agents, when technicians are enroute or onsite at a customer home, part of the awareness they can have is the home readiness state, as determined by iHAT. More deeply in the tools, the sequence of steps to diagnose and fix a MS-related issue should also be available.

Inventory Management: As remediations are made at relatively large scale to remove old drop amps, procurement awareness to the deployment of alternative solutions can ensure that the supply pipeline is tracked and nurtured. This is even more important for proactive amplifier replacement plans, to ensure supply alignment with the plan for the alternative solutions – passive or active.

Data Sciences: As iHAT data is accumulated, new information about the home RF environment – isolation performance, trends over time, and correlations across neighborhoods – can be stored and processed for future optimizations, and to inform future process implications and costs.

Note that the MUSL flow leading into iHAT is from the perspective of a cutover to a Mid-Split-capable DAA node and network, from which flows notifications to activate the spectrum. This includes the instantiation of iHAT, to figure out who can use it at the RPD level. Once this cycle is complete and the Mid-Split node is in operation using the extra spectrum, several reasons were identified to check home readiness status via iHAT on an individual account basis. Noted at the bottom of **Figure 18,** these are referred to as “Asynchronous iHAT Triggers.” They are options that can become part of new Mid-Split operations and maintenance practice, and include:

- New CPE device: XG STB or XB HSD gateway triggers new state-of-home update
- Buy-flow for new product offering that requires Mid-Split: Serviceability tools trigger up-to-date home state
- Premise Health Test (PHT): Tech in the home can trigger from available tools to assess state-of-home locally, as well as to assure that remediation work is completed properly
- Care (E360 tool) – Agent ability to see home state in real-time triage and possibly have access to reverting to Low Split and scheduling remediation
- (New) Home Metrics: Indications that imply a high likelihood that the equipment in the home has changed location or wiring has been changed, such as a persistent DS level change

Lastly, note the hourglass in the middle-right of **Figure 18**. As Mid-Split rolls out in scale, a determination will be made on whether a periodic update of all homes is warranted, and if so, how often. It will be based on empirical data that will reflect findings of just how dynamic the home environment is, and whether it is enough to warrant sweeping all devices on the RPD periodically, or spot checking if the asynchronous triggers do not provide enough off-cycle visibility. The tool will allow for periodic revisiting of each RPD. This eventually could place a lot of additional compute and overhead traffic on the network, storage (and associated cost), and state management resources. For an issue of small enough scale, this might not be warranted – experience and scale will tell.

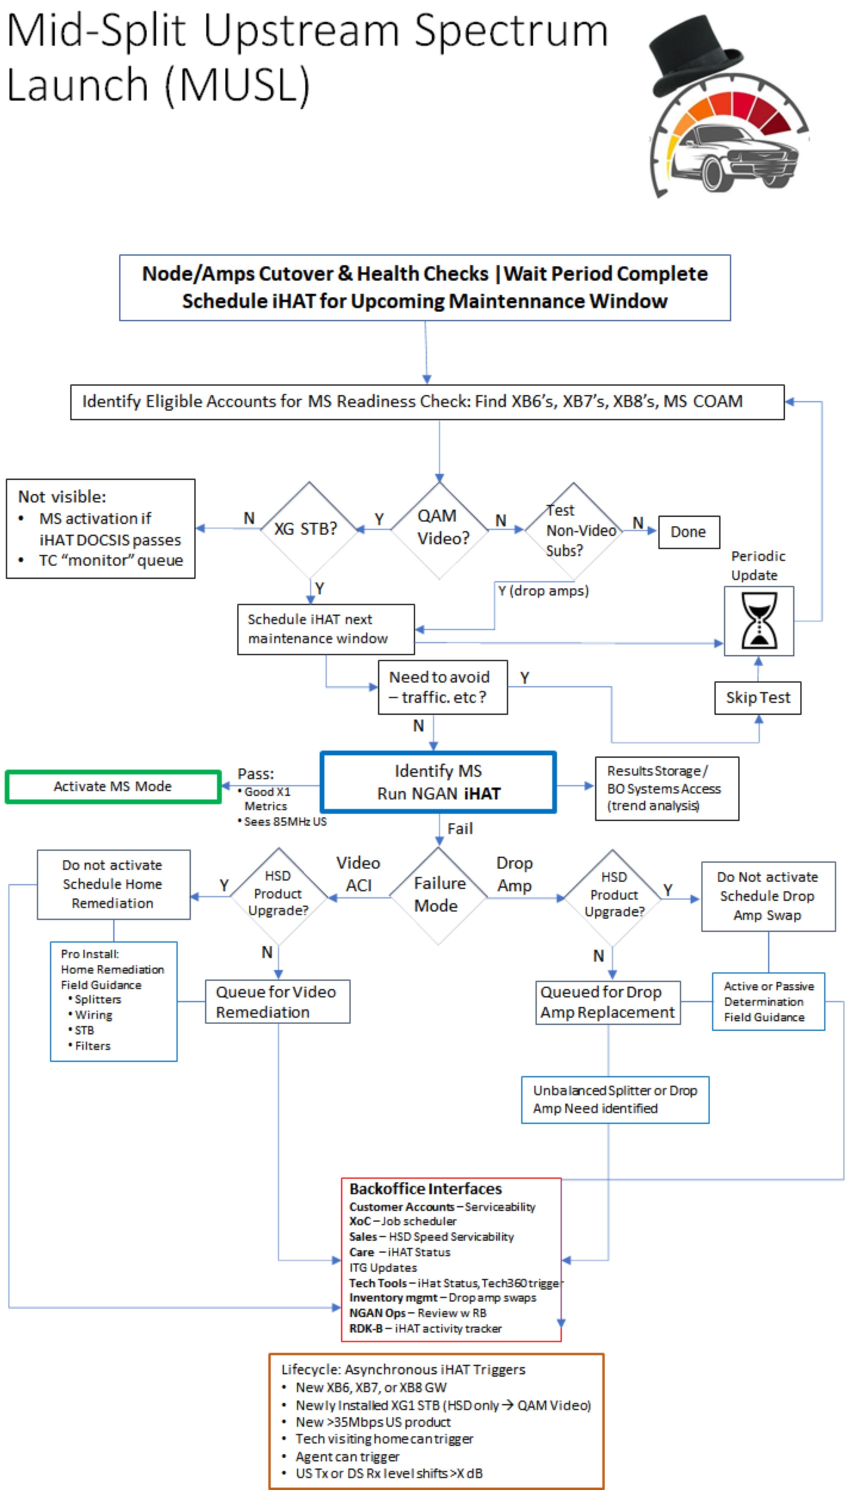

Figure 18 – iHAT as the Engine of the MUSL Framework

8.3. The Care and Feeding of Mid-Split

The iHAT tool described above prepares nodes and homes for the launch of Mid-Split. However, because the home environment can change, and is not under control of the operator, a particular state declared by iHAT is impermanent. New devices can and do come into the home, and customers do change wiring and add in-home passives and actives. Service changes, in particular a broadband speed upgrade for the upstream, increases the likelihood that video service interference could be observed, and exposes any blocking drop amplifiers.

In addition, the video interference potential is statistical in nature: There is an inherent (but small / <1.5%) probability that conditions in the end-to-end system, including OSP and home, shift in such a way that the threshold of interference for visual impairment observed in lab testing is breached.

Finally, the software tool itself will take time to mature, scale, and optimize, and should not be expected to operate perfectly to every potential negative use case and error condition it could encounter at scale.

Because there is a finite probability that a customer could experience video or HSD issues with newly activated OFDMA spectrum, the introduction of Mid-Split spectrum could lead to new inbound call types. As a result, there will need to be process updates such as for ITGs and Line-of-Questioning (LoQ) scripts to diagnose whether Mid-Split is the cause of these issues in the home.

By contrast, a DOCSIS failure is not a statistical phenomenon – either the upstream signal path is blocked by a drop amp, or it is not. As such, there is not very much nuance required around Care processes. The most intricate part of the practice of remediating a DOCSIS failure is two-fold:

1) A home drop amp is replaced with what?
   The knee-jerk answer is another drop amp that supports the expanded frequency split. However, this places a new frequency barrier in place that is likely to be an obstacle in the future, such as for 10G FDX. Best practices and training are being built around a methodology that prioritizes a passive termination at the point of entry, if is not a DOCSIS termination itself (such as for DOCSIS 4.0, in an all-IP home configuration). If a typical splitter implementation is inadequate, a specialized unbalanced splitter may be possible to assure healthy levels at each.

2) Proactive Remediation
   Until HSD products are launched that require the Mid-Split spectrum (upstream speeds of 50+Mbps), the additional upstream spectrum is unlikely to affect the customer experience, one way or the other. To the extent that the added capacity reduces the average upstream utilization and provides a more uniform HSD experience, the customer experience should generally improve with the use of Mid-Split.

   Of course, as mentioned, one of the key benefits of Mid-Split spectrum is the launch of higher upstream speed tiers. As these become available, interested customers who live in a household with a Low-Split upstream limitation will require an additional step. Because faster upstream products are an inevitability, and their penetration will likely grow over time, it makes sense to consider a proactive plan to remove the amplifiers with transactional house calls made for other reasons, and eventually for the specific purpose of pulling the Low-Split drop amplifiers out of the system. This is a business balancing act of operations investment at the right time, to stay ahead of the trajectory of these speed tiers. The alternative is that a percentage of customers who want these speeds will have to await a scheduled truck roll to receive them. However, much can be known in advance with the iHAT tool, so that messaging on the buy-flow front-end can be developed to make that outcome as smooth and efficient as possible for the customer.

For the potential video impact scenario, there is additional nuance and more options to ensure that a quality video experience is maintained. **Figure 19** charts this nuance.

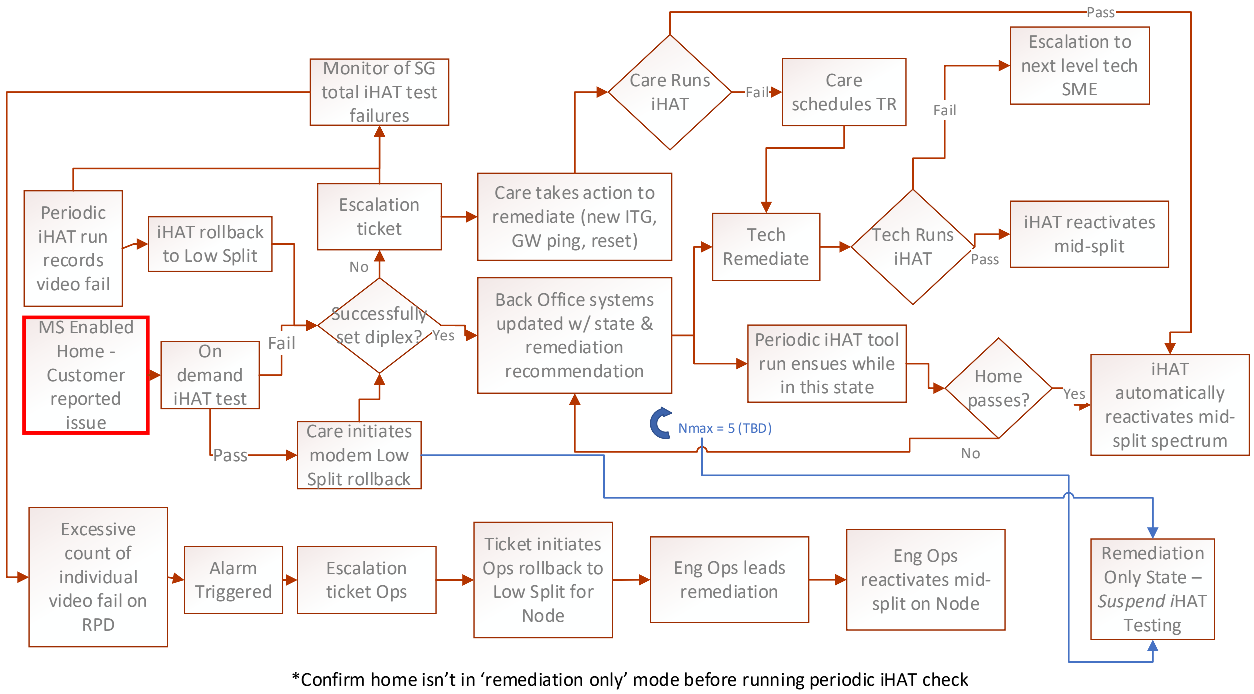

Figure 19 – Care Flow for Support of Mid-Split Related Service Impacts

The flow details are self-explanatory. The diagram assumes that the issue has been diagnosed as likely due to Mid-Split. All other possible causes more probable than Mid-Split have been checked, as they typically would have before reaching this branch of an overall triage flow. In summary form, the fundamental sequence of events in **Figure 19,** after a call to an agent leads to a potential MS diagnosis, are as follows:

1) Run iHAT for an up-to-date state check and disposition of home, compared to prior state
2) Check if iHAT *before|after* state aligns with the TC (iHAT "fail" and device goes from Mid to Low Split). In Low-Split, the video issue will be eliminated, if it was indeed a Mid-Split-related issue. If so, leave the customer in Low-Split mode, schedule remediation, determine the root cause, and reactivate the Mid-Split region for that home, as charted below:

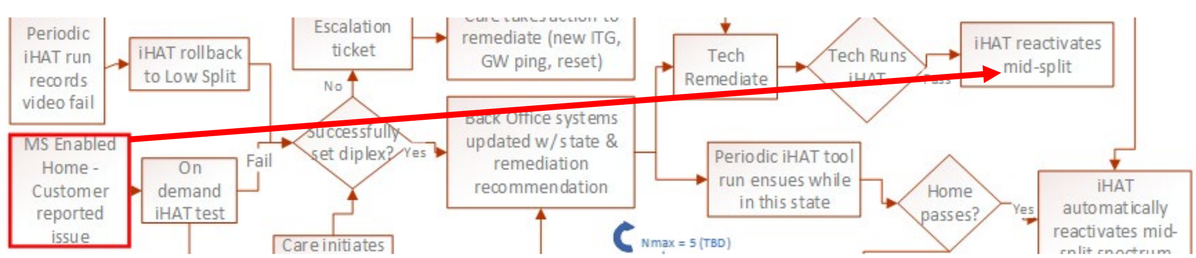

3) If iHAT *before|after* is Mid-Split → Mid-Split, then iHAT believes there should not be a video impairment and yet this is what the customer is experiencing. This does not mean it *is* definitely Mid-Split related. but making this determination on-site is the next step. First, however, the home is manually (via Care directly or via escalation to Operations) reverted to the Low-Split to eliminate the video issue. Again, if this does NOT eliminate the video issue, it is not related to Mid-Split. Assuming this step eliminates the video issue, the account is scheduled for remediation. It is also removed from any further iHAT updates – referred to as the "Remediation Only" queue – that would place it back in Mid-Split mode (because iHAT is giving an erroneous result of "pass" to begin with).

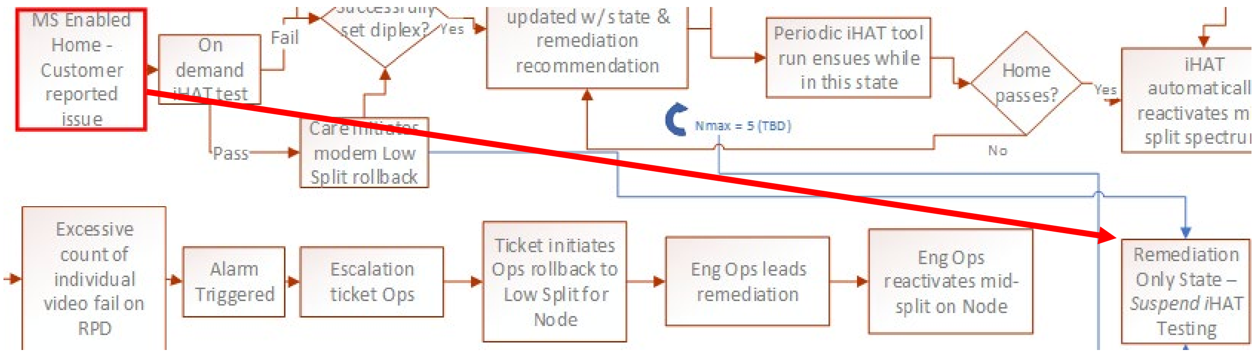

4) A next branch of the flow deals with the case where iHAT does diagnose that the device should be in Low-Split, but it does not properly revert the device to that state. Over time, with code maturity and optimization, we expect this scenario to get diminishingly small. The path first has Care or Eng Ops try to set the modem manually as above, or even factory resetting the modem. If unsuccessful, the path escalates into on-site remediation.

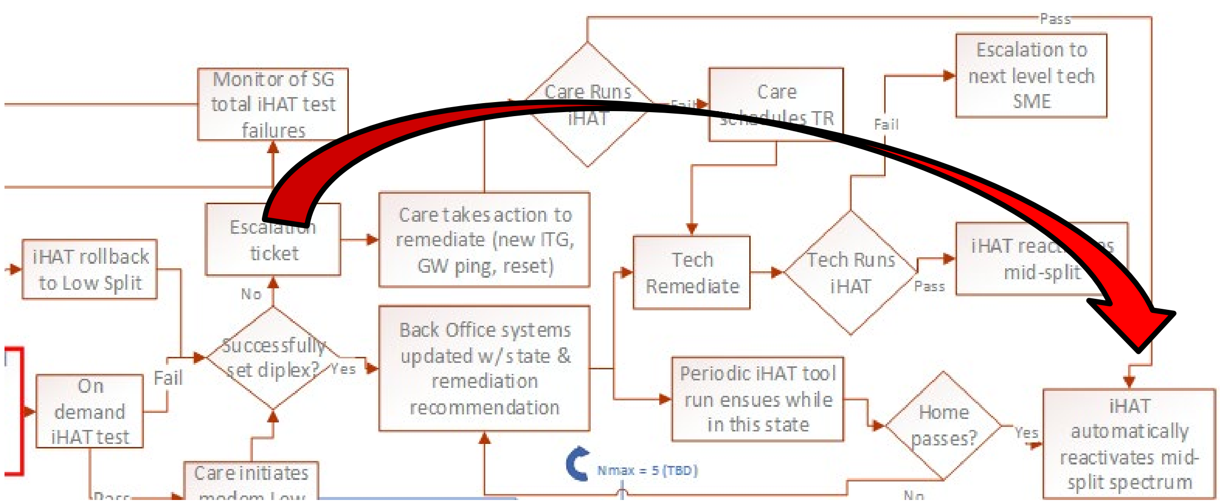

5) A final branch of interest is the scenario where, knowing "typical" metrics for iHAT DOCSIS and video test "fail," a certain threshold of "too many" is set that indicates the issues are probably not on a home-by-home basis, but more systemic. In this case, Engineering Operations is brought in to triage the situation immediately:

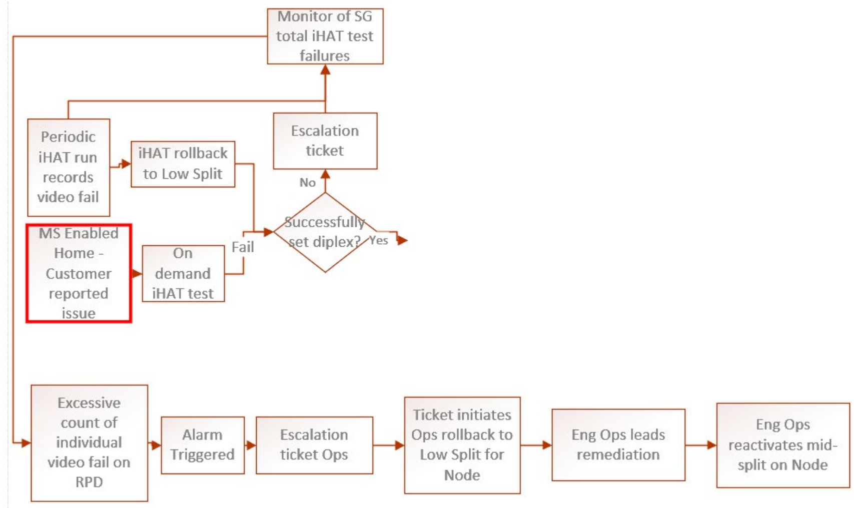

8.4. Scars? What Scars?

In summary, the seamless migration to Mid-Split, from a customer experience perspective, looks like this:

- The full 85 MHz upstream signal can exit the home intact and get onto the network – not blocked by drop amps, filters, or otherwise poor frequency response
- The full 85 MHz signal can be activated, and it causes no QAM video artifacts on any STB in the home
- Any video issues that are encountered by a customer and diagnosed as being caused by Mid-Split can be eliminated remotely and immediately
- Products (speeds) that need Mid-Split spectrum can be delivered to a customer who wants the speeds simply and effectively by any buy-flow means available

We have not detailed further, but these additional components fill out this list:

- In-Home filters are not required to launch Mid-Split spectrum
- Mid-Split can be activated using a self-install kit (SIK) model most of the time

9. Mid-Split → High Split and DOCSIS 4.0

9.1. High Split (5 MHz – 204 MHz)

While this paper focused on a Mid-Split migration scenario, operators are looking also to High-Split and, further out into the future, DOCSIS 4.0. Earlier in the paper we described some of the incremental challenges of operationalizing High-Split, compared to Mid-Split. We left one of these challenges out of that discussion until we were able to go into the Mid-Split details of RF isolation management. Of course, High Split has the overlapping band phenomenon with STBs, only worse. There is much more bandwidth for High-Split that extends into and thus overlaps the forward band. Because of this additional energy that will be launched into the spectrum between 54 MHz and 204 MHz, which currently overlaps the input bandwidth of a STB, there is an even greater chance of ACI interference. Therefore, a migration to a High-Split also includes a migration of the home that is receiving High Split-enabled HSD services (e.g., 1 Gbps symmetric service) to an all-IP configuration – i.e., no QAM video. Without a QAM STB, we can ensure that the video service in the home is not affected by the new service.

However, the relatively loud and wide upstream to 204 MHz has enough energy that this is not necessarily the end of the interference story. A phenomenon that has undergone much study is that of "Neighbor Interference," (NI) whereby the cable "Tap" neighbor of a High Split services user may be "close" enough in the dB sense to have services impacted on the adjacent STB or CM. This scenario is shown in **Figure 20**.

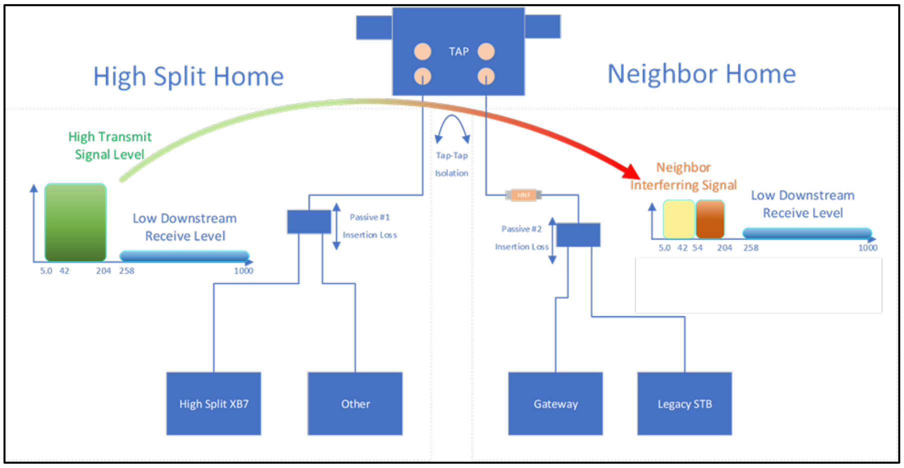

Figure 20 – The Neighbor Interference Phenomenon of High Split

As with the in-home scenario described for Mid-Split, it is also an RF port-to-port isolation and ACI phenomenon with Neighbor Interference, but it is instead the Tap ports that are of consequence. There has been substantial characterization of Tap isolation performance and STB and CM ACI sensitivity recently for this extended upstream band, and the likelihood of this issue has become very well quantified. The MUSL and iHAT tool kit can be applied to a High-Split upstream with these major differences:

1) The OUDP probe signal of interest needs to be modified to one that can be consistently correlated to an extrapolated equivalent of *High-Split* signal energy
2) The potentially at-risk" STB is in a neighboring home, and physical addresses and the relationship of CMs in the field, to Tap ports, is generally not easily known in an automatable way
3) The "victim" device can also now be a non-High-Split CM

Mitigation of NI involves different processes. Visiting a neighborhood home for remediation because a different neighbor on the block upgraded their HSD would be an awkward process on every conceivable level. Thus, the bias for NI would be towards blocking filters in the OSP at the "guilty" Tap port. Documenting filters installed would be an important way to simplify new customer adds going forward. Complete quantification of this phenomenon with respect to the customer experience is difficult until some correlation of High-Split transmissions to video and non-High Split HSD performance can be documented.

9.2. DOCSIS 4.0

Assuming the all-IP home policy that is anticipated for DOCSIS 4.0 homes, coexisting with QAM STBs and legacy CMs in the plant using DOCSIS 4.0 Extended Spectrum technology would take on the equivalent model as High-Split NI, but with the Ultra High-Split bandwidth as the spectrum limit, and STB sensitivity characterization work to be done.

It's worth noting that in DOCSIS 4.0, the very nature of the FDX technology at its core is an overlapping upstream and downstream. As such, it is inherent in the protocol to introduce new technology to manage this overlap. There are two ways this is done, as shown in **Figure 21**:

Echo Cancellation: "Self" cancellation at the PHY level, leveraging knowledge of the transmitted signal at the co-located receiver.

Sounding/Scheduling: Avoids having an FDX CM transmit upstream when it is known that a neighbor could be affected in the downstream that is receiving packets.

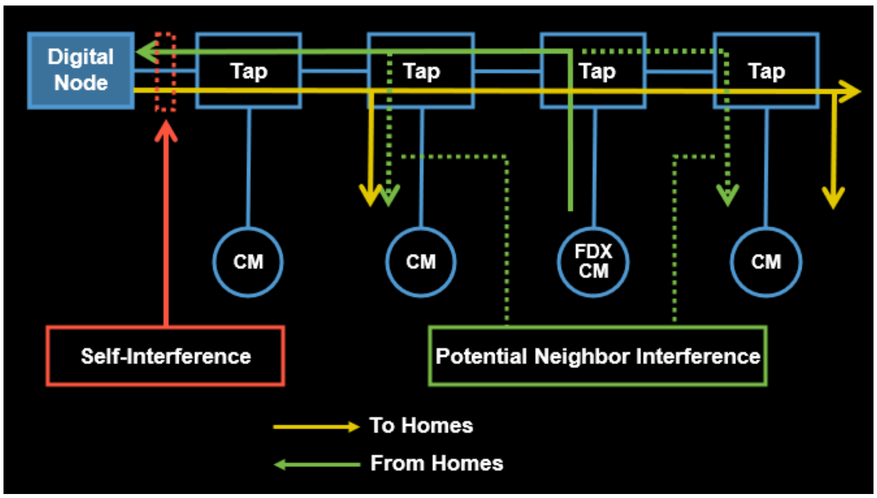

Figure 21 – The Two New Technology Features of DOCSIS 4.0 FDX

For the latter, FDX determines the RF dB isolation relationships among modems to form Transmission Groups (TGs). FDX "sounding" is effectively the built-in, standardized version of NI for the FDX band (108-684 MHz) in FDX systems.

For FDX, the limitation that arises is that these relationships can only be discovered in DOCSIS 4.0 CMs and DOCSIS 3.1 CMs with an FDX-L SW upgrade. FDX-L is a way to make DOCSIS 3.1 CMs aware that they are connected to a DOCSIS 4.0 system, and thereby have their traffic scheduled within the context of the TG assignments. DOCSIS 3.0 CMs cannot participate in sounding at all. The number of DOCSIS 3.0 CMs continues to decline rapidly in the field, but they will not be completely removed from the network before FDX is deployed.

There will be more to come on DOCSIS 4.0 migration challenges and solutions as the 10G technologies continue to be developed [1].

10. Conclusions

Operators are recognizing that, as good as the upstream has been to them since the launch of HSD services, it has given nearly all that it can at this point and needs a spectrum boost to continue to deliver value and support continually growing HSD services, capacity, and speeds. The next step is to add spectrum and launch the next long runway of capacity, with new speed expectations in mind. With the commitment to spectrum comes a commitment to managing it through an HFC lifetime of legacy equipment that is simply not built for it. Building the technology, tools, processes, and practices to enable this transition is a challenge all operators are working through, with a seamless and non-disruptive experience for the customer as the top priority.

In addition, as in any network evolution that touches the outside plant, making sure that enough is done to the network for the longer term, once the commitment has been made to go out and touch it, is an important part of the upgrade. For access network engineers, the billiards analogy is that, as you are lining up the 6-ball at the side pocket, it is important you make sure that after sinking it you've left the cue ball lined up neatly behind the 12-ball at the corner pocket. With the right series of deft maneuvers, the 10G-ball will be lined up to finish out the game.

Acknowledgements

Many thanks to my esteemed peers Leslie Ellis and Larry Wolcott for joining me for this paper, which due to unique circumstances required a full-court press effort to meet the deadline.

Abbreviations

ACI	Adjacent Channel Interference
AGC	Automatic Gain Control
BAU	Business-As-Usual
BG	Bonding Group
CACIR	Carrier-to-Adjacent Channel Interference Ratio
CAGR	Compound Annual Growth Rate
CDF	Cumulative Distribution Function
DAA	Distributed Access Architecture
DSG	DOCSIS Settop Gateway
FDD	Frequency Division Duplex
FDX	Full Duplex DOCSIS
FTTH	Fiber-to-the-Home
HHP	Households Passed
iHAT	In-Home Assessment test
LoQ	Line-of-Questioning
MER	Modulation Error Ratio
MTA	Media Terminal Adaptor
MUSL	Mid-Split Spectrum Upstream Launch
NI	Neighbor Interference
OFDMA	Orthogonal Freqeuncy Division Multiple Access
OOB	Out-of-Band
OUDP	OFDMA Upstream Data Profile
OSP	Outside Plant
OTA	Over-the-Air
PHT	Performance Health Test
QAM	Quadrature Amplitude Modulation
SNR	Signal-to-Noise Ratio
STB	Settop Box
TaFDM	Time and Freqeuncy Division Multiple Access
TCP	Total Composite Power

Bibliography & References

[1] Baumgartner, Jeff, "Comcast Full Duplex DOCSIS trial pumps out 4-Gig symmetrical speeds," LightReading, 4/2/2021.

[2] Howald, Robert, Repair the Ides of March: COVID-19 Induced Adaption of Access Network Strategies, 2021 SCTE Expo, Oct 11-14, Atlanta, GA.

[3] Robuck, Mike, "Comcast notches 10G Milestone with a trial of 1.25-Gig symmetrical speeds in Florida," Fierce Telecom, Oct. 8 2020, https://www.fiercetelecom.com/telecom/comcast-notches-10g-milestone-trial-1-25-gig-symmetrical-speeds.

[4] Thompson, Robert, and Rob Howald, Dan Rice, John Chrostowski, Ronini Vugumudi, Amarildo Vieira, and Zhen Lu, Rapid and Automated Production Scale Activation of Expanded Upstream Bandwidth, 2021 SCTE Expo, Oct 11-14, Atlanta, GA.