



ATLANTA, GA
OCTOBER 11-14



Rapid and Automated Production Scale Activation of Expanded Upstream Bandwidth

A Technical Paper prepared for SCTE by

Rob Thompson

Director, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
robert_thompson6@comcast.com

Rob Howald

Fellow, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
robert_howald@comcast.com

John Chrostowski

Executive Director, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
john_chrostowski@comcast.com

Dan Rice

Vice President, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
daniel_rice4@comcast.com

Amarildo Vieira

Engineering, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
amarildo_vieira@comcast.com

Rohini Vugumudi

Director, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
rohini_vugumudi@comcast.com

Zhen Lu

Manager, Next Generation Access Networks
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
zhen_lu2@comcast.com



Table of Contents

Title	Page Number
1. Introduction	4
2. Background	5
3. New Mid-split Activation Tools	6
3.1. Mid and High split spectrum activation launch (MUSL/HUSL)	6
3.2. Manual vs. Automation Processes – Cloud Infrastructure Choices	7
3.3. In-Home Health Assessment (iHAT) 2.0	8
3.4. iHAT as the Engine for Mid-Split Upstream Spectrum Launch (MUSL)	12
3.4.1. Applying MUSL End-to-End	13
3.4.2. The Remediation Queue	15
3.4.3. iHAT PASS	15
3.5. iHAT Software Solution Implementation	16
3.5.1. Data Collection & Device Interface Platform (Genome)	22
4. High-split Future	25
4.1. Extending iHAT for Neighbor Interference Case	27
5. Conclusion	32
Abbreviations	33
Bibliography & References	35

List of Figures

Title	Page Number
Figure 1 - OUDP Burst Request and Response Parameter Set	9
Figure 2 - Simultaneous MS-CPE OUDP Burst and SS-CPE SNMP Data Collection (incl. spectrum capture)	10
Figure 3 - Spectral View of MS-CPE OUDP Burst, and SS-CPE OOB and Low Frequency Downstream Signals	10
Figure 4 - 30 MHz MS-CPE, SS-CPE Receive Window Overlap	11
Figure 5 - iHAT is the Engine of the Mid-Split Upstream Spectrum Launch (MUSL) Framework	12
Figure 6 - Simplified Mid-Split Activation Flow: Cutover through to Activation	14
Figure 7 - High-Level iHAT Architecture	18
Figure 8 – OUDP/FBC Orchestration Workflow	20
Figure 9 – Genome Architecture	22
Figure 10 – Genome API – NQL Architecture	23
Figure 11 – SNMPv3 USM key architecture	25
Figure 12 - Mid Split Scout Screen Capture with 4 SC-QAMs and 1 OFDMA	26
Figure 13 - High Split Scout Screen Capture with 4 SC-QAMs and 1 OFDMA	26
Figure 14 - High Split Yeti Screen Capture at RPD with 4 SC-QAMs and 1 OFDMA Channels	27
Figure 15 - High Split Yeti Screen Capture showing 2 OFDMA Channels	27
Figure 16 - Neighbor Interference in a High-Split System	28



**UNLEASH THE POWER
OF LIMITLESS CONNECTIVITY**
OCTOBER 11-14 ATLANTA, GA



Figure 17 - Neighbor Interference for STB Test Configuration 29

Figure 18 - Cable-Modem & Set-top Box Interference: Threshold vs Neighboring High Split Throughput
for Several Different Models of Set-tops 30

Figure 19 - High Split Threshold vs Neighboring High Split Throughput Model 31

Figure 20 - Neighbor Interference Notch Filter 85-204 MHz) 32

1. Introduction

Enabling the upstream portion of the access network for Mid-Split or High-Split takes more effort than just configuring the cable modem termination system (CMTS) to activate new DOCSIS carriers. Without considering the potential consequences of in-home networks and their effects on DOCSIS and video services, it may not be as seamless to customers or field operations. Most of the characterization required can be remotely measured and observed through telemetry available in the network. Operators can use algorithms to predetermine if there is work to be done in the field or home to enable new spectrum.

While today's tools are helpful in keeping technicians out of homes to make measurements, they were designed for raising flags, enabling swivel chair dashboards, or troubleshooting specific cases. By contrast, turn-up of network spectrum must initiate action instantaneously at large scale, through multiple scenario permutations, and publish essential data and actions cross-functionally to stakeholder organizations such as Sales, Technical Operations, Care, and even Warehouse Ops.

Instead of reacting to consequences or relying on manual vetting of homes, software running on modernized cloud infrastructure can proactively identify blockers, on an individual customer basis, but in massive scale with machine robustness that is immune to human error. The cloud software creates connective tissue between the upgraded network infrastructure and the incumbent suite of tools currently used by operations. This approach facilitates seamless interaction with the customer experience by maximally enabling new services and identifying blockers to frontline teams to optimize efficiency of support.

The work and development in upgrading the plant and services to mid-split is directly related to future upstream expansion including High-Split. Many of the workstreams will continue and will build off the mid-split effort. Moving the downstream spectrum to allow room for the extended upstream spectrum, the process of upgrading the physical plant with nodes and amplifiers that support different frequency splits, the deployment of Orthogonal Frequency Division Multiple Access (OFDMA) and adjacent channel interference (ACI) -- all correlate to high split deployments. The tools and processes to validate mid-split in-home ACI can be used to evaluate neighbor interference in a high split system. ACI is dependent on in-home splitter isolation and neighbor interference is dependent on tap-to-tap isolation. In both cases, the tools and processes are similar. The mid or high split device can be exercised and the level of interference or effect on the adjacent or neighbor device can be measured prior to permanently moving the customer's device into mid or high split operation. In both cases methods have been developed which are non-service and non-customer impacting.

This paper will introduce the reader to Comcast's cloud software architecture, the high-level algorithms used to "score" a home, and the downstream systems fed by the system that will enable valuable upstream services in scale while minimizing impact to customers and maximizing efficiency of its delivery for the operator.

2. Background

Options for upstream capacity enhancements arrived in DOCSIS 3.0 around the 2006 timeframe [1]. The 42 MHz to 85 MHz or mid-split band could become the first step toward supporting new upstream DOCSIS channels for U.S.-based operators. Since then, new DOCSIS 3.1 and 4.0 standards have been developed to support even more channels, in even larger return bands, like high-split, up to 204 MHz, and Full Duplex DOCSIS (FDX), up to 684 MHz. Those new upstream channels would become part of the larger channel set, enabling operators to increase upstream service capacity appreciably beyond what has been historically constrained to 5 – 42 MHz, known as the standard-split or low-split operating band. During the mid-split introduction, diplexer switching within the cable modem (CM) would also be introduced, to remotely control a new and more flexible diplexer function and thus, the channels over which a CM could communicate. Switchable diplexer functionality enabled operators to seed their network footprint with mid-split-capable CPE, while the rest of the network's technology could catch up to support the activation of enhanced upstream services. What wasn't well understood at that time was the new work needed to enable new upstream spectrum activation -- a topic that will be introduced in this section.

Managing the outside plant to support mid-split has become well understood through ongoing upgrade activities. For example, preparations including vacating downstream spectrum, and relocating the downstream out-of-band (OOB) signal to a higher frequency have become standard operating procedures. Because these outside plant pieces were easily implemented, many operators have activated or at least tested mid-split services, and that experience has proven to be appreciably more challenging for some operators. Two of the challenges, standard-split drop amplifiers (SSDAs) and ACI susceptibility, were recently documented in [4] as home network-based challenges that will block the successful activation of mid-split services.

Mid-split tests at Comcast have helped us learn how to remediate these challenges. SSDAs are ideally removed from the home network for many reasons including that they are a more costly, higher-rate point-of-failure than their passive splitter counterparts. Furthermore, SSDAs used in a home network will have to be removed because of its standard-split diplexer, to make way for mid-split upstream services. Therefore, limiting drop amplifier use, wherever possible, may emerge as one of the preferred remediation approaches.

Remediating ACI susceptibility is accomplished by increasing isolation between the mid-split CPE (MS-CPE) and standard-split (SS-CPE) [4]. The following recommendations have emerged based on our learnings during early mid-split testing, to ensure mid-split connectivity between the home network and the outside plant, while minimizing ACI:

1. If there are few (≤ 4) devices within the home requiring outside plant connectivity, receiving video and/or DOCSIS quadrature amplitude modulation (QAM) signaling, use a higher isolation splitter
2. If there are many (> 4) devices within the home requiring outside plant (DOCSIS) connectivity, use a higher isolation passive network (POE DOCSIS-MOCA Passive or PDMP)
3. If 1 and 2 are not feasible due to severe drop loss, consider the following

- a. Reduce devices requiring outside plant (DOCSIS) connectivity, and subsequent drop loss, by replacing them with Wi-Fi capable devices
 - b. Use a drop amplifier sparingly, which actively adds gain to overcome drop loss
4. If 1 through 3 still do not provide adequate isolation for ACI susceptible devices, use a notch filter

The remediation strategy above can be distilled down to the practice of removing drop amplifiers and swapping splitters --something installers are well-versed in dealing with in their day-to-day. The new work is focusing on when and where to perform the above remediation tasks, since we've already learned that doing them for 100% of the mid-split customers at the time of spectrum activation may not be feasible for some operators [4]. Even though it wasn't called this at the time of publication, in-home Health Assessment Tool (iHAT), as it has since become known, has demonstrated that Remote Feature Control (RFC) and Remote Health Monitoring (RHM) can be used to orchestrate remote identifications of mid-split-capable devices that can be activated without remediation. That's opposed to those that need remediation, which can be placed in a queue for activation when that work completes at some point in the future [4]. Activating mid-split upstream will include performing iHAT tests, and distributing results to teams, including installers, who can perform the remediation tasks described above.

3. New Mid-split Activation Tools

3.1. Mid and High split spectrum activation launch (MUSL/HUSL)

As described, the innovative iHAT tool provides a relatively non-intrusive view into a customer's home and, on a home-by-home basis, making basically a Go-No Go declaration with respect to activating spectrum in the Mid-Split band for that customer's device. The iHAT results score a home's DOCSIS readiness for passing spectrum to 85 MHz, and/or its likelihood of creating video interference.

Of course, the home cannot be assessed, nor the spectrum activated efficiently on a home-by-home basis, by relying on manual processes. 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.

On the input side, iHAT needs information from external systems to identify if a house is eligible for Mid-Split from an equipment standpoint, and thereby worth running iHAT at all. It needs the inventory of Mid-Split capable CPE on a per node basis identified so that it can target those devices to run the test. The output of iHAT – the scoring of a home's DOCSIS readiness or risk of degrading the video experience – is information that multiple other external systems need. For example, technicians in the home need to know what to fix if a problem related to Mid-Split is identified. Care agents need to know how to diagnose a possible Mid-Split related issue, guide a trouble call, and dispatch the proper support. Salespersons need to know if an upstream speed that requires Mid-Split capability is possible. Data sciences team need iHAT results to populate databases to analyze trends to adjust and optimize roll-out processes and operation support. We will look at these in more detail in a subsequent section.

3.2. Manual vs. Automation Processes – Cloud Infrastructure Choices

It is undeniable that a successful business must find ways to scale in order to stay relevant. This is extremely true when applied to the telecommunications industry. The bandwidth requirements of people and businesses, especially in the upstream, have been increasing dramatically year over year due to advancements in technology. To keep up with the high demand for broadband services, innovation is required. There are some problems that can be solved by throwing people at the problem; however, when you have customer devices in the tens and hundreds of millions, this becomes untenable.

In the field of software, the most basic reasons to build an automated solution have always been the need for smart investment on resources, and finding budget efficiencies, instead of spending to do manual work year after year. The manual method is undeniably easy and quicker to build but has its own downfall. Questions raised are 1) What can be automated? 2) What kind of automated solution would work? 3) What defines the “right” solution? 4) What platform is optimal to architect the solution? 5) What factors play a part in choosing and creating the solution? 6) What are the dependencies and integrations? 7) Perhaps the most important question of all is, how much of an automated solution can sustain the business growth year after year, without re-investing and redoing a lot of the work? A companion question: What about component usability, for other parts of the business?

The downside of the automated solution is that its machines must be built and trained in a way that mimics the thought processes for human troubleshooting skills. Outcome planning involves thinking through as many of the business’s use cases as possible, considering regulations and validating as many inputs and outputs as possible. Also, implementing ways to listen and respond back to alerts, and handling notifications to support the operations and avoid potentially impacting customers.

Let’s talk about what factors will make automation an answer for any business, which include but aren’t limited to: Scalability, availability, maintainability, reusability and budget. It is imperative that incoming data feeds are well sanitized before entering an automation solution, so the data that comes out has integrity for better analysis and reporting. Effective automated solutions have a proven record of reducing the cost overtime -- emphasis is on “effective automation solution.”

Taking some time to answer the questions raised above will help to decide upon the right solution. In fact, that’s what led us to choose the solution for iHAT. Like any other software solution, we started with the version 1.0 build, described in [4], to research, implement, test and analyze. This pushed us to build iHAT version 2.0 from the lessons learned. Some of the downfalls of version 1.0 included:

- Manual process
- Scalability
- Speed test failures
- SNMP V3 key reset after every reboot made manual process difficult
- Availability of system depended on human resources
- Time consuming and not easily maintainable
- Handling unexpected exceptions
- Automatic rollbacks
- Having to wait until next day to see the runs, if resources are not available

That list of “cons,” along with other elements we learned with the manual implementation of version 1.0, gave us clarity on next steps. It also elicited a new list of questions and requirements, including: 1) Developing a list of features to be automated 2) Cloud-based automated solution is the preferred option 3) Reduce human dependencies 4) Ensure that the platform is highly available 5) Automation, manageability, high availability, scalability, cost effective 6) Not possible to reduce dependencies, which will make the solution to be focused on understanding data collection and distribution and build reliable integration to pull right data from source of truth (SOT) and be the most reliable source of distribution (SOD). By taking these lessons learned into consideration, and after answering all the questions, a solution can be envisioned and built that encompasses the qualities that sustainably support business.

3.3. In-Home Health Assessment (iHAT) 2.0

iHAT is a method, or in programming terms, a script, whose inputs must include the Mid-Split-capable devices and any set top boxes (STBs) that share a home network connection with them [4] on an operator’s Mid-Split-capable network. iHAT outputs provide operators with results indicating which Mid-Split-capable devices have successfully switched over to Mid-Split and which will require remediation to do so [4]. Remediation associated with Mid-Split was discussed in Section 2. The iHAT PoC from [4] has proven it will work well for small tests with up to 1,000 devices per instance.

Performing iHAT for a larger population (millions of devices) revealed challenges in three key areas:

- Automation – Automation of iHAT, specifically input gathering and output distribution wasn’t described in [4] but is a critical step at the beginning and ending of such a test. Taking steps to ensure the accuracy of the deployment data, which is constantly changing due to business as usual (BAU) activities, including device swaps, can prevent aborted iHAT tests attributable to incorrect device information. Minimizing the time between gathering inputs and executing iHAT will ensure integrity of the inputs and is best accomplished via automated process.
- Service Impact – [4] described the importance of the speed test application for iHAT testing, essentially behaving as a catalyst for determining whether a mid-split based service is going to degrade other customer services, like video. Using a speed test application in this manner is service affecting because it blocks a customer’s use of their CM during speed testing, and therefore is limited to maintenance window activities. Another example of an iHAT service affecting feature is associated with rebooting a CM, which is required when CM switchable diplexer is changed from standard-split to mid-split. Eliminating service affecting components from iHAT process, like the speed test application and reboots, could enable iHAT testing to coexist with customer use.
- Application Dependencies – for many operators, RFC and RHM described in [4] likely translates to multiple application interactions. A sunny day scenario is that all the applications used to support iHAT are responsive, meaning they complete their function in a timely manner. When the sunny day scenario doesn’t occur, iHAT is forced to retry, perhaps multiple times, and ultimately abort a test when dependent application functions fail. One example is when RHM doesn’t return STB signal-to-noise (SNR) data necessary to perform a pass or fail decision on ACI susceptibility. Other examples of application dependencies include the speed test application in [4] and the associated class of service (COS) change application to test at the higher mid-split service rate. Most standard-split COS rates support up to 35 Mbps, but to fully utilize a DOCSIS 3.1 upstream [2], a new COS supporting a much higher rate, like ~300 Mbps, would have to be available for use during iHAT testing. Hardening iHAT for scale operation

requires a careful review of the function calls it makes, an assessment of their ability to fulfill their function in a timely manner, and a consolidation of all functions wherever possible.

As it turns out, there was an opportunity to consolidate application dependencies of the iHAT process involving the assessment of ACI susceptibility. The original version of iHAT described in [4] was successful in testing during worst-case conditions experienced on a home network, which were simultaneous occurrences of (1) fully utilized mid-split upstream and (2) ACI susceptible service, including a customer watching video. There were at least five dependent functions in the original version of iHAT that were required to assess ACI susceptibility:

1. CM COS change to support higher mid-split rate
2. Switchable diplexer state change from standard-split to mid-split
3. CM reset associated with both (1) and (2)
4. Speed testing application to simulate customer activity
5. Telemetry polls to assess against pass/fail criteria

Upstream carriers energized by an upstream speed test can be replaced with a OFDMA Upstream Data Profile (OUDP) burst signal. An OUDP request and response are shown in Figure 1. Using OUDP bursts can be facilitated without originally required COS changes of (1) but more importantly they give greater flexibility than speed test applications for estimating ACI susceptibility. OUDP bursts can be quicker, managed via CMTS scheduling functions, and more benign because their total transmission power can be reduced appreciably using smaller bandwidth, for example 1.6 MHz. This change may set iHAT up for no longer being restricted to maintenance window activities, especially if other service affecting aspects of iHAT, like reboots associated with switchable diplexer changes, are also eliminated.

OUDP Test Request

```
var oudpRequest = {
  transaction_id: tid,
  mac_address: { address: my_cm_mac },
  ofdma_channel: 0,
  iuc: 13,
  center_frequency: 79900000, // 79.9 MHz
  width: 1599999, // 1.6 MHz
  duration: 5000 // 5 sec
};
```

OUDP Test Response

```
{
  "mac_address": {
    "address": "60:3d:26:ab:b5:ec"
  },
  "ofdma_channel": 41,
  "iuc": 13,
  "sid": 8198,
  "tss": "1621964215804",
  "actual_grant_duration_msec": 4999,
  "center_frequency": 79900000,
  "width": 1600000,
  "actual_cm_transmit_duration_msec": 4999
}
```

Figure 1 - OUDP Burst Request and Response Parameter Set

With these changes, iHAT will produce a different view of the ACI susceptibility problem, which will become a spectrum-based view from the SS-CPE point of view. While the OUDP signal is bursting in Figure 2, spectrum captures are performed by the SS-CPE at the highest rate possible, providing a multiple-sample spectral view of both the desired downstream signals and the upstream OUDP signal leakage into the SS-CPE, per Figure 3. From Figure 3, ACI assessment becomes more about visualizing OUDP signal leakage into the SS-CPE receiver, than about driving the SS-CPE into a failure state, like the original iHAT v1.0 using the upstream speed test application in [4].

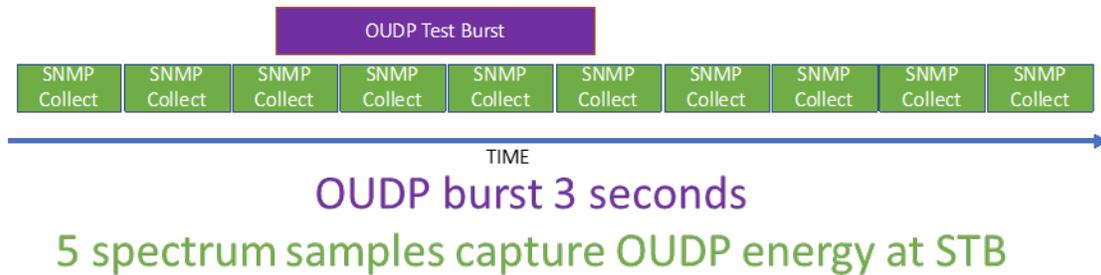


Figure 2 - Simultaneous MS-CPE OUDP Burst and SS-CPE SNMP Data Collection (incl. spectrum capture)

From this new spectrum view, two new parameters can be estimated: MS-CPE to SS-CPE (1) isolation, in decibel units (dB) and (2) interference level, in decibels relative carrier (dBc). When combined with the original metric collection described in [4], the isolation is a simple calculation based upon the telemetry polling of the MS-CPE transmit power in dBmV and the SS-CPE receive power in dBmV. Equation 1 shows this simple relationship, based on decibels relative to one millivolt (dBmV) per 1.6 MHz channels.

$$Isolation (dB) = P_{MS-CPE OUDP TX} - P_{SS-CPE DS RX} \quad \text{Equation 1}$$

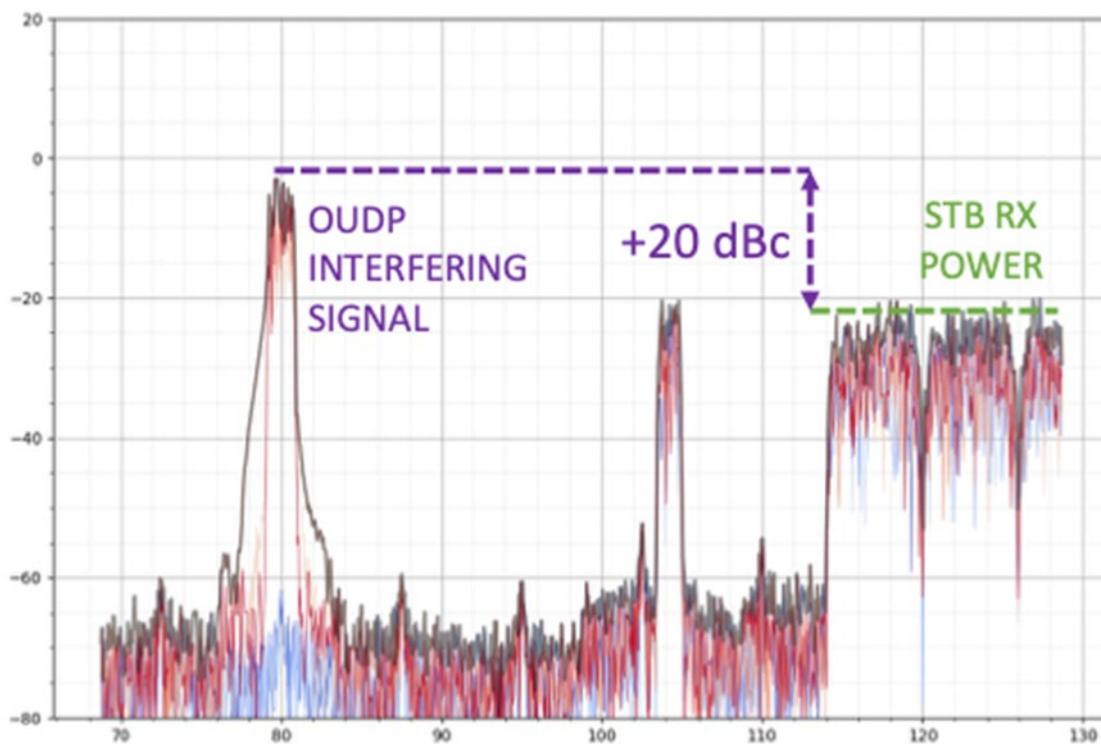


Figure 3 - Spectral View of MS-CPE OUDP Burst, and SS-CPE OOB and Low Frequency Downstream Signals

Adding SS-CPE spectrum data to dependent function 5, “telemetry polls” above and coordinating their capture to occur during the scheduled MS-CPE OUDP burst transmission will enable the estimation of interference level in dBc. Interference level can be estimated via spectrum capture of the SS-CPE per

Equation 2. Multiple captures are required to ensure good estimates can be made for both the OUDP burst and the downstream signal powers. Therefore, the power levels will be a statistical representation of multiple spectrum samples. In particular, the maximum hold, black trace of Figure 3, of the available traces is used to estimate channel power on a 1.6 MHz basis.

$$Interference (dBc) = P_{SS-CPE OUDP RX} \left(\frac{dBmV}{1.6MHz} \right) - P_{SS-CPE DS RX} \left(\frac{dBmV}{1.6MHz} \right) \quad \text{Equation 2}$$

Deciding pass or fail will be different than what was originally shared in [4], where failure assessments were made when SS-CPE MER became worse than a threshold value of 28 dB for example, a point past which degraded video would likely be observed for most STBs. New thresholds targets, like 25 dB for isolation and 20 dBc for interference levels, would be revised as data is aggregated from many iHAT tests. Interference targets are based upon laboratory and other investigations, where greater than 20 dBc interference power would result in a degraded MER and FEC metrics. In other words, the undesired OUDP leakage is 20 dB higher than the desired downstream power observed at the SS-CPE input. This level of interference would put certain makes, models, and vintages of SS-CPE at risk of failure due to ACI susceptibility issues, based on internal testing performed in Comcast labs. Interference thresholds would also consider total power level differences of the OUDP signal (1.6 MHz) versus the overlapping mid-split power of OFDMA transmission in the SS-CPE receive window, see Figure 4. Lastly, process accuracy would be proven over full range of SS-CPE downstream receive and MS-CPE upstream transmit power. The 25 dB isolation target is based upon minimum output-port-to-output-port (OP2OP) isolation requirements specified for Comcast passive components.

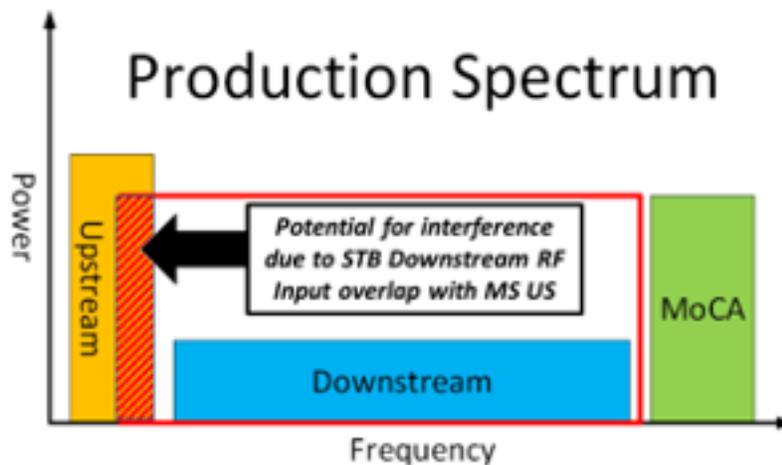


Figure 4 - 30 MHz MS-CPE, SS-CPE Receive Window Overlap

The last piece of scaling mid-split involves configuration of the MS-CPE switchable diplexer and as it turns out, the SNMP set approach discussed in [4] works well for network utilization improvements and optimizing serving group capacity. Another way operators can manage the MS-CPE's switchable diplexer is using bootfiles. However, if there are no COS changes then redundant bootfiles would need to be deployed to have individualized control over a MS-CPE diplexer state. This could be challenging if there are many mid-split-capable platforms with many different COSs. For operators wanting to introduce a higher service tier, like within their commercial service markets, then co-managing the diplexer state with the COS should coexist nicely, and make more operational sense.

3.4. iHAT as the Engine for Mid-Split Upstream Spectrum Launch (MUSL)

As described, 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 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.

Of course, the home cannot be assessed, nor the spectrum activated efficiently 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 Mid-Split Upstream Spectrum Launch (MUSL) ecosystem is shown in Figure 5. As shown, within the MUSL framework, iHAT is the engine. Note that “Tier 1” and “Tier 2” are arbitrary labels meant to represent individualized operator upgrade strategies and illustrate these how these new processes can fit together.

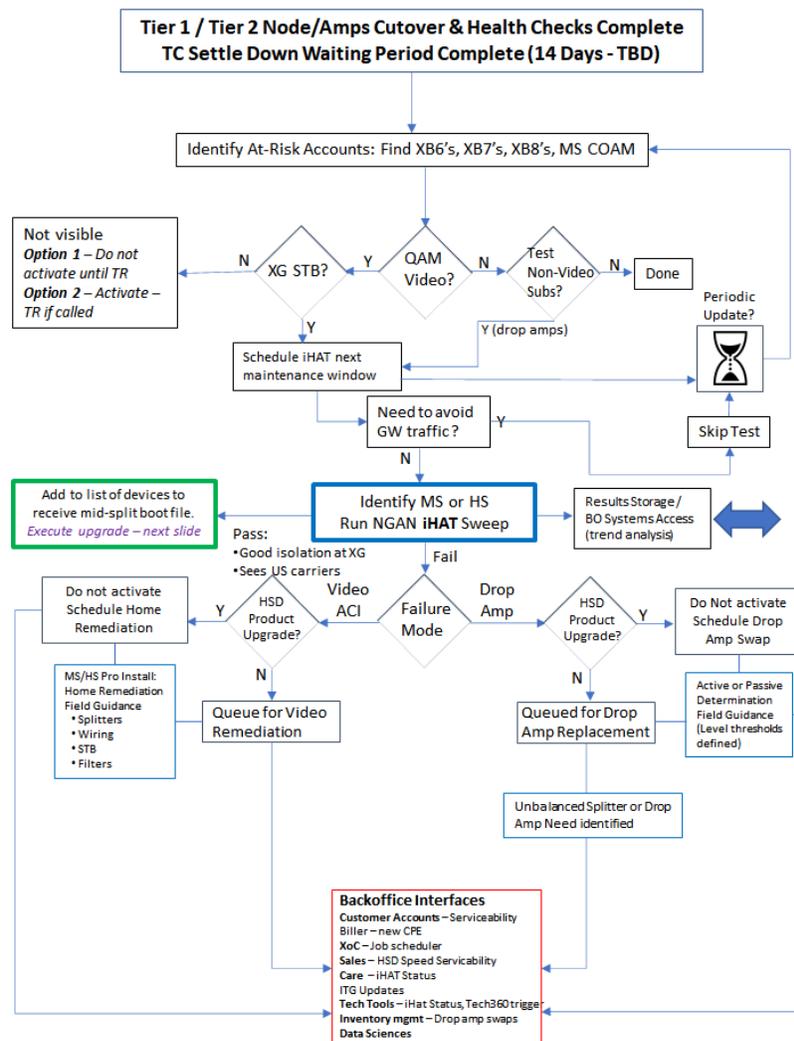


Figure 5 - iHAT is the Engine of the Mid-Split Upstream Spectrum Launch (MUSL) Framework

The interfaces for iHAT for its use in production are highlighted in the red box at the bottom of Figure 5 and briefly described below. These represent interfaces to consider for MUSL to distribute this important information to stakeholders in the success of Mid-Split activation.

Customer Accounts – Serviceability: It is important when there are new upstream speeds that only Mid-Split can provide that the systems to upgrade a customer, whether online or through a service call, recognize if the iHAT status of the home verifies this can be done safely and meet the service speed the customer expects. Alternatively, these tools can trigger an instant iHAT test for an updated result.

Billers – new CPE: When a customer changes CPE, possible iHAT variables that are affected are the device’s 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 risks to the iHAT state as recorded to test (or re-test) the home.

Operations – Job Scheduler: When a home “fails” iHAT, it goes into a remediation queue, with a flag for what needs to be remediated (video or HSD). This allows Tech Ops to plan proactive remediations, occurring routinely and not waiting for a house call to take care of iHAT-known issues.

Sales – Serviceability – Similar to Customer accounts, sales representatives should be able to quickly assess whether a customer, such as an MDU property, is eligible for MS speeds by accessing iHAT status in existing sales tools

Care – iHAT status, ITG Updates: When a TC arrives at an agent, after some amount of ITG-led questioning, the possibility of the issue being MS-related should be considered. A check on the iHAT status of that home, or an instant iHAT test, can help the triage process.

Tech Tools – Tech360: Like Care, when a tech is enroute or onsite to a customer home, part of the awareness the Tech should have is the MS status of the node, and the iHAT status of the customer. Further in the tools, the sequence of steps to diagnose and fix a MS-related issue should also be available.

Inventory Management – Drop Amp swaps: As remediations are made at relatively large scale to remove old drop amps, procurement awareness to the deployment of alternative solutions can ensure the supply pipeline is tracked and cared for

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

3.4.1. Applying MUSL End-to-End

Referring to Figure 6, we can show how iHAT fits within the broader perspective operationally, going from the trigger of a distributed access architecture (DAA) node cutover on the far left to the completion of activation on the far right.

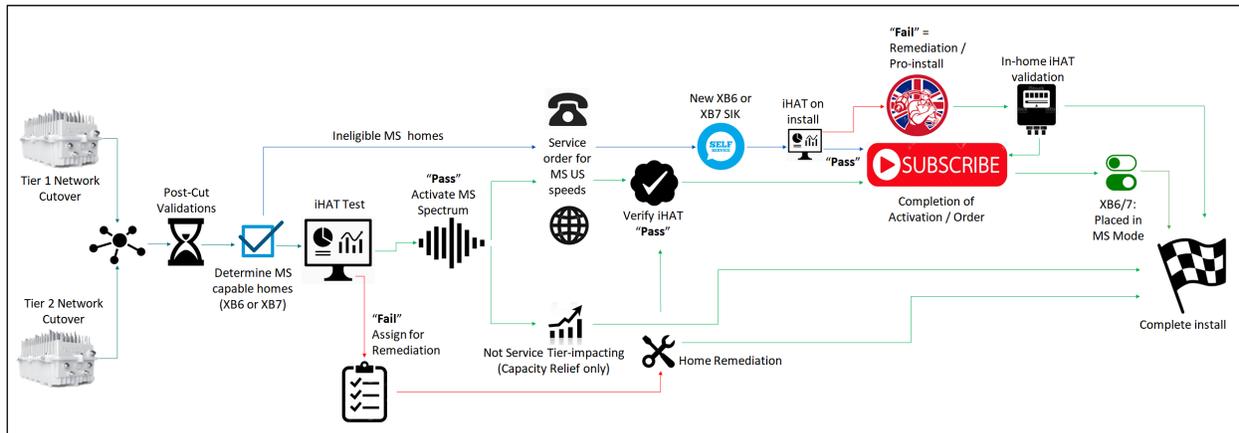


Figure 6 - Simplified Mid-Split Activation Flow: Cutover through to Activation

Beginning on the left, when a mid-split network upgrade occurs, the arbitrarily labeled “Tier 1” or “Tier 2” network cutovers can trigger internal tools, which can notify systems when construction is complete, and the cutover officially closed out. This can then trigger the spectrum activation process. Two things must happen prior to letting iHAT sweep across the node and validate homes where Mid-Split can be turned on. These are

- 1) *Post-Cut Validation* – Make sure that the network has settled to BAU metrics after the cutover. It is not uncommon to have a short period of elevated trouble calls shortly after a cutover, and it is desired to have any residual cutover issues resolved prior to moving to Mid-Split. This can be time-based, or it can be directly associated with trouble call 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 present in 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 6, 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 access the more efficient OFDMA spectrum above 40 MHz.

There is guidance 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 stragglers still in the network to maximize the DOCSIS 3.1 capacity.

Now, as shown in Figure 6, if the customer decides to upgrade their speed tier to one that requires Mid-Split, then of course getting them a gateway capable of that becomes a priority. Also, of course, 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 on iHAT and determine the state of the home for Mid-Split. If iHAT “PASS” is recorded, then the activation process continues, and iHAT sets the device into Mid-Split mode and it becomes capable of using the OFDMA

spectrum, in this case between approximately 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 (nor will they be billed for it!). When the remediation is complete, the technician will validate onsite with iHAT 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 6: “iHAT Test.” Let’s now follow the lower path – “iHAT FAIL.”

3.4.2. 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, and 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 what the failure mode is so that Technicians know what needs to be done. In general, remediation tasks are well-understood and known tasks to technicians – changing out home amplifiers for alternative devices, checking the splitter configuration, model, and wiring to Comcast compliance. After remediation is performed, the iHAT test is run to validate readiness for Mid-Split spectrum, and the activation then completed.

When a home is scheduled for remediation, assuming there is no speed tier ask requiring it, is a business decision with a number of variables having 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. 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 force iHAT to run off-cycle even after the initial iHAT sweep of the node at cutover.

3.4.3. iHAT PASS

The most straightforward flow in Figure 6 is right now 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 ably received at 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 at risk.

Going to the lower green flow down the center of Figure 6, 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 mid split plant upgrade plans, arbitrarily labeled as “Tier 1” and “Tier 2”, are counting on use of this capacity to defer any future network augmentation by many years. So, while it may not be noticeably service impacting to a customer, it is network impacting and indirectly service impacting by lowering the congestion on that node overall.

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 of the fact that an iHAT score is not static, a new iHAT score may make sense to obtain prior to upgrading the customer. The customer expectations for the new service will

be higher, and the awareness 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.

3.5 iHAT Software Solution Implementation

An overview of the innovative iHAT tool process flow is described above. This section contains a detailed deep dive into the automated software solution that was built for iHAT version 2.0. The new version of iHAT has been built exclusively for the devices that exist on a distributed access architecture (DAA) platform only.

There are several aspects of the software that have been prioritized when upgrading to version 2 from version 1.

Invasiveness

The major purpose of iHAT is to test whether a customer can successfully be switched to mid split without impacting their service. To do so, some service test must be initiated and compared to a baseline. Previously, in iHAT version 1, a speed test was run on the customer device. While this provided impressive results and proved to be a robust option, it left something to be desired in terms of efficiency.

The end goal of iHAT, and really any type of automated software that has the possibility of impacting the customer experience, is to lessen that impact as much as possible. To this end, iHAT version 2 has switched to a lightweight OUDP burst in lieu of the more heavy-handed speed test that version 1 relied upon. While running a speed test on customer hardware during peak hours would be out of the question, scheduling a short OUDP burst is not. This opens the possibility of running iHAT tests outside of maintenance windows without impacting the customer experience. The only current barrier to a truly unobtrusive experience is the need for a device reset after the split type has changed on the gateway’s diplexer – a limitation that will hopefully be lifted with a future software upgrade.

Scale

To properly scale out software improvements that can cover the entire network footprint, agile software is needed. Historically, network configurations and optimizations were carried out in an ad-hoc method by network technicians. While there may always be a need for humans to make tough decisions about how to go about solving some problems, there are other problems that are easily solved by simple logical automation.

One case study of this can be seen in the Octave project, where cable modem level data was gathered by the genome service and used to regularly optimize the DOCSIS configurations of CMTSs [6] [7]. Comcast as a company is striving to upgrade the network at a pace that cannot be matched by cherry picking small parts of the network one at a time. Therefore, the system that facilitates such sweeping network changes must be able to operate in a seamless manner.



UNLEASH THE POWER
OF LIMITLESS CONNECTIVITY
OCTOBER 11-14 ATLANTA, GA



In the case of iHAT, this automation comes at several stages. Version 1 had to manually run; this type of model can only be tenable for a handful of node segments/accounts/devices. In Version 2, node segments on the DAA will be automatically registered for the iHAT test pool when they become capable candidates based on network software and hardware criteria. Once a node segment is flagged as capable, automation is set in motion to filter the CPE devices that are serviced by the node segment to those that are iHAT eligible and testing will commence. MUSL will then periodically run iHAT as necessary based on network and account level changes.

Version 2 of iHAT is built to be flat and scalable. All tests are run at a gateway device level leveraging Amazon Web Services (AWS) step functions and lambdas. This type of architecture should theoretically make the iHAT service scale as needed until it can encompass the entirety of the network footprint. The only bottleneck comes from a DOCSIS limitation that only allows one OUDP burst to be scheduled at a time on a remote PHY device (RPD) level basis. This means that the theoretical maximum number of gateways that can be run per RPD is defined by the $(\text{maintenance window time}) / (\text{OUDP burst duration})$ with some padding on either end given for device resets.

Speed

The actual test only takes a few seconds of gathering empirical data and a second or so of computing. The bulk of the time during any given maintenance window is taken up by waiting for device resets, which will ideally be eliminated with hitless split changes sometime in the future.

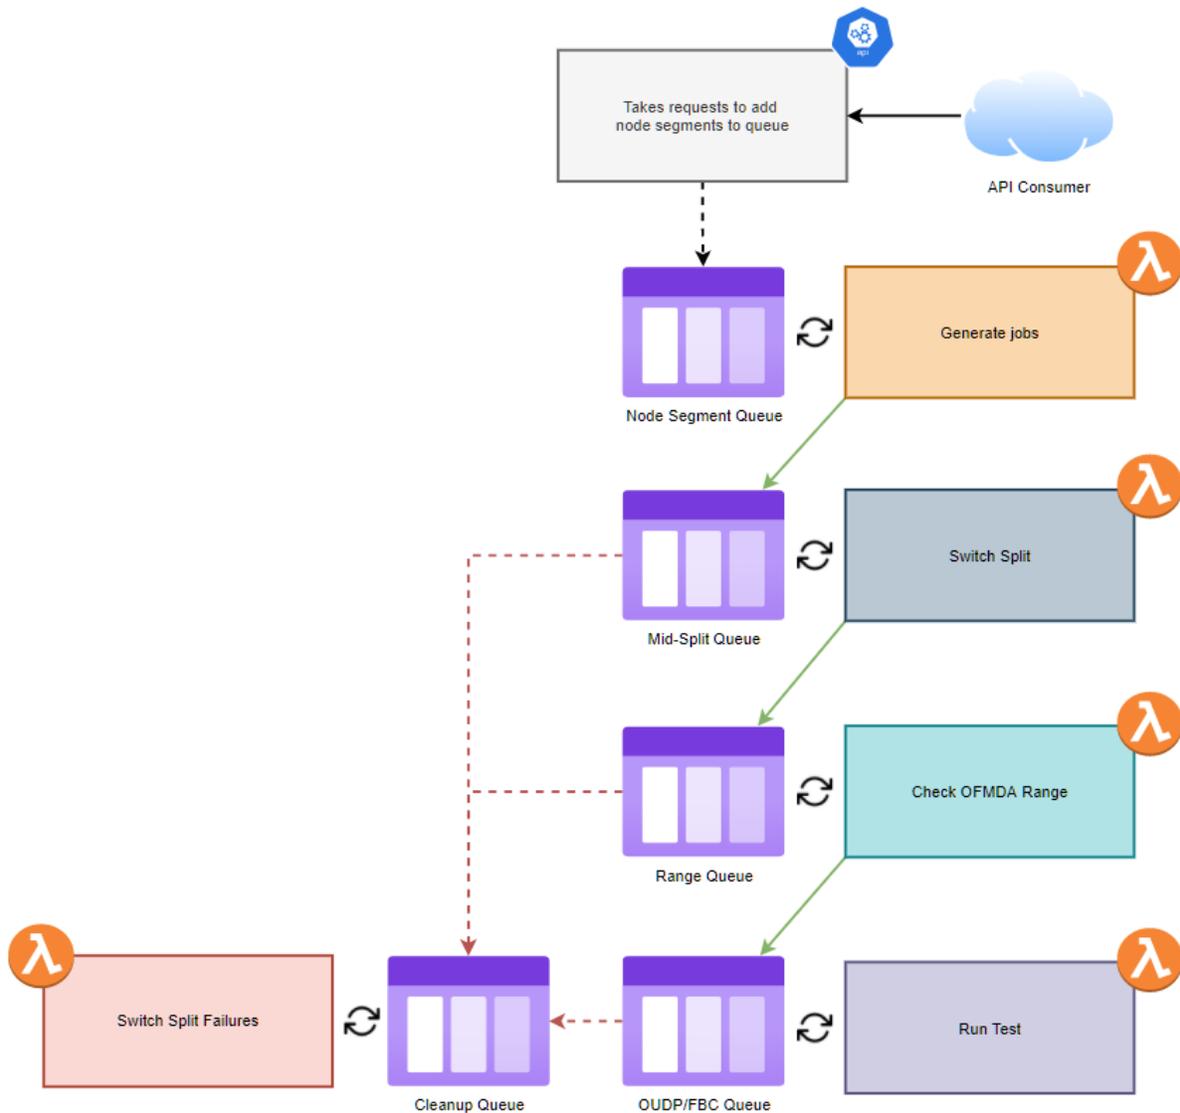


Figure 7 - High-Level iHAT Architecture

Error! Reference source not found. illustrates the high-level architecture that iHATV2 employs to orchestrate mid-split testing at scale. There are 6 distinct phases that are shown here.

Phase 1

Phase 1 is the acceptance phase. This is facilitated by having an application program interface (API) layer that accepts requests for node segment/RPD pairs to be added into the iHAT testing queue to be run in the next maintenance window. This API's sole purpose is to provide an authentication/authorization abstraction layer that wraps the node segment queue.



Phase 2

Phase 2 is the prep phase. This phase starts with spinning up lambda functions using CloudWatch events at the beginning on a maintenance window. Two integrations come into play in this phase. First, the federated data service (FDS) API is called with each node segment name. Node segment names are swapped for MAC addresses that are grouped by service location. At this point, the MAC addresses are opaque, and the genome service must be queried for each MAC address to get device metadata. Genome provides important data such as the model number and current diplexer switch type. The starting switch type is stored so that each device can be placed back in its original state in the case of inconclusive test results.

The model number is then used to classify device types and group devices on an account. Accounts can then be further classified as mid-split eligible based on the kinds and configurations of the devices. Each eligible gateway on an account will generate one “device-level” job. Each device-level job will include one gateway and all set-top-box devices on the account that can produce full band captures (FBCs). Jobs are placed in the switch-split queue for further processing.

Phase 3

Phase 3 is the switch split phase. The phase 3 lambda is responsible for making sure that all modems are switched to mid split mode prior to data collection. If the gateway is already in mid-split mode, then the lambda is responsible for moving the job from the switch-split queue to the range queue, so the cable modem can be immediately processed. If, however, the gateway is not in mid-split mode, then the lambda is responsible for first calling genome to initiate the switch-split into mid-split, and then calling genome again to initiate a modem reset. Once the modem reset has been initiated, the job can be moved from the switch-split queue into the range queue with a 2-minute message timer. This message timer will ensure that minimal resources are spent on waiting for the modem to reset. It is important to note that the rest and delay portion of this lambda may still be required when iHAT switches to a hitless mode of operation.

Phase 4

Phase 4 is the ranging phase. It imperative that the modem is in the correct state before iHAT data collection begins. The range queue that serves phase 4 is the only queue that is not a first-in, first-out (FIFO) queue as it must house two different types of messages. The first type of message represents cable modems that were already in mid split and did not need a modem reset. These messages should appear immediately in the queue as they have no message timer associated with them. The second type of message represents those cable modems that required a modem reset. These messages will only appear after the 2-minute message timer has elapsed. Once a lambda consumes either type of message, it polls genome for a device online status. Once the device is online, then the lambda confirms that the device can range on OFDMA. If the device fails to range, then the modem is sent to the cleanup queue, where it can be switched back into sub-split mode. This type of failure is most likely due to blockage by a drop amp. If, however, everything looks good at the device, --it's online and ranging, the job can be moved to the OUDP/FBC queue with a message ID corresponding to the RPD that the customer gateway is connected to.

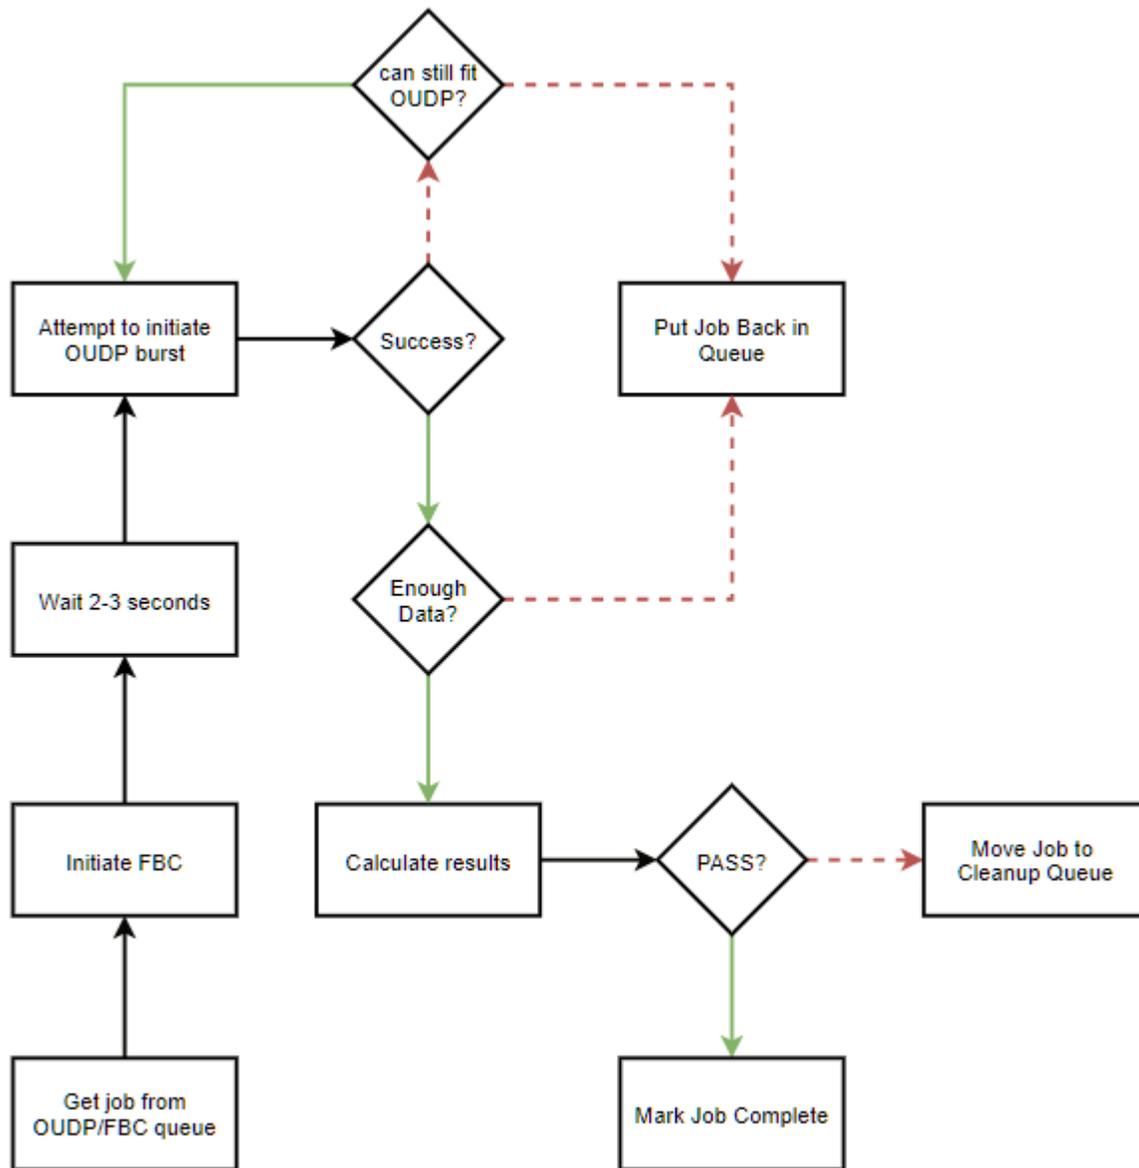


Figure 8 – OUDP/FBC Orchestration Workflow

Phase 5

Phase 5 is where the meat of the iHAT test is performed. A visualization of the workflow is depicted in Figure 8. This phase is slightly more involved than the rest of the phases because careful orchestration of the OUDP burst resource is required. As per the DOCSIS spec, only one OUDP burst on a certain frequency range can occur on a given RPD at one time. This means that one must be very careful when attempting to scale a process that heavily depends on the availability of the OUDP burst service.

The concurrency of this process is controlled by making sure that RPD names are used as message group IDs in the OUDP/FBC queue. In all other queues that exist in the iHAT V2 workflow, the message group ID is the unique value given by the job ID. Amazon simple queue service (SQS) enforces that only one message from a given message group ID is consumed at a given time if the message batch is set to 1. This

means that if all jobs that are associated with a given RPD are associated with the same message group ID, they should execute in serial.

The maximum theoretical concurrency for the data collection process is equal to the number of unique RPDs. Earlier iterations of the architecture tended towards forcing a concurrency equal to the number of unique RPDs in the maintenance window. The fact that a full range of scaling options is available is important, as the iHAT workflow relies on the availability and scale of many other systems to function properly. Being able to calibrate scale to successfully integrate with other systems is a must.

The workflow begins with a lambda firing from the SQS event trigger. First, genome is called to initiate FBC on each associated set top box in the job. This takes 2-3 seconds to initiate, so the lambda sleeps for 3 seconds before attempting to initiate the OUDP burst through the gateway. If the OUDP burst fails, then it is attempted up to 2-3 additional times depending on the FBC duration that was chosen. If it still fails, or if the FBC process does not capture enough data for any of the set top boxes, then the lambda places the job at the end of the queue to be retried later.

If the FBC and OUDP burst orchestration is a success, then the iHAT calculation can commence. This part of the code only takes on the order of a second to complete. If the calculation yields a success, then the OUDP/FBC job is marked complete, and the job has reached the end of the iHAT workflow. If the calculation yields a failure, then the job is moved to the cleanup queue.

Phase 6

Phase 6's purpose is to clean up all failed jobs. It used as the dead letter queue for all queues except the node segment queue. It can also accept jobs from the phase 5 lambda function in the case that a successful test is run, but a failure is noted by the telemetry indicating video service degradation. Categorized failures include, but are not limited to: failure of the device to switch split types, failure of the device to reset properly, failure of the device to range on OFDMA, failure of the OUDP burst to execute properly, and failure due to video noise. It is also possible that the iHAT workflow fails due to underlying software issues, which could be due to network failures or failures of external systems to provide required data.

There are three outcomes of the iHAT workflow. Case 1 is the case where everything is green and the iHAT calculations yield a PASS. In this case, the customer gateway is left in mid-split. Case 2 is the case where the iHAT calculations yield a clear FAIL. In this case, the customer gateway is switched back to sub-split. Case 3 is the case where the iHAT workflow failed and is inconclusive. In this case, the customer device is switched back to its original setting that was recorded in phase 2, whether that is mid or sub-split. Case 2 and 3 may require manual remediation.

The workflow outlined above serves the BAU process of moving customer devices in bulk to mid split and routinely testing them in bulk to ensure network health. There is, however, a separate use case for on-demand testing for customers with a technician in the home. Since the above workflow uses highly modularized components, serving this case is trivial as the individual lambdas can be exposed via API and used by customer care tools. The current use cases are for a technician to move a customer modem from mid to sub-split or vice versa and for an on-demand OUDP/FBC data collection and calculation to yield a score.

The iHAT project requires several integrations with external systems. The one that is used most heavily is the integration with the genome service. The iHAT tool would be much more complicated if it had to deal with authentication and SNMP, but the genome service abstracts those technical details away, allowing for a much cleaner and elegant solution. The next section describes in detail how the genome service achieves this level of abstraction at scale.

3.5.1 Data Collection & Device Interface Platform (Genome)

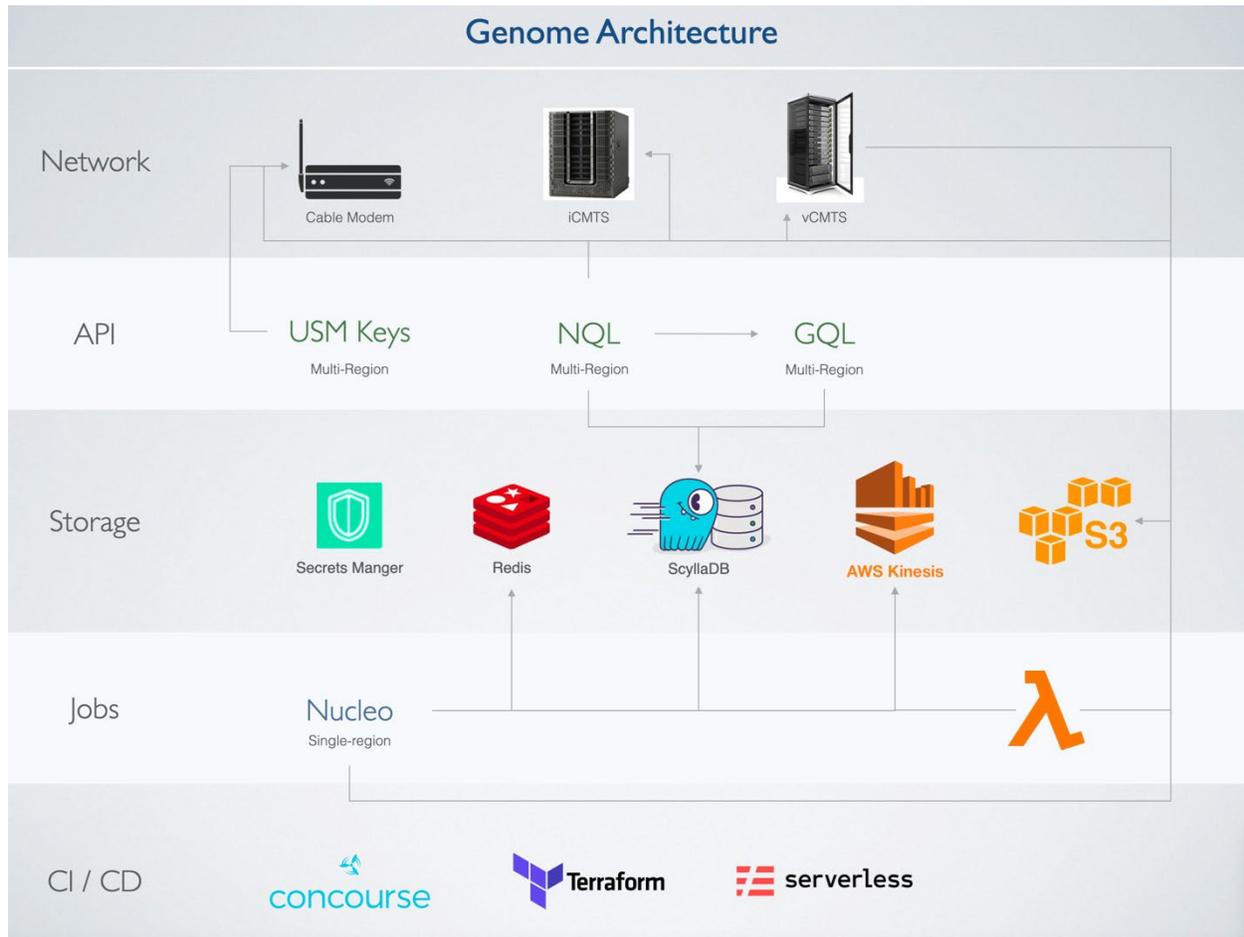


Figure 9 – Genome Architecture

The Genome platform in Figure 9 was originally built to serve the Octave Profile Management Application (PMA) initiative but has grown into a platform of its own since inception [6] [7]. The main purpose of Genome is to actively poll and cache data from devices such as cable modems and CMTSs in a scalable and configurable way while offering consumers the ability to analyze the cached data or get live data in a seamless fashion (on-demand data polling). This positions Genome to be the sole data provider, SNMPV3 secure key collection, data validator and device interface layer for cable modem and CMTS data, eliminating inefficiencies around oversaturating CMTSs or cable modems with multiple connections and over fetching duplicative data. In the context of iHAT, Genome is responsible for verification of eligible accounts and the devices that has been provided by business and their current status, running spectrum analysis while the OUDP has been triggered.

Genome is made up of two layers, the poller layer, and the query layer. A “poller” is simply a piece of software that is responsible for scheduling, collecting, and standardizing a set of data from devices. Pollers are built to be lean, modular, and extensible. The query layer for Genome is called Network Query Language (NQL), which exposes a declarative API service through which consumers request live or cached data. All pollers are primary consumers of NQL to collect live data, while also offering external consumers to do the same. NQL’s goal is to offer a declarative abstraction layer for edge network devices, allowing

consumers to query using standard hypertext transfer protocol secure (HTTPS) instead of simple network management protocol (SNMP), trivial file transfer protocol (TFTP), and various other network communication protocols.

Genome requires a list of eligible devices to poll the necessary data for iHAT. iHAT integration layer works as the mediator between business system which holds the list of eligible Nodes/Accounts and verifies the OFDMA being active, along with mid-split capability and eligibility, then feeds Genome with this eligible device list to check the device status to make sure they are in the right state to run iHAT. The polling could be done at any configured frequency or on-demand.

In addition, Genome must ensure that data collection can scale when devices are added over time such that data collection and aggregation can be achieved in a given polling window. It must also ensure that the ingested data is validated and cleaned before it is cached. As the amount of data is large, especially in the case of cable modem data, Genome must also manage data retention policies in order to reduce cost.

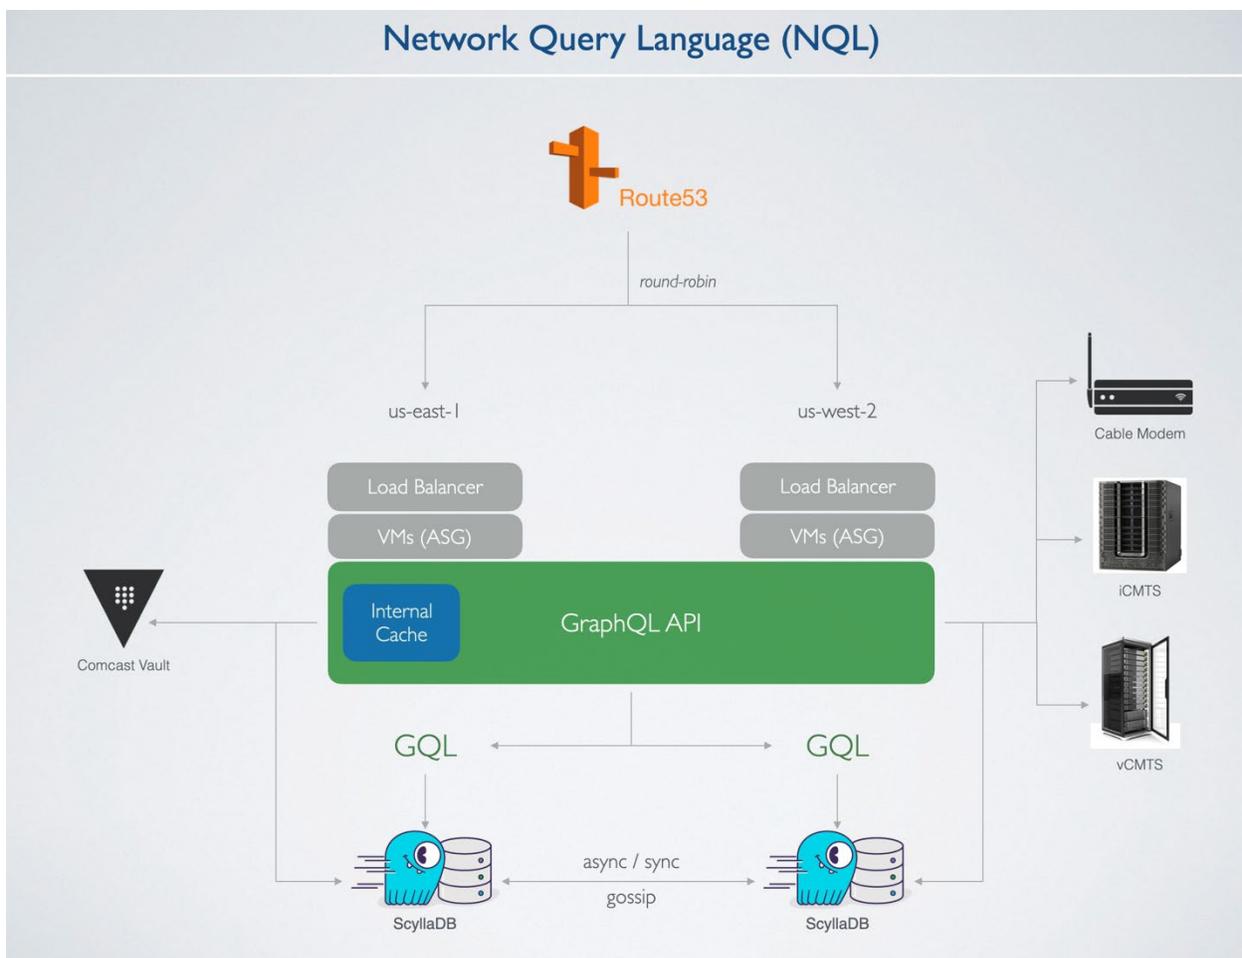


Figure 10 – Genome API – NQL Architecture

NQL’s journey began in the early days of Octave development as well. The goal was to create an API service which abstracts away all different protocols around networking and let end users interact through HTTPS. The first iteration came in the form of a representational state transfer (REST) API which accepted object identifier (OIDs) parameters and returned the output through HTTPS. While this iteration ran in

production for several months, several shortcomings were brought to light. For example, use cases involving many data points per CMTS or CMs, necessitated making multiple requests to the API. Each request would connect to the CMTS, which resulted in opening and closing sockets many times over multiple requests. Furthermore, the work of encoding and decoding large amounts of javascript object notation (JSON) data was repeated for each request, degrading the overall performance of the service through repetition of work. Graph query language (GraphQL) was a natural option that allows data schemas to easily evolve and can handle complex relationships between data sets. NQL architecture in Figure 10 was born.

As NQL offers an abstraction layer for connectivity to both cable modems and CMTSs, it must therefore manage all writes as well as reads to each network device. In order to accomplish such a feat, it supports both IPv4 and IPv6 communication protocols. NQL is built primarily using the GraphQL specification and has been optimized at node level. In some cases, consumers may require a request-response type handshake, where the consumer keeps a socket open until the requested data is returned. However, many queries may be long running, and it may not be practical or possible for a consumer to keep a connection open for the duration of the process.

Historically, the team used Nodejs as a declarative jack-of-all languages. However, in this use case, performance is king. As Golang is routinely touted as a performance-oriented language geared towards networking applications, it was a simple choice after orchestrating some internal benchmarking. Genome as of today supports 3+ billion transactions a day with ability to scale up as we need. Simple GraphQL APIs were built in both languages to test a small fraction of our total scale.

NQL's secure shell (SSH) feature allows consumers to log into a supported remote device through HTTPS rather than SSH. This allows NQL to abstract away details around authentication, authorization, and managing the underlying SSH connection. Users are able to connect to a host, run multiple commands, and get output back for each command. NQL uses GraphQL to define the API, which allows developers to develop powerful features and evolve our API without breaking the world. NQLs SSH is a simple layer around SSH, all the logic around what commands to run in what order and what to do with the output is handled by the client itself. For example, configuration management can apply configurations to a CMTS and then analyze the output to see if everything went well. All this is done through HTTPS using NQL. NQL has evolved from a standalone service running on an elastic compute cloud (EC2) virtual machine (VM) to where it could be leveraged anywhere within our codebase and able to run in any container or lambda alongside our codebase.

SNMP v2/v3 capability and the user-based security model (USM) are shown in Figure 11. Genome by default uses SNMP V3, and if the device is not V3 enabled, the fallback option is V2. The role of USM is as follows:

- Get public and manager key from cable modem
- Use manager key and V3 security name as input to get vault secret
- Use cryptographic hash function on vault secret and public key and compute auth priv keys.

With authentication in place, collection of the SNMP data can be performed on an as needed basis and streamed to Kinesis, which will be used by Analytical Engine from PMA and other consumers of the data including iHAT. iHAT uses NQL layer as the interface to do the SNMP set for moving the device from sub to mid-split.

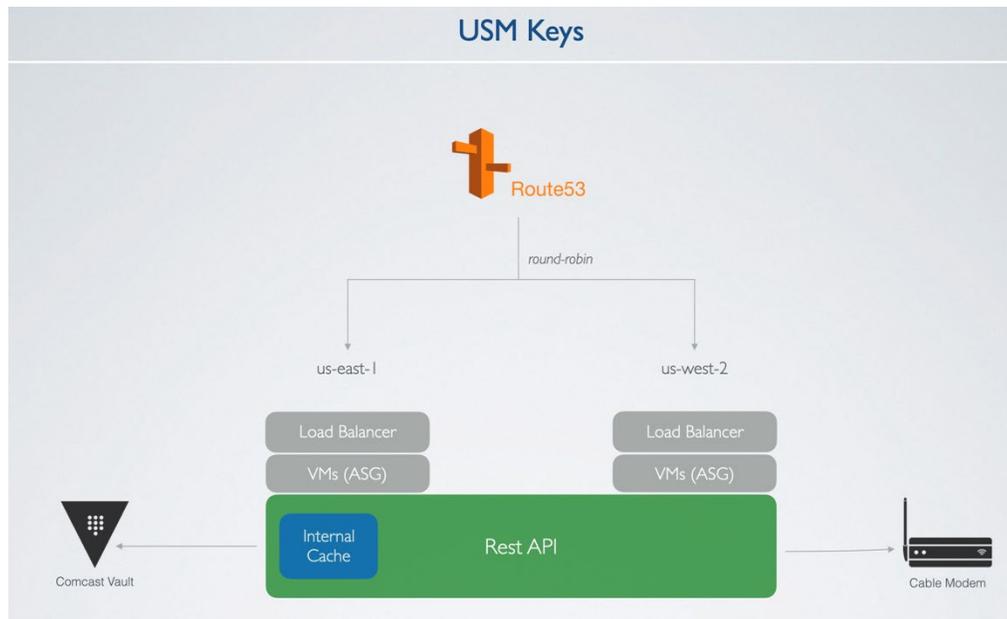


Figure 11 – SNMPv3 USM key architecture

This section has provided a detailed view of the iHAT software architecture covering genome, NQL and USM. Overall, the architecture is like previously presented PMA architecture [6] [7], with many software components being reused to support iHAT functionality. With an understanding of the iHAT operation for mid-split, we would like to now discuss future enhancement considerations for expanding upstream path to high-split and what new functionality would need to be considered for future generations of iHAT software.

4. High-split Future

Plant upgrades with mid-split are just the beginning of extending the upstream spectrum and increasing the upstream throughput capability. Mid-split spectrum activation adds needed capacity in the upstream and extends the life of node splits. To achieve 1 Gbps symmetrical services significantly more upstream spectrum is needed. High-split development and deployments are in process and FDX is right around the corner to achieve one Gbps and greater symmetrical services.

Throughout the development of mid-split, there were many aspects of the access network which were evaluated and many challenges which were overcome. Many of the workstreams and lessons learned in mid-split extend to high-split development and deployment.

OFDMA activation, node and amplifier upgrades, the shift of the downstream spectrum to allow room for the extended upstream spectrum, adjacent channel interference, plant design levels, CPE devices, in home amplifiers and splitting networks were all extensively evaluated and validated for mid-split deployments. The learnings from the development were and continue to be used for high split.

Comcast's network upgrades have predominately been N+0 and recently include amplifiers (N+X) for mid-split deployments. The process of splitting and upgrading a standard-split analog node to mid-split digital nodes is well understood and completed many times a day throughout our networks. Creation and push of the downstream channel map, upstream channel configuration and upgrade of plant equipment; nodes and amplifiers all happen simultaneously and with minimal disruption to our customers. This same

process can be used for mid to high split upgrades. Mid split nodes amplifiers are 1.2 GHz, so there's plenty of spectrum available to move the downstream spectrum above 258 MHz prior to high split equipment installation. Tools, like Comcast's Scout Monitoring, have been developed to monitor and validate many physical layer parameters, including OFDMA in the upstream, and these have been extended to 204 MHz in the upstream. Examples of these tools are shown below. Figure 12 and Figure 13 show the upstream DOCSIS 3.0 and 3.1 channel status via Scout.

Mid-Split
4 SC-QAM
1 OFMDA

Device Health	
Registration State	8 (Online)
Down Rx Power	-2.4 -2.6 -2.3 -2.2 -2.3 -2.4 -2.2 -2 -2.1 -2.1 -1.9 -1.8 -1.8 -1.8 -1.7 -1.7 -1.6 -1.6 -1.4 -1.1 -1 -1.1 -1.7 -1.7 -1.5 -1.6 -1.7 -1.6 -1.6 -1.7
Downstream SNR	45.7 45.6 45.5 45.7 45.5 45.5 45.6 45.5 45.5 45.6 45.7 45.8 45.7 45.7 45.7 45.7 45.6 45.7 45.9 45.7 45.6 45.3 45.3 45.3 45.3 45.3 45.2 45.2
Upstream Tx Power	44.3 44.5 45.3 46
Upstream SNR CM	40.5 41.2 42.7 43.1 42.3
Upstream Rx Power	0.0 -0.25 0.0 0.0 0.5
US RX/WO Padding	0 -0.2 0 0 0.5
Upstream SNR Ch	40.2 41.6 42.0 42.2 42.0
Upstream Ranging	4 (Success) 4 (Success) 4 (Success) 4 (Success) 4 (Success)
Upstream ICFR	
Upstream Ripples	For more detailed analysis, click on the Flux icon below: 
Upstream Distortion	
T3 / T4 Timeouts	null / null 0 / 0 0 / 0 0 / 0 0 / 0
Resets / Lost Syncs	434 / 0

Figure 12 - Mid Split Scout Screen Capture with 4 SC-QAMs and 1 OFDMA

High-Split
4 SC-QAM
2 OFMDA

Device Health	
Registration State	9 (Online)
Down Rx Power	-1 -1 -0.9 -0.7 -0.7 -0.7 -0.7 -0.5 -0.7 -0.2 -0.2 0 -0.2 0 0.2 0 0.4 0.2 0.2 0.7 0.7 1 1 0.7 1 1 1 1.4 1.2 1.2
Downstream SNR	40.3 40.0 40.3 40.3 40.3 40.9 40.3 40.9 40.9 40.9 40.3 40.9 40.9 40.3 40.9 40.3 40.9 40.3 40.9 40.3 40.3 40.9 40.3 40.3 40.3 40.3 40.3 40.3
Upstream Tx Power	45.5 46 46 46.5
Upstream SNR CM	45.0 45.3 45.3 45.3 46.3 46.3 46.3 46.3
Upstream Rx Power	0.0 0.0 0.0 0.0 0.0 0.25
US RX/WO Padding	0 0 0 0 0 0.2
Upstream SNR Ch	45.2 43.1 43.1 42.8 43.3 43.3
Upstream Ranging	4 (Success) 4 (Success) 4 (Success) 4 (Success) 4 (Success) 4 (Success)
Upstream ICFR	
Upstream Ripples	For more detailed analysis, click on the Flux icon below: 
Upstream Distortion	
T3 / T4 Timeouts	null / null 1307 / 0 1690 / 0 1961 / 0 1843 / 0
Resets / Lost Syncs	0 / 0

Figure 13 - High Split Scout Screen Capture with 4 SC-QAMs and 1 OFDMA

Figure 14 and Figure 15 show the upstream spectrum as received at the RPD via Yeti for mid split systems to 85 MHz and high split systems to 204 MHz. Yeti is a “real time” upstream spectrum analyzer tool used to evaluate the upstream channels and also the underlying upstream noise.

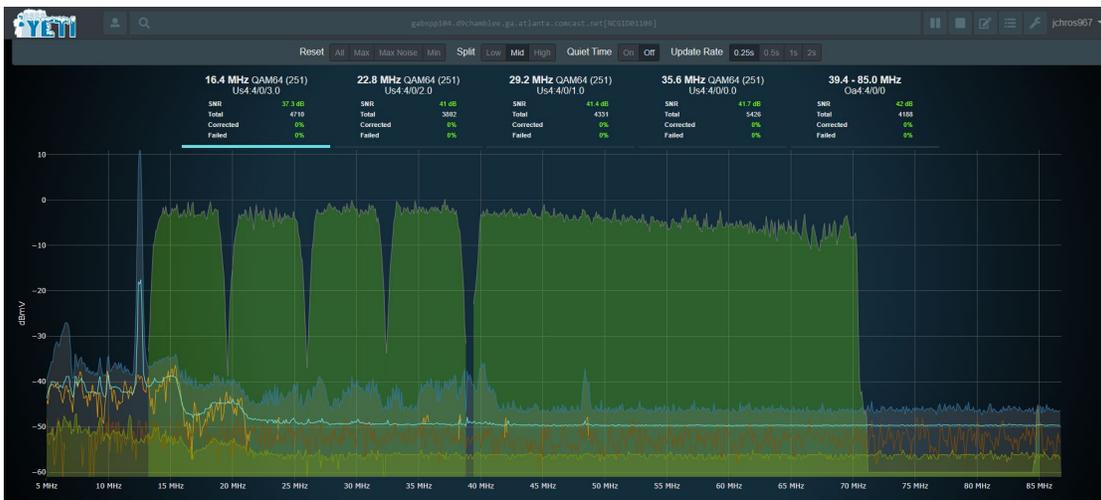


Figure 14 - High Split Yeti Screen Capture at RPD with 4 SC-QAMs and 1 OFDMA Channels

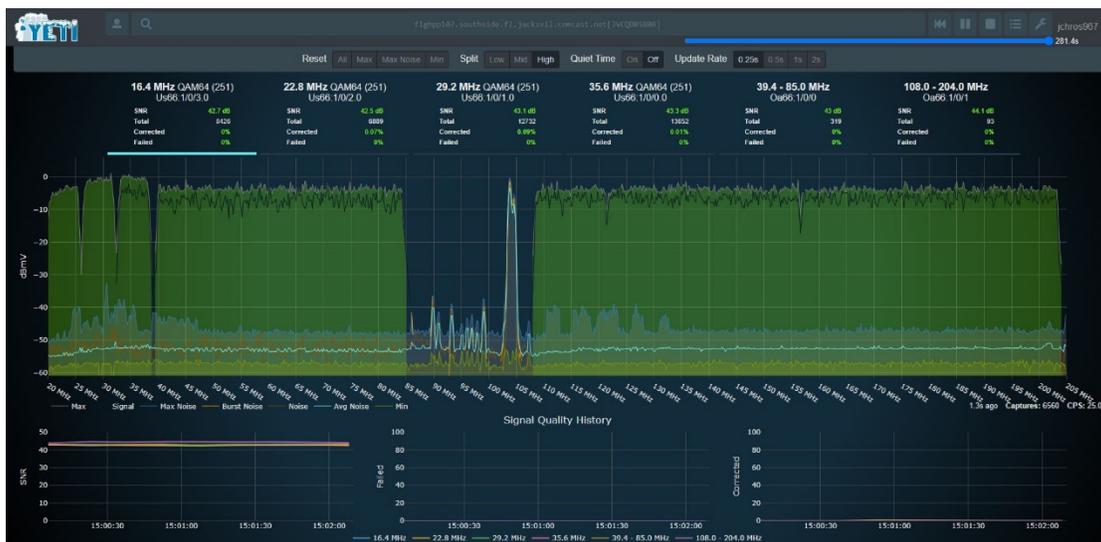


Figure 15 - High Split Yeti Screen Capture showing 2 OFDMA Channels

Note that the lower spectrum edge starts at 20 MHz due to current FFT limitations

4.1. Extending iHAT for Neighbor Interference Case

ACI has been studied extensively for mid-split plants. Thresholds of interference for our set-top boxes are well understood and as discussed iHAT development is well under way for mid-split deployment to analyze ACI with no customer impact. iHAT first started out exercising the full mid-split upstream band with a speed test and the adjacent in-home set-top box performance was monitored during the speed test. This was a very effective way to measure the impact of the mid-split spectrum on the adjacent in-home set-top boxes. The speed test could interfere with customer high speed data usage and interference was possible on the set-top box, if the in-home isolation between the gateway and set-top box was inadequate.

During the development of high-split, a method of evaluating leakage was needed where the high split gateway in the home creates a leakage tone. DOCSIS 3.1 supports OUDP signals and APIs have been developed to have the vCMTS command a cable modem to create an OUDP burst to be used for leakage detection. These same OUDP bursts are being used to evaluate adjacent channel interference. The OUDP burst is created by the mid-split gateway and a simultaneous full-band capture is taken on the adjacent in-home set-top box. Through extensive testing the thresholds for the set-top boxes for interference are well known and the measurement of the delta between the OUDP burst as measured on the adjacent set-top box and the downstream signal is measured to evaluate the potential for ACI. If the delta is less than the threshold the account can support mid-split spectrum activation. If the delta is more than the threshold, the account is targeted for remediation.

This same process can be used for high split systems to evaluate the potential for high split neighbor interference. In the mid-split plant, there is enough tap-to-tap isolation and the interfering bandwidth is smaller, so neighbor interference is not a concern. In a high-split system, the extended upstream bandwidth coupled with lower tap-to-tap isolation with higher frequencies creates the potential for neighbor interference. See Figure 16.

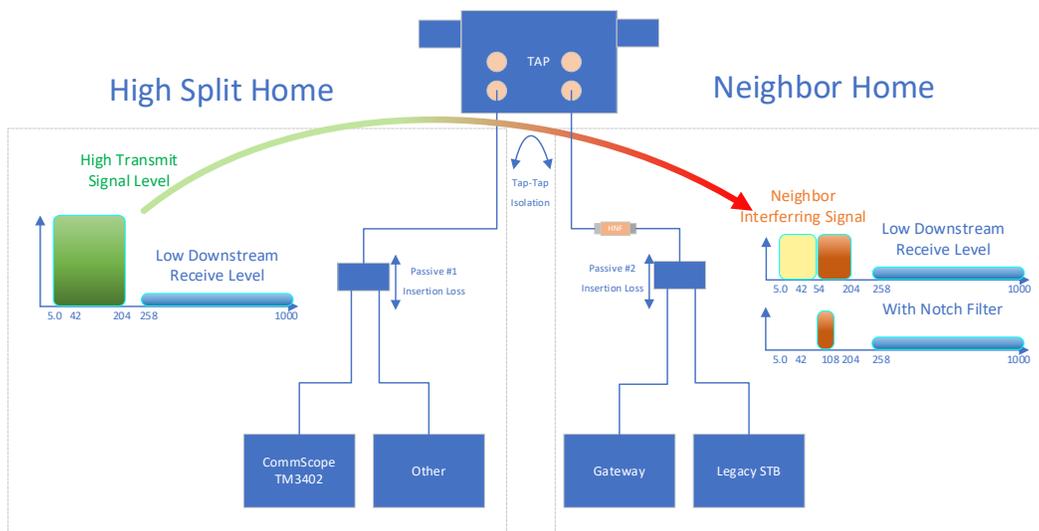


Figure 16 - Neighbor Interference in a High-Split System

The same process the iHAT tool uses with the OUDP burst can be used to evaluate high-split neighbor interference. When high split modems are first deployed and before service is activated, iHAT can be used to create the OUDP burst on the high split device and the neighboring standard and mid-split devices can be evaluated using the same full band capture technique used for mid-split ACI. Additional system and plant information is used to determine which devices are connected to the same tap as the high split device and these devices are monitored.

In a Mid-Split system, ACI and iHAT are focused on adjacent set-top boxes in the same home as the mid-split gateway or cable modem. In a high-split system, neighbor interference can occur both on neighboring set-top boxes and neighboring gateways and cable modems. As in the mid-split case, extensive testing has been completed to understand the interference and threshold levels for high split interference both on the set-tops and gateways.

Testing was completed on both set-tops and gateways with a “Neighbor” high split cable modem passing traffic under various upstream throughput scenarios. This setup for neighbor interference testing is shown in Figure 17.

For neighbor video interference testing there are two main test scenarios

1. Video: system connected to DAC for the STB testing
2. Data: connected to CMTS for CM testing

The same set-up was used for testing both Data and Video (QAM) CPEs. The first vCMTS is connected to the HS node to generate the upstream load to the HS modem is configured for interference as follows

- 4 x 6.4MHz SC-QAM channels (for the SS CMs only)
- 48 MHz OFDMA (up to 85MHz, MS)
- 96 MHz OFDMA (108-204MHz, HS)
- US Traffic and Level adjusted to produce the desired interference signal

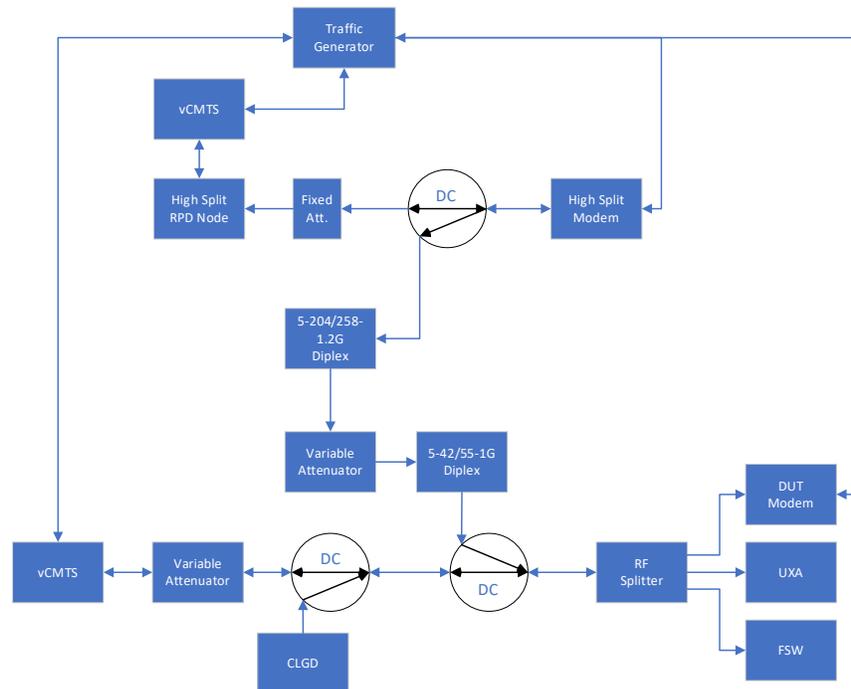


Figure 17 - Neighbor Interference for STB Test Configuration

The DUT modems are configured as follows:

- CM DS Load (258 – 1002MHz):
- OFDM (96MHz, 792 to 888 MHz)
- SC-QAM (603 to 789 MHz, Symbol Rate: 5.360537MSym/s, 256QAM)
- SC-QAM above and below those frequency ranges from CLGD
- Add noise to the DS to decrease the received MER to 38 dB
- A traffic generator is used to produce a fixed DS traffic for the DUT (300Mbps data rate) while the HS data rate was varied from 1 to 980 Mbps

- STB DS Load (258 – 1002MHz)

The Interference Level is defined as the ratio of the high-split Interference power spectral density (PSD) (dBmV/6MHz) with respect to the device under test (DUT) downstream PSD (dBmV/6MHz). It is important to note that the high-split Interference ratio takes in consideration the occupied bandwidth for each high-split interfering traffic rate which is CMTS scheduler dependent. The methodology for determining the High-Split threshold interference was based on a measured DUT DS codeword error rate (CER) above a threshold of $1e^{-6}$ for the Data scenario while an observer noticing video artifacts was used for the Video scenario. The variable attenuator controlling the high-split Interference level applied to the DUT is adjusted for each particular interference high split traffic rate until the CER threshold is reached or video artifacts are observed.

When testing interference low upstream throughputs correspond to low OFDMA utilization and/or a low duty cycle for the OFDMA channels occupying the band from 39.4-204 MHz. As upstream traffic rates are increased, the OFDMA channels are utilized more and the duty cycle is reduced, up until the point that the full spectrum from 39.4-204 MHz is being utilized 100% of the time. The threshold of interference under this testing shows that under very low utilization the interfering threshold is high. As traffic increases, there is a point where the interference threshold is very low and then as the traffic is maximum, the upstream spectrum is at a steady state and the interference threshold rises. This can be seen in both the interference of set-tops and in gateways, as shows in Figure 18.

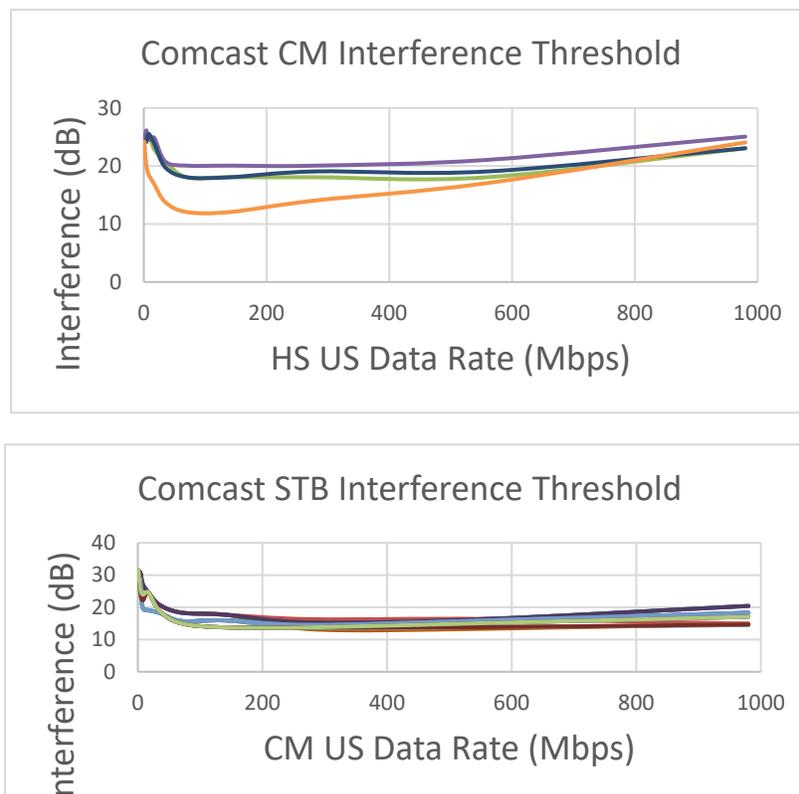


Figure 18 - Cable-Modem & Set-top Box Interference: Threshold vs Neighboring High Split Throughput for Several Different Models of Set-tops

A mathematical model has been developed to better understand the DUT front-end behavior and identify opportunities to improve the CM/STB performance. The goal is to capture the impact on the CM front end including front end overhead, automatic gain control (AGC) behavior, Downstream/Upstream power ratio, high-split upstream occupied bandwidth, and duty cycle. A typical CM front-end behavior is well described in reference [5].

Figure 19 shows the comparison of the model with data measured for a Comcast CM. It also shows the main regions of a typical CM or STB behavior.

- Low Energy Region: range with low interfering signal utilization (occupied BW and duty cycle) requiring higher levels of interference signal to impact the victim CM front end
- Maximum Interference Region: range where the increased occupied BW and duty cycle produces the highest levels of interference and corresponds to minimum interference thresholds
- Quasi Steady-State Region: range where the high interfering signal utilization gets close to a continuous mode requiring higher levels of interference signal to impact the DUT CM front end

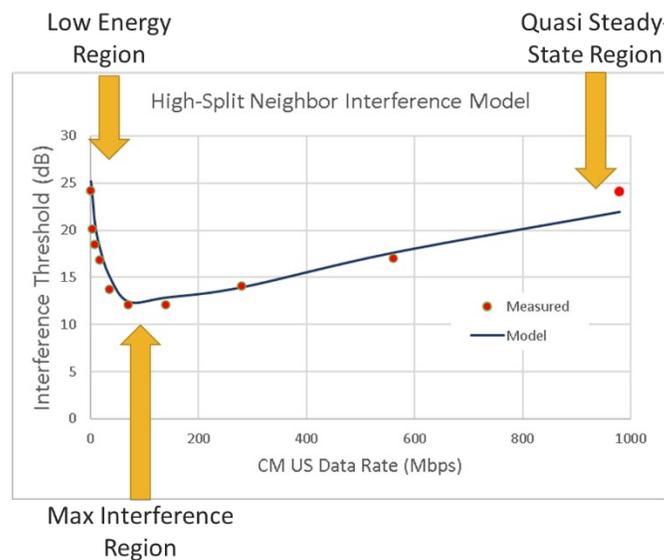


Figure 19 - High Split Threshold vs Neighboring High Split Throughput Model

It is important to note that the maximum interference frequency range is dependent of the CM front-end and CMTS scheduler characteristics. In the case of the CMTS scheduler, grant allocation implementation, which translates into a certain upstream occupied bandwidth and duty cycle, defines the interaction with the CPE front end.

Additionally, for system design purposes, the minimum threshold level is used since the upstream data demand varies all the time. Comcast measurements have indicated a minimum threshold levels in the order of 12 dB which is a little better than the “Maximum average power of carrier input to CM, within any 6 MHz channel from 54 MHz up to 1002 MHz” and “256-QAM Image Rejection Performance” PHYv3.0 DOCSIS specs, which limit the interference level to 10 dB. New case scenarios could be developed and proposed to the CPE chip-set manufacturers to further improve the CPE performance, particularly as the industry starts to get ready for FDX. However, tools like iHAT are critical for a smoother transition.

If high OUDP signal to downstream levels are measured and potential neighbor interference is identified, several remediation techniques can be implemented. One of the potential solutions is the use of notch filters that can be added to neighbor drops and suppress the interfering signal. Although not operationally desirable, these filters can be designed to notch out the complete spectrum from 54 to 204 MHz or to be compatible with mid-split systems, notch filters have also been designed and tested to pass up to 85 MHz and notch the spectrum between 85 and 204 MHz. Figure 20 shows the frequency response of a notch filter used to prevent the impact of the High-Split neighbor interference. Filter suppression higher than 20 dB would suffice for most case scenarios.

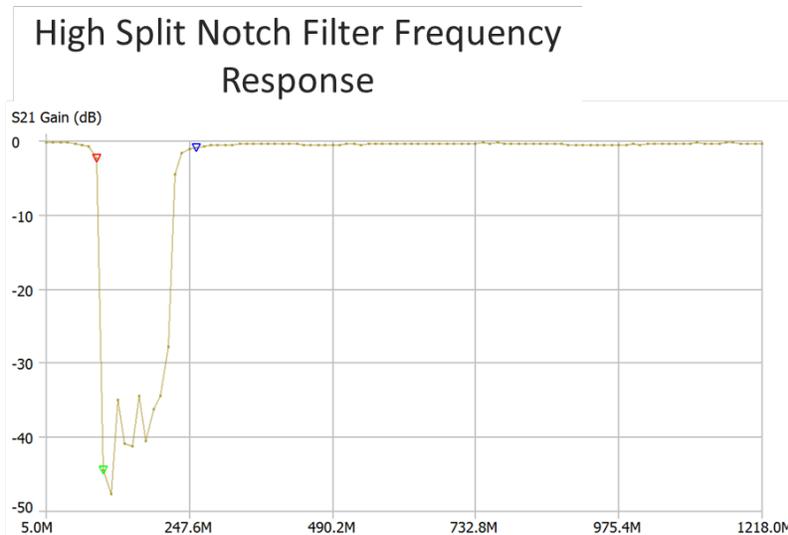


Figure 20 - Neighbor Interference Notch Filter 85-204 MHz)

This allows mid-split and high split devices to co-exist and minimizes device types deployed in the field. New tap plates can also be installed. Newer taps are specified with better tap to tap port isolation to help minimize neighbor interference.

With minor updates to the existing iHAT tools for ACI, iHAT can be adapted to solve the high-split neighbor interference. The threshold of the OUDP to downstream signal can be adjusted to match the measured interference levels for high-split interference, and the back-office tools modified to run/measure OUDP on different accounts. The frequency of the OUDP burst and duration can be optimized for certain areas of the band based on the isolation measurements and the duration may also be optimized.

iHAT and OUDP measurements are an ideal tool to validate the high-split system prior to activation to optimize the customer experience.

5. Conclusion

New work and learning have continued since the introduction and subsequent rebranding to iHAT this past year. Aside from the cool name, improvements upon the process of evaluating ACI using spectrum-based methods have been made that will one day enable OUPD leakage detection to occur in a non-service affecting manner, harmoniously with regularly scheduled traffic. As we work toward scaling iHAT, our intent is to focus on making this tool work as quickly as possible, while offering robustness of carrier grade software at scale using proven tools readily available in the marketplace. A well-defined

object model acting as the glue between iHAT process and downstream APIs can lead to the scale operations needed at Comcast, and reliable distribution of results into the hands of frontline teams who can act in a more proactive manner to address remediation issues. With an eye on the future, including high-split, iHAT has room to grow into a toolset that will one day make it easier for operators to operationalize 1 Gbps upstream speeds. It is our hope that these steps in upstream evolution will help us to embrace new DOCSIS 4.0 FDX sounding technology when the time is ready, and work to harmonize new technologies while always maintaining support for legacy service.

Abbreviations

ACI	Adjacent Channel Interference
AGC	Automatic Gain Control
API	Application Program Interface
AWS	Amazon Web Services
BAU	Business as Usual
BW	Bandwidth
CER	Codeword Error Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
COS	Class of Service
COVID-19	Corona Virus Disease 2019
CPE	Customer Premise Equipment
DAA	Distributed Access Architecture
dB	Decibels
dBc	Decibels Relative to a Carrier
dBmV	Decibels Relative to one Millivolt
CER	Codeword Error Ratio
DAA	Distributed Access Architecture
DAC	Digital Addressable Controller
DOCSIS	Data Over Cable System Interface Specification
DS	Downstream
DUT	Device Under Test
EC2	Elastic Compute Cloud
FBC	Full Band Capture
FEC	Forward Error Correction
FDS	Federated Data Service
FDX	Full Duplex DOCSIS
FIFO	First-In, First-Out
Gbps	Gigabit per Second
GraphQL	Graph Query Language
HS	High Split
HSD	High Speed Data
HTTPS	Hypertext Transfer Protocol Secure
HUSL	High-Split Upstream Spectrum Launch
iHAT	in-home Health Assessment Tool



**UNLEASH THE POWER
OF LIMITLESS CONNECTIVITY**
OCTOBER 11-14 ATLANTA, GA



ITG	Interactive Troubleshooting Guide
JSON	Javascript Object Notation
MER	Modulation Error Ratio
MDU	Multiple Dwelling Unit
MHz	Mega Hertz
MoCA	Multimedia over Coax Alliance
MS	Mid Split
MS-CPE	Mid-Split Customer Premise Equipment
MUSL	Mid-Split Upstream Spectrum Launch
NQL	Network Query Language
OFDMA	Orthogonal Frequency Division Multiple Access
OID	Object Identifier
OOB	Out-of-Band
OP2OP	Output Port to Output Port
OUDP	OFDMA Upstream Data Profile
PDMP	POE DOCSIS MoCA Passive
PHT	Performance Health Test
PMA	Profile Management Application
PoC	Proof of Concept
POE	Point of Entry
QAM	Quadrature Amplitude Modulation
PSD	Power Spectral Density
REST	representational State Transfer
RF	Radio Frequency
RFC	Remote Feature Control
RHM	Remote Health Monitoring
RPD	Remote PHY (DOCSIS physical layer) Device
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SIK	Self Install Kit
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratio
SS-CPE	Standard-Split Customer Premise Equipment
SSDA	Standard Split Drop Amplifier
SSH	Secure Shell
SOT	Source of Truth
SOD	Source of Distribution
SQS	Simple Query Service
STB	Set Top Box
TC	Trouble Call
TFTP	Trivial File Transfer Protocol
US	Upstream
vCMTS	Virtual Cable Modem Termination System
VM	Virtual Machine
XOC	eXcellence Operations Centers



UNLEASH THE POWER
OF LIMITLESS CONNECTIVITY
OCTOBER 11-14 ATLANTA, GA



Bibliography & References

- [1] Cable Television Laboratories, Inc., Data Over Cable System Interface Specifications, DOCSIS© 3.0, Physical Layer Specification, <https://specification-search.cablelabs.com/CM-SP-PHYv3.0>, December 7th, 2017
- [2] Cable Television Laboratories, Inc., Data Over Cable System Interface Specifications, DOCSIS© 3.0, Physical Layer Specification, <https://specification-search.cablelabs.com/CM-SP-PHYv3.0>, December 7th, 2017
- [3] Cable Television Laboratories, Inc., Data Over Cable System Interface Specifications, DOCSIS© 4.0, Physical Layer Specification, <https://www.cablelabs.com/specifications/CM-SP-PHYv4.0>, April 29th, 2020
- [4] L. Zhou, “A Proactive Network Management Scheme for Mid-Split Deployment”, SCTE Virtual Expo, 2020
- [5] S. Shulman, “Operating Legacy Cable Modems in an FDX Environment”, 2019 Fall Technical Forum, SCTE-ISBE CableTec Expo, 2019
- [6] R. Vugumudi, “Dynamic Data Collection and Configuration Management”, SCTE Virtual Expo, 2020
- [7] M. Harb, “Lessons Learned from Deploying the Profile Management Application System at Scale and Considerations for Expanding the System Beyond OFDM”, SCTE Virtual Expo, 2020