Operationalizing the Grey Optics Architecture:

An Update One Year After

A Technical Paper prepared for SCTE•ISBE by

Venk Mutalik
Executive Director, HFC Architecture
Comcast
1401 Wynkoop St, Denver, CO
+1 (860)-262-4479
Venk_Mutalik@Comcast.com

Dan Rice, Comcast
VP, HFC Architecture
1401 Wynkoop St, Denver, CO
Daniel_Rice4@Comcast.com

Bob Gaydos, Comcast
Comcast Fellow
1800 Arch Street, Philadelphia, PA
Robert_Gaydos@Comcast.com

Doug Combs, Comcast
Consultant Engineer
1401 Wynkoop St, Denver, CO
Doug.Combs@Comcast.com

Pat Wike, Comcast
Sr. Director, Access Engineering
1401 Wynkoop St, Denver, CO
Patrick_Wike@Comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

High Speed Data (HSD) customer growth and capacity requirements continue to increase year over year. A Distributed Access Architecture (DAA) provides enhanced upstream and downstream capacity to match demand. Last year, we covered a new approach in DAA called Grey Optics Aggregation (GOA) that provided significant savings and enables DAA deployment at scale regardless of plant density.

This year, we report on the lessons learned during our GOA trials and deployments of the architecture. In addition to proving out the technology and cost benefits in the field and observing its stability, deploying the architecture also enabled us to interconnect and create an ecosystem of remote monitoring and control for the various GOA elements. We will report on our ability to continuously monitor Grey Optical Terminating (GOT) nodes, detect cable modems connected to specific GOT nodes and remotely identify and mitigate ingress in the GOA domain.

Many of the remote techniques were used during the COVID pandemic and can potentially be generalized further to enable remote monitoring and mitigating capabilities across the network. Furthermore, we report on recent work on bi-directional Coherent Optical modules that enable efficient extension of the architecture, which, over a period of time, can create cost effective convergent and scalable “Switch On A Pole” (SOAP) access networks.

2. Architecture Recap

Last year, we presented a new Grey Optics Aggregation (GOA) architecture that can reduce cost in the outside plant, reduce cost and space in primary and secondary headends, and provide substantial upgrade opportunities to increase capacity and enable new Switch On A Pole (SOAP) type architectures. The GOA architecture works especially well in medium and lower density areas and can exist side by side with traditional nodes equipped with Remote PHY Devices (RPDs). As such it helps to accelerate the virtualization of the CMTS (vCMTS) by enabling it to be deployed across complete primary and secondary headends.

Figure 1 shows a simplified DAA deployment of vCMTS, DAA Switch (also called DAAS, connects vCMTS data to RPDs) ports and RPDs in the field. Recall that each RPD requires its own DWDM SFP (Dense Wave Division Multiplexer Small Form Factor Pluggable transceiver), multiple such SFPs can be optically multiplexed and connected to individual DWDM SFPs in DAAS ports through optical de-multiplexers. Multiple DAAS ports are aggregated and connected to the vCMTS core via optical methods common to long haul networks (~100G -400G).

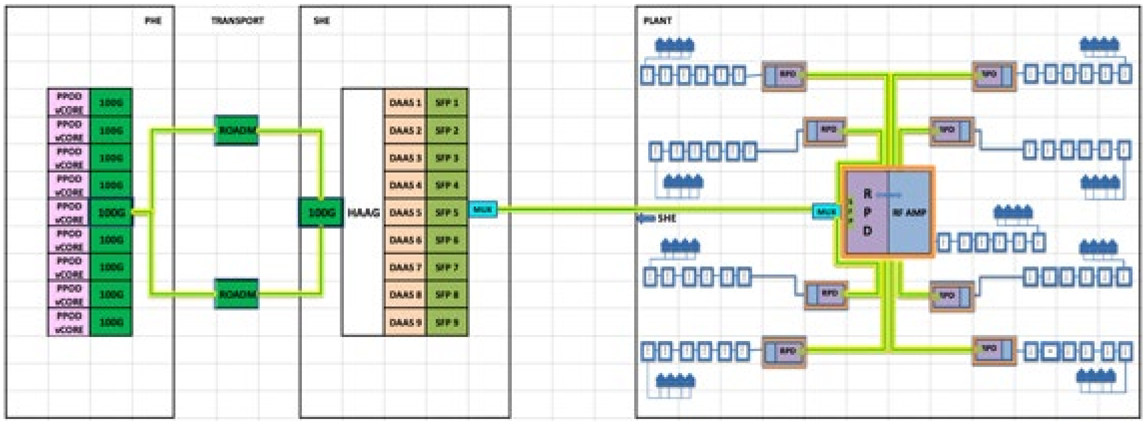

Figure 1 – Traditional DAA in Primary and Secondary Headends and Plant

This architecture works well when home/neighborhood densities are high and as such the cost of each RPD and the associated optics at the headends can be spread across a large base of household passings (HHPs). However, if the area is sparsely populated, this may be difficult. Even when a high-density residential area is selected, it can contain pockets of lower density. Addressing them cost-wise might be the proverbial tail wagging the dog.

Presented in Figure 2 is the GOA Architecture that provides a cost effective and pervasive way to deploy virtualized CMTSs across the whole DAA footprint.

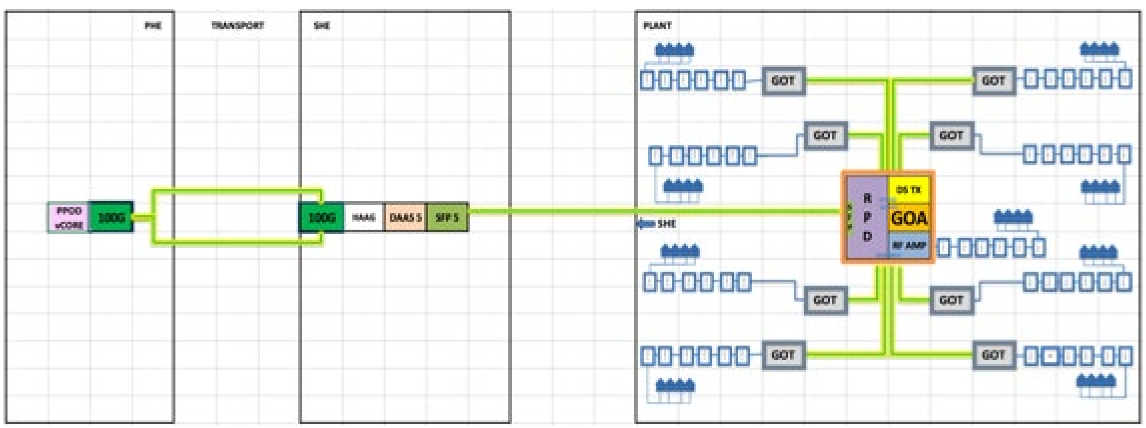

Figure 2 – Illustrating GOA Architecture

In this architecture, a prime node contains the RPD and is called the GOA (Grey Optics Aggregating) node, since it is populated with an analog forward transmitter and return receiver modules for aggregation. The downstream (DS) transmitter output is then split and connected to multiple (up to 8) nodes, called Grey Optics Terminating (GOT) nodes. Each of the GOT nodes contains a downstream receiver module and distributes the RF output to its four constituent RF blocks. The RF levels of the GOT node both for downstream and upstream (US) traffic are identical to regular RPD nodes. For the US, the RF signals get combined in the GOT RF tray and fed into a digital return transmitter (DRT). From there, the signals are sent to the GOA node, where they are fed to the aggregating receiver module, one each for each receiver. Then they are RF-combined and fed to the RPD, thus completing the signal circuit.

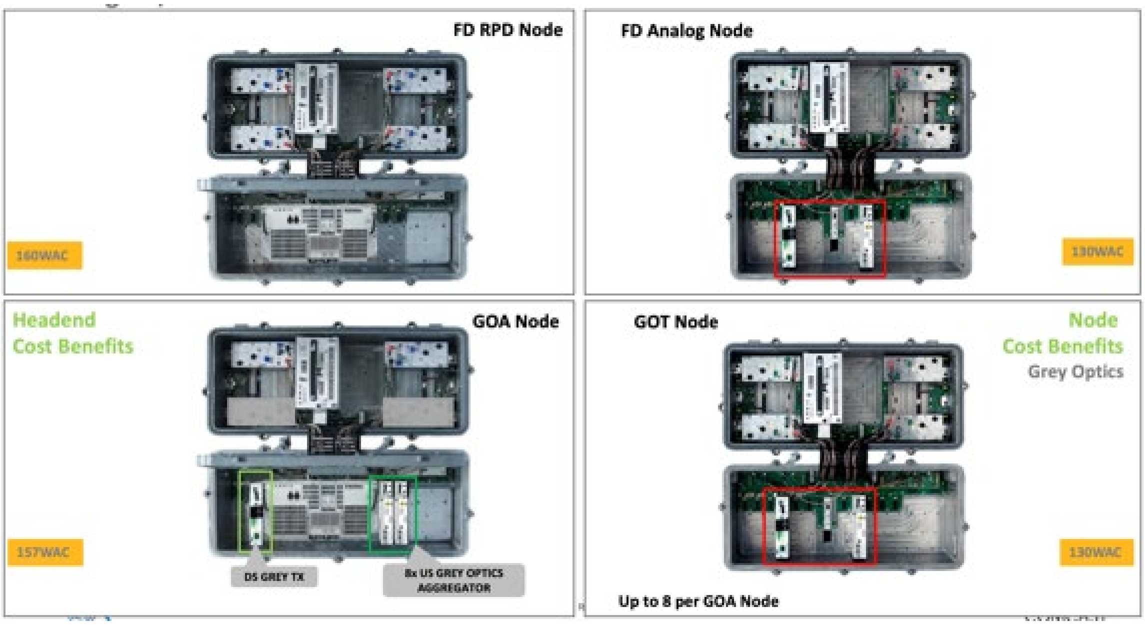

Figure 3 – Illustrating all Nodes and their Power Dissipations

As can be seen in Figure 3, all nodes are the same physical size. And it is the case that all of the above nodes have identical RF US and DS characteristics. The GOT node has an optical AGC receiver that always adjusts the RF level of the receiver to be identical to the RPD output. The DRT to the aggregate return receiver module gain is set to unity, so that the combination of thee DS receiver and the US DRT are an almost perfect replacement for the RPD RF levels. In fact, it is for this reason that replacing the DS receiver, the US DRT and the node monitoring card together with an RPD would instantly transfer the GOT into a full RPD node when and if needed.

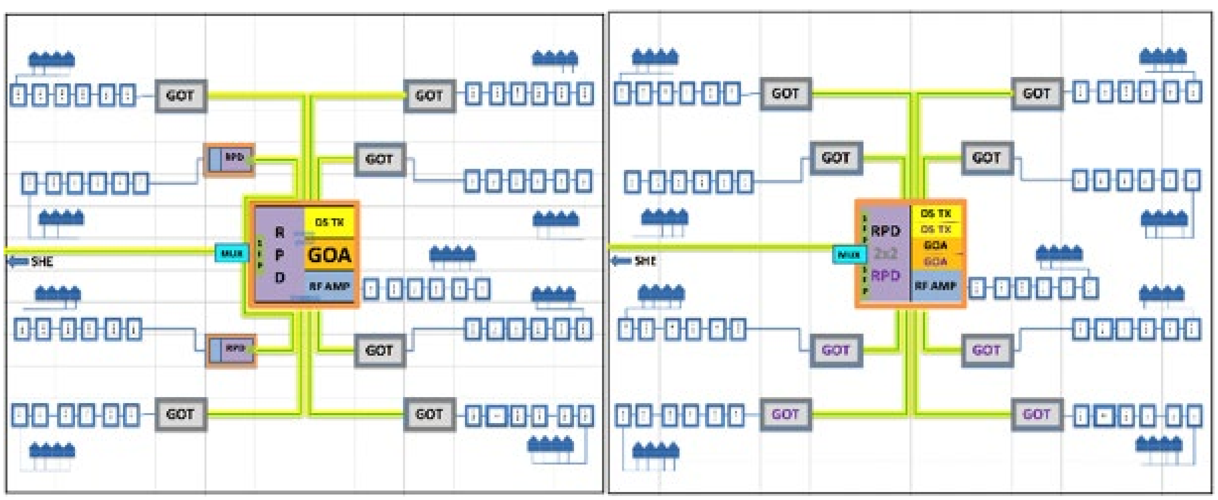

Figure 4 – Upgrade Models: Localized Growth (Left) and Uniform Growth (Right)

Figure 4 illustrates one of the major advantages of the GOA architecture: Not only does the GOA allow for cost effective fiber deployments across multiple areas, but it also allows for easy upgrades, should the need arise. Generally, the DAA deployments are projected to be stable and address growth needs for a substantial period of time. However, if there is localized growth at any one of the GOTs that defied earlier predictions, the move from a GOT node to the GOA node is easy, from a field perspective, because it just replaces the analog receiver and DRT with the RPD and connects the two ends to the optical passive. Furthermore, the new RPDs that would need to be deployed would be more computationally capable, consume lower power and occupy less space (as described by Moore's Law, Dennard's Law and Koomey's Law, respectively) and cost less as well. This will, in all likelihood, enable us to use the uniform growth model (shown in Figure 4) to deploy a 2x2 RPD that provides increased capacity across the board, when the time comes thus requiring far fewer site visits for network upgrades. Furthermore, a move to Full Duplex DOCSIS (FDX) or Switch on a Pole (SOAP) that relies quite heavily on RPD, Node and vCMTS enhancements would also be well served by the GOA architecture.

For all of the above reasons, the GOA architecture was deployed on a trial basis and its performance, robustness, tech friendliness and customer satisfaction have been established. With this information, deployments are now occurring in additional areas.

In this paper, we discuss some of the lessons learned during the deployment. In addition to the above features, several innovative features suggested themselves. One was an intelligent use of the aggregation modules to associate cable modems (CMs) not only to the RPD and to the vCMTS core, but also to each of the GOTs. In addition, the same type of intelligence enables us to identify ingress or noise in each of the GOT nodes or on the RF blocks in the GOA node. These features and other interesting applications of harnessing the monitoring information flow between the GOTs and the GOA will also be discussed.

3. Initial Deployment Site

After discussions, a suburban area in one of the Comcast divisions was chosen for initial deployment. This area already had just been built with traditional RPD nodes. Figure 5 is a description of the chosen area, which served ~550 HHPs with 21 RPD nodes. Some nodes served as few as 15 HHPs while the maximum HHP count was 40, with an average count of 25 HHP / RPD.

	HHP	NODE
240 HHP	19	**GOA A**
	18	GOT A1
	27	GOT A2
	20	GOT A3
	40	GOT A4
	38	GOT A5
	33	GOT A6
	23	GOT A7
	22	GOT A8
202 HHP	24	**GOA B**
	26	GOT B1
	39	GOT B2
	15	GOT B3
	29	GOT B4
	19	GOT B5
	29	GOT B6
	21	GOT B7
109 HHP	24	**GOA C**
	31	GOT C1
	28	GOT C2
	26	GOT C3

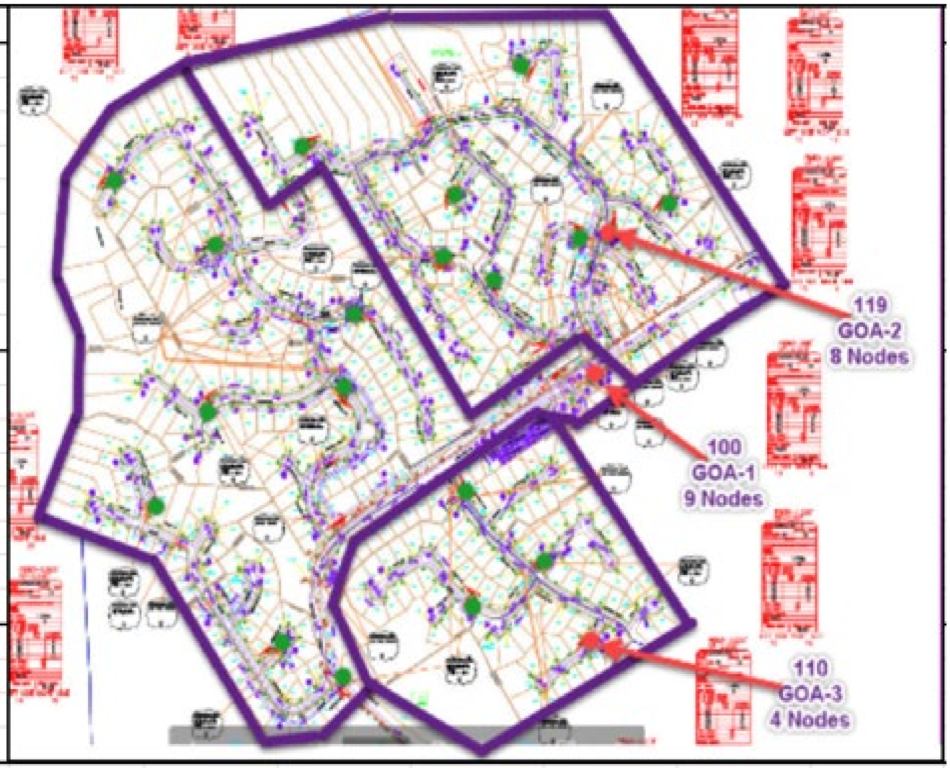

Figure 5 – Technical Details of Initial Deployment Area

As shown in Figure 6, the selected area was a regular subdivision, as would be expected off of a major metropolitan area. What we found was that the common RF build practice of running RF through the front of each property meant that the number of HHPs that could be served per node was rather limited. This is especially true given that DAA builds follow a node plus zero amplifier (N+0) design, thus do not have amplifiers to extend RF runs.

Figure 6 – Geographical area illustrating the Physical Subdivisions

Because of this, it is easy to see that a build that spans across a SHE location may have on average have high HHP/DAA Node density but could still have several nodes with much lower HHP/DAA node densities. We have investigated many ways to enhance the HHP/Node, including a thorough analysis of the drop levels needed. A judicious application of those design rules does extend the number of HHPs/Node and are defined as acceptable in newer playbooks.

In this case, taking a look at the geography enabled us to see that the 21 RPD nodes could be reduced to 3 GOAs and 18 GOT nodes. This represents a 7x reduction in amount of SFPs, DAAS ports and vCMTS microservice instances required. The GOA locations are indicated in Figure 5, along with the GOT nodes and their HHP counts.

4. Initial Deployment

To prep for the transition from RPD to GOA, we set up a team of analytics and traffic experts. We then charted all the CMs associated with each of the nodes and noted the cumulative traffic in each RPD. Then the team threaded the total traffic and time series, aligning them to estimate the cumulative traffic, were the nodes converted to the GOA architecture.

A word about HHPs and CMs. While there is wide variation in respect to broadband penetration rates and number of DOCSIS devices per HHP, in general, when both are taken together, the number of CMs in a given area is approximately equal to the HHPs in that same area (this was certainly true for our case here). With this in mind, we will use HHP and CM count interchangeably in this paper.

We did this to get a baseline view of the traffic before the RPD nodes would be aggregated with the GOA architecture. This laborious process was then automated, and we found that as the aggregated node size increases, the cumulative traffic increases, but at a slower rate than the increases in HHP. And, as expected, we find that the traffic, which is lumpy with small numbers of CMs, gets more blended with larger numbers of CMs. The idea was to ensure that none of the GOAs would be overwhelmed with traffic -- and none were.

We also put together a dashboard that threaded the US RF output levels and input levels, SNR, PER/FEC rates and RF frequency per CM, as well as DS input levels and MER for each frequency, from multiple sources. Some data were directly from the combined vCMTS/RPD and others from the CMs, as they reported back. All of these were time stamped and reported daily as Max, Min, Mean and 95th percentiles. This enabled us to be able to see each CM and all associated metrics simultaneously to see the impacts in almost real time. Over time, this was very useful.

Once the decision to move ahead was taken, the division leaders took over the process of combining nodes in an appropriate manner and drawing up a schedule for node cutovers. The leadership, engineering talent and commitment of the division was unquestionably paramount in the move forward. As can be seen from the figures above,

- GOA A had 240 HHP, with 8 GOT nodes
- GOA B had 202 HHP, with 7 GOT nodes
- GOA C had 109 HHP, with 3 GOT nodes

These were selected in this fashion to ensure that we had sufficient opportunities to verify performance with larger and smaller HHP groups and GOT node aggregation. In parallel, we also performed lab tests to determine that the RPDs could handle the 240 HHPs that were likely to be added. In reality, the RPD and vCMTS combination can handle much more in some of the newer virtualized systems.

Figure 7 – GOA and GOTs assembled in the headend, Closed (Left) - Open (right)

Once this was decided, we drew up entry and exit criteria and a backout plan. Part of the plan was to train our field technicians on the new products, and also to set up a full GOA-GOT system in the headend, to ensure that we had a working system for burn-in, and to work out any kinks. Setting up the GOA-GOT system allowed us to verify some of the very preliminary software features then available. These were road-mapped to improve over time and were deployed and tested on the headend system.

When it came time for construction, in a maintainence window, the first GOA was deployed. After a day of observing its performance, the GOTs were then cut in during daytime. This is because the RPDs have a turn-on time, while the GOT nodes (being analog) are turned on instantly with no additional settle-in time needed. All in all, the 21 RPD nodes were turned in to 3 GOAs and 18 GOTs with careful planning, diligent effort led by the division, and periodic checks on customer call logs and performance verifications, using available internal tools, which will be described in detail later in Section 7.

Overall, the construction of the GOAs and GOT (i.e migration from RPDs to GOA) was uneventful, thanks in large part to the planning and efforts described above. The performance of the units was well within what was predicted, there were no customer impacting events, and the customer experience was maintained well. Next we will characterize the many additional software and hardware innovations that were developed and deployed to enhance the useability of the platform.

5. Heart to Limbs

Construction of the GOA architecture follows a “Heart to Limbs” strategy. By that we mean that the GOA node is first deployed. Its functionality is established, the local splice enclosure is primed and then the GOT nodes are activated, one after another. Since the RPD in the GOA node feeds all other nodes, this strategy will require that the plant that feeds all other GOT nodes stays online even after the GOA is deployed. This is not the case for RPD deployments, since each RPD node is individually feeding its own limited set of HHPs. To achieve good visibility and to speed up the construction process, the team came up with an innovative splicing matrix approach, depicted in Figure 8.

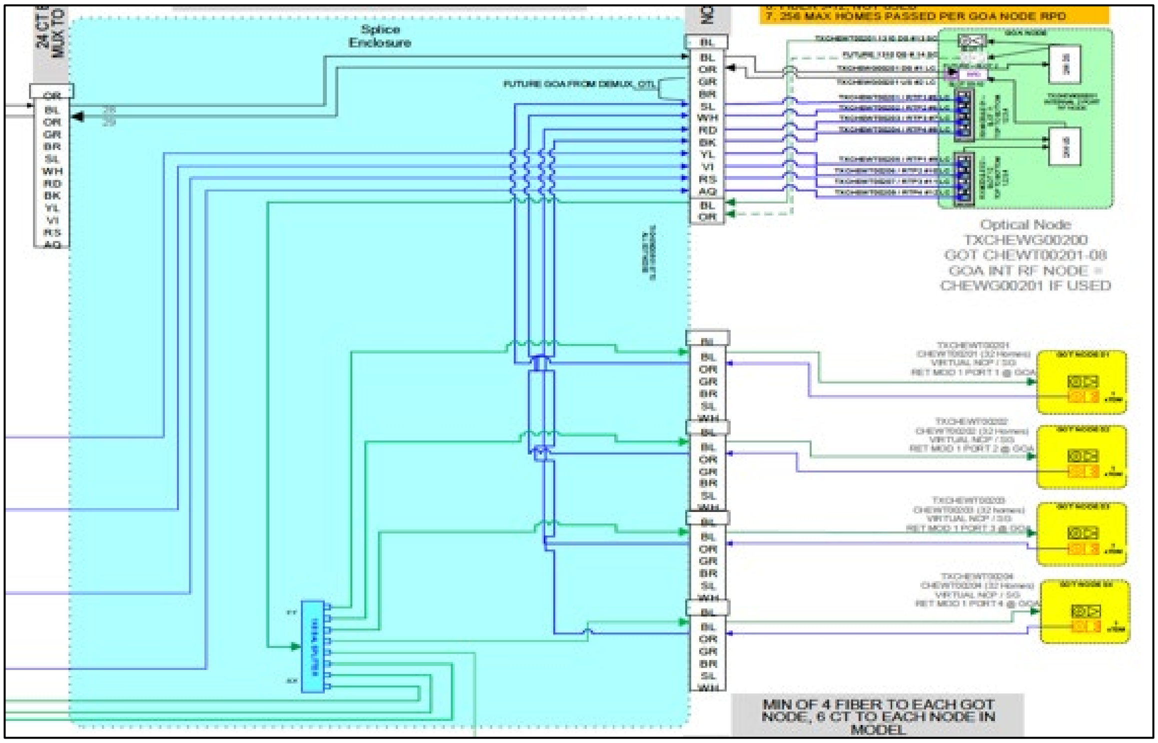

Figure 8 – Splice Matrix and the Wiring Diagram

In the case of the trial deployment, the fiber plant was already designed with RPDs in mind. In cases where GOA is the preferred deployment, we allocate a minimum of 4 fibers per GOT (typically 6 fibers). The GOA is allocated a 14-fiber pigtail. The standard 12 colors, from Black to Aqua and Black and Orange, repeated. For any GOT, the two fibers that are spliced are color coded so that the Black is always connected to the DS receiver and the Orange is always connected to the DRT. On the GOA side, the GOT designated GOTs 1 – 8 follow the Slate – Aqua colors of the standard Fiber Color scheme. For the GOA RPD, the Black is connected to the Rx of the SFP, and the Orange is connected to the Tx of the SFP. The other Black is connected to the DS Tx and goes to the 8-way DS splitter and then on to the GOT's DS Receiver.

We also followed a minimal design for optical passives that connect to the GOA RPD. We generally note that a parent node rarely has more than 3 GOAs. Therefore, the DAA Mux has been standardized to the 12 port Mux (ITU 61 – ITU 50) and exclusively uses sub-band tunable SFPs for the DAAS ports and for the GOA RPDs.

The GOA architecture, in addition to performing well as a cost reduced DAA option, also contributes to the elevation of DAA architectures as a whole. The GOA node has an I2C bus (this is a bus that enables protocols that can monitor or control various elements inside the node) that connects the RPD to the aggregating Receiver modules, and reports the state of DS receiver power, the US transmit power, the US receive power and the DS transmit power (more on this in Section 10.) But it also connects the RPD to the node itself and helps report on the RF amplifier blocks and the overall power and current usage of the node. Furthermore, each of the US aggregating receivers can either attenuate their own RF out by 6dB or turn themselves off. This helps identify ingress sources amongst each of the GOT nodes. This same feature is available in each of the RF blocks, as well, where the RF block can attenuate the US by 6dB or turn it off. This feature is called the Ingress Control Switch (ICS). As an aside, we currently do not use these features in regular RPD nodes. The primary reason was the lack of available remote software control or monitoring. With the software development that took place for the GOA architecture use of I2C/ICS can now be used in regular RPD nodes as well in monitoring and ingress detection on the node ports. With this understanding, most of the software that is described in the paper can apply equally well to GOA architecture or to normal RPD deployments.

6. Deployment Guidelines

In terms of plant makeup, roughly 60% of our plant is aerial and 40% underground. While DAA is a preferred architecture for high growth tier 1 markets, and it helps us stay ahead of impending node splits and also prepare the plant for further DS and especially US capacity, it takes a fair amount of work to lay the fiber cable, optimize the RF cabling, condition the taps and light up the nodes. In recent times, we have made good progress on optimizing the various parameters that go into designing the network, to improve its efficiency and to select areas especially well fitted for DAA. As reality goes, this is easier said than done – which in turn makes the GOA architecture a good fit for enabling the DAA to be more successful in moving to the vCMTS architecture. To this end, a set of rules have been established for the GOA deployments:

- The GOA node has 2 RF ports and operates at standard DAA RF levels (73.6 dBmV TCP for 1.2 GHz load for the DS and 8 dBmV for the 85MHz US)
- The GOT node has 4 RF Blocks and operates at standard DAA RF levels (73.6 dBmV TCP for 1.2 GHz load for the DS and 8 dBmV for the 85 MHz US)
- The GOA the GOT and the RPD nodes use the I2C and the ICS functionality to communicate/monitor and identify/mitigate ingress
- The GOA RPD has a maximum of 256 HHP, this is the sum total of all passings of the GOA and all associated GOT nodes (in rare instances, an additional 30 HHP may be added to the total count if it can avoid having to add a GOA altogether)
- No more than 8 GOT nodes may be subtended by a GOA node, and it is highly preferred that the GOT nodes be contiguous (in other words, a GOA and its constellations should be in proximity rather than be interspersed with GOTs of other GOA nodes)
- A power interruption extension circuit must be deployed with the GOA RPD (this is a recommended for all RPDs, not just a GOA RPD), and no such power pack is required for the GOT nodes (since they are instant-on nodes)
- 160W is designated per node for dissipation, even though GOT nodes dissipate considerably less. This is to be prepared for when GOT nodes need to be upgraded to an RPD node.
- As discussed, the optical passives deployment uses 12 wavelength devices (ITU 61-ITU 50) and use the sub-band tunable SFPs for GOAs and the DAAS ports. The optical splitter will always be an 1x8 device
- Regular DAA nodes and GOA nodes can be used in the same build footprint at convenient locations, as needed, and the DAAS ports and the vCMTS continue functioning with these devices

These guidelines were presented to our design partners, and with their help we have elected to deploy the GOA architecture in two of our divisions this year (this is in addition to the GOA deployment discussed in section 2), covering multiple DAA builds. "Fingers are crossed" to achieve this, given the COVID pandemic and its impact on construction schedules. It is recognized that underground construction is especially challenging in the current environment and further thought is being directed towards this end.

7. Ingress Identification and Mitigation

As discussed earlier, all GOA and GOT nodes have the I2C connectivity to monitor and control many node components. The RPD itself is controlled by Command Line Interface (CLI) commands. There is a secure way to access the RPD to be able to determine how the GOA is configured and display it.

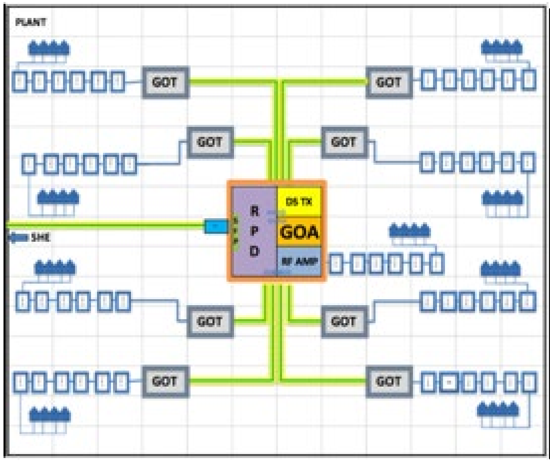


```
*********************************************
*          CLISH (see-lish)                 *
*                                           *
*      WARNING: Authorised Access Only      *
*********************************************

Welcome admin it is Mon Dec 16 12:01:44 -06 2019
# show goat
12/16/2019 12:01:47
--------------------------------------------------------------------------
          DSTX dBm  USRF1 dB  USRF3 dB  USRF4 dB  USRF6 dB  Serial Number
--------------------------------------------------------------------------
GOA           10.1         0        NA         0        NA      519301882

--------------------------------------------------------------------------
          DSRX dBm  USTX dBm  USRX dBm   USRF dB  Node Lid  Serial Number
--------------------------------------------------------------------------
GOT 1         -0.6      -2.6      -3.9         0    Closed      419243508
GOT 2         -0.9      -2.6      -4.1         0    Closed      419243512
GOT 3         -0.2      -2.6      -4.6         0    Closed      419243495
GOT 4           NA        NA        NA        NA        NA             NA
GOT 5           NA        NA        NA        NA        NA             NA
GOT 6           NA        NA        NA        NA        NA             NA
GOT 7           NA        NA        NA        NA        NA             NA
GOT 8           NA        NA        NA        NA        NA             NA

#
```

Figure 9 – Illustrating the Basic CLI and how it summarizes the GOA Constellation

Access to the CLI for a GOA RPD displays the entire constellation of GOT nodes connected to the GOA. All the GOT nodes that are connected to the GOA have a one-way EMS connection to the GOA via the DRT. Therefore, the GOA receiver is aware at all times about the GOTs connected to it. The GOT node sends its state periodically to the GOA, which then updates this information and passes this on to the vCMTS. In addition, this information is also available at the RPD and can be uncovered via the CLI described above. Section 10 will describe how the information that arrives at the vCMTS can be used, but for now we describe two important aspects of what the GOA can provide which could be of great value.

The entire GOA constellation is represented in Figure 9. The output power of the DS transmitter (DSTX) is displayed and it is nominally 10 dBm. Next, for each of the GOT nodes, the DS receiver power (DSRX) is displayed, which is nominally 0 dBm to -1 dBm, because the DS light passes through an 8-way splitter and traverses some fiber. The US Transmitter (USTX) power is shown next, which is nominally -3dBm (this is a lower-cost Grey 1310 nm SFP). The US Receiver (USRX) power reported by the GOA is shown next, and is nominally -4 dBm, because the US light does not go through any splitter or combiner.

In addition to the light levels, we mentioned earlier that there is an ICS; each of the states of each ICS are reported as well. Thus, the USRF ports refer to the attenuation setting of the ICS of each of the RF blocks in the GOA (recall that we have 2 RF blocks per GOA). Recall that the GOA US Digital Receiver has the capability to attenuate the RF output of each individual GOT US as well. That state is also represented in Figure 9. Furthermore, each GOA and GOT serial number is reported as well. As a neat feature, the GOT node lid is reported as open or closed.

With these settings established, we can now remotely monitor and control the whole constellation. We will show in Section 8 how we have automated this, and how a ticketing structure was developed to automatically indicate major and critical alarms and identify escalation and mitigation options.

```
# show goat
11/20/2019 16:24:47
-------------------------------------------------------------------------------
          DSTX dBm  USRF1 dB  USRF3 dB  USRF4 dB  USRF6 dB  Serial Number
-------------------------------------------------------------------------------
GOA           10.0         6        NA         0        NA      519301888

-------------------------------------------------------------------------------
          DSRX dBm  USTX dBm  USRX dBm   USRF dB  Node Lid  Serial Number
-------------------------------------------------------------------------------
GOT 1           NA        NA        NA        NA        NA             NA
GOT 2          0.2      -2.7      -9.0         6    Closed      419243545
GOT 3           NA        NA        NA        NA        NA             NA
GOT 4         -0.2      -2.6      -9.1         0    Closed      419243566
GOT 5           NA        NA        NA        NA        NA             NA
GOT 6          0.2      -2.6      -8.9         0    Closed      419243513
GOT 7           NA        NA        NA        NA        NA             NA
GOT 8         -0.7      -2.6     -10.5         6    Closed      419243477
```

```
# show goat
11/20/2019 19:21:35
-------------------------------------------------------------------------------
          DSTX dBm  USRF1 dB  USRF3 dB  USRF4 dB  USRF6 dB  Serial Number
-------------------------------------------------------------------------------
GOA           10.0         0        NA         0        NA      519301888

-------------------------------------------------------------------------------
          DSRX dBm  USTX dBm  USRX dBm   USRF dB  Node Lid  Serial Number
-------------------------------------------------------------------------------
GOT 1           NA        NA        NA        NA        NA             NA
GOT 2          0.2      -2.6      -9.0         0    Closed      419243545
GOT 3           NA        NA        NA        NA        NA             NA
GOT 4         -0.3      -2.6      -9.0         6    Closed      419243566
GOT 5           NA        NA        NA        NA        NA             NA
GOT 6          0.2      -2.6      -8.6         0    Closed      419243513
GOT 7           NA        NA        NA        NA        NA             NA
GOT 8         -0.5      -2.6     -10.4         0    Closed      419243477
```

Figure 10 – CLI with Various Settings

Figure 10 shows how these values can differ. Notably, the USRX values are nominally at -9 dBm because each of the US links have a ~6dB optical pad. However, since the return path is digital, the optical attenuation does not affect the RF levels at all. The only thing that affects the RF level is RF attenuation, which is shown here in Red. On the top, GOTs 2 and 8 are attenuated as is the GOA RF Block 1. In the bottom picture, only GOT4 is attenuated.

When GOT 4 output level is attenuated in the GOA, two things happen:

1. The noise and the signal level of the GOT4 node is attenuated by 6 dB. Therefore, if GOT4 was the ingress culprit, we would have ingress relief of 6 dB
    a. Since all the RF was combined together, any reduction in ingress will beneficially improve the US SNR of the entire system
    b. If the aim is to identify the ingress point, that is now accomplished. If the aim was to mitigate ingress, then the attenuation stays on until technicians can trouble shoot and eliminate the ingress
2. In addition, the long loop AGC administered by the combined vCMTS/RPD kick in and the RF output levels of all CMs connected to GOT 4 are increased by 6 dB
    a. When this happens, the CMs connected to the GOT 4 benefit and their SNR improves as well
    b. Some CMs that are on the high edge of USTX could reset. This is dictated by the duration of the attenuation and how close they are to the edge

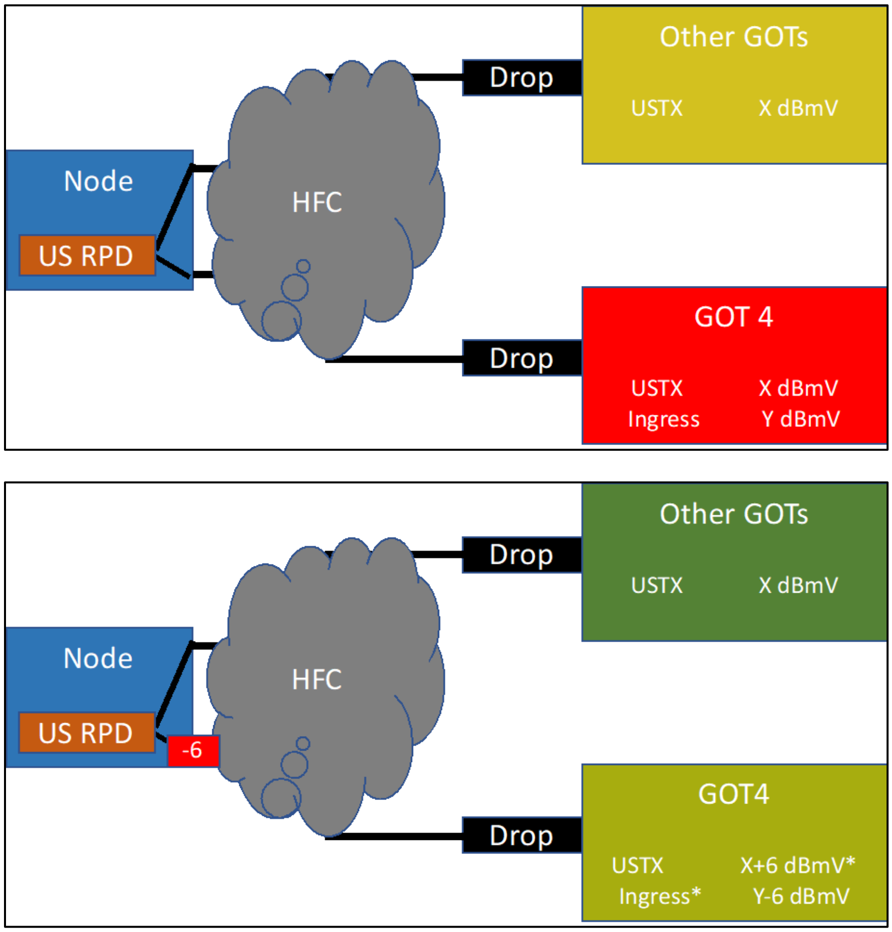

Figure 11 – Highly Simplified Ingress Illustration and the Role of US RF Attenuation

The above points are illustrated in Figure 11. Roughly 90% of upstream ingress accumulates at the home, as characterized in an excellent paper written by Larry Wolcott [5] who explains that minimal additions to it occur in the RF plant -- but the plant does funnel the ingress, which compromises the entire plant's SNR. At the extreme it is possible that an aggressive ingress source from a single household can impact the SNR of the entire system. Therefore, it is certainly possible that a single GOT can impact the SNR of the entire system. According to experts whose job it is to keep the upstream signal path functional, around 80% of trouble reports and technician time is spent on ingress identification and mitigation. Since ingress is generated at the home, it can occur in DAA plant and in GOA plant just as it can occur in a regular HFC plant. The ability to peer into the network remotely and identify the ingress-impacted GOT nodes or RF ports, and reduce its impact until it is all improved, is a potential game changer for the future of the RF plant. This capability, although developed for GOA, can be extended to any RPD node, whether it is deployed for DAA or HFC applications.

8. Tooling, Eventing, Layering … All in a Sprint

Using a Command Line Interface (CLI) for quick checks is one thing, but using it to continuously monitor and control a large number of GOA RPDs is not a good idea. This is because repeated logging into the RPD (using a scheduling utility such as a Cron job) has a way of destabilizing the network. In addition, RPD logging is taken very seriously in Comcast and the security steps make the process much more difficult.

Fortunately, we have a set of software that uses TLVs (Type Length Values that are used in data communication protocols) that specify information of interest and "broadcast" these to those interested. Harnessing the TLVs by appropriately defining them and connecting them to the various pieces of information will then enable us to "listen in" on them, and create dashboards that enable us to look at various network elements in real time.

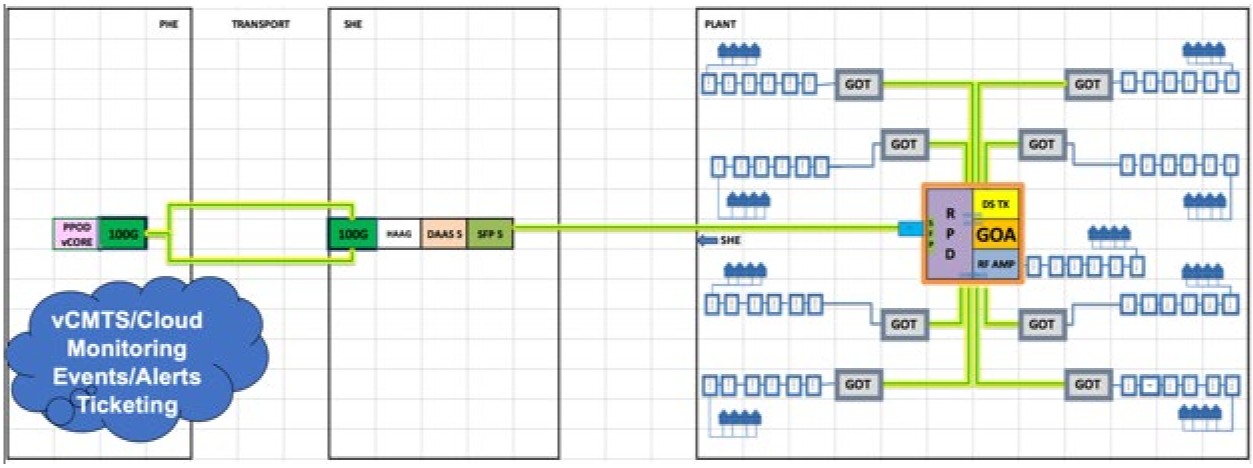

Figure 12 – The Tooling Diagram – Cloud Based Provisioning/Monitoring/Tooling

Within Comcast, a project or a product is never considered finished until the "tooling, eventing and layering" tasks are complete. This is a special lingo that takes a bit of time to absorb, particularly for a hardware optical engineer such as this author. But after a bit of learning, we focused on defining the TLVs for the various elements of the CLI table defined in Figure 12. Thus, all of this information would be continuously available (every 15s in the limit) and the entire constellation of the GOA, including all the GOTs and their US and DS levels and RF attenuation state.

This information is available to Grafana, an analytics and visualization tool, for read-only graphing and history purposes for anyone in the organization. This is especially useful to the engineers, and also to the technicians, division engineers/leaders and also to operations centers as necessary.

We also use a piece of software called the RPD Life Cycle Manager (RLCM). This is secure software that enables one not only to view all aspects of the RPD, but also to control aspects of it without additional security. Thus, the ability to manipulate the RF attenuation values are listed in the RLCM.

By now, it is clear that there is ample information available within the company to understand all aspects of our network. We continue to make strides in making all of this data actionable, with increasing application of machine learning and artificial intelligence algorithms. In fact, several of these aspects were used during the COVID pandemic to not only improve reliability but also increase capacity.

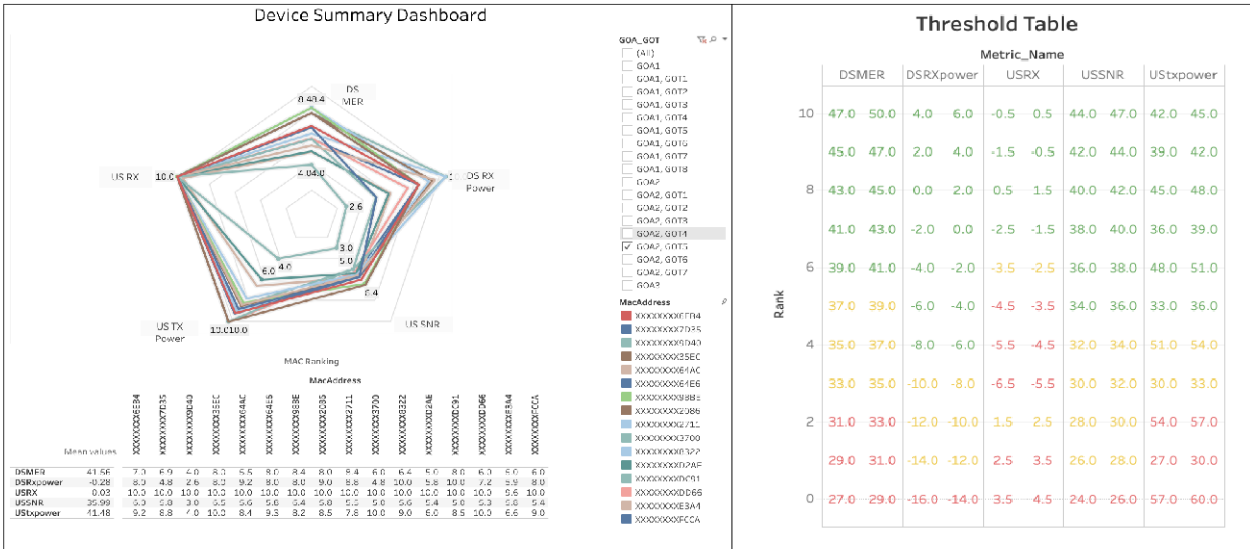

Figure 13 – Visualizing All US and DS Metric at a Glance Across the Node

For measuring performance of the network, however, we have so many metrics that it becomes difficult to look at all simultaneously, yet, looking at them is unsatisfying and probably dangerous as the limited information could lead us to incorrect remediation.

Consider a spider chart of the kind depicted in Figure 13. Here, 5 of the most important parameters have been coded to the table alongside and plotted. Since we have data on each modem in real time, these data are available for view on a modem by modem basis and available for display. It is easy to see here that for this node, the US RF RX is always close to ideal (0 dBmV in real life and 10 on the spider chart), but one of the CMs has a rather poor time of getting all other parameters out (DSRF MER is between 35-37 dB, DS RFRX is -8 to -10 dBmV, USRF SNR is 30-32 dB and the US RFTX is 51-54 dBmV). Meanwhile, all the other CMs are in much better shape. This allows us to know, at a glance, that this modem has a low DS RFRX and a high US RFTX value (probably due to a high tap value in an in-home network) which is why the DSRF SNR is low, so, improving it might serve all the parameters well. In general, the wider the "Spider Eye" is, the better the system is, and the eye opening may be used to verify overall performance in a non-invasive way as an added verification of the Orchestration process.

Together the RLCM and the RPD feed events to the internal National Watch Tower (NWT) engine. NWT is a powerful visualization and event reporting and ticketing tool that is one of our internal operational staples. For the architecture to function smoothly and be a template for the future crop of DAA RPDs, all these features must work seamlessly, and are next described in detail.

9. Provisioning the GOA-GOT system

We have now entered the Software World, where verbs function as nouns and all specific actions are gerund-sounding … starting with the provisioning process.

Provisioning an RPD is usually done by scanning the QR code of the RPD and matching the IP address, MAC address and the node name together as a first step. The case is similar with a GOA RPD. However, for the GOT node, there is no MAC address, nor is there an IP address. We solved this lack of information by using a piece of software we developed called the RPHY Scanner App.

The RPHY Scanner App provisions an RPD and recognizes when the RPD is a GOA RPD. This can happen because the RPD is connected to a Receiver aggregating module and has a DS Transmitter in the node. Therefore, the App automatically opens up a GOT tab. Recall that we always go from the 'Heart to Limbs' for the GOT deployment. Since the GOA RPD is already live, when the fibers (of specified colors as described before) are connected, and the DS light enters the GOT, the US light leaves the GOT and enters the GOA node, the node is automatically given a name, which is a single digital extension of the GOA node name. For example, if the GOA were called NGACD00110, then the GOT connected to GOA receiver 1 would just be NGACD001101, and the one connected to receiver 2 would be NGACD001102, and so on.

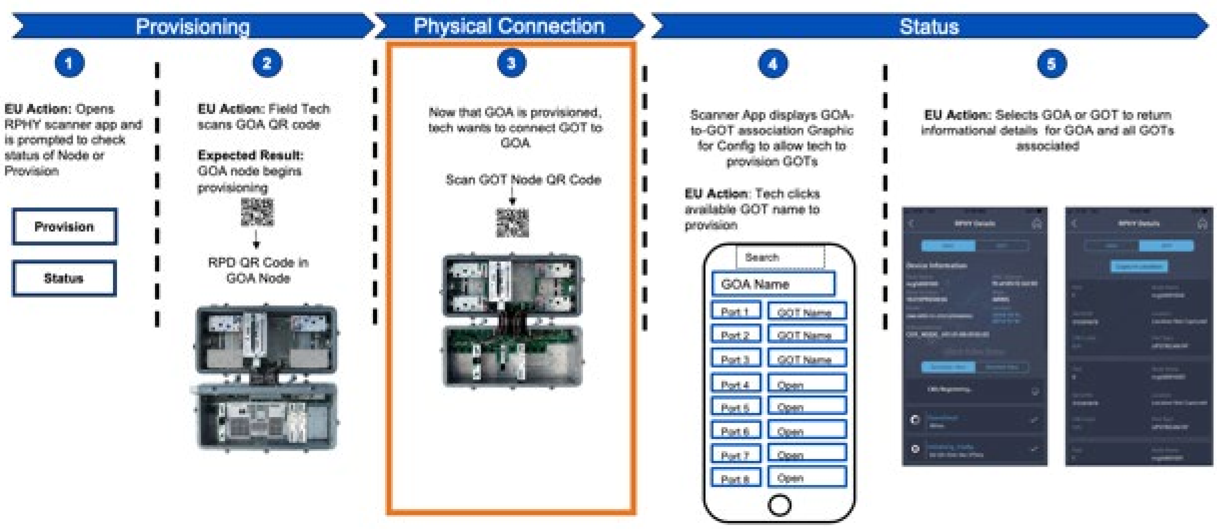

Figure 14 – Provisioning Application Flow

With a name available before the GOT is provisioned, the technician is prompted to take a picture of the GOT QR code. When that is done, the serial number of the QR code is matched against the serial number that is sent via the TLV. Assuming they both match, the Latitude/Longitude of the GOT position is stored and the node is declared provisioned. Thus, each GOT node is provisioned, its Lat/Long position set, its serial numbers verified, and it can start as a fully functional node in the GOA constellation.

10. "Grafana-ing" the Constellation

Grafana is a formidable tool that enables one to graph the performance of any device over time. As discussed, we get a range of US and DS RX and TX information along with data on the RF attenuations, in real time, every 15s or so. It should therefore be possible to slice that information and present it in a graphical form. Such a task is achieved in the 'GOA Constellation' part of Grafana as a read-only feature. This is especially helpful to see trends, identify intermittent issues, or perform historical comparisons (such as day to day or week to week).

Figure 15 – Grafana and Real Time and Historical Constellation Information

In Figure 15, the following can be seen:

1. Each GOA can be individually selected for viewing based on the PHE location or the node name The timeline can be varied as well.

2. The top portion of the dashboard shows the DSTX level in real time as a graph. The value is nominally 10 dBm, as we have seen before. Each of the DSRXs at each GOT are shown as well, and we know that the nominal values are -0.5 dBm. If, for any reason, the DS fiber were cut/disconnected, this would show as a disconnect and point to the issue.

3. The middle of the dashboard shows the USTX and the USRX data for each GOT. Here it is seen that the nominal value of the DSTX is ~-2.5 dBm and the nominal value of the USRX is ~-4 dBm, as can be seen in the drop-down list.

4. In the bottom part of the dashboard, one can see the RF attenuation experienced by the various ports. Recall that the GOA has 2 RF Ports, and each of the GOTs is treated as an RF port for purposes of attenuation, to identify and mitigate ingress. Here one can see the 8 GOTs undergoing RF attenuation, one after the other. Even though the attenuation is applied for a uniform time, the graph is limited by the resolution of the tool and the display and the time it takes for all the information to be collected.

This tool has been of utmost importance, as we have used it extensively to study the robustness of our constellation, especially as COVID raged across our country, with its unquestionable impacts on available capacity. Understanding these issues was impossible with direct observation. The vast amount of historical data is a perfect complement to the CM data and traffic data that is generally available and stored in other portions of the Grafana tool. This portion of tooling is vastly complex and not discussed further.

11. Orchestration and Association

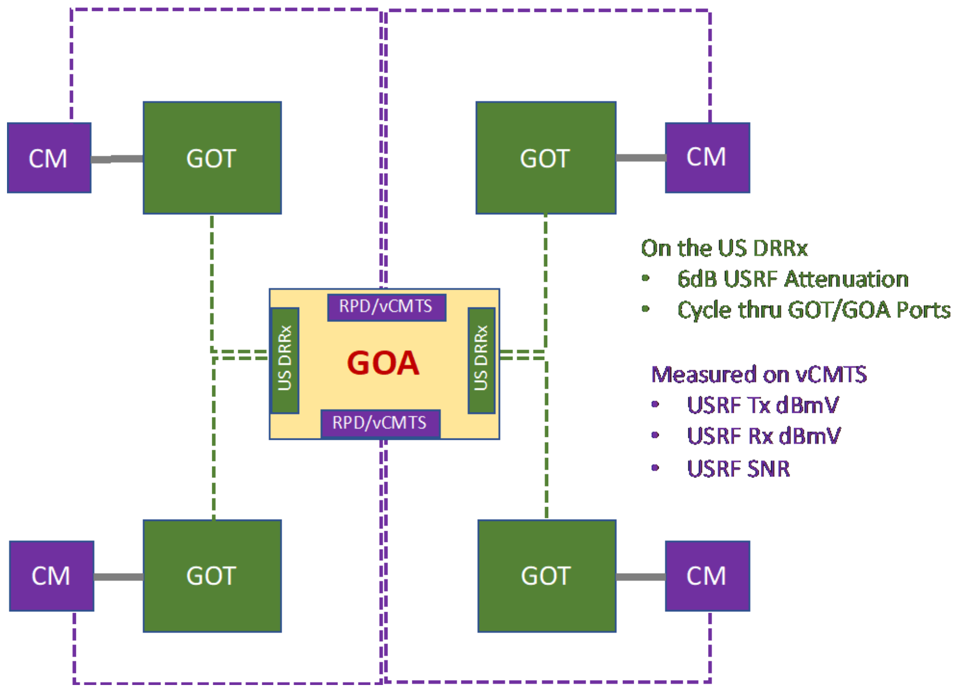

Figure 16 – Illustrating the concept of Association and Orchestration

It is common for the CMTS to know all the CMs connected to it, as well as the modem RF transmit power, the modem RF receive power, and also the RF receive power at the CMTS input. In this case, the vCMTS is expected to keep track of these parameters in concert with the RPD (since features of the CMTS have been distributed between these two by principles of virtualization described in the beginning). We have further seen that the GOA node knows about the existence and state of all the connected GOT nodes, and is aware of the optical DS input power, the optical US transmit power and the US received power. But since the US is a DRT/DRRx connection, changes in optical power do not affect RF power received power as long as the link is functional. Of course, we would not like the optical power to change because this will cause events. We will discuss how these are flagged and ticketed in Section 12.

In Figure 16, the situation described above is represented by the purple dashed lines connecting the RPD/vCMTS to the CMs, to signify that the US RF TX and US RF RX power, as well as the US RF SNR is known. In addition, the DS RF SNR and the DS RF MER is known. The green dashed line represents the connection between the GOA and GOTs via the constellation program described above. What we do not have is the knowledge of how the CMs are connected to the individual GOTs. Nor do we know if ingress exists, and if so, where it might be originating. Note here that even though the GOA itself is also connected to the plant, we have not depicted it, as a way of simplification.

The ability to map the individual CMs to the individual GOTs is called Association. To do this, we orchestrate 6 dB RF attenuation one after another on the various GOTs in the DRRx in the GOA node. This way, the RF levels of the CMs that are connected to the specific GOTs are raised by 6 dB, due to the long loop AGC of the CMTS in response to the attenuation. By keeping track of the US RFTX, US RFRX and the US RFSNR, we can pinpoint the CMs associated with individual GOTs, and also, by tracking the SNR improvements, we can pinpoint the magnitude and location of any ingress sources.

By way of explanation, we keep track of the baseline of all US RFTX, US RFRX and USRFSNR for each of the CMs connected to the GOA RPD. We then put in 6 dB attenuation on one of the GOAs, which will cause the RF levels of all CMs to rise by 6dB. Furthermore, we track if the SNR of the system has improved, and the extent of any improvement. All CMs that reacted by increasing their RF levels are associated with the GOT node. If the SNR improved, then said GOT is identified as an ingressor. We then rate this on a scale of 1-6, with 6 dB being the most egregious ingressor.

With proper verification, we have been able to secure 100% accuracy of the CM associations, in real time, remotely. Because DAA deployments tend to have RF transmit variables well within limits, we have seen very few (~1%) of CMs reset with changes in RF levels, and they come back online within the duration of the test, once the attenuation has moved away. We have also been able to track ingress events remotely, which was immensely helpful in the COVID environment. This is because technicians can approach the exact node and troubleshoot for ingress in the rare events that presented themselves during the deployment. This reduces their time in the field and reduces any associated virus exposure in customer homes. How exactly this is implemented securely, remotely and automatically is described next.

12. Secure, Remote and Automatic: RPD Life Cycle Manager

Enabling DAA is a complex operation. It not only required the vCMTS scheduling effort, but a similarly complex back office system to monitor and control various parameters that weren't normally visible or useable in real time. We described the provisioning application, as well as the traffic and connectivity managing application; next we described the process of using a complicated real time learning app that helps the RPD monitor and control the various parameters of devices connected to the RPD. We have previously indicated that the use of the I2C bus in the GOA node allowed us to be able to surface all aspects of the node to the RPD monitoring system. This system already exists with RPDs and has been updated to include the GOA-GOT system.

The RPD is controlled remotely via the RPD Life Cycle Manager Application (RLCM). This application identifies the RPD as a GOA. This then opens up a table with the various GOTs and their serial numbers and their latitude and Longitude information (recall that this is secured when the GOTs are provisioned). The App further maintains a list of all current CMs, and their state, as is expected of a CMTS. This is based on direct observation by the RPD/vCMTS and not based on inference of any sort.

The App has a soft button for association, which, when pressed orchestrates the 6dB attenuation periodically, using machine/logic in the cloud, and establishes the CM association and the ingress rating. A button allows one to download the data for further analysis. Currently this is an individual operation. The thought is to expand it for periodic data harvesting, and to use machine learning to understand the nature, frequency and magnitude of ingress. This will provide a lasting improvement to any cable system.

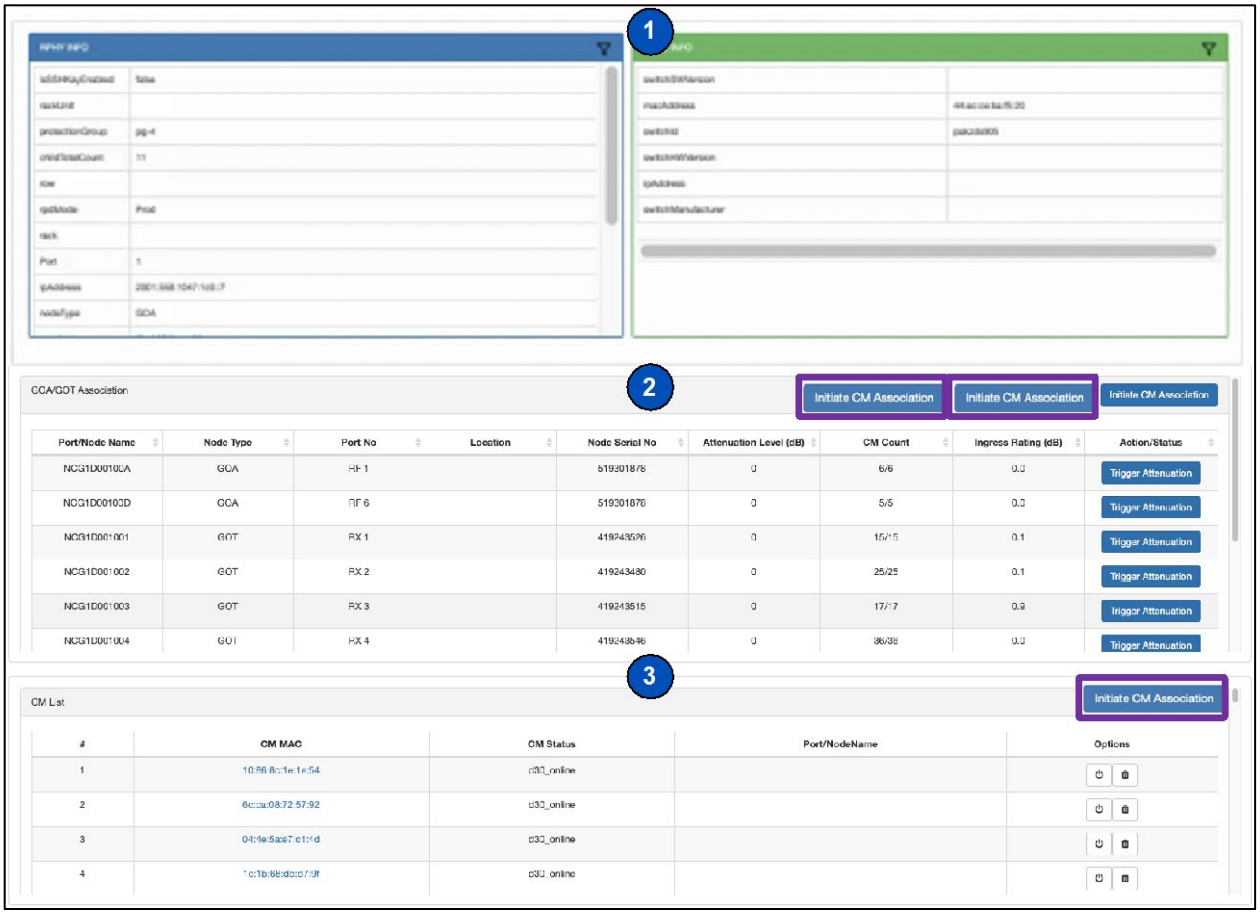

Figure 17 – Association and Orchestration in the GOA RPD

13. Eventing and Ticketing

The optical connectivity between the GOA and the GOTs is normally quite reliable, given that the fiber connections are typically less than a km. However, due to natural disasters and other events it is possible that the connectivity would get disturbed. If the DS and/or US fiber is cut, or if, for some reason, the GOT should lose connectivity to the GOA, it is imperative to identify the impacted GOA and GOT with a ticket to flag them for remedial actions.

It is, however, the case that information reaches the decision engine to call an event in multiple forms, all with delays peculiar to their content. If the fiber between the GOA and the GOT gets cut, the CMs go offline and this information reaches the decision engine via the vCMTS path. The fiber cut will also result in the GOA US RRX missing light on a previously lit receiver, which will then propagate this information to the decision engine. There may be outages attributable to power which could take down multiple nodes; this information reaches the decision engine via the powering network. Since outage information from various sources is arriving with various minor delays, if the decision engine begins reacting immediately, it will have reacted with insufficient information. Thus, to avoid this "race" condition, the decision engine "soaks" or stores all incoming information with various minor delays to look at it holistically.

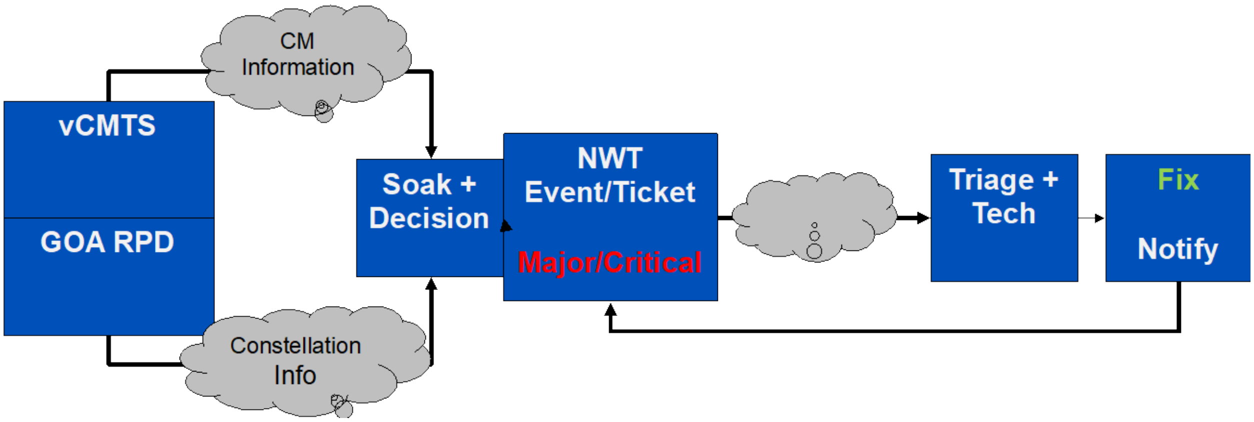

Figure 18 – Automatic Eventing and Ticketing

Having complete information from all sources, by inference and direct observation, the engine declares an informational/major /critical event, and an application such as National Watch Tower generates a ticket and sends it to the concerned division/headend/operations center. At the center, it is triaged and allocated to the technicians. All GOA-GOT outages are identified by a GOA and GOT node name and also the Lat/Long locations of both devices. This is to give as complete information as possible, to not only identify the devices, but also point out the likely fiber route that may have been taken. In the NWT, the GOAs and GOTs are layered, so that it becomes easy to match the outage to the map area.

Once the technician has fixed the problem, whether fiber cut or node maintenance or ingress event, the event is cleared by the technicians themselves. This way the node is brought back to an event-free state and normal operation resumes.

14. Future Steps

This year, we reported on the lessons learned during Comcast GOA trials and deployments. In addition to proving out the technology and cost benefits in the field and observing its stability, deploying the architecture also enabled us to interconnect and create an eco-system of remote monitoring and control of the various GOA elements. We will continue to report on our ability to continuously monitor Grey Optical Terminating (GOT) nodes, detect cable modems connected to specific GOT nodes and remotely identify and mitigate ingress in the GOA domain.

Many of the remote techniques were used during the COVID pandemic and can potentially be generalized further to enable remote monitoring and mitigating capabilities across the network. Furthermore, in the future we will report on recent work on bi-directional Coherent Optical modules that enable efficient extension of the architecture that, over a period of time, can create cost effective convergent and scalable ‘Switch on a Pole’ access networks.

15. Acknowledgements

It is our pleasure to acknowledge the entire team within Comcast who has been working directly and indirectly on the GOA project. A project of this magnitude benefits from the dedication of diverse expertise, from the hardware and software of the RPD team, the reliability and functionality testing of the Physical and Environmental team, the ticketing and alerting of the NWT team and the functional rules for deployment from the access engineering team for both ISP and OSP deployments. Special thanks to our vendor partner CommScope for your insights. We sincerely thank the Senior Leadership Team at Comcast NGAN in supporting this project.

Abbreviations

AGC	Automatic gain control
CLI	Command Line Interface
CM	Cable modem
CMTS	Cable modem termination system
DAA	Distributed access architecture
DAAS	Distributed access architecture switch
DRT	Digital Return Transmitter
DS	Downstream
DSRX	Downstream receiver
DSTX	Downstream transmitter
EMS	Element Management System
FEC	Forward error correction
GOA	Grey optics aggregation
GOT	Grey optics termination
HHP	Households passed
HSD	High speed data
ICS	Ingress control switch
MER	Modulation Error Ratio
NGAN	Next Generation Access Network
PER/FEC	Packet Error Rate/Forward Error Correction
RF	Radio frequency
RFSNR	RF Signal to Noise Ratio
RLCM	RPD Life Cycle Manager
RPD	Remote PHY device
SOAP	Switch on a pole
SFP	Small form-factor pluggable
SNR	Signal to noise ratio
TLV	Type length value
US	Upstream
UPRX	Upstream receiver
UPTX	Upstream transmitter
vCMTS	Virtual cable modem termination system

Bibliography & References

1. *Fifty Shades of Grey Optics: A Roadmap for Next Generation Access Networks*, Venk Mutalik, Bob Gaydos, Dan Rice, Doug Combs, SCTE Expo 2019

2. *Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access*, Nagesh Nandiraju et. al., SCTE EXPO 2016

3. *When Wavelengths Collide, Chaos Ensues: Engineering Stable and Robust Full Spectrum Multi-wavelength HFC Networks,* Venk Mutalik et. al., SCTE Cable-TEC EXPO 2011

4. *Cable's Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture,* Venk Mutalik et. al., SCTE EXPO 2016

5. *A Comprehensive Case Study of Proactive Network Maintenance,* Larry Wolcott, John Heslip, Bryan Thomas and Robert Gonsalves, SCTE EXPO 2016