Accelerating the Virtualization:

Introducing Hybrid Fiber Shelf into the Mix

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

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

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

Jorge Salinger, Comcast
1401 Wynkoop St, Denver, CO
Jorge_Salinger@Comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Abstract

Broadband access networks continue to experience customer growth and higher capacity demands year over year. Early and ongoing analysis of these demand trends enabled Comcast to develop a Distributed Access Architecture (DAA) that enabled us to keep up with capacity needs and enhance customer satisfaction with appreciable capital economics.

Successfully virtualizing complex portions of our network, including the CMTS, resulted in network simplification and the harmonization of multiple purpose-built platforms into one common entity. However, DAA deployments require headend and field modifications as part of the deployment. Even before the 2020 spike in capacity demands created by the COVID pandemic, and its consequent work from home requirements, efforts were underway to accelerate the virtualization, and associated benefits, of our networks, and to separate headend-centric innovations from those in field construction. These parallel efforts enabled us to press forward more rapidly in the adoption of virtualized CMTS technology.

A new concept was developed called a Hybrid Fiber Shelf (HFS) that integrates with virtualized CMTS/DAA Switches at one end, and with transmit/receive analog optical signals on the other end into the outside plant. This innovation provides rapid and economically sustainable increases in capacity, by independently accelerating vCMTS integrations while DAA construction proceeds in other areas of the network. HFS improves critical infrastructure in headends by saving wasted space, power and time inherent in RF combining and splitting circuits. Extending HFS into secondary headends (SHEs) also enables significant fiber reclamation. Locating these assets in secondary facilities closer to the traditional HFC nodes improves performance in ways that translate into enhanced capacity with upstream profile management tools also presented at this SCTE.

In this paper, we begin with a description of the optical and RF innovations that enabled a hybrid fiber shelf concept, and its theory of operations. We then describe the end performance and the impressive improvements in headend critical infrastructure. We finally describe the commonalities and nuances of this approach and its fit within the overall DAA architecture.

2. Introduction to Virtualization

A typical Cable System uses Cable Modem Termination Systems (CMTSs) to provide high availability upstream (US) and downstream (DS) signals between IP-connected devices (cable modems, gateways, set-tops) and IP-delivered services in digital Quadrature Amplitude Modulation (QAM) format. These QAM signals are RF-combined with other video sources, and then provided as inputs to high performance analog transmitters. Multiples of such analog transmitters' output wavelengths are then multiplexed into one fiber and sent to multiple Secondary headends. At the secondary headend, these signals are either demultiplexed and retransmitted, or amplified and distributed to individual nodes. After the node, multiple RF amplifiers in cascades amplify this signal to reach homes. In the reverse path, signals from the home arrive back at the RF amplifiers, which are bidirectional, and then reach the SHE from where they are often aggregated and sent back to the PHE to be terminated in the CMTS to complete the signal circuit.

The traditional CMTS is a formidable equipment type, with purpose-built hardware and significant scheduling and modulation software that has soaked over decades. All high-speed data (HSD) traffic flows through and terminates at the CMTS, making it a gatekeeper of broadband traffic for the MSOs.

Each time additional data was needed in the plant, typically in chunks of 50 MHz, the CMTS capacity would have to be augmented by combining additional ports, which creates capital costs.

This type of traffic infrastructure was fine while the HSD was a sliver of the total cumulative spectra delivered to the home, which was dominated by MPEG video, pay-per-view and VOD and fit the overall criteria of "pay as you grow" dictum of the MSOs. Readers may remember that long span of time when two 6 MHz channels served most of the voice and Internet services for most MSOs. However, when HSD usage started to grow exponentially, with CAGRs of ~35% year over year, and the CMTS ports started piling up faster and faster, the model of a traditional CMTS started becoming unwieldy. Something had to change, when HSD traffic dominated the US and DS spectra.

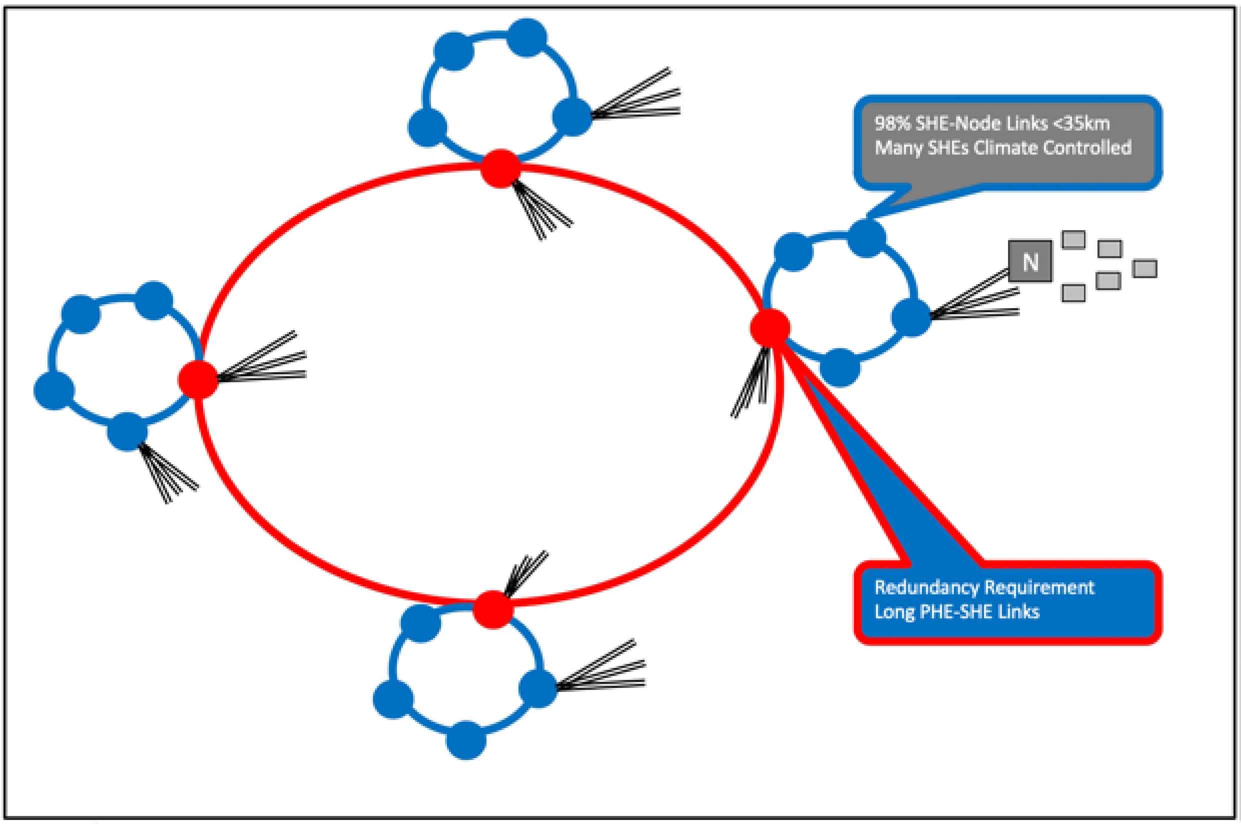

Figure 1 – Simplified Connectivity Diagram. Red=PHEs, Blue=SHEs, N = Fiber Nodes

It was around this time when we began to look at a new paradigm of network architecture. Distributed Access Architecture (DAA) envisioned the disaggregation of the CMTS into its component scheduling (aka MAC) and modulation (aka PHY) parts, accomplished by placing the PHY parts within the optical fiber node. Doing so would enable the headend to only concentrate on scheduling issues, while the PHY layer would concentrate on the modulation issues and thus distribute computational resources across the network, improve performance and reduce overall cost. The devices inside the nodes were called Remote PHY devices (RPDs) and thus was born the term "RPD node."

For the RPD nodes to function well, the CMTS functionality in the headend must be allowed to scale for traffic from day one. Because the RPD has just one 10G optical connection, we would ideally like to have all traffic flow to it from one source, without any RF combining of other formats that would mitigate the benefits afforded by the RPD.

This is what prompted us to rethink the CMTS as a virtual machine: One that depends less upon purpose-built hardware that controlled the RPD, and thru it the Cable Modems (CMs) in the house. By creative use of software structures, such as Containers/Kubernetes, the entire CMTS MAC was thus virtualized and distributed across the network. The success of this effort is of breathtaking importance, because for the first time it allows MSOs to increase traffic on the networks without the additional CMTS port license fees for each upgrade.

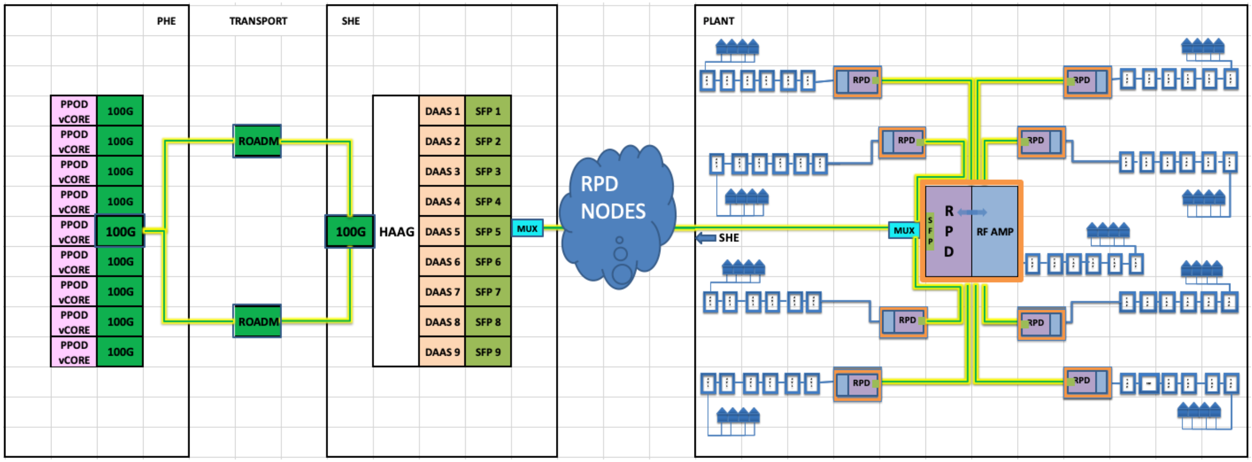

Figure 2 – End to end DAA Architecture DIagram

Once the CMTS is virtualized, the core CMTS functionality, known as the virtual core (vCORE) or interchangeably the virtual CMTS (vCMTS), is located in the primary headend (PHE), while multiple Secondary Headends host what are called DAA Switches (DAASs). These DAAS ports are connected to the vCOREs thru high speed, sometimes coherent links, while the DAASs themselves are connected to the RPDs via 10G DWDM links. The RPD node then converts the 10 Gbps baseband data into QAM format and then distributes it into the networks. At the home, the CM receives this data, and sends its US data to the RPD, which converts that into baseband data and sends them to the DAAS ports in 10Gbps format. The DASS then sends it back to the vCORE via the high-speed links, completing the circuit.

DAA is just one type of architecture in market in Comcast. In reality, we support more traditional HFC nodes than the DAA kind. Unfortunately, for the traditional nodes, the vCORE-RPD node connectivity architecture does not work, because there is not RPD in the picture. Therefore, some of the best benefits of the vCORE are not available to parts of our footprint.

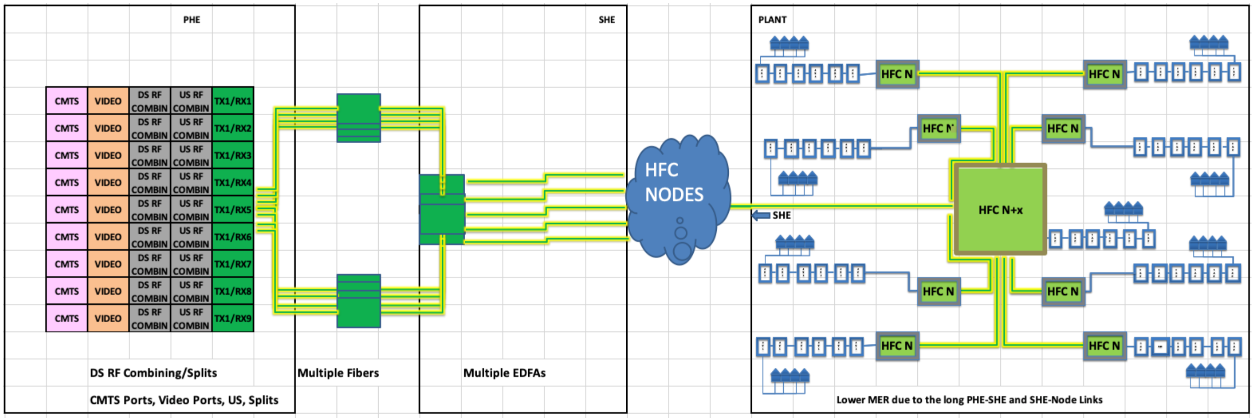

Figure 3 – HFC Architecture Diagram

Fortunately, there are ways in which this can be accomplished and in this paper we will discuss the basic idea of unleashing the vCORE onto a vast majority of our plant. Some of the ancillary benefits that accrue because of the vCMTS involve the significant critical infrastructure savings in the PHEs and SHE. Additionally, there is a potential for performance improvements over the existing link. Finally, we discuss a common wavelength plan that enables us to converge coherent, analog and 10G wavelengths onto one single fiber.

3. Architectures and Virtualization

The key to the success of any standard is its ability to fit into and enhance multiple diverse architectures. Such is the case with virtualization. Virtualization of the cable modem termination system (CMTS) began as a way to enhance Distributed Access Architecture (DAA) strategies. However, as it turns out, virtualization is a great fit for multiple architectures. Accelerating the reach of virtualization enhance all these architectures.

We have already seen that in the DAA architecture, the RPD node is a perfect fit for vCORE, however we have a preponderance of the original HFC nodes. Therefore, it is essential to provide an option for using the vCMTS architecture, so as to realize its benefits in regular HFC nodes as well. It may well be asked why we would not go about and change the entire architecture to DAA; however, that is not such an easy task. A typical DAA architecture can also be a "node plus zero" (N+0) architecture, meaning no amplification/amplifiers after the node. A that goes into making a system conform to N+0, not the least of which is the ability to secure enough HHPs per RPD. While this might happen easily in high density areas, such is not the case for lower density areas. For those areas, Grey Optics Aggregation (GOA) architectures will come in handy to shore up the number of passings per RPD. However, in current environments, and especially with COVID resulting in a lot of work from home households, the capacity needs for both US and DS are quite substantial, and capacity must be improved on a timelier basis. This is one of the reasons to leave node placements and RF amplifier placements as is.

It may further be argued that at least the nodes could be replaced, leaving RF amplifier cascades alone. While that is a better option, even it will take time. RPD nodes are physically bigger and will almost certainly need a complete node re-splicing. As well, a training regimen will be needed for field technicians who may not yet be trained on RPD devices. In some cases, it will continue to make economic sense to leverage existing investments in analog nodes and optics, especially in cases where additional capacity can be leveraged. In any case it is quite hard and inconsistent to migrate RPDs in the field.

For these reasons, the best approach for some of the network segments is to move the RPD into the headend and co-locate these with the DAAS ports in the SHE. The RPDs are then connected directly to the DAAS ports with active optical cables (AOCs), and the output of the RPDs are then fed to DS analog transmitters (TXs). These TXs feed typical HFC analog optical nodes, which reach household CMs. The return signals from the CMs arrive at the node and get transported to the SHE, where they are fed to the return receivers, then connected to the RPDs, which are then fed back to the DAAS ports, completing the signal circuit. These returns from analog nodes can represent typical digital or analog signals.

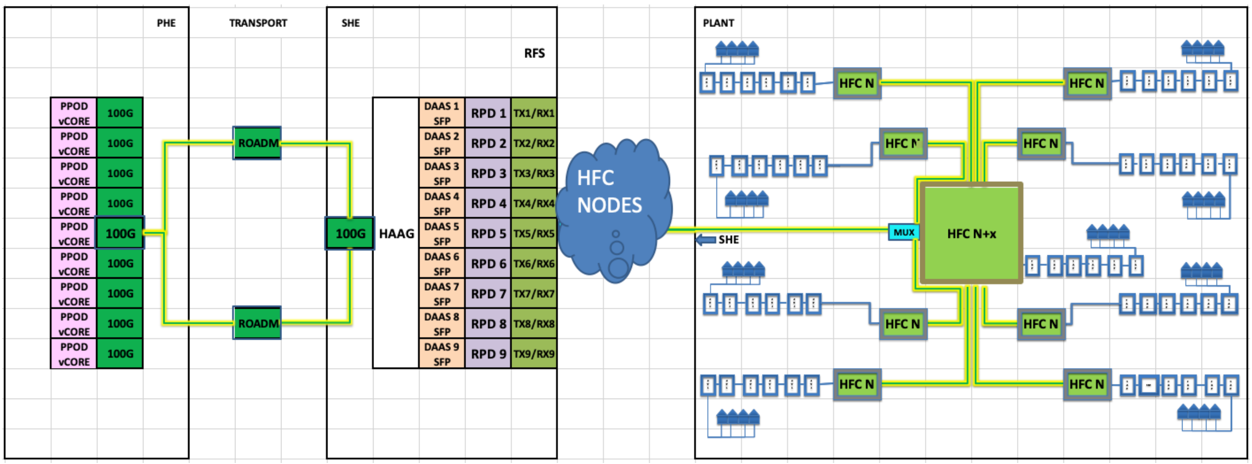

Figure 4 – RFS Architecture DIagram

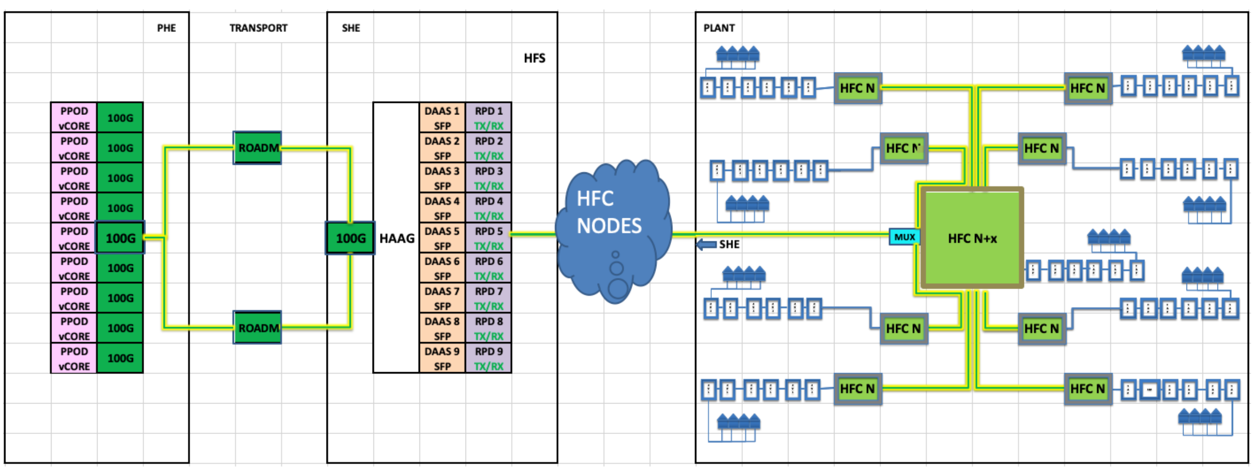

Figure 5 – HFS Architecture DIagram

This new approach to accelerate the rate of virtualization is called as the RPD Fiber Shelf (RFS) and Hybrid Fiber Shelf (HFS).

Figure 3 shows a set of details about a PHE and SHE within one of our regions. A general connectivity diagram for this was presented in Figure 1Basic Connectivity Diagram. We see here that the PHE hosts the CMTS and video traffic, along with the associated RF combiners. This primary headend also houses the analog transmitters that are multiplexed and sent over redundant links to the secondary headend. At the SHE, and opto-mechanical switch resolves the redundancy, and, after amplification, passes it over to the optical node.

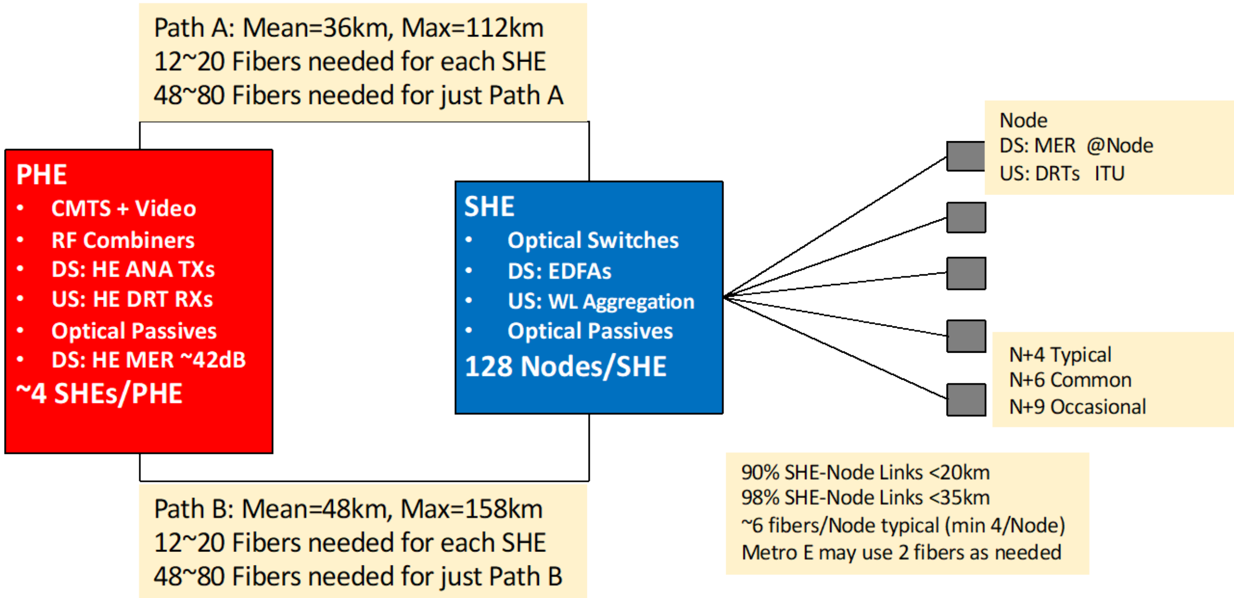

Figure 6 – An example of Traditional CMTS and PHE to SHE to Node connectivity

It must be noted here that due to the long reach from the PHE to the SHE and then to the neighborhood node, the Modulation Error Ration (MER) at the node itself is rather limited. Furthermore, the number of fibers need to traverse from the PHE to the SHE is substantial, because to avoid the deleterious fiber effects, analog transmission allows for 16 full spectrum wavelengths, which will be discussed in Section 4. In many cases, the rather large difference in primary and redundant routes creates a substantial MER delta when the redundant route becomes switched -- not to mention CMs generally reset themselves when there are substantial temporal changes caused by flight delay changes to the CMTS. The MER further reduces when we move from the node through the RF amplifiers.

In the return path, baseband digital returns have a rather large dynamic range, which gives them long reach capabilities. In this case, they are all combined together, and their light is shipped back to the primary headend. In any case, the number of fibers used from the PHE to the SHE are considerable, and any additions in the number of nodes to accommodate node splits, for example, could easily overwhelm the fiber infrastructure. Adding fibers to the PHE to SHE infrastructure is "easier said than done" due to cost of the large number of needed redundant links and also of the time associated with acquiring permits and the subsequent construction.

Fortunately, a move to HFS solves all these issues. In the diagram below, when HFS is applied, the PHE "only" has high speed coherent links of ~100-400G that traverse to the SHE over redundant routes. Since these are a well-established technology, and higher capacity, they require only a small fraction of PHE to SHE fibers. At the SHE, we connect up the coherent links to the DAAS ports. Typically, up to 48 DAAS ports can be lit with a 1RU DAA switch box. The DAAS ports are connected to the HFS and the analog link, then connected to the fiber node. Note here that the analog link traverses only the link from the SHE to the node.

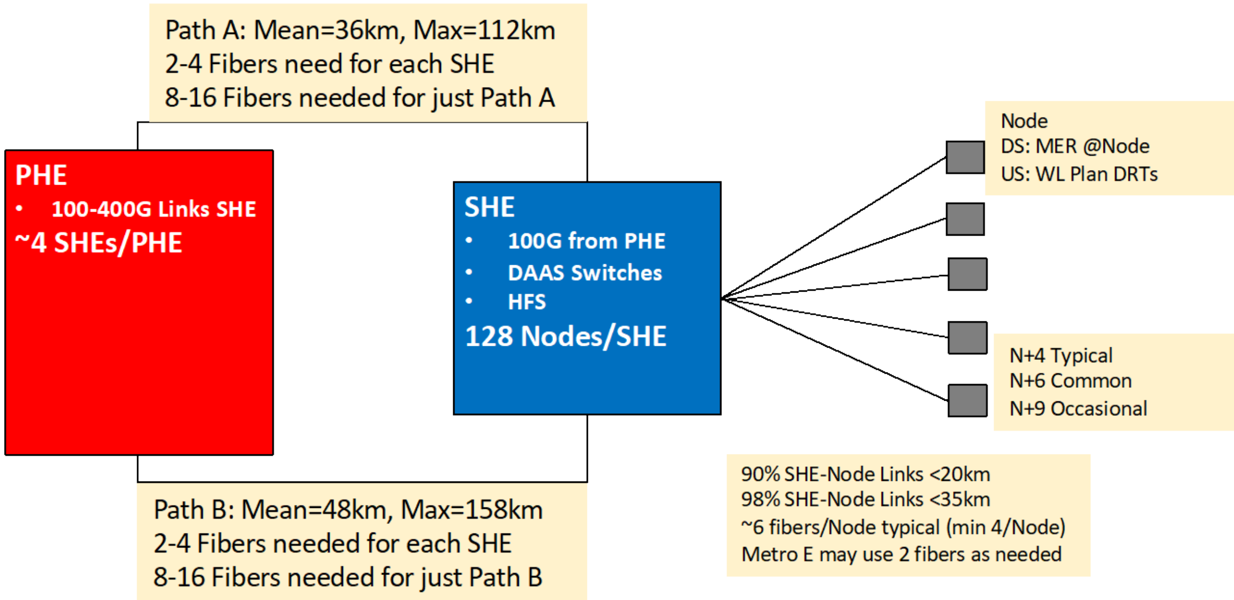

Figure 7 - An example of Traditional vCMTS and PHE to SHE to Node connectivity

By way of example, comparing the two figures immediately above, we see that there is a net savings of 10-16 fibers per PHE to SHE hop, and since there are typically 4 SHEs per PHE, this results in 40-64 fibers saved per SHE cluster. There is an equal number of secondary fibers that can be saved, as well. Some of our bigger markets may have as many as 100 SHEs, in which case the numbers of fibers saved on the (very hard to get) HE-SHE link would be up to ~1500 to 3000 when the primary and redundant routes are counted!

It is this huge savings in PHE/SHE fibers, and the associated reduction in critical infrastructure pressure at the PHEs, that accounts for a push towards HFS. But for this to happen, the HFS should be able to fit into a secondary headend – and some of them can be small and overpopulated. Furthermore, the link lengths from the SHE to the nodes are short enough for the HFS to offer good performance, even accounting for any RF cascades after it. All of these points are discussed in this paper, but we begin with a survey of our fiber links.

The graph shown in Figure 5 is the result of a survey of >100K geographically dispersed fiber nodes within our U.S. footprint. The average link length between the secondary headend to the fiber node is ~10 km; 98% of nodes within 30 km. This survey spans 16 major markets and all of our regional divisions.

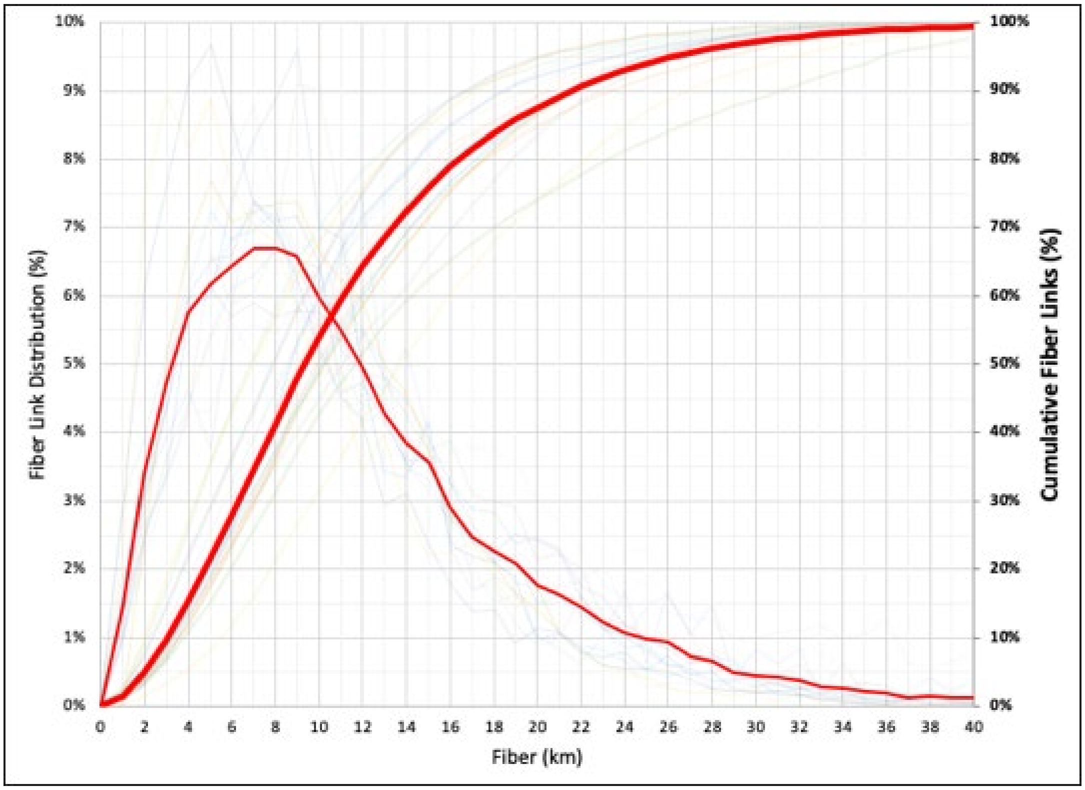

Figure 8 – Analyzing ~100K SHE-Node Fiber Links Across Our Footprint

With this visual, it can be seen that the MER of the node for a signal generated at the SHE and fed with an analog transmitter could have a much higher MER than that of a long PHE to SHE link followed by a fiber run to the node, which will be shown in Section 6.

To recap, a move to the HFS will result in a significant reduction in fibers from the PHE to the SHE. With fiber lengths that are more modest, there also a substantial improvement in MER at the node. The adoption of a Hybrid Fiber Shelf also results in critical infrastructure benefits, including substantial powering and space reductions in secondary headends. These savings, combined with the speed with which HFSs can be deployed, and the resultant uniformity of vCMTS usage across the company, are perhaps its most attractive features.

4. Converging the Headend and the Fiber Plant

We have seen in earlier sections that the primary headend/PHE hosts vCOREs and coherent optics, and the secondary headend/SHE hosts the same high-speed optics, as well as the DAAS ports and analog optics. In many cases, the PHE also supports nodes directly, which are (aptly) known as "direct fed nodes." In this case, the conventional CMTS is replaced with the vCORE; the DAASs are also connected to the vCORE and fed to the analog transmitters as described earlier. In this instance, all forms of optical conversions take pace directly within one converged headend.

Similar to the case of a converged headend, we will encounter a substantial period of time when we might expect cases of converged fibers. By converged fibers, we mean fibers with bi-directional transmission of analog and digital wavelengths on the same single fiber. Consider the case shown in Figure 6, where a simplified diagram is presented.

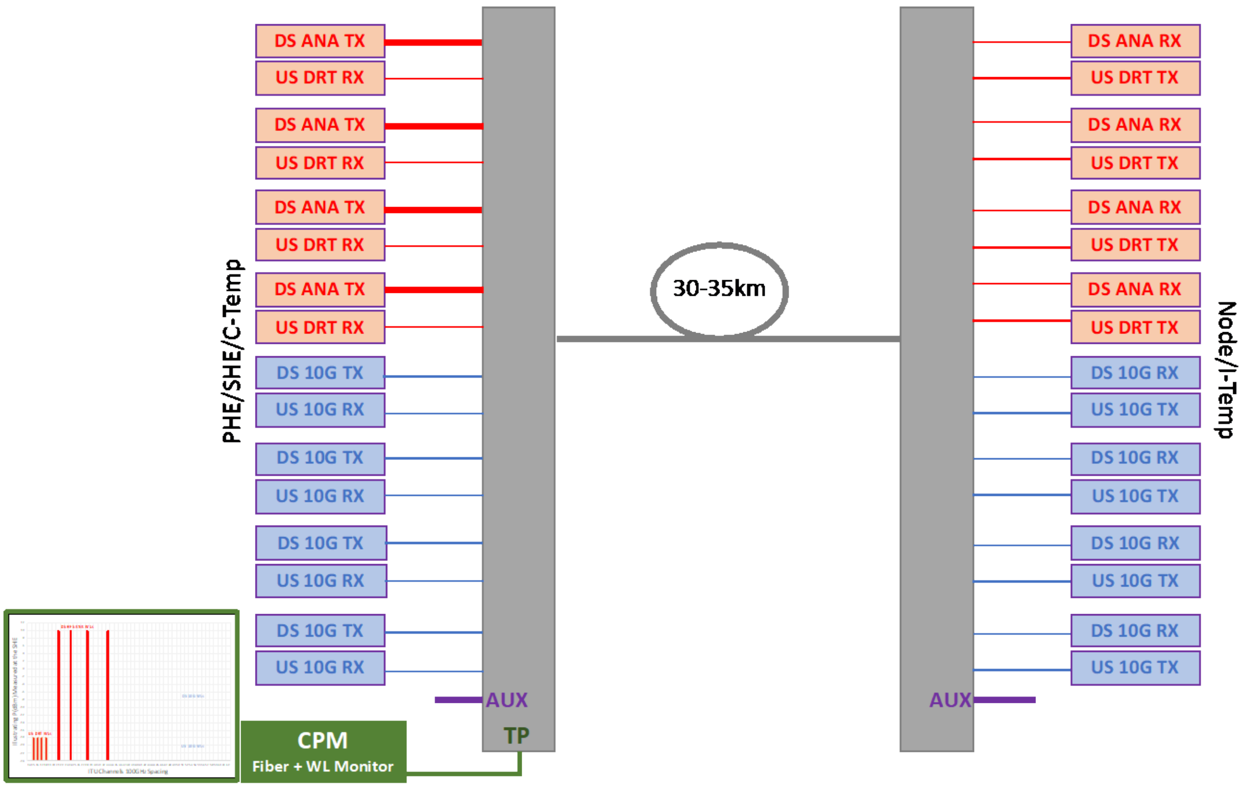

Figure 9 – Converged Fibers

On the same fiber that hosts analog DS wavelengths, the US digital return transmitter (DRT) wavelengths may be hosted as well. Furthermore, 10G wavelengths that support DAA nodes and some other 10G links may need to support MetroE links and be hosted as well. For DAA operation, today we use 10G tunable SFPs.

Considering Figure 6 as part of the SHE, we can see that there are up to 4 DS analog transmitter wavelengths and up 4 US DRT wavelengths, and up to 4 bidirectional 10G wavelengths circuits that could support any combination of DAA or MetroE circuits. A detailed description of the wavelength selection process is given in Figure 7. But here, we note that additional wavelengths that are part of a well thought out plan may still be added on the auxiliary port. In addition, the entire fiber and all sets of wavelengths are monitored over the consolidated test port by a continuous and pervasive monitor that is described in a different paper in the 2020 Cable-Tec Expo program.

There will be cases when over time, 100G/200G coherent wavelengths may need to be hosted on the same fiber, in order to enable Switch on a Pole (SOAP)-type architectures (this point will be discussed briefly in Section 5) In general, dual fiber options are not favored because the fiber coexistence use case presupposes a fiber scarcity that argues against this option. To enable multiple wavelengths of differing power levels into the fiber, a comprehensive understanding of optical effects and non-linearities is needed.

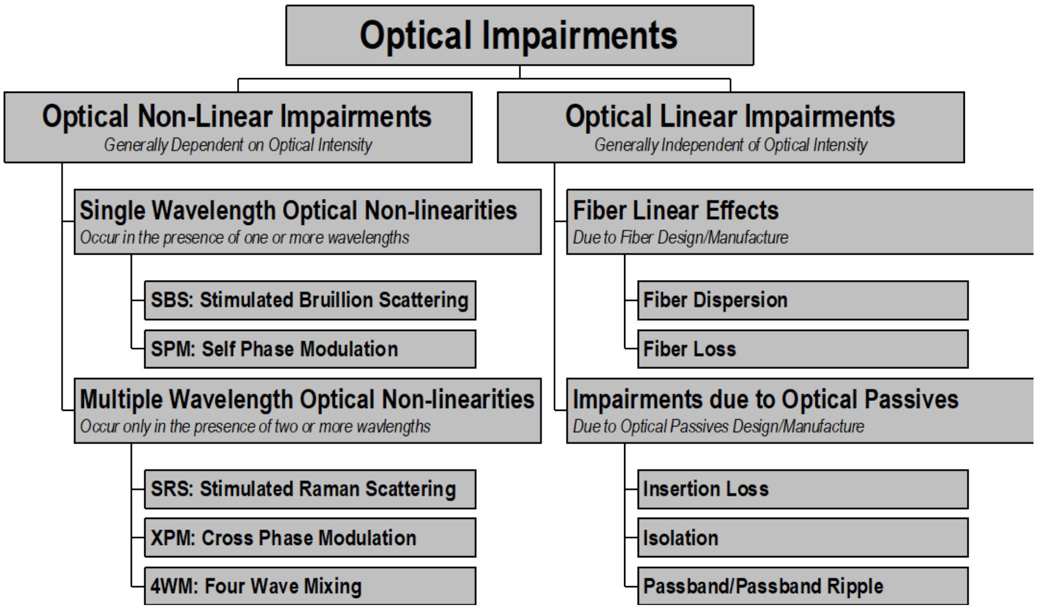

Figure 10 – Taxonomy of the Optical Effects and Non-linearities

Presented in Figure 7 is a taxonomy of optical impairments. We begin by dividing impairments into linear (those that do not generally depend upon the optical power level) and non-linear (those that generally become worse with higher optical levels) impairments. Linear impairments can be attributable to the fiber itself, such as dispersion and fiber loss, which have to be dealt with in all of the signal types. For example, the use of externally modulated transmitters heavily reduces the effects of dispersion, since they are almost chirp free. For baseband signals, propagation length is inversely proportional to the square of the dispersion. Therefore, a 10G link that can go over 80km can only go over ~15km with 25G line rates, and less than a km at 100G with the same Non-Return-to-Zero (NRZ) signals. Transmitting the signals in the 1310 nm region, where the dispersion is close to zero, would work, but here we run into the other impairment -- the fiber loss -- which is considerably higher at 1310 nm than it is at 1550 nm, with no recourse for easy optical amplification.

Add to this the fact that optical passives also add to analog system impairments. Directly modulated lasers (DMLs) can be affected by passive ripples (undulations in the passband), however, externally modulated lasers do not have this effect because they are chirp free. For digital systems, the passband of the filters should be large enough and the adjacent channel isolation deep enough for successful transmission. We currently use thin film filters, as they are well suited to indoor and outdoor plant, and because we also have high and low wavelength counts distributed throughout the infrastructure.

A common theme here is that that for analog transmission, externally modulated transmitters seem to be a good fit. They are able to handle dispersion and optical passives better. In fact, well designed externally modulated lasers can have much better MER than DMLs could -- over longer reaches and over a wider variety of optical passives. The industry has used EMLs for quite a long time, but in recent years their prices have compared favorably to DMLs. As such they represent a great alternative for the HFS and provide appreciable performance benefits.

We now come to non-linear signal performance. These generally present as single wavelength impairments which will need to be optimized. For example, while EMLs behave well with dispersion and over the passives in the linear domain, they have very aggressive single wavelength non-linearities due to stimulated Brillion scattering (SBS) which must be overcome. Fortunately for us, in HFS, we have modest fiber links, power levels are modest, and these have been overcome and are on par for the course. For digital transmission, these effects are minimal.

Finally, when we put multiple wavelengths on the same fiber, multiple nonlinear effects can take place simultaneously, such as Stimulated Raman Scattering (SRS), Cross Phase Modulation (XPM) and 4 Wave Mixing (4WM). SRS is mitigated by packing wavelengths closer; XMP is mitigated by moving wavelengths far apart; and 4WM is mitigated by avoiding all beat products that would produce significant SNR degradation by the effects of optical beat interference (OBI). OBI is a known issue in RF over Glass (RFoG) systems and has the ability to singlehandedly shut an entire system down.

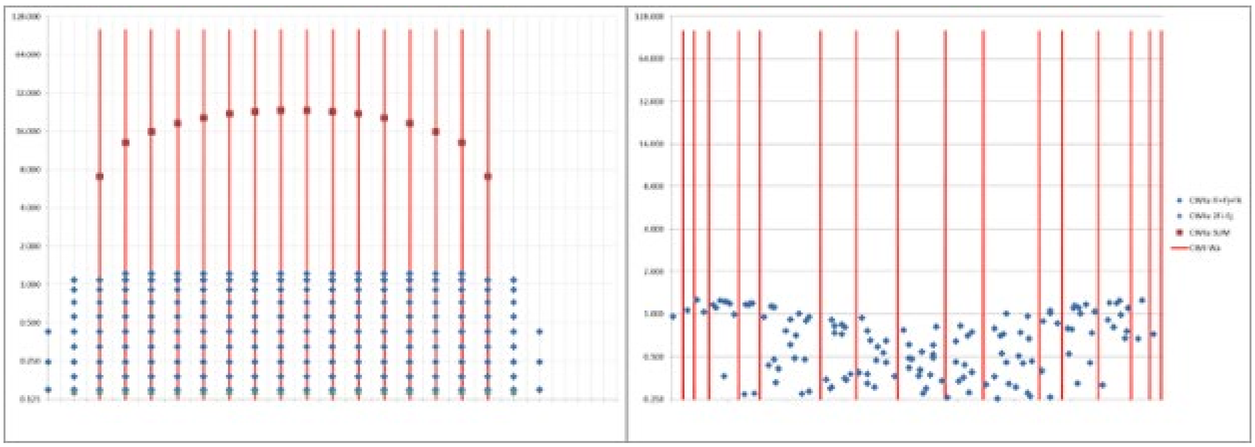

Figure 11 – Full Spectrum Wavelength Planning

To make matters worse, 4WM is intermittent, seeing as how it depends upon beats lining up with carriers. Therefore a standard uniform wavelength spacing is not appropriate. A wavelength plan that spreads the beats around is the only effective way to eliminate 4WM, while securing an effective, robust and high-power transmission in the optical fiber. Such a wavelength plan, as described above, can be enhanced with wavelength offsets, which are especially important for the EML lasers. This is so that they can accommodate the multiple lobes associated with SBS suppression circuitry.

Typically, the recommended maximum total launch power for a well-designed FS wavelength plan is 19 dBm. For 16 wavelengths, light may be launched with up to 7 dBm/wavelength of power launched into the fiber, and for 8 wavelengths light may be launched at 10 dBm/wavelength by the same token. Higher light may be launched for 4 wavelengths, but one must be careful to spread out the wavelengths within accepted guidelines to reduce the effects of SRS and XPM. Since analog wavelengths operate at high power levels, we use the wavelength plans noted above -- and therefore the links between some PHEs and SHEs are congested.

We have seen that the Analog FS wavelengths utilize the wavelength plan described in Figure 8. It is possible that the 'gaps' between the 16 wavelengths can be used by upstream Digital Return Transmitter (DRT) wavelengths. In fact, this is a way to enhance fiber efficiency, because the bi-directional transmission of light utilizes more wavelengths, without the consequent increases in optical intensity at any one end. Thus, we could have 16 WLs in the DS and a like number in the upstream with no significant degradation to the signal integrity in either direction.

The desire to use the same fiber for analog and digital (baseband) operation is an old one -- we first wrote about this in 2006 (SCTE ET). Some of the recommendations there still hold true (although that was about 1 GbE and for CWDM operations). The main idea is to ensure that the analog and digital (in this case 10 Gbps) wavelengths behave as good neighbors and do not influence each other. This is possible when we maintain a substantial optical level differential.

As mentioned, analog transmission is generally at very high levels since the performance needed is quite high. Not so for digital transmission, because the optical input requirement is really in the -21/-22dBm range for signal recovery. This being the case, there is really no need for high power transmission especially if the reach is limited (and we have seen that our links are maxed out at 35km). Therefore, one could launch light at much lower power, which would also reduce the contribution to optical non-linearities. For this reason, if we maintain a substantial optical level differential between the analog and digital transmission, the two formats would behave as good neighbors and not mutually affect each other.

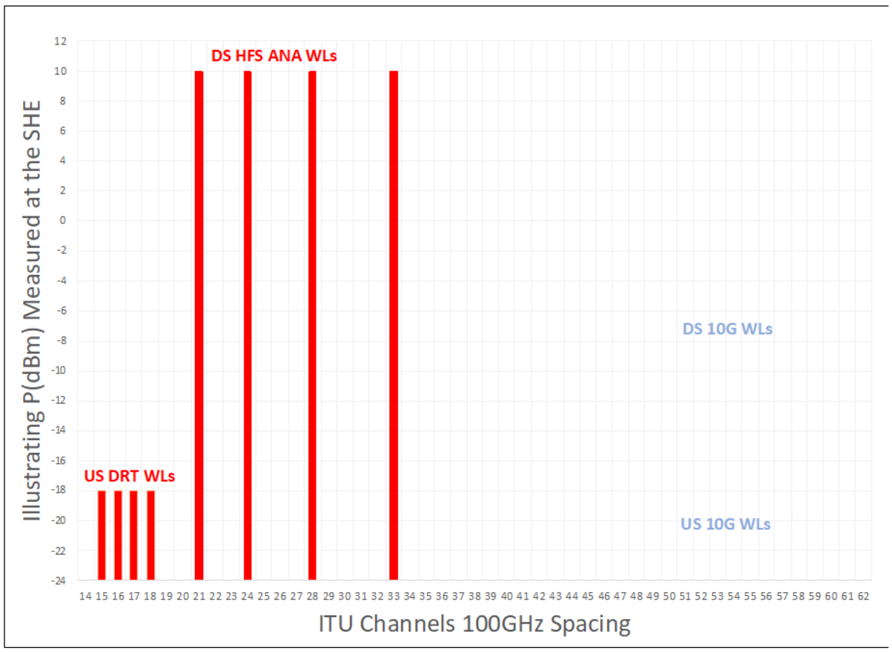

Figure 12 – Illustrating Analog and Digital Coexistence on the Same Fiber

Figure 9 shows a representative wavelength plan as seen on the CPM at the SHE, that minimizes non-linearities and enables 4 DS and US wavelengths along with 4 10G circuits on one single fiber. This type of wavelength selection and optical level differential will allow for several digital wavelengths to co-propagate with analog wavelengths. The counter propagation is considerably easier because the return wavelengths are all baseband digital and require very little power compared to the downstream.

5. A Word about Coherent Optics

Wavelength planning for accommodating Coherent Optics on the same fiber is an ongoing activity at Comcast and around the industry.

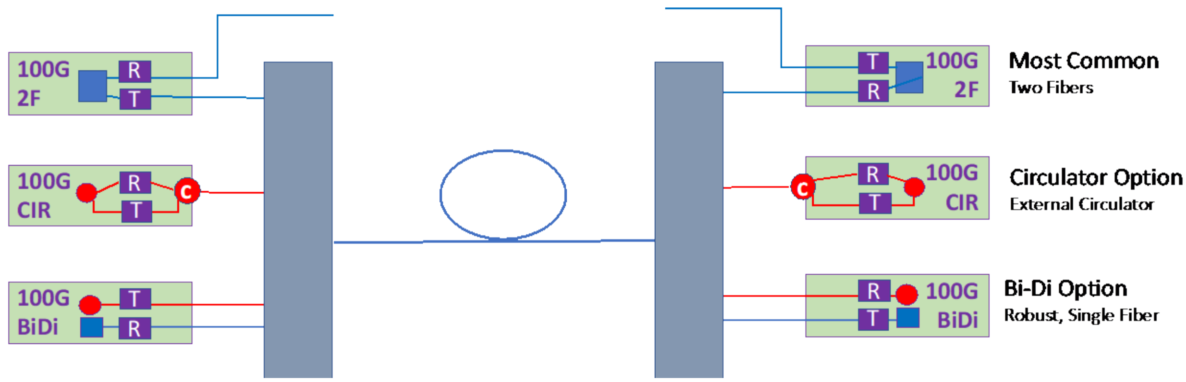

Figure 13 – Coherent Optics Options

Dual Fiber Single Laser Coherent Pluggable: Traditional Coherent Pluggables are built with an optical engine that has only one laser. The same laser pulls double duty as the transmitting light source that then passes through a nested modulator to create complex light forms to transmit over the fiber. At the far end, the laser receiver adjusts its wavelength precisely to the incoming light, and the heterodyne receiver deciphers the incoming complex signal and send it over to the DSP. The processor then decodes the signals and streams them in either 10Gbps or more likely 25Gbps binary signals to the outside world. The other half of the aforementioned laser's light is sent over to the nested modulator. Its light is then sent over a separate fiber to the first module, where the lasers will heterodyne and decipher the incoming light. This is perhaps the only place where the OBI is intentionally used to enhance the communication system. It is what makes Coherent systems so interesting and special. As mentioned in Section 4, coexisting in a single fiber means that there is not a second fiber available to be to be used, so this option may not work in those cases.

Single Fiber Single Laser Coherent Pluggable: For single fiber co-existence, it is proposed to use an optical circulator to reuse the same wavelength. In this case, the light in an optical circulator can move only one way, therefore the transmit light can only move out of the module while the received light might move only inside the module. With this arrangement the wavelength may be reused. The circulator has good but imperfect isolation and is prone to losing isolation at temperature extremes. While this may not be an issue at the PHE or the SHE, it may be an issue at the node, where temperatures can cycle between -40C to +60C. Furthermore, reflections in the system can also impinge on the available dynamic range and reduce the effectiveness of the coherent receiver, thus limiting the fiber reach or speed.

Single Fiber Dual Laser Coherent Pluggable: A dual laser design for the coherent pluggable has two separate lasers, one for light transmission and another to receive light. The corresponding transmitting and receiving lasers, in two separate pluggables across the fiber mutually lock on to each other. Because their sources are independent, they deliver true light bidirectionality over different wavelengths. As a result, this arrangement is not prone to any of the reflection issues described earlier. Previously, these types of devices were expensive and had higher heat dissipation. However, the advent of 7 nm DSP technology has propelled efficient drivers and with efficient lasers and cost-effective optical amplifiers these issues have been resolved. These modules are also now very cost effective and have been realized in (C-Formfactor Pluggable 2) CFP2 form factors.

This discussion on coherent optics is intended to prepare the reader for the notion of a "Switch on a Pole" (SOAP) type of an architecture, that would, with the help of a coherent termination device, transform the capacity of access networks dramatically. In this context, we gratefully note the efforts of CableLabs and its P2P link and Coherent Terminating Device (CTD) specification efforts. There are also multiple efforts underway towards 25G NRZ transmission, based on a 10G SFP platform, and also of a new type of point-to multipoint (P2MP) architecture using subcarriers. All are trying to bring higher speed optics closer towards the access space.

6. Hybrid Fiber Shelf

A typical headend that has CMTS is a pretty big facility. Generally, it holds the high-speed circuits coming in to connect to the internet, and video sources coming in either from a satellite dish farm or through terrestrial fiber. Multiple modulators groom the video to be on specific outbound frequencies. Any on-demand servers are connected to modulators as well. The CMTS itself is processing, scheduling and getting traffic, and modulating it. Finally, all of these signals are RF-combined and fed into an analog transmitter.

Each RF combining attenuates the output signal, so a very common requirement is for high RF output power as mentioned in DRFI (DOCSIS RF Interface) Specification to burn through all the of the combining network. Generating high output RF power of sufficient quality requires RF amplifiers and a high-power dissipation to power them up. Of course, all of this is frittered away in the RF combiners until it reaches the transmitter. If there is cabling in a PHE, then the cable loss and its tilt will have to be taken into account. Typical transmitters today require 10 dBmV/6 MHz (in days past they would have required 15 dBmV/6 MHz) but it is common to have individual per channel modulator outputs as high as 60 dBmV/6 MHz.

As we have discussed, there are two main problems with this scenario. For starters, as HSD traffic grows, so too will the CMTS ports or the needed capacity per port. At the limit when HSD has subsumed most of the traffic needs, then the RF can all be potentially available at one port. This will also require the maximum CMTS licenses (discussed in Section 2), but until that time, the entire RF combining infrastructure is carried. However, once the vCOREs come online, and the entire CMTS is virtualized, the timing is perfect to dismantle the RF combiners and recover the space. This is possible because the entire RF spectra is populated by the vCMTS output port and can be connected via a grey 100G connection to the DAAS switch. We note here that the DAAS output port of 10G puts out the complete RF spectra that may be generated at the RPD.

The DAAS may be connected to the headend RPD device thru 10G AOCs, which will have fairly low power consumption. If the RPD is co-located with the VCMTS then this might even be connected via DAC cables that are even more power conserving and at very low cost.

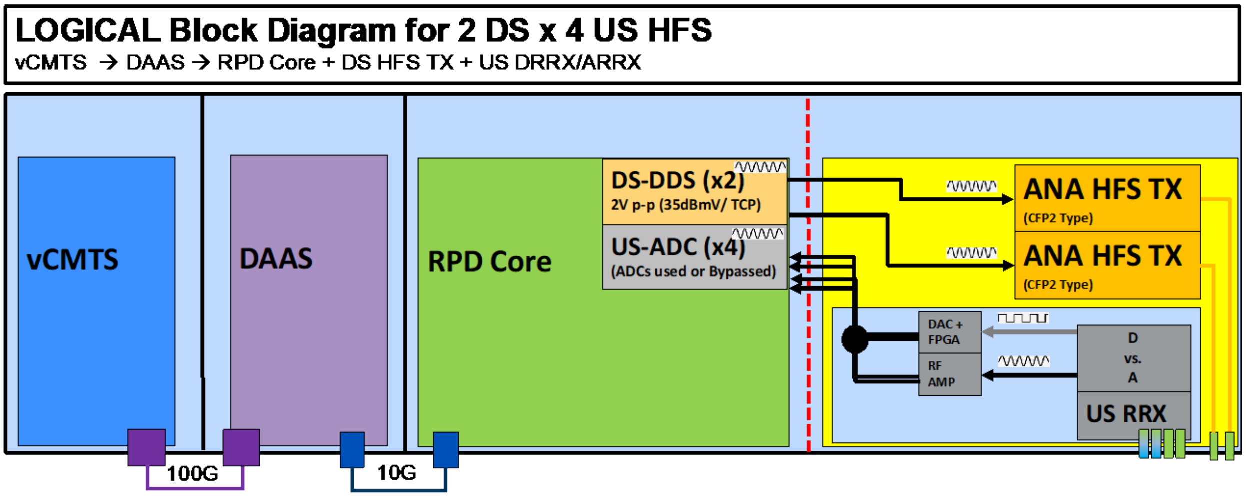

Figure 14 – Logical Block Diagram for HFS

A typical RPD without its post amplifier puts out 2V p2p. This is roughly 35dBmV of RF power (for distortion-free performance). Typically, the RPD MER is better than 50 dB, on average.

If a chassis has only RPDs in it, we call it an RFS (for "RPD Shelf") to distinguish it from a Hybrid Fiber Shelf (HFS), which would have had the transmitters and return receiver integrated. For our purposes, the red line shown in Figure 13 is the demarcation point. We will use the acronym "HFS" generically, to mean that the RPD Shelf and the optics shelf is either integrated or co-located – either way, the RF combining circuits are eliminated or heavily reduced. While the RFS is already of material benefit to us as an operator, the HFS enhances that benefit in cost effectiveness for critical infrastructure.

Years back, the analog transmitter was a device that took considerable space and power and used high performance Directly Modulated Laser (DML) technology. State-of-the-art transmitters are now built using Externally Modulated Laser (EML) technology within the size of a CFP2 form factor typically used for 10G transmission. These comprise an EML, the linearization circuit for it and any SBS suppression circuitry. Typically, we need only around 10-11 dBm of SBS suppression to cover the 30km or so of the SHE to node link. A typical HFS transmitter would have around 4W of dissipation and around 11 dBm of output optical power. As is oftentimes the case, EML transmitters have lower optical modulation depth (OMI) than do the DMLs, but their chirp free dispersion performance and lower RIN make up for it and generally provide a modicum of advantage. The best feature of these is that there is no reason now to declare the fiber lengths.

Using DMLs would have required us to declare fiber distances, because the electronic dispersion compensation would need distance to cancel fiber dispersion-based distortions generated in the fiber. Even though this is a one-time provisioning issue, it causes problems when a redundant route gets activated. As it happens, there is not redundancy from the SHE to the node, but still, eliminating this extra step is a welcome feature. So is the feature of being able to use common optical passives for 10G, coherent and/or analog, which was harder to do with DML transmitters.

In another paper this year, we have described a pervasive monitoring tool to track all wavelengths. That tool also needs a description of the fiber length. One thing to keep in mind here is that the needed number of analog wavelengths per fiber is going to be rather limited in this architecture. In this architecture we typically would have had each SHE feed one node, on average. This is because the link would have started at the PHE, then thru the SHE and then all the way to the node, and we have seen that long-distance transmission of analog wavelengths is limited.

In fact, while multiwavelength analog transmission is limiting, single or limited wavelength analog transmission over modest fiber distances with EML transmission is has high performance. Figure 15 shows a measured result of an RPD followed by EML transmitter over 38km of fiber. This shows a node MER of ~45dB, which quite handily beats the 4K QAM requirement of 41dB.

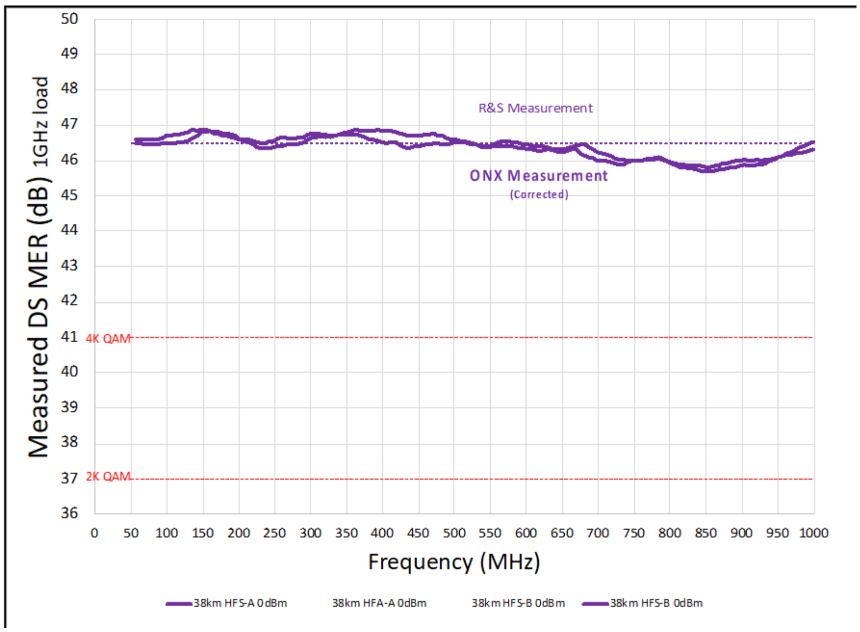

Figure 15 – Measured RPD and EML Transmitter MER

However, it is often the case that multiple amplifier cascades are in place after the RPD. In Figure 16, we can see the performance of an amplifier cascade that has been added to the RPD MER. Assuming each RF amplifier at around 50 MER each, this indicates that the cumulative MER performance is also dictated by the RF amplifier cascade as well as that of the RPD and of the HFS EML transmitter. Figure 16 shows a simulated performance of the RF cascade, beginning with an MER appropriately reduced to account for unit to unit variation and temperature performance. The cumulative performance is then very close to ~41dB MER, even for N+6. This performance complements the efforts of DS profile management application (PMA) that automatically adjusts the modulation complexity to available SNR and maximizes capacity deployed at Comcast. With smaller cascades, the performance is even better.

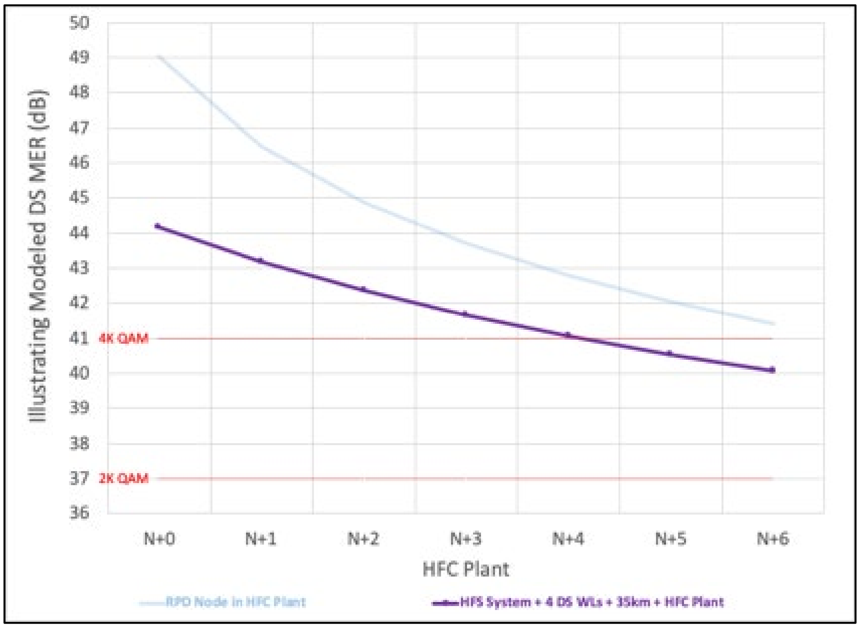

Figure 16 – Node and RF Cascades

Figure 17 summarizes the benefits of the Hybrid Fiber Shelf.

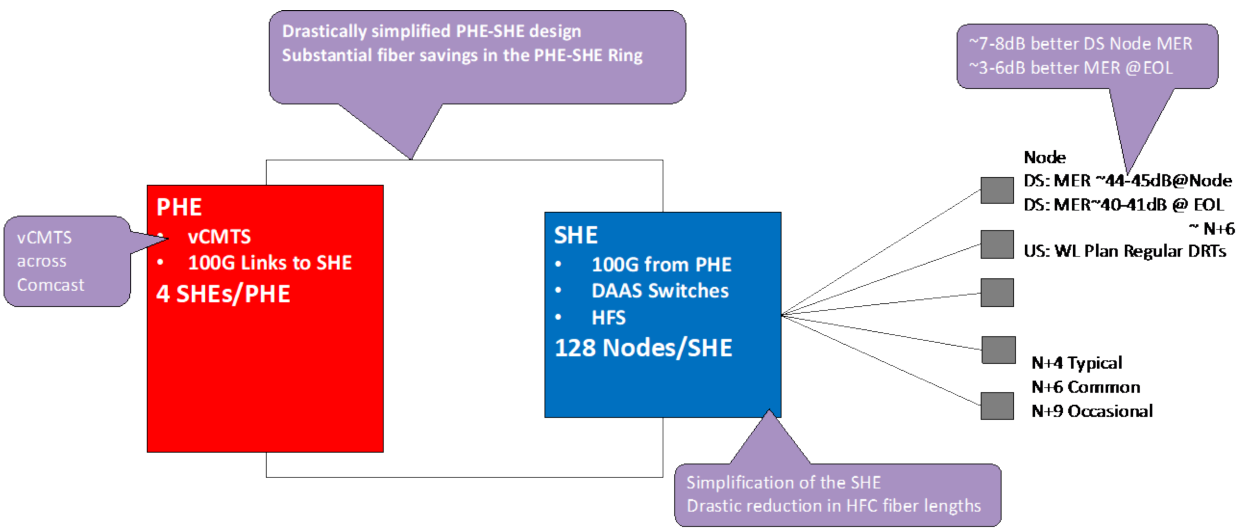

Figure 17 – Summary of HFS

7. Critical Infrastructure

In this section we talk about how both HFS and RFS can potentially save on Critical Infrastructure (CI) in the PHE and the SHE. As we have indicated before, a regular CMTS occupies several rack units (RUs). Further space is occupied by the video generators, and the most space is occupied by the RF combining circuitry. Overcoming the excess loss of the RF combiners requires that the CMTS and other video equipment put out a rather large RF level. This RF Level is then split and combined to incorporate the various signals and then provided to the Analog Transmitter at an appropriate level. A typical CMTS-based system connected to 192 service groups can take five 7-foot racks worth of valuable headend real estate. Some of the older CMTSs and the existing RF combining equipment take up even more space, because the CMTS and the DS transmitters are not co-located in the building, so RF cables are run across the headend to connect them. This has the unfortunate result of needing to have slope and amplitude correction circuitry, on a transmitter by transmitter basis, to groom the input RF to the transmitter for distribution. Added to this, because RF combining is complex and folks will not want to touch it once installed, there are numerous test points distributed across the RF cables which provide the convenience of test points but at added power and cost. The same process is repeated for the US as well. The US RF out of the US RX comes in at a high level, is split multiple ways, and combined. This conserves CMTS ports, with enough test ports to track RF levels along, and is then connected up to the CMTS.

With a system described above, accommodating necessary node splits will require additional headend space that is hard to get. Adding to a headend is quite expensive and may not even be possible, in the time needed for the node splits, and after accounting for permits and construction.

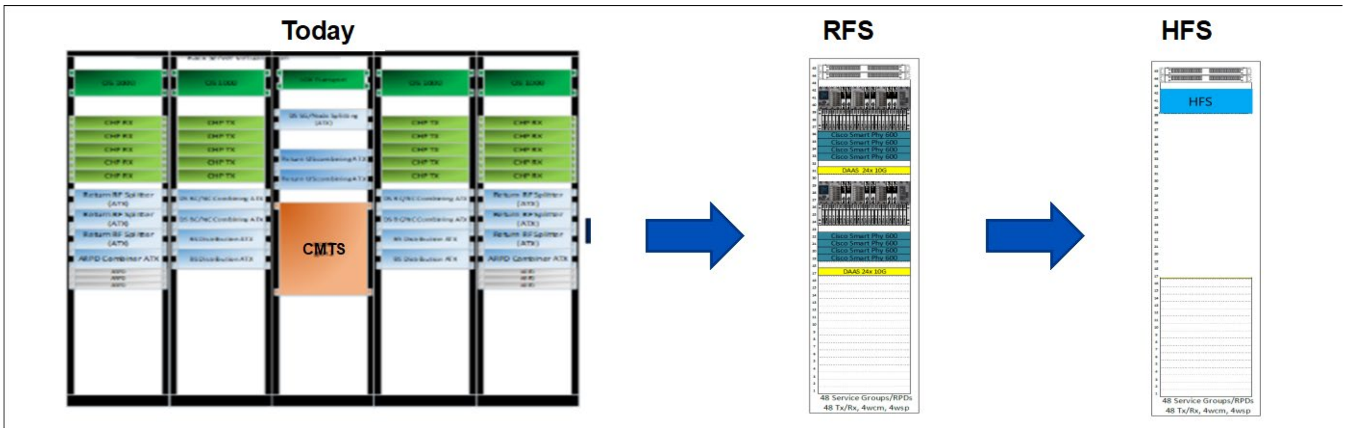

Figure 18 – HFS Critical Infrastructure

With an enterprise wide vCMTS/vCORE one can now compress all of the HSD, video and other assets into one converged platform and have a single RF connection. In Figure 18, we have indicated that the RFS is co-located. Here, the RF combining is totally replaced with just a DAAS port that is connected to the RFS via AOCs or DACs. Then, there is a modicum of RF circuitry before the output level is connected to the analog chassis. Thus, the space and power savings are considerable, as shown above. Even though the RFS is very compact, it still has a higher output, because we may need to feed existing analog transmitters which may require higher outputs. Sometimes the analog gear may not be co-located (especially in the PHEs) and therefore it may need to be connected to long sections of RF cable. For this reason, the RF output of the RFS is also higher and will require the consequent grooming. One additional fact will need to be considered as well: any time a new node is commissioned, it is common to perform what we call "First Node Provisioning" (FNP). This is done by connecting a CM directly to the chassis and ensuring that the CM can turn up and stay connected. This is done to ensure that there are no loose ends in the node commissioning process. It requires test points for the US that will be able to not only display the ingoing RF levels, but also provide a path for the CM signal to be injected. For this reason, the modicum of RF combining is also maintained.

A study by our Critical Infrastructure group on one of the PHEs found not only substantial space savings, but also that there would be 23% reduction in powering. A 23% reduction in powering is quite a major achievement, because it will also contribute to lower air conditioning costs and our corporate sustainability goals. There is a substantial overall cost reduction with a move to the vCMTS/RFS combining, as well, but that is out of scope for this paper.

The real breakthrough, in hindsight, was the move to a HFS: Integrating the whole Remote PHY device (RPD), downstream transmitter and upstream receiver, all in the same module and installed in the same chassis. When done this way, there is no RF combining, by definition, the test points are conducive not only to verify RF levels, but also to do FNP. Also, the RF level out of the RPD is exactly equal to the RF level needed by the DS transmitter, and the RF level out of the US RX is also exactly equal to the RF level needed for the RPD. In this case, the power requirement is much lower, while the overall density is higher. Note that the RFS will have higher density, but when combined with the need to add in the DS TXs, US RXs and the RF combining, it still results in an overall lower density than the HFS. Furthermore, since the transmitters are very low in power consumption, similar analysis indicated above shows a further 40% reduction in power and space. This makes it very attractive and enables us to start placing the HFS in secondary headends and other such “hard to fit” areas. This is where the maximum benefit of this architecture comes into effect.

It is important to realize that both RFS and HFS have their place. For example, in a PHE, where there is a preponderance of existing optical transmitters and there is a modicum of space, many of these could be repurposed and either co-located or distance located from the RFS. But if the issue is to move closer to the nodes or if node-splits are to happen and new transmitters and receivers are needed anyway, then a move to SHE may be warranted and here the HFS would be a far better fit.

It may be asked that if additional nodes are to be added, why would we not just deploy RPD nodes. Indeed, that may be a legitimate option, and 1x2 and 2x4 RPD nodes are more available now than they were a couple of years back. However, if the aim was predominantly to eliminate traditional CMTS and standardize on the vCMTS, replacing of all nodes would be too onerous of a task. A HFS would be a better choice and enable use of existing nodes with available segmentation capacity. We have further shown that the performance is quite good and holds itself well in a N+x type system, where there is a fair contribution from the RF cascade as well.

To summarize, a combination of DAA RPD nodes, 1x2, 2x4 HFC RPD nodes, RFS and HFS are a set of tools to accelerate the virtualization of CMTS and usher in uniformity and standardization within our access networks.

8. Conclusions

We began this paper with a recap of our work to virtualize what are some very complex portions of our network, including the CMTS, which resulted in network simplification and harmonization of multiple purpose-built platforms into one common entity. However, DAA deployments require headend and field modifications as part of the deployment. RFS takes great step forward by accelerating virtualization with a design of RPDs in a shelf, but still requires some RF combining and a separate optical chassis.

HFS takes that concept further, integrating with virtualized CMTS/DAA switches at one end and with transmit/receive analog optical signals at the other end, into the outside plant. This innovation provides rapid and economically sustainable increases in capacity by accelerating integration with the vCMTS, as DAA construction proceeds in other areas of the network. HFS improves critical infrastructure in headends by saving wasted space, power and time inherent in RF combining and splitting circuits. Extending HFS into secondary headends also enables significant fiber reclamation. Locating these assets in secondary facilities, closer to the traditional HFC nodes, improves performance in ways that translate into enhanced capacity.

9. Acknowledgements

It is with gratitude that we acknowledge the team within Comcast who has been working on the RPD and RFS projects. We also thank the Critical Infrastructure team for their analyses, and the Access Engineering team for their support on fiber convergence and planning efforts.

Abbreviations

4WM	Four wave mixing
AOC	Active Optical Cable
CAGR	compound annual growth rate
CFP2	C-Formfactor Pluggable -2
CI	Critical infrastructure
CM	cable modem
CMTS	cable modem termination system
CPM	Continuous Pervasive Monitor
CTD	Coherent Terminating Device
DAA	Distributed access architecture
DAC	Digital Access Cable
DAAS	Distributed access architecture switch
DRFI	DOCSIS RF Interface Specification
DRT	Downstream Return Transmitter
DML	Directly modulated laser
DS	Downstream
DSP	Digital Signal Processor
EML	Externally modulated laser
FNP	First node provisioning
FS	Full spectrum
GOA	Grey Optics Architecture
HFC	Hybrid fiber coax
HFS	Hybrid Fiber Shelf
HSD	High speed data
MAC	Media access control
MER	Modulation error ratio
MPEG	Moving Pictures Experts Group
MSO	Multiple systems operator
NRZ	Non Return to Zero
OBI	Optical beat interference
P2P	Peer to peer
P2MP	Peer to multi-peer
PHE	Primary Headend
PHY	Physical
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RFoG	Radio Frequency over Glass

RFS	RPD Fiber Shelf
RIN	Relative intensity noise
RPD	Remote PHY device
RU	Rack unit
Rx	Receiver
SFP	Small form-factor pluggable
SHE	Secondary Headend
SOAP	Switch on a pole
SRS	Stimulated Raman Scattering
US	Upstream
vCMTS	Virtual cable modem termination system
VOD	Video on demand
WL	Wavelength
XMP	Cross phase modulation

Bibliography & References

1. *Operationalizing the Grey Optics Architecture: An Update a Year After*, Venk Mutalik, Dan Rice, Bob Gaydos, Doug Combs and Pat Wike, SCTE EXPO 2020

2. *It is 10PM: Do you know Where Your Wavelengths are? Continuous and Pervasive Monitoring of Optical Assets in the Access Domain*, Venk Mutalik, Dan Rice, Rick Spanbauer, Simone Capuano, Rob Gonsalves and Bob Gaydos, SCTE EXPO 2020

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

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

5. *Gigabit Ethernet and Analog/QAM Traffic Compatibility in HFC Networks: A case for Physical Layer Convergence,* Shamim Akhtar, Doug Weiss, Venk Mutalik, Marcel Schemmann, David Heisler, Down Wesson, Liyan Zhang, SCTE ET 2006

