Photon Avatars in the Comcast Cosmos:

An End-to-End View of Comcast Core, Metro and Access Networks

A Technical Paper prepared for SCTE by

Venk Mutalik
Fellow
Comcast
1800 Arch Street, Philadelphia, PA 19103
+1 (860) 262-4479
Venk_Mutalik@Comcast.com

Steve Ruppa, Senior Principal Engineer

Fred Bartholf, Senior Principal Engineer

Bob Gaydos, Fellow

Steve Surdam, Vice President

Amarildo Vieira, Principal Optical Engineer

Dan Rice, Vice President

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

As the largest broadband company in the US, Comcast serves millions of customers and businesses with a reach that stretches coast to coast. All of this is the result of a large optical network that spans core, metro and access layers with multiple intersecting points all intended to increase capacity, reduce latency, and enhance reliability.

In this paper, we describe for the first time an end-to-end view of our optical network including the core, metro, and access layers. At the core and metro, we increase capacity with a move towards flexible 400G connections and reduce latency and enhance reliability with an infrastructure that meshes color-less, direction-less, and contention-less reconfigurable multiplexers thru to each of our headends. At the access layer that connects these headends to customers and businesses, we discuss capacity increases with our move to all-digital fiber links and the distributed access architecture paradigm. Of note is a cost effective environmentally hardened dual laser bidirectional coherent 100G system and the converging all bidirectional access transmission formats on one single optical fiber. Reliability enhancements accrue with real-time continuous and pervasive optical monitoring of all these access assets. We then briefly describe the infrastructure that helps provision, visualize and event these layers. Finally, we will venture into the future of optical technology at Comcast and its positive impact on network robustness and enhancing the customer experience.

2. The Big Picture

Photons flood into the Comcast backbone network from giant Internet routers and reach the various metro routers. At the metro center, photons reincarnate and course thru the highly meshed Converged Regional Area Network infrastructure that terminates in the thousands of our headends. At the various headends, photons reincarnate again, traverse access fibers and light up the many homes, businesses and fiber nodes eventually completing their journey in the downstream. A similar process ensues in the upstream where photons transmigrate thru the access system to the various headends and mesh metro circuits before making their way back thru to the internet.

2.1. Backbone, Metro and Access

Presented below is a picture that illustrates the reach of our backbone and access systems in Comcast [1].

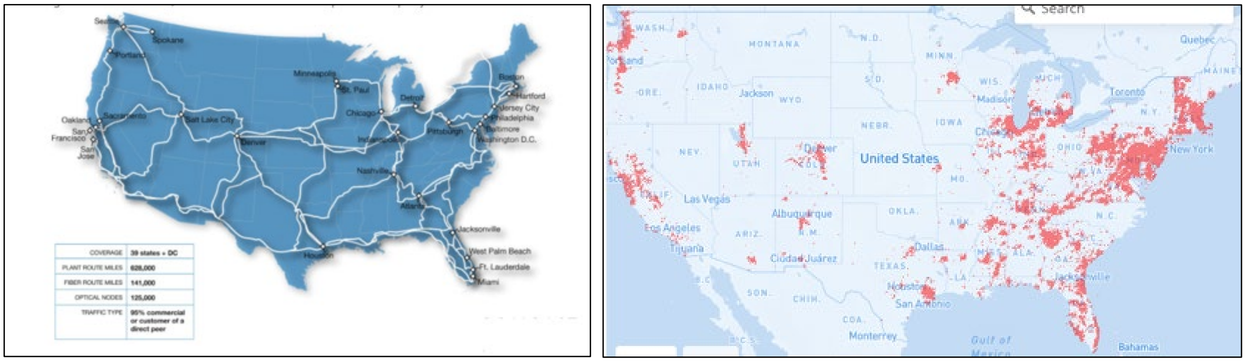

Figure 1 – Comcast Backbone and Access Networks

With the breadth that reaches coast to coast connecting important Internet Data sites and the depth of reaching 60M homes and businesses, the Comcast communications system comprising the Backbone, Metro and Access sub-systems today is now all under one single group in Comcast.

Internet sites located at Chicago, Atlanta, Ashburn and others are connected to the Backbone and via route redundancy reach all our Metro sites that interconnect hubs and headends within individual cities. At the extremities of the Metro network, residential access network interconnect with Metro and via vCMTS (virtualized Cable Modem Termination System) connect to Remote PHY based digital fiber Nodes. Commercial/Business services which were always closely connected to Metro are now converging with Residential access networks and are sharing common connectivity and fiber assets.

The residential access network is orchestrated by Comcast own provisioning, monitoring and fix agents and provide an end-to-end view. The Metro and Backbone orchestrations for provisioning, monitoring and fix are more closely related to the vendors of choice.

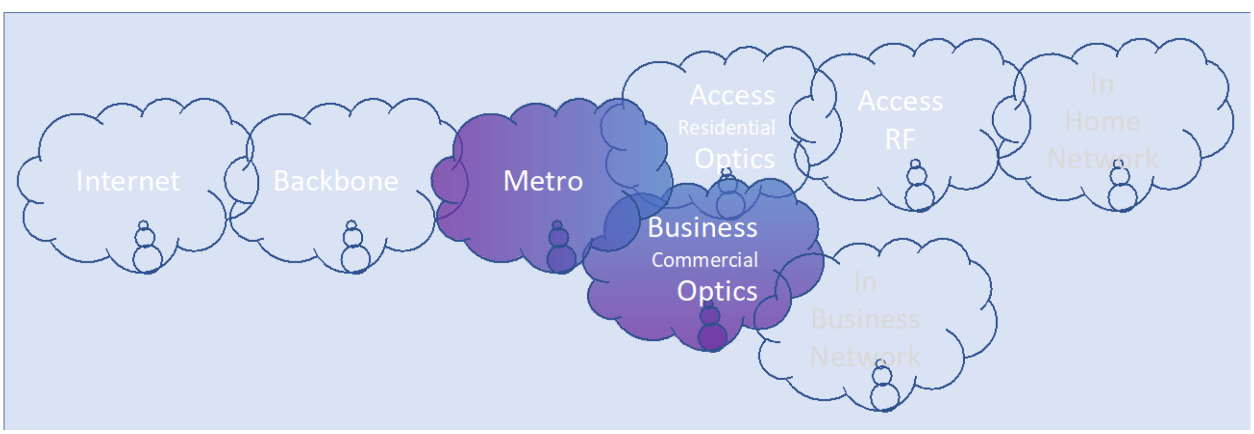

Figure 2 – Comcast Architecture

With this understanding in the next few sections, we describe a more detailed view of each of the subsystems along with a brief history evolving to the current state and potential innovations that we are looking into for meeting capacity and performance and transform Comcast into the leading communications provider in the country.

3. The Backbone

The early days of the Comcast network began as 28 separate networks or islands connected using Transit connectivity and leveraging local access interconnects. Comcast's network needs dramatically changed in the early 2000's. The company was looking to centralize video downlinks in order to distribute them terrestrially, which coincided with the initial roll out and testing of Broadband cable modems. Between the needs of nationally distributed video, transit for high-speed internet access, and the birth of data centers for VOD / back-office systems, there was an obvious need for a national network. The decision was made to build a National Backbone that enabled more control over the products and services Comcast would deliver by acquiring the Indefeasible Right of Use to a 2-fiber national footprint. The footprint touched all 28 Comcast metro networks and had a presence in the majority of all the tier 1 cities that enabled easy connectivity to our transit provider(s). Since the fiber passed close to Comcast buildings, but not through Comcast buildings, short fiber laterals would need to be constructed to attach metros to the Backbone.

3.1. Comcast Backbone History

Each of the 28 metros chose 2 locations for Backbone connectivity, and core locations were also strategically selected in major cities. Each core location was given a core router for Comcast "internal video" and a core router for "external" or "internet" services. Although multiple "networks" were built from a routing perspective, fiber availability, and cost, drove the need for a single transport platform to carry all traffic and lines of business. The original Backbone was designed and built as a greenfield taking advantage of new MEMS based reconfigurable add, drop multiplexers (ROADM) technology. ROADMs were placed at each metro aggregation points, core location, and fiber branch. This created an 88 channel, fixed 50GHz, meshed network, that allowed channels to be turned up between ANY add/drop locations. A large portion of the footprint contained amplifier sites that didn't serve any add/drop purpose. These sites were built as amplifier only sites, unless optical regeneration was needed in which the site was built as an add/drop.

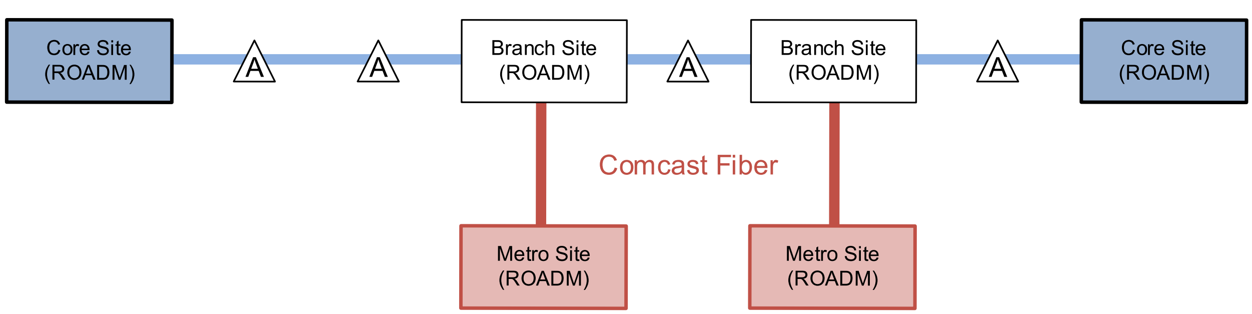

Figure 3 – Comcast Backbone Architecture - 1

Traffic was then turned up using 10G DWDM transponders. Initial calculations arrived at each metro requiring a 10G channel for "internal" Comcast video traffic, and a 10G channel needed for "external" data/internet traffic. To support this in the core, all core routers were also connected via 10G channels per supported service.

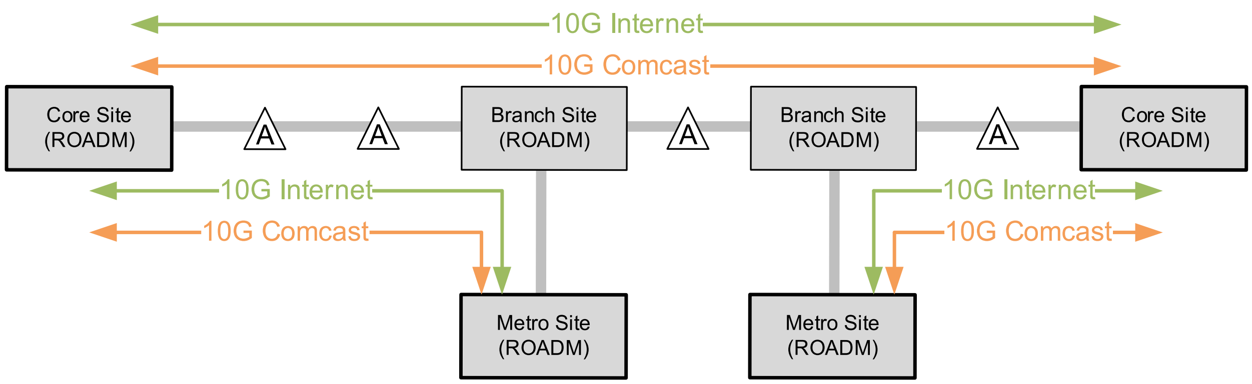

Figure 4 – Comcast Backbone Architecture - 2

Before the entire greenfield was even completed, it became VERY obvious that the exponential growth of the internet would quickly drive 2 changes.

First, Comcast needed to transition from leveraging a transit provider, to becoming a full-service provider. Core sites needed to move to locations that were more conducive to Settlement Free Interconnects. "Carrier Hotel" sites were chosen, and "Edge" router ports now faced offnet Settlement Free Interconnects and customers.

Second, the Compounded Annual Growth Rate (CAGR) growth far exceeded Comcast's ability to deploy technology such as 10G transponders. We quickly realized that in order to support our growth that we would need to drive Optical Technology roadmaps to support higher Optical speeds. The move to 40G (2006) and then to 100G (2011) was required to manage our optical spectrum efficiently to avoid exhaustion on the original line system.

By 2015, the original fixed grid system was full and unable to utilize the new 400G transponders that required flex grid. Comcast chose to acquire more fiber and construct a 2nd line system instead of trying to upgrade the existing system. The 2nd backbone optical system was completed in 2018 and was almost exclusively 400G from the start with built in capabilities to expand to 800G and beyond.

3.2. Optical Innovations in the Backbone

The tremendous CAGR growth realized over the last 15+ years, drove many optical industry innovations in an effort to keep pace. In addition, these innovations had to be economically viable driving Watts per Bit efficiencies.

The primary obstacle that separates backbones from other optical networks is simply the fiber distance. Fiber Distance drove the initial 10G and 40G NRZ deployments to install over 40 optical regen sites across the country. The first big break in technology was the arrival of the

40G coherent modem. It not only quadrupled speeds per channel, but just as importantly, it more than doubled the reach. Cost per bit dropped, and with regens reductions dropping into the teens, network costs were reduced significantly. Within a few years, Comcast was then able to augment longer spans with Raman amplifiers, move to the industry's first 100G coherent modem, and reduce Optical regen sites to 4.
The second backbone built in 2017-2018 was built utilizing colorless direct attach (CDA), flex grid ROADMS, with a flexible add/drop structure. Combining low loss fiber and 2nd Gen Raman amplifiers, Comcast is now able to cross every link without regen at a minimum of 200G and maximum of 800G. Additionally, all shelves were deployed with L-band amplifiers, so L-band channels can be added in the future without any service interruption.

Finally, the integration of an OTDR into the optical line system must be mentioned. When your network spans 18,000 route miles, knowing exactly where an outside plant problem occurs is invaluable with the hopes of reducing the Mean Time to Mitigate. (MTTM)

3.3. Capacity of the Backbone

The most valuable commodity in National Optical Networks is long haul fiber and the efficient use of optical spectrum management. Fiber longevity within the access and metro networks is also important, but typically more feasible and economical than acquiring and deploying national fiber where such options are very limited. Any fiber, if available, usually comes at an exorbitant cost.

Comcast has historically been a proponent of partnering with the Optical Industry to drive key technology innovations and continues to collaborate on next generation Optical capacity and efficiencies innovations.

Table 1: Fiber Capacity

Year	Line System	Modem	Fiber Capacity
2005	50GHz Fixed	10G	880GB
2009	50GHz Fixed	40G Coherent	3.5TB
2013	50GHz Fixed	100G Coherent	8.8TB
2018	CDA Flex	400G Flex	12.8 - 25.6TB*
2021	CDA Flex	800G Flex	16.8 - 33.6TB*

* C-band only (all fibers are L-band capable which should double capacity when modems are available)

3.4. Future of the Backbone

There is no denying the need for long haul networks to interconnect metro networks and peers, both in footprint and offnet. Even with the overall CAGR slowdowns after the Pandemic, the need to scale efficiently and economically will not disappear anytime in near the future. Moving forward, it is critical to drive software, tools, automation and integration in addition to the speeds and feeds of the typical hardware innovations. It is important that as an Industry we continue to partner and drive these Optical technology evolutions in support of future network products and services.

3.5. Purpose Built Networks

Comcast traffic contains 2 basic types. The majority (>95%) is very predictable internet destined traffic between the metro aggregation shelf and the edge. The smaller portion is very unpredictable traffic that has varied source and destination points. The large portion emphasizes the need for lowest cost per bit, while the other demands flexibility over cost. A point-to-point network of a flexible CDA line system with state-of-the-art modems would really push capacity up, and costs down, but with almost no flexibility. Where a CDC network can move traffic between any location, but at a lower capacity, and at a higher initial cost. Since the original Backbone was full and reaching end of life, a replacement was needed. Comcast utilized the luxury of having more than 1 fiber pair and decided to build 2 purpose built networks instead of one Swiss Army knife of a network. This allows the cost-effective network to grow quickly, while preserving the life of the costlier one.

3.6. Power Consumption

Power consumption is an industry wide concern. However, this is amplified in peering locations (BB Core sites) specifically. Peering locations are needed to interconnect with other service providers, and they are independently owned by facility managers. This forces Comcast and other service providers to rent space and power. Network growth has forced these locations to fill and make space and power very scarce. If a space fills, a migration to a new facility is needed. To avoid (or delay) the high cost and complexity of moving a core peering location, driving power consumption down is a must. Comcast has seen watts per optical bit fall from 7.3 in 2005 to below 0.5 today. Unfortunately, growth has outpaced the reduction in power consumption, so this continues to be a hot topic.

3.7. IPoDWDM and Alien Waves

IP over DWDM, primarily led by ZR+ plugs, is getting serious testing in Comcast. Although they are not as applicable in the Backbone as the metros, there is still some key use cases. The Comcast Backbone network tried this in 2006-2007 with a 40G optic being placed in our core routers. Engineering limitations were understood and overcome, but operational limitations ultimately ended the experiment. The same operational short comings were identical to alien wave issues. Comcast continues to seek operational solutions both internal and external in hopes of taking advantage of the possible cost reductions of alien transponders and plugs.

3.8. Automation

Turn up of new bandwidth and services across large national networks can seem like it happens at a snail's pace. Equipment and resources tend to be needed in several locations, often very far apart. Planning, deploying, testing, and turn up requires a lot of time. This time is stretched even longer as they sit on long waiting lists seen in very large networks. Comcast has been able to streamline and automate some of the predictable, business as usual (BAU), bandwidth growth, but much more is needed. Commercial and wave services will, due to their unique and custom designs, amplify this need. Comcast continues to drive automation from planning to turn up and utilize more flexible equipment to drive down equipment installation times.

4. The Metro Network

Designed to support multiple lines of business (LOB) with a variety of bandwidth needs and traffic patterns, however, to say designed is a bit of a misnomer. This would imply the metro was conceived and deployed from day one in a greenfield manner to meet all the needs of today. What had actually happened is that over the last 25 years these networks came to be through mergers of smaller cable television (CATV) providers into multiple system operators (MSO). There were technology evolutions from passive transport systems with amplifiers (AMP) and multiplexers (MUX) to what is present today with reconfigurable add, drop multiplexers (ROADM) colorless add / drops, and flex grid technology. The type of data and rate has also dramatically changed. The earliest deployments were to support video to areas that were beyond the reach of long coaxial (COAX) trunk runs, followed by hybrid fiber-coaxial (HFC) and 1 gigabit ethernet (GbE) commercial customers. Today 100G to 800G wavelengths are the norm! The more than 60 metro networks at Comcast have now either been upgraded or planned to go to colorless, dispersion less designs with flex grid ROADMs optimized for coherent wavelengths. Additional paths have been added between sites (referred to as degrees) for resiliency and latency reduction as well. The Metro networks are unique in that they are intentionally not purpose built to support a single demand but are LOB agnostic. The Swiss Army Knife of transport at Comcast!

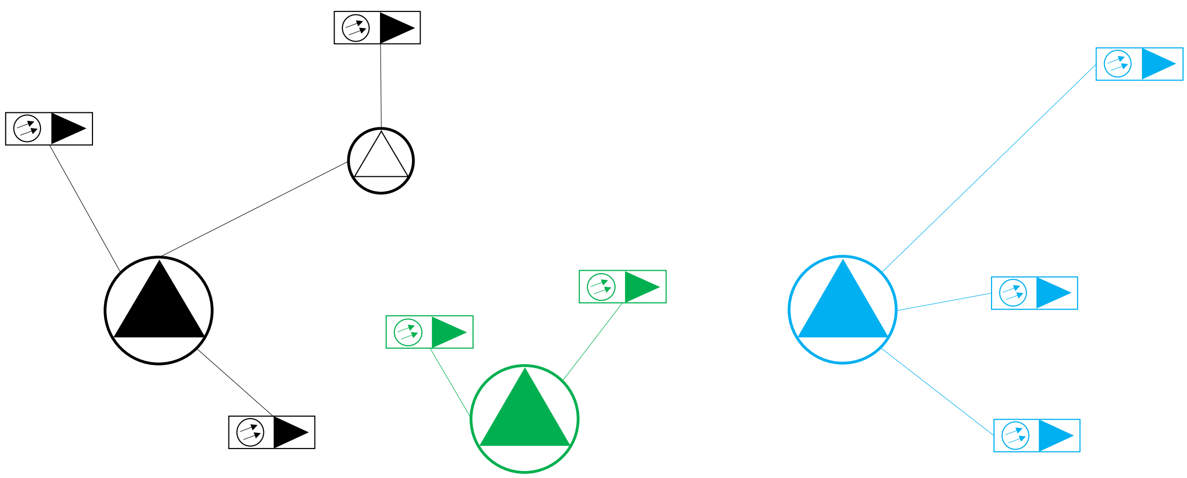

Figure 5 - The beginning of the Metro network. 3 independent CATV providers

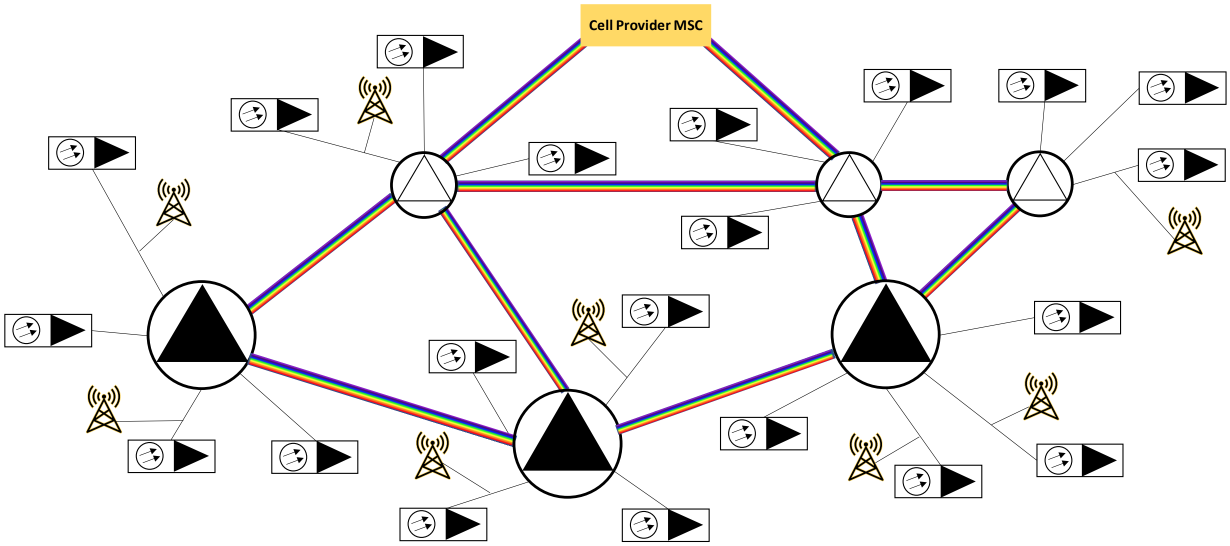

Figure 6 - Metro network today. One MSO. Flex grid and ROADMs. The Swiss Army Knife

4.1. The 28 CRANs of Comcast

The largest consumer of metro transport bandwidth are the Comcast regional area networks (CRAN). These networks are the core of the internet protocol (IP) infrastructure consisting of a pair of large aggregate routers (AR), residential U routers (RUR) and commercial super U routers (SUR). Each edge site is U ring with a pair of RURs with one router connecting to one market AR and the other connecting on diverse path to the other AR. If there are commercial services at the site, the same topology is applicable but with SURs in addition to the RURs. Each CRAN has two ARs which are connected to two backbone connections on diverse routes. This topology ensures failures of equipment or fiber will not impact customers. They are called U rings because the shape of the letter U describes the topology. The top of each leg of the U is an AR location and the bottom being the RUR / SUR site.

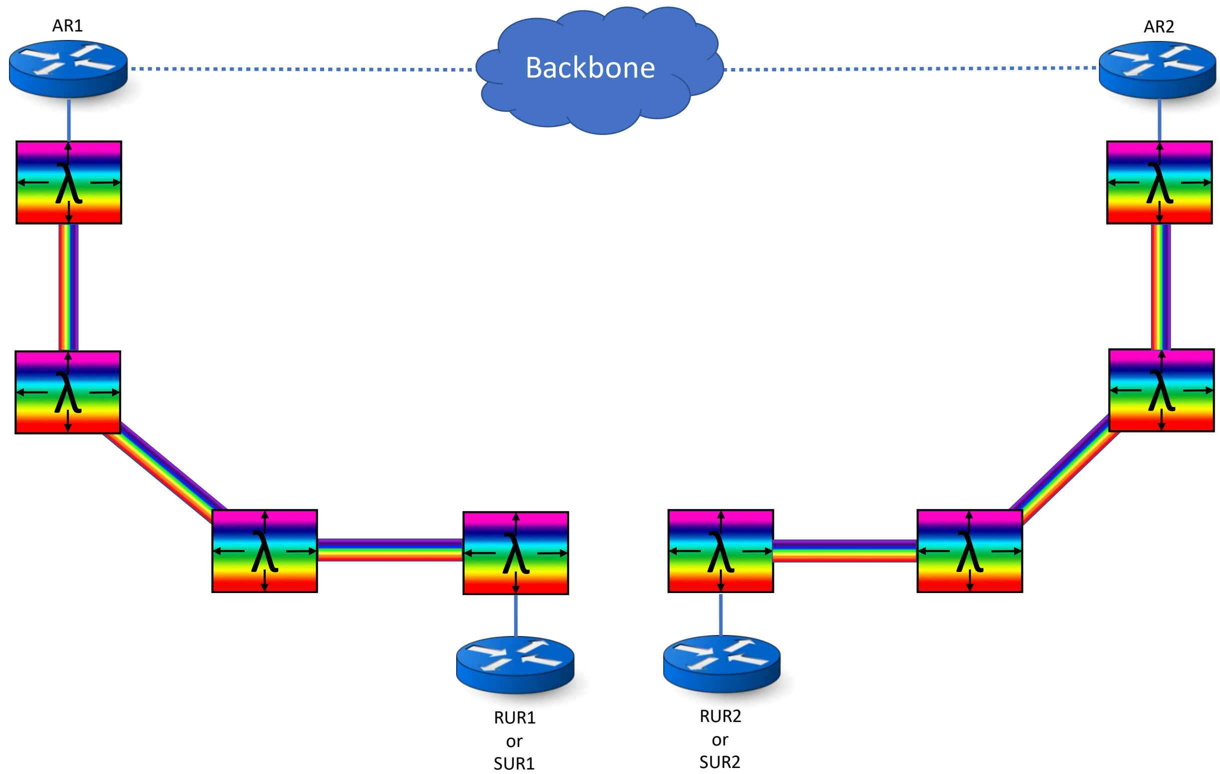

Figure 7 - U Ring Topology

4.1.1. In the beginning. Video on demand (VOD)

The Metro Transport Network tends to evolve to meet the needs of the routed network. The original iteration was dense wave division multiplexing (DWDM) and coarse wave division multiplexing (CWDM) uni-directional 10GbE optics placed directly in routers with MUXs and AMPs as needed. This was to support VOD and was also the birth of the U ring at Comcast. The routers we aptly named U routers (UR). A major difference in the topology compared to today is that that these URs to support VOD were configured hop by hop. Each router was connected in a chain. Since the circuit never had to go more than one span this allowed the use of point to point amplified optical links referred to as passive transport

4.1.2. Next up! High-speed internet (HSI).

If all routers need the same amount of capacity the hop-by-hop VOD architecture worked fine. With the introduction of high-speed internet (HSI) not only did the traffic need to become bi-directional but also the capacity needs of the routers began to become more fluid based on consumption of the product. For example, if there were five routers in the leg of a U all at 10GbE of demand and the fourth router in the chain grew and needed 20GbE then all the spans between the AR and the fourth router need to be augmented to 20GbE. This applied to both sides of the U. This was not a scalable model to throw away bandwidth along the way to reach a downstream router.

This led to the installation of active transport with ROADMs which allowed each router to get a direct connection to the ARs bypassing sites along the way optically. This first generation of transport had dispersion compensation, fixed grid wave selective switches (WSS), fixed grid MUXs and was optimized for 10GbE. Depending on the era this equipment could be either 50 gigahertz (GHz) or 100GHz spaced. This spacing is not critical for the 10G but would play a role in future evolutions.

As ROADMs were being installed and traffic could be steered at site optically it made sense to start breaking up large rings into smaller sub rings by building degrees between sites and including more remote sites into the network. These new shorter paths began to open the door for commercial services like cell backhaul (CBH).

4.1.3. 10 x the bandwidth and then some!

The next substantial change was to 100GbE interfaces on routers and the 100G coherent wavelengths required on transport. This was a huge leap! Fortunately, the metro transport was able to carry this traffic without upgrading the line system. A typical 100G coherent wavelength is ~37.5GHz wide. With most of the technology at this time being deployed with 50GHz spaced ROADMs and MUXs the 100G wavelength fit without issue. In fact, up to a 200G wavelength can fit in these networks.

At this point the transport was one step of ahead of the routers. With 400G line rates right around the corner and the router interfaces to follow the transport needed to maintain this lead and be ready for 400G as soon as possible.

In an interesting twist some of the oldest ROADM transport networks were the first ones capable of supporting 400G. As mentioned above both 50GHz and 100GHz systems were deployed with the 100GHz typically the first generation. With a 400G wavelength being ~75GHz wide it fit in the oldest technology deployed allowing another entire evolution of bandwidth increases to be deployed without an upgrade to the photonics. This is however the end of the road. 800G is already being deployed and with its ~112GHz wide wavelength it can only be deployed using flex ROADMs and a colorless add / drop structure. Given the fixed 50GHz systems are not capable of 400G and the networks that can are the oldest, the standard has been set to only deploy new networks as colorless and Flex grid to ensure any size future wavelength would be supported. Having also learned that there can never be enough available capacity a decision has been made to design all new networks to be C+ L band capable as well. Adding the L band doubles the capacity of the fiber! A final benefit to deploying all colorless networks going forward is the availability of Layer 0 (L0) control plane. This sets the stage for further resiliency by offering protection and restoration at the optical level.

4.2. Optimizing further - Capacity vs. Connectivity

Now the network is well prepared for bandwidth expansion by being able to support any width wavelength and the addition of L band has doubled the capacity. The construction of new degrees and creation of sub rings improves resiliency, latency and provides new routes for commercial customers. This architecture is the new performance baseline. Can it be optimized further though?

Of the traffic on the CRANs the largest and most predictable is residential. RUR U rings are high capacity with many being over 1 terabit (Tb). Although the physical flow of the traffic is passing through several ROADMs the logical circuit looks like point-to-point. With the large and predictable growth of these U rings does it make sense to go back to the original VOD architecture and build large, simple point to point networks? The answer might be yes, but not exactly.

4.2.1. Internet protocol over dense wave division multiplexing (IPoDWDM)

Much of the focus on technology advancements has been on creating larger wavelengths that consume less spectrum, however the transport is now outpacing the IP platforms with up 1.2TB line rates while the routers are just starting to adopt 400GbE. To best utilize and align the IP and optical platforms the answer could be IPoDWDM. Simply put, this is making the source of the DWDM wavelength a pluggable optic and placing it directly in the router, thus eliminating the transponder. Sound familiar? This is exactly how the first generation of transport to support VOD was deployed as described previously. This may make sense again but with several important caveats.

The optic that has been developed for this application is called a ZR+. It has a 400GbE rate and like other 400G transport interfaces has a spectral width of ~75GHz. It does have roadblocks to overcome before being deploying at scale.

4.2.1.1. Form factor and transmit power

The first two issues are related. They are form factor and optical transmit power. There are two current versions of the optic. C form factor pluggable (CFP) and quad small form factor pluggable (QSFP). The form factor is directly related to optical transmit power. The CFP is larger and can house the components required to achieve ~0 decibel-milliwatts (dBm) optical transmit power commonly seen in transponders. However, IP host platforms that use CFP are no longer common at Comcast. They have been replaced by devices that use the much smaller QSFP optic. The issue is given the smaller size of the QSFP optic it can only produce ~ -10dBm transmit power with the current technology. This is as much as 14dBm lower than transponders! To allow the high-power transponder and low power optic to coexist in the same ROADM they need to have a similar input. There are two options. Either turn down the transponder and greatly reduce its performance or amplify the low power ZR+ optic. Neither of the solutions are acceptable in a brownfield environment. Fortunately, both problems are remedied by high power ZR+ optics in a QSFP form factor becoming available in late 2022. In fact, Comcast is collaborating with an optical partner to be one of the first to get this optic in the lab in advance of the general availability release.

4.2.1.2. Reach and ROADM limitations

The third issue is the reach of the optic. The ZR+ optic not only has distance limitations but every ROADM it passes through adds penalties and decreases the reach. In the end all the ROADMs that were added to increase flexibility in steering traffic and reducing latency now could be preventing the use of these optics. One solution being explored is to build ROADM bypass express paths to decrease the number of ROADMs and increase the reach.

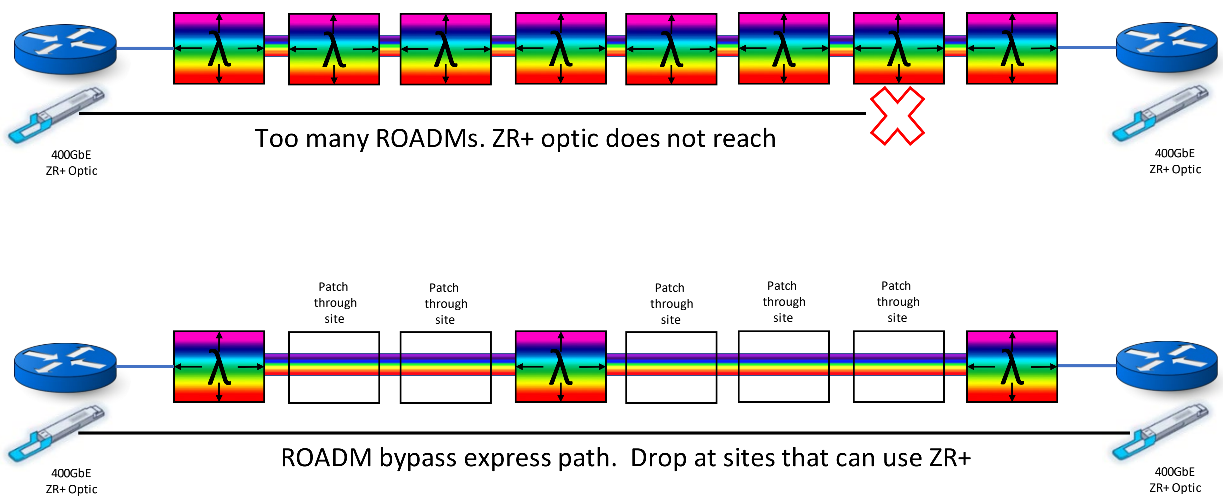

Figure 8 - ROADM impact on 400GbE ZR+ optics

4.2.1.3. *Operational support*

Finally, the largest issue is the operational and engineering support of these optics. By disaggregating the DWDM source from the optical line system the topology connection between the DWDM wavelength source and add / drop structure is no longer known. It is not possible to know what the pluggable is connected to. This splitting of the optical network into two domains essentially makes the link an alien. This was the case also in the VOD era however with only a handful of links to support per market it was manageable with static documentation. Today with thousands of links this is no longer possible. Comcast engineers and software developers and the industry in general are currently working on solutions to operationalize aliens which would make the use of ZR+ optics possible.

If these challenges can be overcome ZR+ optics do not replace transponders everywhere nor do the ROADM bypass express paths displace the multiple path options of many meshed degrees. Instead, it becomes yet another layer, optimizes further and solidifies the Metro Optical Networks role of being the do it all, Swiss Army Knife at Comcast.

4.3. Leaveraging access and metro together to support commercial services

Where the metro ends the access starts. Unlike other networks that are completely independent and do not converge the metro and access are different. They can and do and have a symbiotic relationship.

Comcast combines the metro optical network, optical transport network (OTN) tails over the access fiber, small optical shelves at customer sites called network terminating equipment (NTE) and dedicated commercial shelves at hub sites called wave integration shelves (WIS) to support multiple types of commercial business.

As seen in the above network evolution diagrams the metro transport has become more meshed and better able to support an any to any traffic pattern. The access has also gone through evolutions but what has stayed the same is the fact that it stops at the hub site. What has been found is that most commercial customers can connect to the access and be brought back to a hub site where metro is present, but how can they be connected to each other if their fiber terminates at different hubs? Or in the case of ethernet dedicated internet (EDI) how do they get to the router that provides internet access? This is where the access and metro can work together by using OTN tails.

OTN is taking an Ethernet (or other protocol) payload and encapsulating it in a wrapper to be carried across an optical transport network. It is completely transparent, and a standard governed by International Telecommunication Union (ITU) G 7.09. An OTN tail is transmitting that wrapped payload over dark fiber or a channel on a MUX over fiber. Using this tail and transponders (or Muxponders) in the NTE and WIS facing each other can extend the metro core over the access without using ROADMs or even AMPs.

This tail is then transmitted over the metro by a second card in the WIS with an optical transport unit (OTU) client to client connection to the card facing the NTE. This back-to-back configuration allows a DWDM OTN signal to face the NTE and the add / drop structure of the metro. The client-to-client connections have been made in the past as ethernet where a similar design was used for regens. In this case, keeping them OTN so that the circuit remains transparent from end to end has been selected as the standard. There is no conversion to ethernet and back to OTN. One advantage of this is that a trail trace identifier (TTI) remains intact from end to end even if the circuit crosses to another network or even to another vendor. You can transmit a trace on the A end and read it on the Z end. This is particularly useful for trouble shooting and topology verification.

This Unified Optical Architecture to Support Wavelength and High Bandwidth Ethernet Services is now a common way the metro and access work together at Comcast and is fully detailed in another paper being presented at the SCTE Cable-TEC Expo 22. [1]

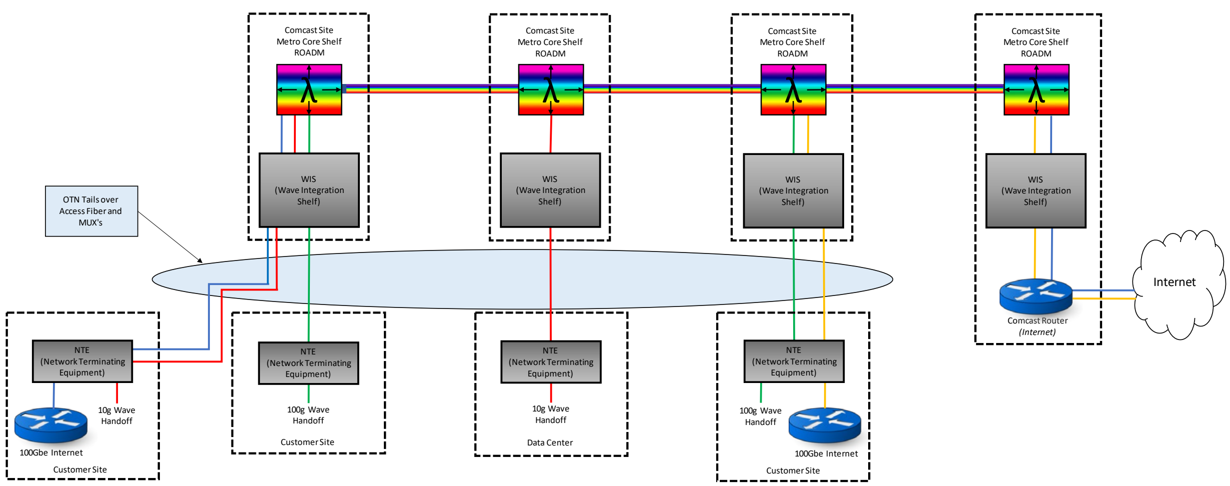

Figure 9- Metro and Access used together to support commercial customers

5. The Access Network

The Comcast Access network consists of a myriad set of architectures some organic, some acquired and many augmentations over the years. Some of these architectures are the traditional 1310nm/1550nm analog and digital/analog return nodes connected to traditional CMTS. Others are DWDM QAM Overlay solutions and yet others are the latest distributed access architecture (DAA) based nodes.

Comcast has pioneered the DAA that has virtualized the CMTS functionality and separated the PHY layer from the CMTS and distributed it out into nodes in the field. These are now called the Remote PHY Device (RPD) nodes. In Comcastvirtually all of the traditional nodes and CMTSs will be replaced by the DAA in the next couple of years. For this reason, this paper will concentrate of the DAA system but will address other systems along the way as needed.

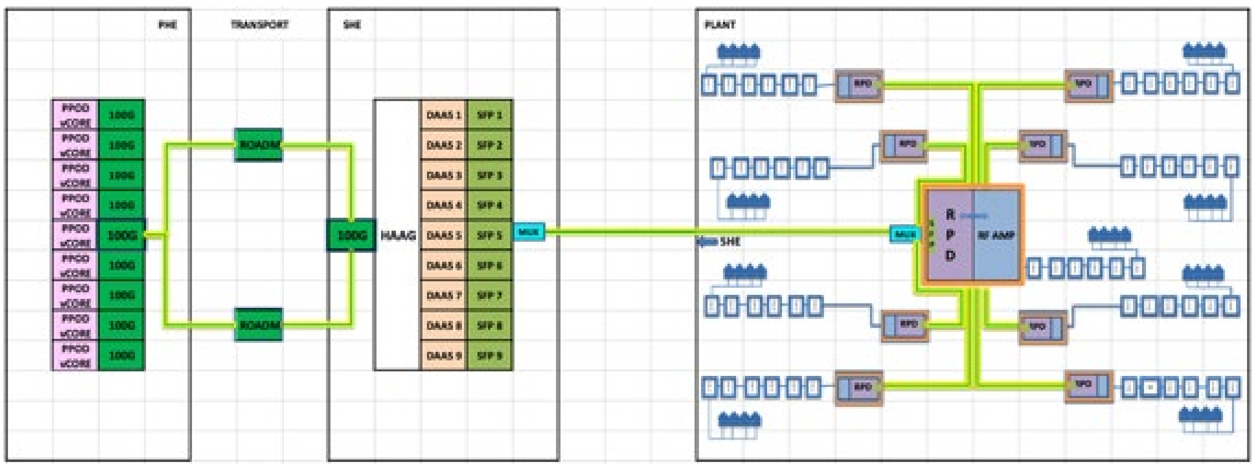

Figure 10 – Illustrating the DAA Network of Comcast

Shown above is a simplified version of the DAA system. vCMTS cores from the Primary headend are connected to DAA switches (DAAS) in the secondaries. Note here that these connections are typically done via the Comcast's own Metro networks described in the above sections. While the Primary to Secondary connections are based on 100G or higher capacity coherent optics, the connections from the DAAS on wards in the access domain are always 10Gbps DWDM based signals. The aggregating switches called HAAGs typically combine US signals and de-combine DS signals to the respective DAAS ports.

Once out of the DAAS, both the US and DS 10G wavelengths are multiplexed on a single strand of fiber and sent over into the field, where they are demultiplexed at an OSP enclosure location. For there, dual strands of fiber carry the 10Gbps signals over to the various RPDs in the immediate location. The distance from the Secondary to the field OSP mux is called the trunk fiber, while the fiber from the OSP Mux location to the various RPDs is called the distribution fiber. On the trunk fiber Comcast has a 100GHz ITU WL Plan standard that has adjacent odd and even pairs of ITU channels serve each node beginning at ITU 61/60 pair to ITU 15/14 pair. Thus, all in all, 24 pairs of wavelengths are available on each trunk fiber capable of serving 24 different remote devices.

While distribution fibers are new fibers sometimes put in explicitly for newer nodes, the trunk fibers are in extremely short supply as they were installed decades back and augmented over the many years to serve nodes that sprung up as population centers sprung up. It is for this reason that all DAA trunk fibers are bi-directional as that relieves pressure on the trunk fibers and promotes fiber efficiency across the organization.

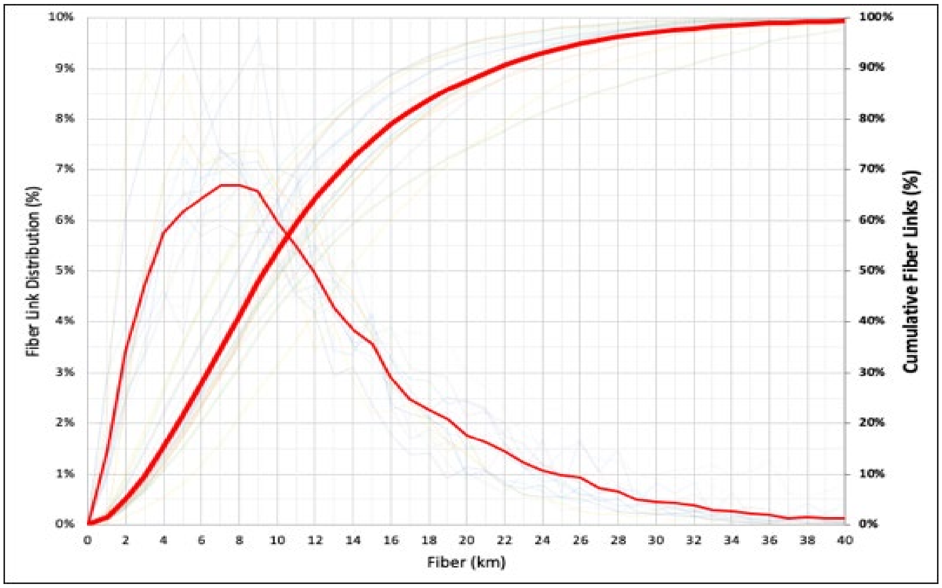

Figure 11 – Typical Secondary to Node distances in the Access Plant

Generally, primary headends also supply node signals directly from their location to the local nodes as well. It is for this reason that the distance between a headend and the node is around 10km, but a fair number of nodes extend up to 40km. This is illustrated in a Comcast survey of optical links from a couple of year back [2]. In general, distribution links are on average 0.5km and as discussed construction crews may build this fiber for some of the newer nodes which arise due to node splits or fiber extensions from the OSP Mux. For this reason, there are multiple fibers connecting OSP Mux to fiber nodes. It must be mentioned here that most of access field infrastructure requires Industrial-Temp operation, this means that all optics actives and passives must be qualified up to 85C. Therefore, the field trunk fibers are standardized to 100GHz spacing utilizing thin-film based optical filtering and all the SFPs used in remote devices are Industrial-Temp qualified tunable SFPs.

Our fiber to the home (FTTH) offering comprises initial RFoG deployments but that is expected to move over to 10G EPON offerings for all the newer builds. The 10G EPON DS is at 1577nm while its US is at 1270nm. With the advent of SFP based OLT and ONUs, it is expected that current RPD nodes that have the capability to received 10G signals will receive 10G signals from their respective DAAS ports and convert via the SFP based OLT to the 10G EPON standard and serve the FTTH customer base. A lot of attention has been paid to enable the convergence of conventional RPD devices that can serve residences via standard HFC plant AND the FTTH customers thru the remote OLT (also called virtualized broadband network gateway or vBNG).

The vBNG line puts out 1577nm in the DS and accepts 1270nm in the US one single fiber and thus enables centralized or decentralized splitting as it serves homes via single fiber stands.

5.1. Converging the Access Network

The previous has been a simple description of Comcast network, but while the direction of digitization and distributed access is clear Comcast still has various older architectures comprising analog DS wavelengths and potentially analog or digital US wavelengths and they must all be accommodated as the DAA transition takes place. In addition, one interesting aspect of access plant is that most of the businesses that require separate optical wavelengths are also interspersed among the residential customers. The ability to have analog and digital WLs coexist on the same fiber has been described in earlier papers and is summarized here. Briefly, full spectrum analog WLs in the C-Band require a comprehensive WL plan to minimize optical non-linearities such as 4WM, XPM, SRS and overcome dispersion and imperfections of the optical passives. For this reason, they are not on a uniform spacing, but rather non-uniformly spaced. In Comcast, we have carved out Analog and 10G WLs in DS and, 10G and DRT WLs on the US on the same single strand of fiber.

Our current business base is a mix of single and dual fiber CWDM and DWDM connections that reflects the years of organic and acquired growth. Many of our current business customers that require fiber to their premises require 1G to 10G services (discussed in the Metro section of this paper) and could be in the vicinity of existing OSP Mux enclosures. For a go forward plan, since these are digital signals, these customers can be served by the same strand of trunk fiber as would connect analog and digital nodes. This is an important way to relieve stress on our fiber assets by enabling a more effective use of our fiber assets. In addition, it enables a common way to continuously and pervasively monitor optical assets which is described in a subsequent section.

But what is interesting in the figure below is that we have now been able to incorporate Coherent optics in the access plant in much the same way as direct detect has been deployed. This type of convergence requires game changing technology of bidirectional dual laser pluggable optics conceived and implemented in Comcast in recent years and is described in this section.

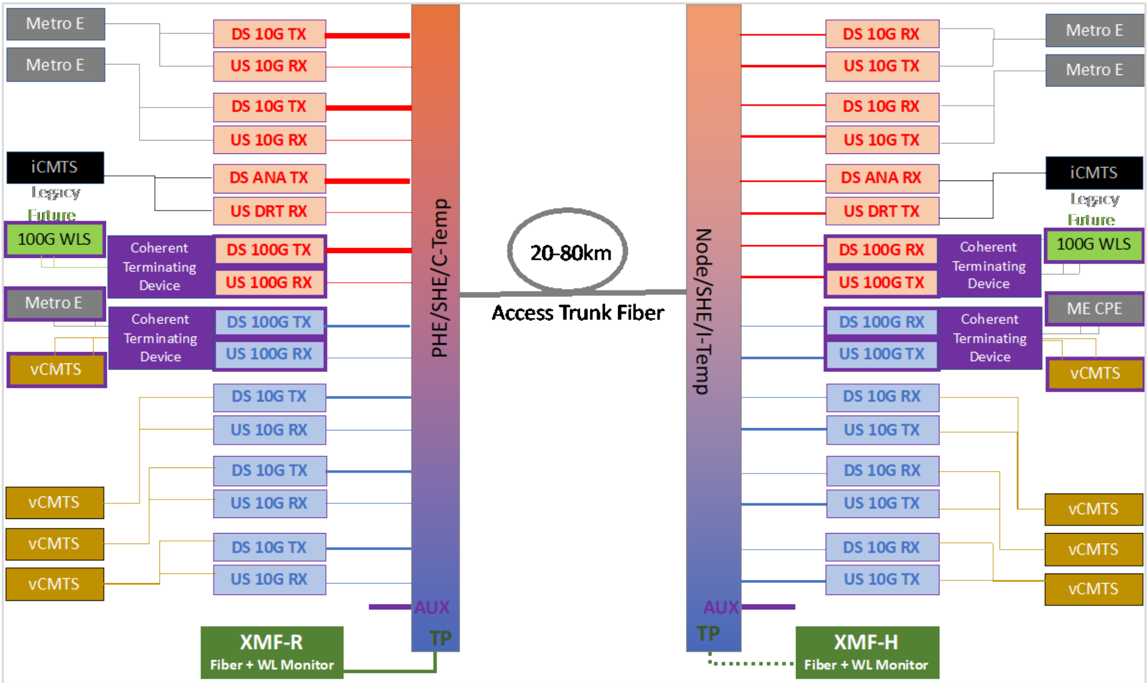

Figure 12 – Illustrating the Converged Comcast Access Network

Since ROADM cascades are not used in our DAA access networks due to modest fiber lengths, the design of high-capacity (100G – 800G) solutions could be elegant since there is no degradation due to these devices. Furthermore, these links can support bidirectional transmission quite easily since there are no non-reciprocal elements in the fiber path. Indeed, the relaxed relatively modest fiber reach required in DAA enables us to innovate on optical devices as well as reduce cost, oftentimes by innovating on the trailing edge of technology, which is discussed next.

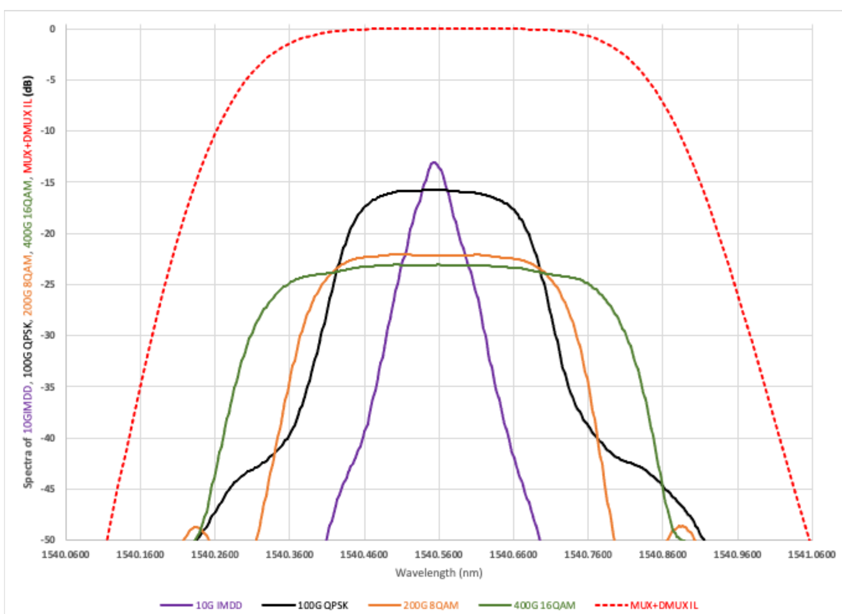

Figure 13 – Converging Direct detect and Coherent bi-Directional Wavelengths

Presented above is a measured plot of the various direct detect and coherent optical signals overlaid with the optical passband of typical filters being deployed in Comcast [3]. It is seen here that 10Gbps direct detect as well as 100Gbps QPSK Coherent all the way to 400Gbps 16QAM Coherent optical signals fit within the optical passband of deploying optical filters. Although there is little concern for data rates up to 400Gbps, we still need to further investigate the case for 800Gbps where, even with higher QAM modulation orders and with higher OSNR requirements, it is expected that the occupied bandwidth will start to encroach into the WDM passband and impose power penalties that need to be considered, particularly over temperature. More generally, a skillful use of optical passives express/upgrade ports with a contiguous C-Band can easily be dedicated for higher capacity/bandwidth modulations or for flex grid type applications.

5.1. Dual Laser Bidirectional Coherent Transmission

Coherent optics relies on having a laser at the transmitting end send out phase and amplitude information in the form of QAM constellations to a receiver. At the receiver, the incoming signal is ‘beat’ up against a laser that has in effect the exact wavelength of the transmitting laser and it teases out the phase and amplitude information out. Since information is coded in phase and amplitude, as opposed to only amplitude in direct detect systems, the information carrying capacity of the coherent system is vastly superior to that of direct detect systems for the same amount of bandwidth available.

To reduce cost and also to ensure that the wavelengths are identical, most of the plugs available in the market use just one laser in each plug and split its light to do double duty, in effect to act as the encoding transmitters and the decoding receiver. This is illustrated in the top half of the picture below.

The industry has gotten quite good at this type of design and has successfully used it to provide 100s of Tbps over 1000s of kms in core and metro networks. Unfortunately, this approach always requires two separate fibers, one for downstream and another for upstream information. When a fiber can be fully loaded with data, say for 100s of Tbps, this approach seems reasonable. But when the fibers are unable to be fully loaded, like what happens in access networks, an approach of this kind is very inefficient with fibers and can cause severe fiber exhaustion.

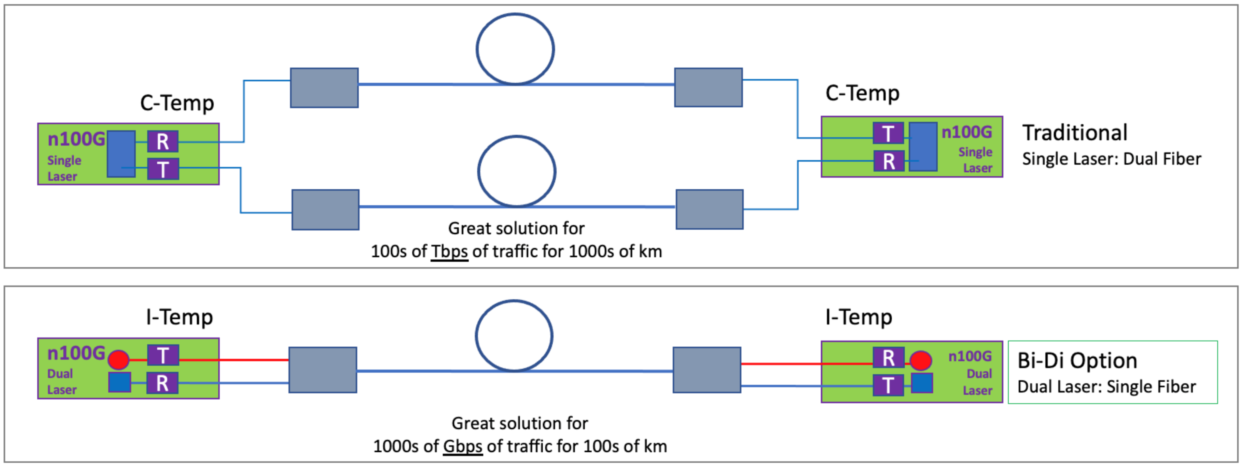

Figure 14 – Illustrating Bi-Directional Coherent systems

It is for this reason that a new approach had to be thought of for accommodating Coherent optics in the access domain and enable high-capacity wavelengths to feed our residential DAA systems and as standalone high-capacity wavelengths for our business customers.

We note here that the use of optical circulators is a valuable option and could potentially allow the same wavelength to traverse in opposite directions, but that approach can be prone to reflections in the fiber plant and also to fiber Rayleigh backscattering that could limit the link budget. A more robust solution is needed for wide bi-directional deployments that envision the use of existing fibers with other bi-directional signals already running on it.

Such a robust solution to this problem is to have a separate laser for the transmitter and one for the receiver at each end. Doing this would enable the two wavelengths to be multiplexed (or combined) together and transmit bidirectionally (Bi-Di) on the same single strand of fiber. In fact, many Coherent optical wavelengths could be multiplexed on the same fiber along with direct detect wavelengths thus leading to a converged system that was described in the previous section.

Interestingly, low earth orbit (LEO) satellite communications systems have also come up with the same requirement of dual laser bi-directional transmission. In free space, coherent optical modules communicate with adjacent satellite and hop signals from satellite to satellite to minimize the need for ground stations. But the vast speeds of the LEO satellites and their relative speeds to other satellites create the familiar doppler shift in frequency sufficient to thwart the use of a single laser design. So, in this effort of defining a dual laser plug, we have had the unusual conjoining of space requirements as well.

With this in mind, Comcast have specified dual laser bidirectional fully C-Band tunable industrial temperature CFP2 plugs for use in access networks much in line with the CableLabs point to point Coherent specification [4]. Current systems support 100Gbps rates with QPSK modulation at 32GBaud with a 25dB link budget, enough to support point-to-point links up to 100kms. As mentioned previously there are no ROADMs in the access plant and there are no non-reciprocal elements in the system, this allows for relaxation of certain specifications and needs for chromatic dispersion compensation that enable these new innovative devices to be cost effective as well.

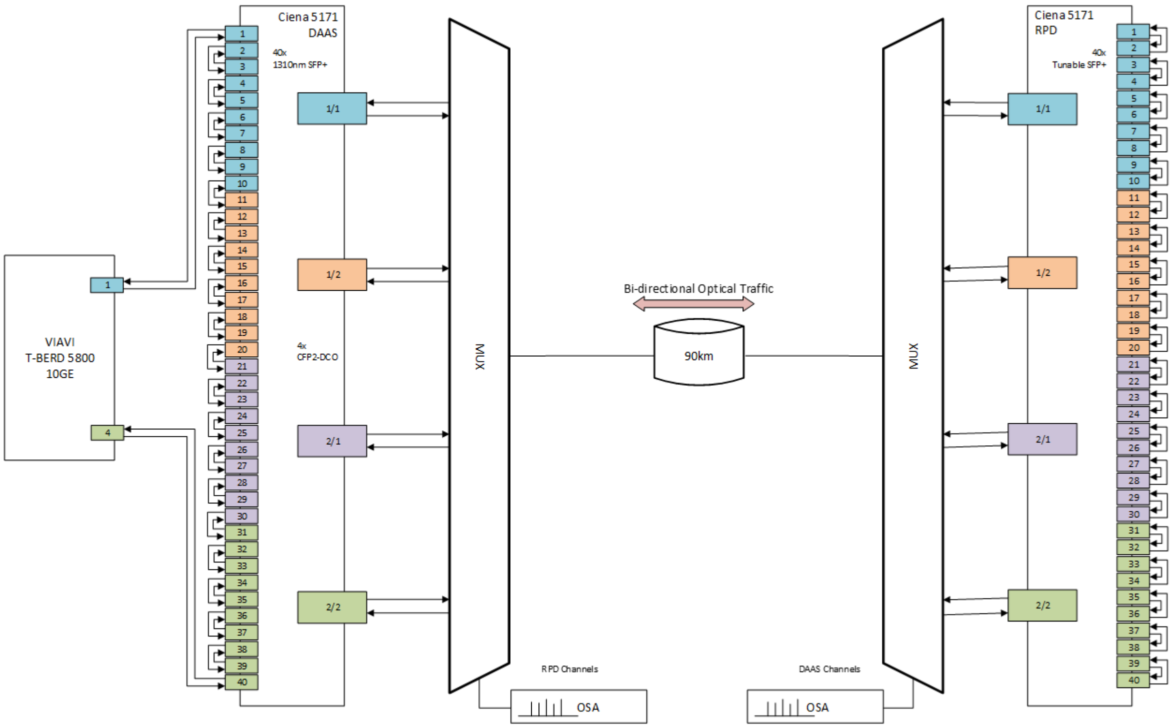

Figure 15 – Hybrid Loop testing the Coherent Optics in the Access Plant

While 100G coherent systems can exist on their own and support individual wavelength services for our commercial customers in the converged access architecture described in the previous section, a particularly important application of the dual laser bi-di system is to support DAA systems. Often these DAA systems have to operate out of smaller OTN or secondary cabinets that make direct metro transport system connections difficult due to limited critical infrastructure. And adding fiber to the trunk line is an expensive and time-consuming effort. In such cases, we have developed a system that takes in 10 of the 10Gbps streams and electronically converts to 100Gbps stream that modulates the CFP2 which is then multiplexed to a single fiber and sent over to the secondary. At the secondary the 100Gbps stream is split back to its constituent 10Gbps streams and serves DAA nodes from there.

To test this system for its robustness, we set up a hybrid loop test where 10Gbps signal was injected into one of the 10Gbps port that then entered the 100Gbps CFP2 which then traversed 90km of SMF fiber and was converted to its constituent 10Gbps stream and looped back via the CFP2 and traversed the 90km. In this way the traffic traversed 90km*2*40 = 7200km before closing the loop. In our tests this was done with zero errors and with latency that was predominantly dictated by the optical fiber time of flight.

The system designed and shown above can support up to 40 10Gbps streams each way that modulate 4 of the 100Gbps streams in each of the two locations. In effect, 6 such systems can be optically multiplexed together and serve up to 240 10Gbps Nodes all on a single strand of fiber. As we have said before these our efforts at innovations will be to extend to higher capacity systems cost effectively over the coming years.

5.2. Convergence Continuous Pervasive Monitoring

Our current business base is a mix of single and dual fiber CWDM and DWDM connections that reflects the years of organic and acquired growth. Many of our current business customers that require fiber to their premises require 1G to 10G services (discussed in the Metro section of this paper) and could be in the vicinity of existing OSP Mux enclosures. For a go forward plan, since these are digital signals, these customers can be served by the same strand of trunk fiber as would connect analog and digital nodes. This is an important way to relieve stress on our fiber assets by enabling a more effective use of our fiber assets. In addition, it enables a common way to continuously and pervasively monitor optical assets which is described in a subsequent paragraph.

In the figure below we have shown a converged system that is a combination of residential and commercial services, incorporating coherent and direct detect systems, all capable of running on a single fiber. Since all services can be multiplexed on a single fiber, the current series of optical passives in Comcast have consolidated test ports that provide access to view forward and return wavelengths of live links. Furthermore, the test ports also allow unrestricted 1611nm access for use of OTDRs to shoot over live links and provide fiber impairment and cut information.

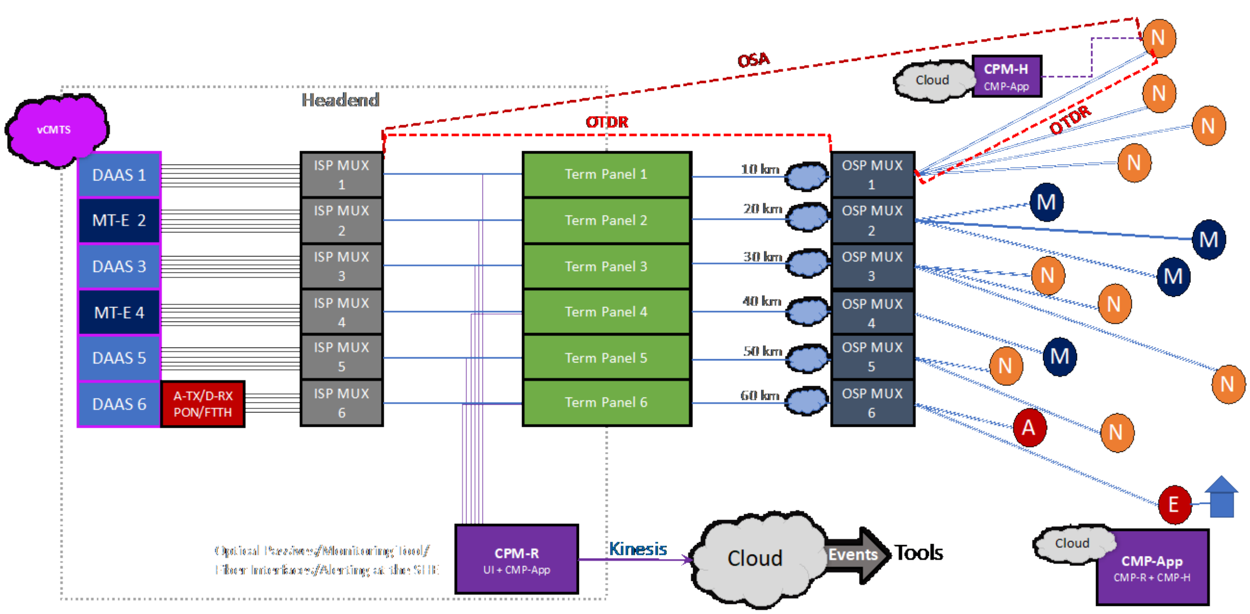

Figure 16 – Illustrating Continuous Pervasive Monitoring

Detailed description of the continuous and pervasive monitoring paradigm in Comcast was described in earlier SCTE papers. By way of summary, all optical passives are connected to a continuous monitor comprising an OSA and OTDR and an optical switch. This arrangement continuously monitors all connected fibers in a round robin fashion and generates fault information comprising either fiber impairments or individual wavelength impairments on average within 90 seconds. The entire information set is sent into the cloud and alarms and events are sent to fix agents.

To promote a common language between headend and field technicians, a version of the monitor is available as a handheld device that has the same cloud connectivity as the headend version. Connecting the headend unit or the hand-held unit or both to the cloud now enables the entire region, division or company tie in and view all events in real time and strive to achieve closure on important customers impacting issues. This is a real game changer with the net effect of providing continuous monitoring on residential, commercial, direct detect and coherent optical signals across the access network.

6. Conclusion

In this paper, we described for the first time an end-to-end view of our optical network including the core, metro, and access layers. A brief history of the organic and acquired properties and technologies at play in Comcast was described. At the core and metro, we have increased capacity with a move towards flexible 400G connections and continue to reduce latency and enhance reliability with an infrastructure that meshes color-less, direction-less, and contention-less reconfigurable multiplexers thru to each of our headends. At the access layer that connects these headends to customers and businesses, we discussed capacity increases with our move to all-digital fiber links and the distributed access architecture paradigm. The vision shared here is part of a larger commitment that leverages Comcast's technical and business expertise to deliver cutting edge and reliable services to our residential and commercial customers combining a robust and scalable core and metro network with a converging range of access networks ranging from wavelength services to traditional HFC to FTTH solutions.

Abbreviations

4WM	Four Wave Mixing
AMP	Amplifier
AP	access point
AR	Aggregate Router
bps	bits per second
FEC	forward error correction
HCF	Hollow core fiber
Hz	hertz
K	kelvin
SCTE	Society of Cable Telecommunications Engineers
AMP	amplifier
AR	aggregate router
CATV	cable television
CDA	Cororless Direct Attach
CBH	cell backhaul

CFP	C form-factor pluggable
CMTS	Cable Modem Termination System
COAX	coaxial
CRAN	Comcast regional area network
CWDM	course wave division multiplexing
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
dBm	decibel-milliwatts
DRT	Digital Return Transmitter
DS	Downstream
DWDM	dense wave division multiplexing
EDI	ethernet dedicated internet
FTTH	Fiber To The Home
G	gig
GbE	gigabit ethernet
GHz	gigahertz
HCF	Hollow Core Fiber
HFC	hybrid fiber-coaxial
HSI	high speed internet
IP	internet protocol
IPoDWDM	Internet protocol over dense wave division multiplexing
ITU	International Telecommunication Union
L0	layer 0
LOB	line of business
MEMS	micro-electromechanical system
MTTM	Mean Time To Mitigate
MSO	multiple system operators
MUX	multiplexer
NRZ	Non-Return-to-Zero
NTE	network terminating equipment
OLT	Optical Line Termination
OSP	OutSide Plant
OTDR	Optical Time Domain Reflectometry
OTN	optical transport network
OTU	optical transport unit
PHY	Physical layer
QAM	Quadrature Amplitude Modulation
QSFP	quad small form factor pluggable
RFoG	Radio Frequency (RF) over Glass
ROADM	reconfigurable optical add drop multiplexer

RPD	Remote PHY Device
RUR	residential U router
SRS	Stimulated Raman Scattering
SUR	super U router
TB	terabit
TTI	trail trace identifier
UR	U router
US	Upstream
vBNG	virtualized Broadband Network Gateway
vCMTS	Virtualized Cable Modem Termination System
VOD	video on demand
WIS	wave integration shelf
XPM	Cross Phase Modulation
WL	Wavelength
WSS	wave selective switch

Bibliography & References

[1] 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, Bob Gaydos, SCTE 2020
[2] Accelerating the Virtualization: Introducing Hybrid Fiber Shelf into the Mix, Venk Mutalik, Bob Gaydos, Dan Rice, Jorge Salinger, SCTE 2020
[3] Approaching the Universal Speed Limit: Introducing Hollow Core Fibers for Low Latency and High Capacity, Venk Mutalik, Amarildo Vieira, Bob Gaydos, Elad Nafshi
[4] Cable Labs point to point Coherent Optics specification

