Unified Optical Architecture to Support Wavelength and High Bandwidth Ethernet Services

A Technical Paper prepared for SCTE by

Stephen Ruppa
Sr Principal Engineer
Comcast
2001 York Rd, Oak Brook, IL 60523
847-489-7953
Stephen_ruppa@comcast.com

Table of Contents

Title **Page Number**

1. Introduction 3
2. Network evolution 3
    2.1. Many small CATV providers 3
    2.2. Acquistions and HFC 4
    2.3. Networks merge, advanced services, cell backhaul and ROADMS 5
    2.4. All ROADMs, Flex Technology and one vendor 5
3. Products 6
    3.1. Wavelength Services 6
    3.2. High Bandwidth Ethernet Services 7
4. The Building Blocks 7
    4.1. Network Terminating Equipment 7
        4.1.1. NTE Remote Management 8
    4.2. Wave Integration Shelf 9
        4.2.1. Capacity Planning 10
5. Leaveraging Access and Metro together: OTN Tails 11
    5.1. Access MUXs to support 400 Gbps 12
6. Complete end to end archicture 13
7. Benefits 14
    7.1. Engineering and Service Delivery 14
    7.2. Procurement and Deployment Engineering 14
    7.3. Operations and Service assurance 15
8. Conclusion 15

Abbreviations 15

List of Figures

Title **Page Number**

Figure 1- Three independent CATV providers 4
Figure 2- Two neighboring networks continue to grow 4
Figure 3- Single network with CBH, redundant hubs, some ROADMs and various transport vendors 5
Figure 4 - Meshed Metro Optical Networks with ROADMs and Flex Technology 6
Figure 5 - LR to DWDM conversion 8
Figure 6 - Management IP for NTE being distributed using GCC0 9
Figure 7 - Muxponder. All circuits share A and Z termination 10
Figure 8 - OTN aggregation allowing different A and Z termination 10
Figure 9- Recolor with back-to-back cards in WIS 12
Figure 10- Spectrum consumption by 200 Gbps, 400 Gbps and 800 Gbps Wavelengths 13
Figure 11 - Unified Optical Architecture 14

1. Introduction

The optical networks present today for cable television (CATV) systems look nothing like their predecessors first deployed with the birth of fiber nodes. Over the years, as technology evolved, these networks were augmented and upgraded as customer expectations also grew requiring a more reliable, higher-quality experience. As many smaller CATV companies were operating with various technologies and architectures, mergers and acquisitions continued to take place creating larger multiple system operators (MSOs). Having this variety was not inherently an issue with the residential services the networks were designed to support, which were primarily one-way plants providing linear video. With the advent of two-way plants and high-speed Internet (HSI) access, it was still not an issue to have separate optical networks. Other than backbone, services either originated or terminated on devices in a local headend. There was not a requirement for an optical service to continue beyond this termination. Even as commercial services were originally productized to support cell backhaul (CBH), and soon after retail metro Ethernet (ME) customers, the same access fiber infrastructure used to support hybrid fiber/coax (HFC) architectures fit the bill. Once the fiber circuit reached a headend, if it needed to continue to another location it became an Ethernet service and leveraged the routed Internet Protocol (IP) network. The introduction of two new commercial offerings, wavelength and high bandwidth Ethernet services, has made using the existing architecture a challenge.

This paper demonstrates how the metro optical network, optical transport network (OTN) tails over the access fiber, small optical shelves at customer sites, and dedicated commercial shelves at hub sites can be combined to support multiple types of commercial business. This enables new features like remote management, performance monitoring (PM) data, alarming, and a full end-to-end circuit view including the customer site. In addition to these operational benefits there are other efficiencies seen by using the same hardware and software as the rest of the core network.

2. Network evolution

To understand what the new architecture means it is important to look in a little more detail as to where the metro and access networks started, how the current state was reached, and where they are seen going.

2.1. Many small CATV providers

Some of the earliest deployments of fiber in CATV were to support one-way, video-only nodes, typically using low-count fiber cables feeding large pockets of customers from a headend, and areas that were beyond the reach of long coaxial trunk runs. These were single-threaded and there were some, but not many, hub sites. At this time it was not uncommon to have several small CATV operators with shared borders. Figure 1 represents this with three independent providers and a limited amount of fiber, hubs, and fiber nodes.

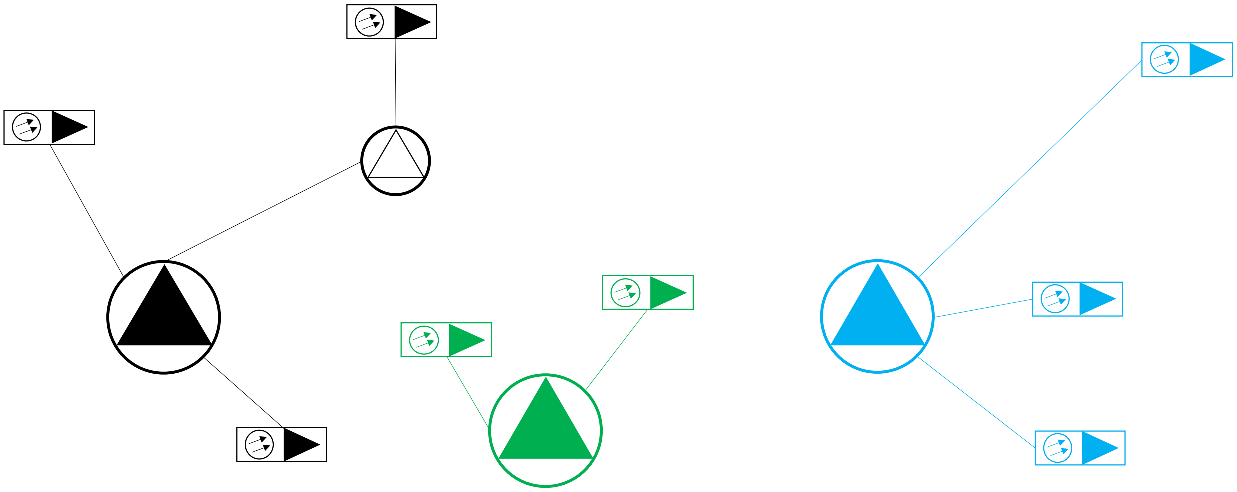

Figure 1- Three independent CATV providers

2.2. Acquistions and HFC

With the upgrades to two-way plant and HFC architectures many new nodes were added as operators began creating hub sites with diverse routes and transport. These were mainly passive, point to point systems with amplifiers (AMP) and multiplexers (MUX) but no reconfigurable optical add / drop multiplexers (ROADM). Mergers and acquisitions continued over the years resulting in larger cable companies. This is shown in Figure 2. The three previously independent CATV networks illustrated in Figure 1 have grown and merged and are now owned by a single operator.

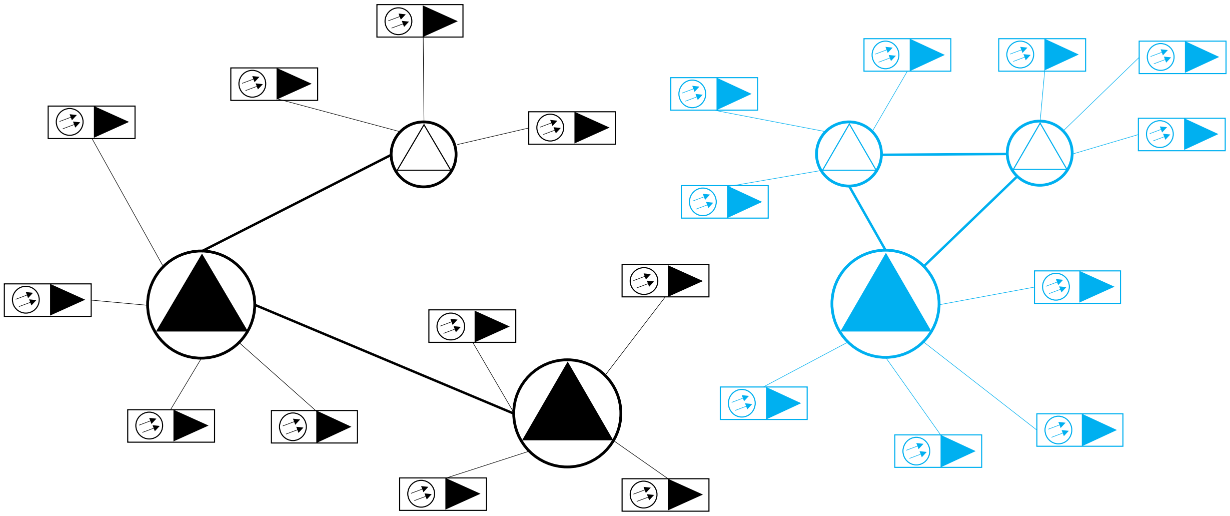

Figure 2- Two neighboring networks continue to grow

2.3. Networks merge, advanced services, cell backhaul and ROADMS

With advanced service deployments like HSI and video on demand (VOD) in full swing, the idea of placing routers in the hubs with redundant feeds to regional aggregate routers and video content servers became common. Systems continued to merge and some of the headends became hub sites. Each previously separate network was connected with fiber and active transport. ROADMs become more prevalent. This is also when the first commercial services over fiber were introduced, consisting primarily of CBH connecting cell towers to carrier's mobile switch centers (MSC). This made sense as the ~~access~~ fiber placed during HFC builds passed close to many towers. The routed infrastructure was getting more robust and resilient. At this point, there were many transport vendors and technology types. While the geographic borders between networks are blurring, transport networks that do not interop stop in the same hub keeping the optical networks still mostly segregated.

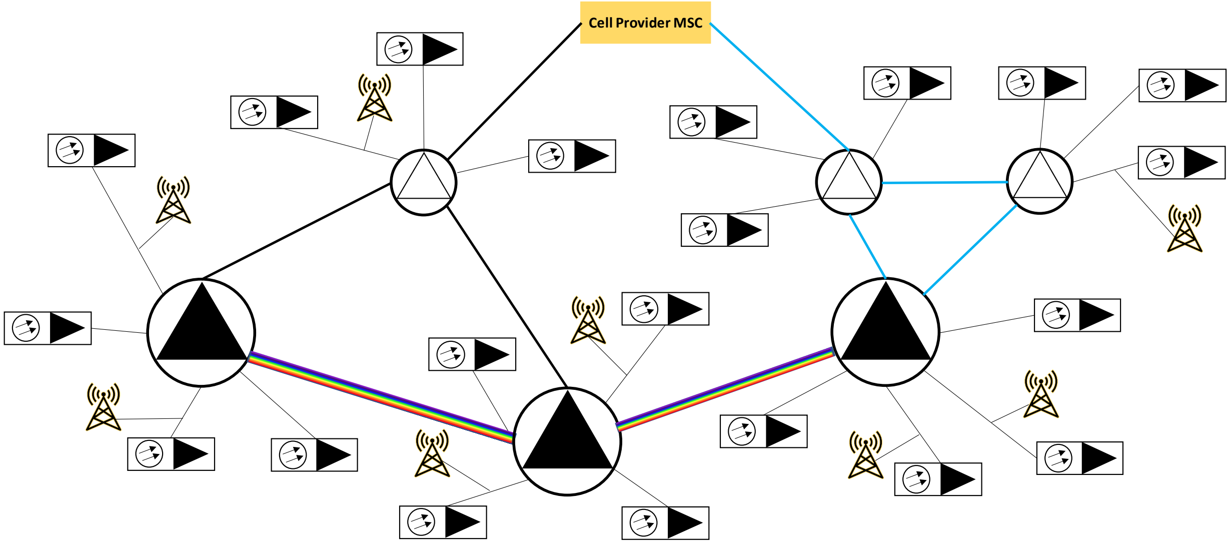

Figure 3- Single network with CBH, redundant hubs, some ROADMs and various transport vendors

2.4. All ROADMs, Flex Technology and one vendor

Over time, bandwidth needs became much greater for residential and commercial needs. Many systems within a region became part of the same network. The same technology type and vendor of transport are being deployed and meshed in the metro optical network. The access still stops at the metro hub.

From here forward all networks are being deployed with flex technology with an initial focus on supporting 400 gigabits per second (Gbps) wavelengths.[1] This paves the path for the IP infrastructure to use 400 gigabit Ethernet (GbE) native interfaces. Flex technology is also known as a colorless

---

[1] To better understand sections in the remainder of this paper there are important definitions when referring to wavelengths that are dependent on context. Wavelengths can be defined by the center frequency that aligns with an International Telecommunication Union (ITU) standard channel. The width of the wavelength can also be referenced in gigahertz (GHz). And lastly while referring to a network facing wavelength the transmission rate can be reflected in Gbps.

system. Instead of a traditional fixed wavelength grid add and drop structure with either 50 GHz or 100 GHz channels, flex technology allows 6.25 GHz slices of spectrum to be combined into any width wavelength. Doing this allows more efficient use of spectrum and future proofs the network against new technology that may require wider wavelengths than the fixed grid systems can support. At the same time, Comcast has decided to add two additional optical design requirements: Layer 0 (L0) control plane capable which allows protection and restoration of optical paths through wavelength switching and L band. These both make the network more robust by adding resiliency and capacity.

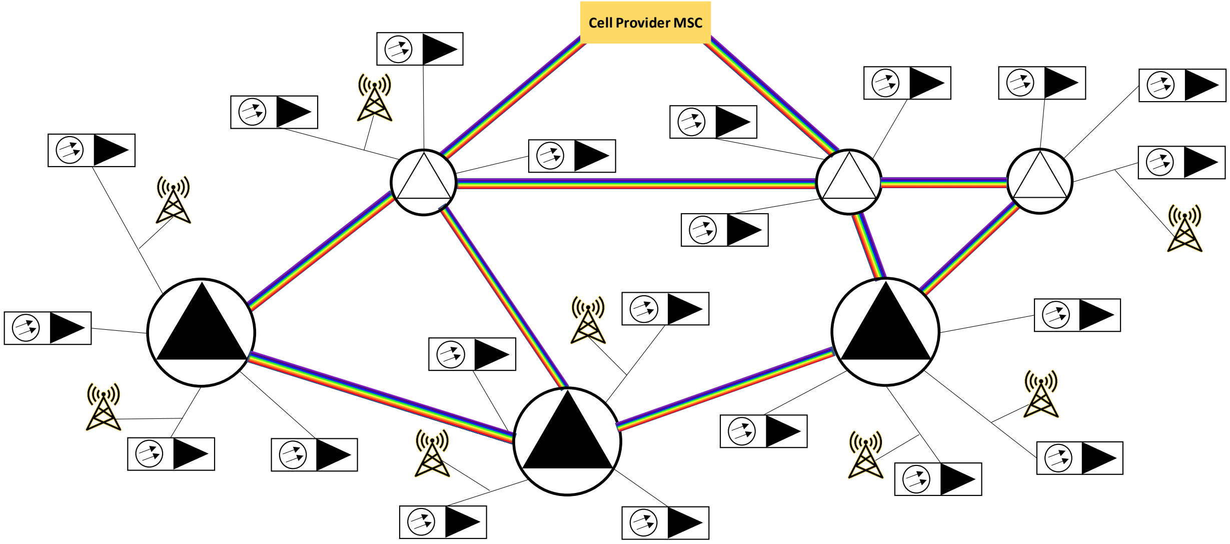

Figure 4 - Meshed Metro Optical Networks with ROADMs and Flex Technology

3. Products

Equally important to understanding the underlying fiber and transport infrastructure is knowing the requirements and challenges of the two products that drove the creation of the unified architecture.

3.1. Wavelength Services

Wavelength services is a product that delivers a single or diverse point-to-point, optical (Layer 1) 10 Gbps or 100 Gbps circuit over the Comcast dense wavelength division multiplexing (DWDM) network. There is no Layer 2 (L2) switching, Layer 3 (L3) routing, or any IP customer premise equipment (CPE) provided by Comcast.

Simply stated, wavelength services are point-to-point, transparent pipes over long distances with extremely low latency. The customer can connect their own IP devices between sites and as far as they are concerned, it is like a fiber jumper, is transparent, and lower latency than Ethernet services.

It is different than dark fiber in that PMs are available to observe for degradation or other issues. It can also go much farther because it uses optical transport.

3.2. High Bandwidth Ethernet Services

This product differs from wavelength services in that there is a Comcast owned and managed CPE on the end of the circuit and can be either point-to-point between customer sites or Internet access, referred to as Ethernet dedicated Internet (EDI). To achieve this at 10 GbE and lower rates, the CPE are connected over the access network to Comcast routers using coarse wavelength division multiplexing (CWDM) or DWDM optics, and leverage the IP infrastructure to transit to other customer sites or the Internet. With customers now requesting much larger bandwidth (up to 100 GbE) this cannot always be supported without extensive augments to the IP network and platform upgrades to support 100 GbE interfaces. This can lead to long deployment times and be cost prohibitive.

To solve this issue the access and metro optical networks can be used together with the same unified optical architecture as wavelength services to backhaul the high bandwidth Ethernet services to locations that the IP infrastructure can support.

4. The Building Blocks

The unified optical architecture started as a solution for wavelength services. This product needed to be launched in several markets that had different transport vendors and technology types. All tools for order processing, design, service delivery, and service assurance were centralized and needed standards applicable across the board to successfully provide the experience customers expect. It would not be possible to handhold each service and use development cycles on several iterations of the same product.

As seen above in the network evolution, the metro optical network is starting to be well-defined and positioned to carry any line of business (LOB) to any location. While the access is also maturing to support distributed access architecture (DAA) technology, it does remain segregated from the metro. This is not necessarily a bad thing provided the two networks can work together when needed.

With the metro being able to carry any LOB agnostically there needed to be a way to provide demarcations for engineering, service delivery, and service assurance. This was done by subtending independent devices dedicated to commercial services under the metro shelves. These needed to have the same role and functions regardless of the vendor. Thus, the network terminating equipment (NTE) and wave integration shelf (WIS) were born.

4.1. Network Terminating Equipment

The NTE is a small optical transport shelf that today can support 10 GbE and 100 GbE services over up to 400 Gbps wavelengths. It resides at a single customer site or data center to serve multiple customers and assumes the role similar to a CPE or edge gateway (EG) would for Ethernet-only services. A challenge when developing wavelength services was deciding how to hand off to the customer. Being an all-optical product, a point of demarcation was required to be the handoff. When only a demarcation was needed, a fiber patch panel could be used. A fiber patch panel directly connecting a customer to the network was not an option because there also needed to be a way for the wavelength and transmit power to be controlled.

A major roadblock for using high bandwidth Ethernet services was that anything over 10 GbE of committed information rate (CIR) would require the use of a 100 GbE interface. The use of 100 GbE optics in IP platforms is not new. This has been widely adopted by the industry for years. 100 GbE has become the new 10 GbE. However, these are all 1310 nm, long reach (LR) optics. The only way a CPE could be connected to the IP core with these optics would be through a dedicated pair of fiber with no MUXs assuming the customer was within ~10 km of the Comcast router. At the time of creating this architecture, coherent DWDM optics at 100 Gbps or greater for the routers and CPE were in their infancy and were not within the project timeline. Even as this paper is being written these optics and their host platforms are still being developed and are not ready for production at scale.

The solution for this is to use optical transport cards that support DWDM optics at 100 Gbps (or greater). This allows taking a 1310 nm, LR signal from the router and CPE, and converting it to a DWDM wavelength for use over the access MUXs as shown in Figure 5.

To deploy this LR to DWDM conversion a place was needed to house the transport cards. The NTE fit this role perfectly at the customer site.

An additional benefit to placing the NTE is now the device can be remotely managed and enrolled in the transport network management system (NMS). This allows alarms to be monitored, PM data is available, loopbacks can be placed remotely for troubleshooting, provides an end-to-end circuit view and other features useful for supporting the service.

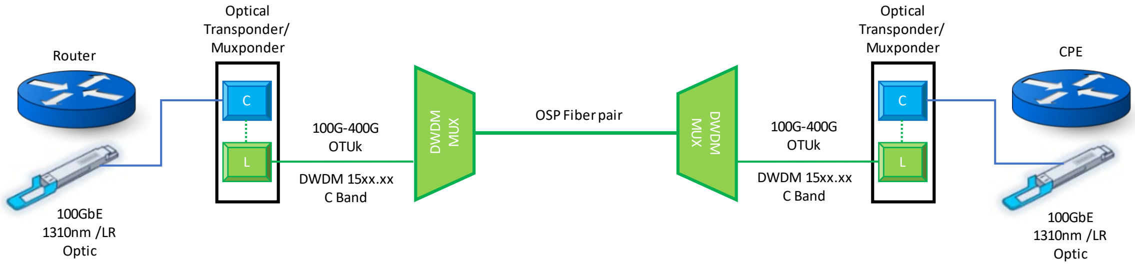

Figure 5 - LR to DWDM conversion

4.1.1. NTE Remote Management

With many proposed solutions come new challenges. To enroll the NTE in the transport NMS and be able to reach it remotely it needs a management IP address. In most cases the IP for transport devices is provided by directly connecting to a co-located switch. With the NTE being in a customer site there is no switch present. The closest switch is typically in the hub at which the fiber from the customer site terminates. The solution is to use the general communications channel (GCC). The GCC is bytes embedded in the overhead of transport links and can be used for small amounts of information. There is GCC0, which is two bytes within the optical transport unit (OTU) overhead, and GCC1, which is two bytes in the optical data unit (ODU) overhead. GCC0 was selected because the OTU is on the line (or network facing) interface. GCC1 is associated with ODU which is tied to the client payload. While both are transparent to the customer it was decided to keep the management on the network facing interface partially due to more hardware supporting GCC0.

One caveat when considering hardware in this architecture is that not all interfaces support GGC0 or GCC1. Due diligence was required to select the right transport cards and shelf for the NTE role to ensure it supported this feature.

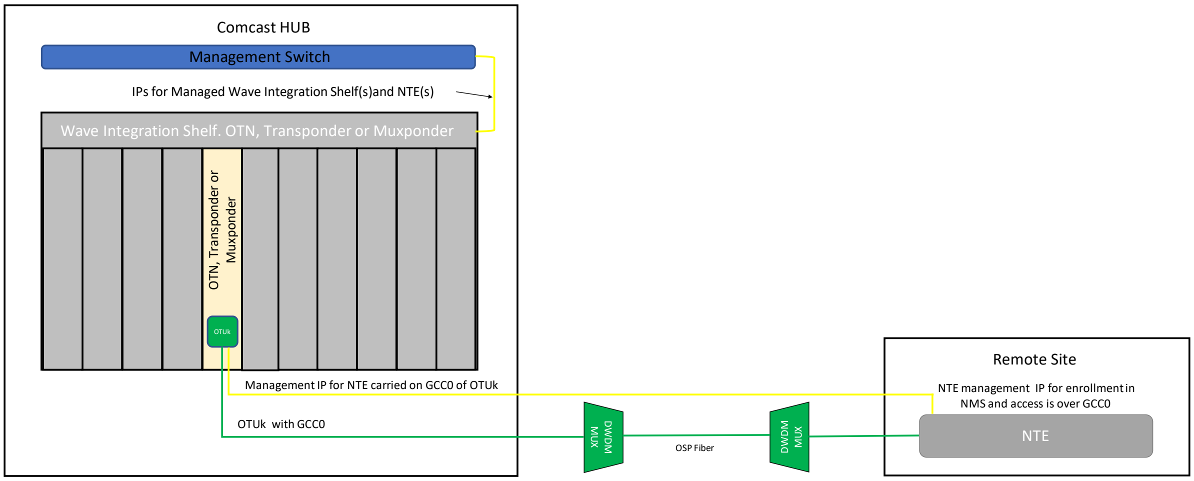

Figure 6 - Management IP for NTE being distributed using GCC0

4.2. Wave Integration Shelf

In addition to needing a device in customer sites another role was needed in the hubs and headends. It serves a similar point of demarcation function as the NTE. It also functions as an optical aggregation point. Other than more effectively using wavelengths through aggregation this also solves another issue. As optical networks are upgraded to flex technology they are optimized for coherent channels. In doing this the dispersion compensation modules (DCM) are removed. Without dispersion compensation 10 Gbps non-coherent channels will not work on most spans. By making the WIS an aggregation point for 10 G channels they can be combined onto 100 Gbps to 400 Gbps coherent wavelengths for transit across the flex optical network. There was a decision to be made here regarding the type of aggregation technology to use.

Traditional muxponders could be used to combine the nx10 Gbps circuits to a single coherent trunk. As seen in Figure 7, all traffic must have the same A and Z ends. Analysis showed it would not be common for multiple 10 Gbps circuits to be going to same end point except in the case of dense data centers.

A better solution was to use OTN switches. This allows traffic to be assigned to a slot on the trunk that can be picked out at a location along the path between A and Z without impacting the pass-through services. This is particularly effective when you expect a common A end with different Z ends along a path as is often seen with single customer sites that need to reach a common data center.

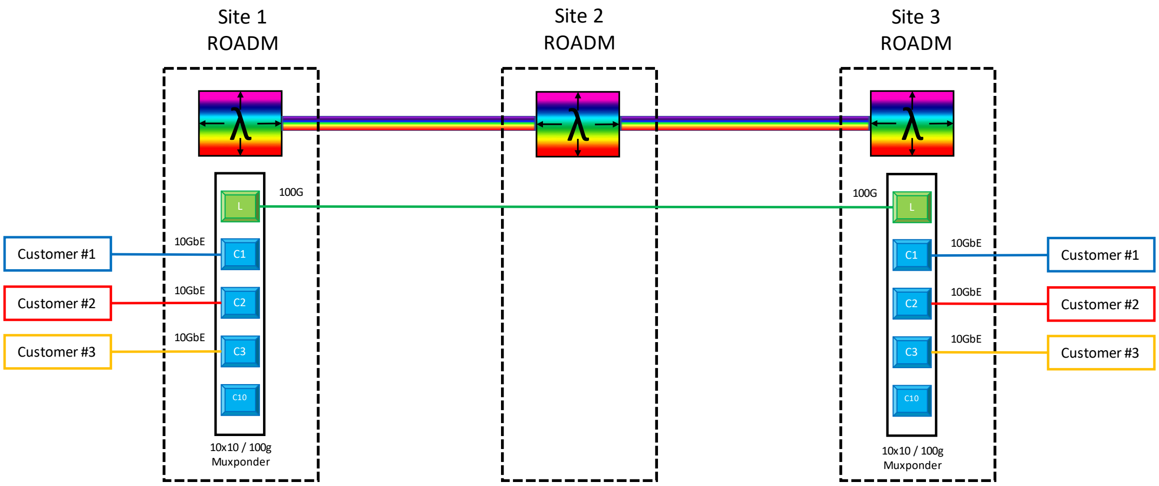

Figure 7 - Muxponder. All circuits share A and Z termination

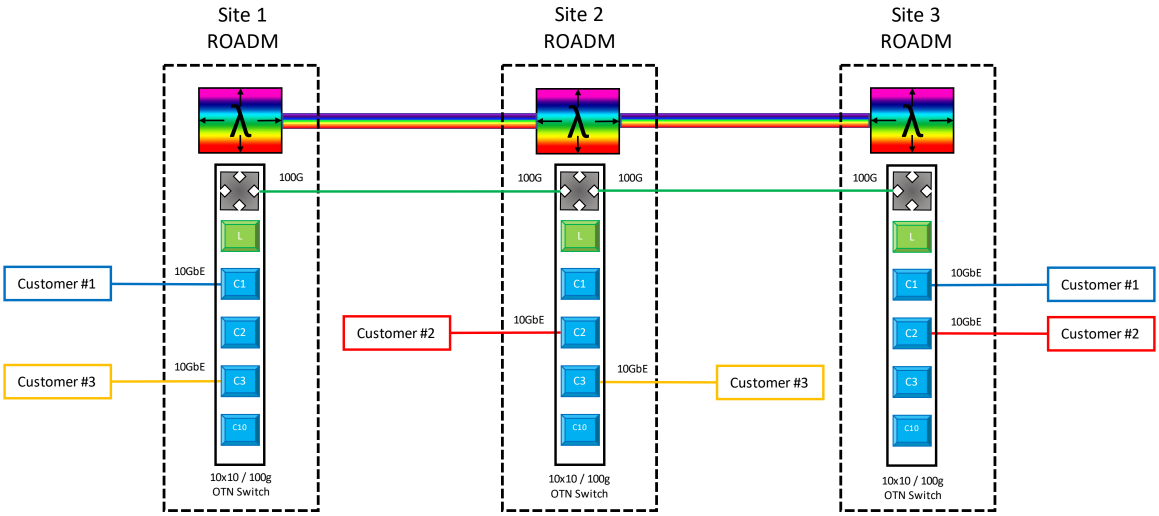

Figure 8 - OTN aggregation allowing different A and Z termination

4.2.1. Capacity Planning

Much of the focus on creating the unified optical architecture has been on engineering and operations but another important piece is planning, specifically capacity planning. Commercial services have sales forecasts, but they are very fluid. Residential networks have well-known, proven compound annual growth rates (CAGR) and can be planned with good accuracy. Since the metro optical network supports all LOBs to effectively manage slot capacity, each demand, or group of demands, benefits from having its own client shelf to support the cards needed. The WIS fits this role perfectly. It provides a shelf that will only support commercial services. Here is an example to illustrate the benefit.

Assume a shelf in a headend that supports both residential and commercial had five open slots, and the CAGR for residential traffic dictated that over a year one slot per quarter would be consumed.

That leaves one slot to support a commercial sale. Initially, that is not a significant issue. With commercial being so fluid though, sales can come in at any time. If more than one was to come in before the planned residential circuits were turned up, the commercial asks would consume the slots planned for residential. Everything could be oversized to make sure there is enough slot capacity but with space and power being a premium it does not make sense to do this “just in case.”

The WIS solves this issue by giving the commercial cards a place to reside other than the well-known static needs of the residential network.

This benefit proved itself with the capacity explosion during COVID-19. With no risk of a commercial demand unexpectedly consuming a slot in the residential shelf, engineers were able to execute augments to keep up with bandwidth needs on the residential network without delay while at the same time being able to provide high bandwidth commercial services to business customers to support remote employees needing to reach their infrastructure.

5. Leaveraging Access and Metro together: OTN Tails

Now that the hardware and roles have been defined to solve those challenges, we can address the separation of the access and metro networks and how they can be merged when needed to create a single end-to-end circuit.

As seen in the above fiber evolution maps, 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. It has been found that most customers can connect to the access and be brought back to a hub site. But the remaining question is how can they be connected to each other if their fiber terminates at different hubs? Or in the case of 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 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 core by a second card in the WIS with an 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. Client-to-client connections have been made in the past as Ethernet where a similar design was used for regens. In this case the decision was made to keep them OTN so that the circuit remains transparent from end-to-end. There is no conversion to Ethernet and back to OTN. One advantage of this is that the trail trace identifier (TTI) trail trace 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 troubleshooting and topology verification.

The signal flow described above is applicable to 100 GbE services using a 100 Gbps to 400 Gbps trunk.

10 Gbps services differ slightly at the WIS. They still use an OTN tail over access from the NTE to the WIS but instead of back-to-back cards the 10 Gbps is aggregated on the OTN switch card.

Using the back-to-back cards or OTN switch card has also solved another issue with using MUXs in the access. By having the ability to tune to any wavelength facing the NTE or the metro it is not required to maintain the same wavelength end to end from NTE to the metro, across the metro, and to the other NTE. An effective way to illustrate this is to imagine many MUXs being deployed in the access with ITU channels (Ch) 20-59. The metro uses the same channels and, in some cases, adds another set of 50 GHz spaced channels the access does not use. If Ch 20 is used on the A end without this recoloring (that is, changing wavelength) in the WIS you would need to have Ch 20 open on your metro and the other end of the circuit over the access. Once you use Ch 20 for one circuit in the access or metro you cannot use it again. Back-to-back cards allow any channel that can reach the hub site where the WIS is located to be recolored to any open channel across the metro.

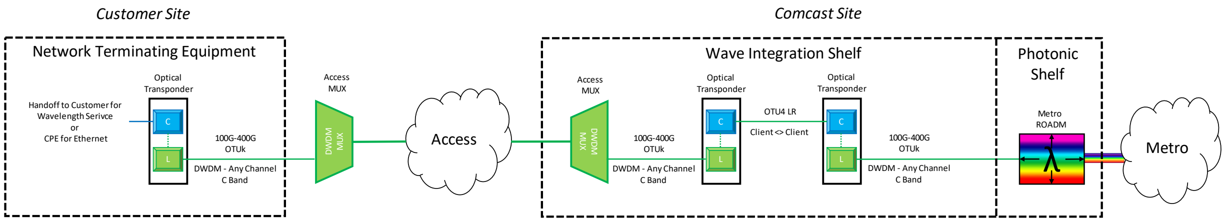

Figure 9- Recolor with back-to-back cards in WIS

5.1. Access MUXs to support 400 Gbps

With the metro being a flex system, it can combine slices of spectrum in 6.25 GHz increments to create the width needed per channel. Legacy fixed grid transport used two sets of 44 channel 50 GHz MUXs. This provided up to 88 channels but with a maximum per channel width of 50 GHz. This would support up to 200 G channels that are typically ~37.5 GHz wide.

The Access followed the same fixed grid pattern with one important distinction. The MUXs used were 100 GHz spaced. Why is this important? The last major technology advance that has been heavily adopted is the use of 400 G line rates. Depending on vendor and specific technology type, the width of this is ~75 GHz. With the access having up to 100 GHz of width available per channel, 400 G wavelengths can be deployed on a system that was first created and optimized to run 10 G. This is a fantastic use of the available spectrum; however, it does appear to be the end of the road. The next technology being adopted by the industry, including Comcast, is 800 Gbps wavelengths. As can be seen in Figure 10, an 800 Gbps wavelength occupies ~112 GHz of spectrum and will not fit in the 100 GHz MUXs used in the access.

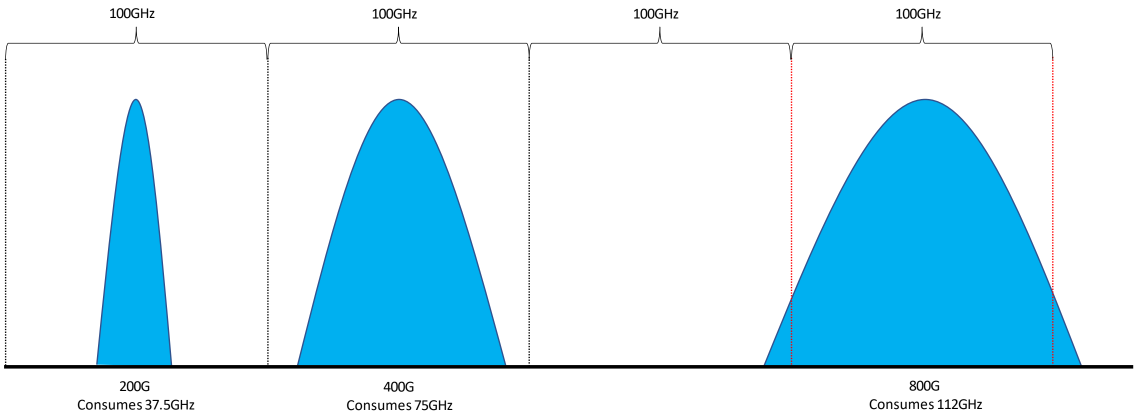

Figure 10- Spectrum consumption by 200 Gbps, 400 Gbps and 800 Gbps Wavelengths

There is one particularly important caveat when using 100 Gbps and greater wavelengths over the access. Today's coherent optics by nature require both the transmit and receive frequency (wavelength) to be the same. This requires the use of a two-fiber MUX. A single-fiber MUX uses different frequencies for the transmit and receive.

Being that single fiber MUXs are the standard deployment at Comcast for all new DAA builds this puts the unified optical architecture at odds with DAA. However, work is being done in the industry to create bidirectional coherent optics to use for this application that would operate on a single fiber MUX.

Comcast Fellow and industry leader in this effort Venk Mutalik says, "Many businesses are within a short distance of optical fiber nodes, and with the expansion of DAA, we are going to have more of them out in the field. So, a really great way to use fiber assets already available is to converge business services on the same fiber as the residential services. This is done rather easily with 10 Gbps services, but with the ability to do the same with 100 Gbps and even 400 Gbps coherent services on the same fiber, access convergence is a game changer for the industry!"

6. Complete end to end archicture

With the building blocks established and the access and metro working together, a complete end-to-end architecture can be established. In Figure 11, each colored line represents a different commercial service type and rate.

All services, regardless of type, originate or terminate on an NTE at the customer site, use an OTN tail over access and are added into the metro optical network via WIS.

In Figure 11 the following types of services are shown.

- Blue – 100 GbE EDI
- Red – 10 Gbps wavelength service, customer site to data center
- Green – 100 Gbps wavelength service, customer site to customer site
- Orange – 100 GbE EDI

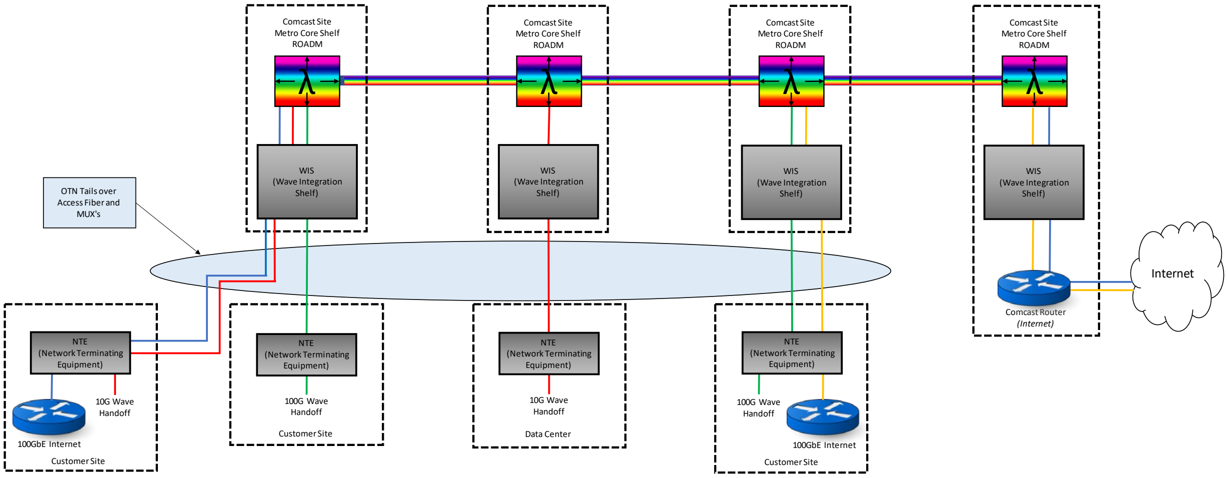

Figure 11 - Unified Optical Architecture

7. Benefits

All the work is now done. Building blocks have been assembled, technology challenges have been overcome, and a unified optical architecture to support multiple lines of commercial services has been created. The benefits are simple and clear. By developing this vendor-agnostic, unified optical architecture, commercial circuits can be deployed and supported the same way every time regardless of the type of service or traffic pattern, while increasing efficiency and maintaining or improving customer satisfaction.

The overarching theme is sameness. As explained by Shane Portfolio, Comcast Senior Vice President of Network Services, and leader of the Comcast One Network Initiative, “The value of sameness should not be underestimated. It allows for speed to market, greater scale, increased customer felt reliability, costs reductions, and enables next generation technology like machine learning, automation, and artificial intelligence to come to life because it is a homogeneous system that can be built to leverage these technologies in ways doing things differently simply cannot. It’s a game changer.”

7.1. Engineering and Service Delivery

By using the same architecture regardless of vendor, engineers can streamline processes and become deeper subject matter experts. The designs will always include the basic building blocks of the NTE, WIS, access and metro core. They will always follow the same device and port naming standards.

7.2. Procurement and Deployment Engineering

Using the same equipment as the metro networks benefits procurement by reducing the number of unique items in the ordering systems. A basic tenant of economics, volume buying power, also applies here.

Deployment engineering benefits from the ability to shuffle hardware to the highest priority project. Using the same hardware for multiple demands makes this possible. This again was a valuable option to meet demand and bandwidth shifts during COVID-19.

7.3. Operations and Service assurance

Possibly the biggest benefactors are the operations and service assurance teams. They are tasked with keeping the circuits and network healthy and code up to date. The ability to see PM data like optical levels and errors on an interface as well as place loopbacks at the customer site remotely is invaluable. By using the access to create OTN tails and merging them into the metro, the power of the NMS can be used to see a circuit from end to end including a simple green is good, red is bad, graphical representation. Alarms are also collected from all the optical devices and centrally managed in the NMS. This can dramatically cut the time to resolve an issue. Using the same hardware as the metro for NTEs and WIS also streamlines code certifications, deployment, and development efforts by reducing the variety of equipment deployed.

One of the greatest accomplishments this architecture has made possible is the development of software that merges customer account information, provisioning details, and correlates L2 and L3 ports to their connected transport clients. This also shows a hop-by-hop trace of the circuit from the NTE over the access and metro. Work was even done to stitch together circuits on multiple networks provided by different vendors. This creates a true end-to-end view with all the information an engineer needs to understand the circuit in one place. Comcast and a partner company co-developed this software.

8. Conclusion

All the work that was done to create a unified optical architecture may seem obvious. Why wouldn't the same equipment and standards already be used everywhere? The reality is that with the speed at which networks have merged and evolved they have been in a perpetual brownfield state with a variety of architectures. If everything were a greenfield build with no existing network to support, this would certainly be a lighter lift. In addition, optical technology has gone from 10 Gbps to 1.2 terabit per second (Tbps) wavelengths over the last decade. Simply keeping up has required the full attention of all teams. To put everything together has taken several years and a dedicated effort beyond keeping the lights on. In the end, though, by creating this unified optical architecture the goals to continue to deliver or improve the experience customers are accustomed to, offer new technology, and increase efficiency have been achieved. By not only solving the challenges present today but looking ahead to the future while creating this architecture, the networks are well-positioned to deliver the next generation of 400 Gbps wavelengths and 400 GbE to the customers with no changes.

Abbreviations

AMP	amplifier
CAGR	compound annual growth rate
CATV	cable television
CBH	cell backhaul
Ch	channel
CIR	committed information rate
CPE	customer premise equipment
CWDM	coarse wavelength division multiplexing
DAA	distributed access architecture

DCM	dispersion compensation modules
DWDM	dense wavelength division multiplexing
EDI	Ethernet dedicated Internet
EG	edge gateway
Gbps	gigabits per second
GbE	gigabit Ethernet
GCC	general communications channel
GHz	gigahertz
HFC	hybrid fiber/coax
HSI	high-speed Internet
IP	Internet Protocol
ITU	International Telecommunication Union
L0	Layer 0
L2	Layer 2
L3	Layer3
LOB	line of business
LR	long reach
ME	Metro Ethernet
MSC	mobile switch center
MSO	multiple system operator
MUX	multiplexer
NMS	network management system
NTE	network terminating equipment
ODU	optical data unit
OTN	optical transport network
OTU	optical transport unit
PM	performance monitoring
ROADM	reconfigurable optical add drop multiplexer
SCTE	Society of Cable Telecommunications Engineers
Tbps	terabit per second
TTI	trail trace identifier
VOD	video on demand
WIS	wavelength integration shelf