Capacity and Technology Considerations in DAA Backhaul Deployment Strategies

A Technical Paper prepared for SCTE•ISBE by

Fernando X. Villarruel
Architect
Cisco Systems
5030 Sugarloaf Pkwy, Lawrenceville GA 30044
770-236-1385
villarf@cisco.com

Martin Mattingly
Technical Solutions Architect
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30044
770-236-1338
mattinm@cisco.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

This white paper compares the benefits of several architectural options for the Distributed Access Architecture (DAA) backhaul in the context of bandwidth growth over time. Some of the specific topics covered include the networking and optical implementations needed to address DAA backhaul, routing and TDM framing in hubs and HFC nodes, methodology for estimating needed capacity, concurrency, implementation of direct detect and coherent optics, and cable's new-found synergies with standard bodies. After reading this white paper, the cable operator will be able to compare the various architecture options and decision-making process for the deployment and long-term evolution of DAA.

Background

In the world of DAA there are at least two types of digital endpoints, remote PHY (RPHY) and remote MAC-PHY, (RMAC-PHY.) From the networking perspective, they both have the same function and they are both point to multipoint signaling that terminate at 10 Gbps endpoints. With a closer look, there are certain efficiencies concerning multicast that favor R PHY. Thus, in this paper we use RPHY as our main example, pointing out relevant differences with RMAC-PHY when necessary.

Figure 1 shows a summary for the RPHY architecture. It also helps us to define the Converged Interconnect Network (CIN) as the Ethernet/IP network between the packet cores and the RPHY nodes, drawn as the shaded cloud. We note that DOCSIS and other packet core payloads are encapsulated within an L2TPv3 pseudo-wire and then in IP and Ethernet layers. This process makes any subscriber specific data effectively invisible to networking elements thus making it fully addressable via switching and routing principles. It is also important to note that the packet cores and RPHY nodes have network facing Ethernet client interfaces, which mean that the most direct path for transmission is via client-to-client connections, not unlike what could be used in non-access networking such as in home or data centers applications.

The DAA backhaul is a subset of the CIN, in particular it is the physical link and method for aggregating multiple DAA endpoints. Backhaul methods then are distinguished by connectivity, packet processing and transport options. In practical terms, it is the way we leave the hub or headend and connect to a networking element that combines the signaling to and from RPHY nodes.

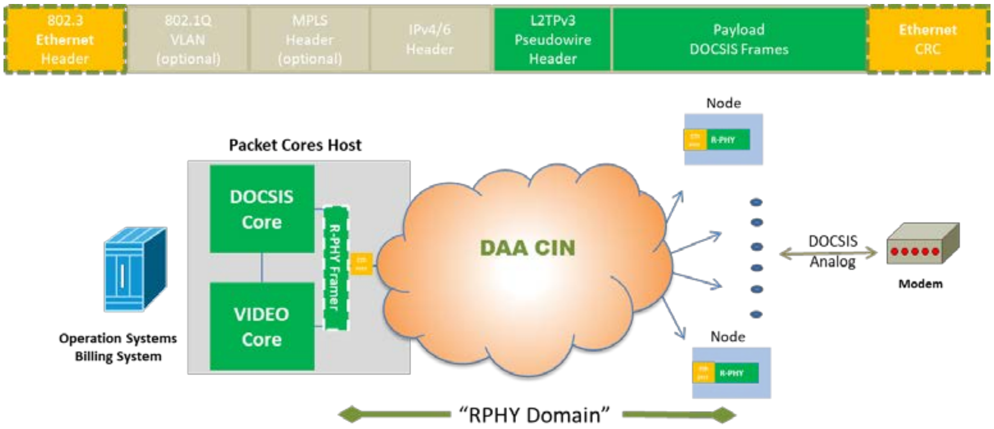

Figure 1 - High Level Remote PHY Architecture

Architecture Comparison

There are many variations in how to design the DAA backhaul. In Figure 2 we present several generalized options. Focusing on commonalities, we note that the northbound elements are a collection of packet cores, which execute subscriber management policy and create the data plane necessary. These can include the DOCSIS packet core, multiple video packet cores including broadcast, video on demand and switched digital video. There are also support packet cores that include such elements as out of band signaling and HFC RF monitoring tools. Generally, there could also be other service cores such as broadband network gateways (BNG) for PON or mobile. These packet cores can create point to point or point to multipoint sessions where signaling could be unique or shared as it passes through an initial series of hub routers towards their final destination, RPD's in the field. In particular from Figure 2, we note the existence of a layer of “core routers” whose job is to coalesce the signaling directly from packet cores. This routing layer uses 100 Gbps connectivity, with typical forwarding capacity nearing 1 Tbps with typical port count of 36 ports that can be used as either uplinks or downlinks. The packet cores themselves have direct 100 Gbps connectivity or are facilitated by an extra layer of routers with 1 or 10 Gbps connectivity for their uplink and 100 Gbps connectivity in their downlink.

Following the 100 Gbps connectivity router there is the existence of a dense 10 Gbps connectivity aggregation router. This router has 100 Gbps connectivity on the uplinks and 10 Gbps connectivity on the downlinks. Typically, these routers have two or four 100 Gbps connections along with 40 or 48, 10 Gbps connections. We call this the aggregation router because this router is the last logical connection at the hub or headend as the signal enters the outside plant. This multilayered approach of the remote PHY CIN is very similar to the spine – leaf architectures that are now prevalent in data centers and useful to facilitate the evolution to virtual packet cores. We note that between the 100 Gbps and dense 10 Gbps routers there could be photonic network as the packet cores might not be located at distribution hubs. Photonic equipment allows for a TDM aggregation of signals for transport of large bandwidths. We will discuss this type of transport gear again later.

These items account for commonalities. We now look at the differences in the options for the backhaul of Remote PHY Devices (RPD's).

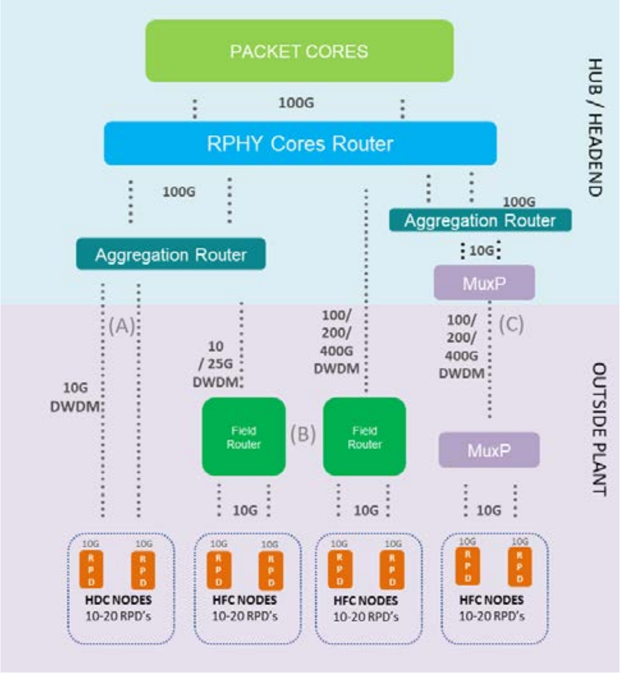

Figure 2 - Remote PHY Backhaul Architecture Comparison

1. Backhaul Option (A), "Direct Connect"

The connection labeled (A) in Figure 2 is part of an architecture description generally called "Direct Connect." We call this Direct Connect because there is a direct connectivity between the dense 10 Gbps aggregation router and the 10 Gbps RPD endpoints. The main quality of this option is its simplicity and availability. In most cases, a collection of signals will be incident on one fiber along with an accompaniment of a Mux/Dmux to manage them. This architecture leverages 10 Gbps DWDM ZR optics, which is a description for optics that are wavelength specific within the 100GHz channels, as described by (G.694.1). The one variance is that the ZR optics used have to be thermally hardened to exist within the RPD enclosure that exists in the outside plant. This adds some complexity and some cost over the otherwise commodity structure of 10 Gbps ZR optics.

In situations where fiber is taken deep into the outside plant, the number of 10 Gbps endpoints can be 10 to 20 times the number of nodes that exist without DAA. As such, the number of DAA endpoints translate to a corresponding large number of DWDM wavelengths per trunk fiber. In Figure 3 we cover in more detail the evolution of the trunk fiber and DAA endpoint in a way that would be typical in end to end architectures.

It is worth noting that after the transition to DAA, we typically see a large disparity between the provisioned bandwidth per DAA endpoint and the actual utilization of bandwidth by the end user. In instances where multiple DAA endpoints are logically clustered together to create "service groups", we often see that the 10 Gbps connection to each DAA endpoint is significantly under-utilized. In a later section we will also explore how obtaining a better understanding of provisioned versus utilized bandwidth can be an effective tool when estimating backhaul capacity.

2. Backhaul Option (B), Field Aggregation Router (FAR)

In Figure 2, we show two connectivity options labeled (B). Both have in common the introduction of an active networking element in the outside plant, a Field Aggregation Router, (FAR). The FAR has the task of facilitating several packet-processing functions to its subtending RPD's. From a topology perspective the FAR fits well at the same location where once was an analog node that now spawned multiple connections to RPD's. Southbound, there is 10 Gbps connectivity from the FAR to RPD's that are typically no further than 1 or 2 km away. This creates an opportunity for using lower cost 10 Gbps LR optics, which could be further de-rated for less than 2 km. These optics can be as low as one fourth the cost of the 10 Gbps ZR discussed in the previous section. Also, note that because of the short distance and the low link budget the necessity for thermally hardened optical components changes in scope and should not be a considerable cost adder.

Northbound of the FAR this architecture brings several dimensions of flexibility, and in order to appreciate the flexibility we explore the connectivity options available to the FAR by design. Routers are built around silicon chips whose inputs and outputs (I/Os) are well defined but with inherent flexibility. For example, a typical router in the service provider space might have a collection of 10 Gbps I/Os or 25 Gbps I/Os, all with equal access to the forwarding plane. How these I/Os are used is up to the designer of the router, with many variations possible. These transmission lines can be combined or down rated, when accompanied with the right media access control, (MAC). Four 25 Gbps lines can be combined to facilitate a 100 Gbps signal, or a 25 Gbps line can run 10 Gbps, or a 10 Gbps line can run 1 Gbps. In the case of the FAR, this allows for a range of options in northbound connectivity. We will see in a later section that pairing this flexibility with the expected capacity over time allows for pay as you grow scenarios.

2.1. FAR, Initial Uplinks

In Figure 2, the left portion of (B), we see that the FAR can have uplinks in speeds of 10 or 25 Gbps that leverage the existence of first generation hub aggregation routers. This allows a fine-tuned way to address the capacity needed for the uplink over time. In detail, this means purposing some of the 10/25 Gbps transmission lines of the FAR for the purpose transmitting the uplink, where the rest of the 10 Gbps, or 10/25 Gbps lines can be used for downlinks. Note, this type of uplink connectivity only makes sense to do with a few ports and maybe only at the beginning of the lifetime of the FAR, which should be a decade or more. This also makes sense as a transitional step if initially investments have already been made for aggregation routers in a hub. In this case, the introduction of the FAR enables fewer ports on the hub aggregation routers, as we will see in the backhaul capacity section.

2.2. FAR, Long Term Uplinks

In Figure 2, the right portion of (B), we see the FAR can have uplinks in speeds of 100 Gbps, and above. These uplinks represent the natural evolution of optics for higher transmission bandwidths and allows for the long-term transmission of signals to the FAR. These uplinks are made to address the challenges mentioned earlier in the direct connect section, where many lambdas were needed to address a collection of endpoints. In this case, the full bandwidth needed for all the RPD endpoints subtended by the FAR can be addressed with one lambda. From the perspective of the FAR, the usage of 100, 200, or 400G on one lambda can still employ the same routing fabric, if sized accordingly and accompanied with the necessary collection of transmission lines in the design. For example, 4 x 25 Gbps transmission lines facilitate a 100 Gbps optical PHY interface. If available, 8x25 Gbps transmission lines can facilitate a 200 Gbps optical PHY interface. Moving beyond first product implementations, the expected eventual addition of 50 Gbps transmission lines will facilitate higher bandwidths with even more simplified connectivity.

3. Backhaul Option (C), Muxponding

In Figure 2, option C, we see the insertion of a TDM framing layer between hub networking gear and the RPHY endpoints. Note that this practice would be new in the cable access but has been practiced in long distance optical transport for a long time. In the case of the cable infrastructure for DAA, the insertion of this transport mechanism makes perfect sense when hubs are collapsed to a more centralized location and the needed bandwidth between the hubs and head ends is on the order of many hundreds of Gigabits. Transport platforms currently have of up to 1 Terabit, with interfaces of up to 400 Gigabit, (Microsemi, 2017).

The common framing mechanism used for TDM solutions is called Optical Transport Network (OTN). This solution in essence takes in different client signals, possibly even with varied rates, and without any examination or manipulation of packets it stitches these different signals in the time domain and puts them on a signal at a much faster speed, a process that is reversed on the other side. We call the platforms that execute this function a muxponder. In practice, a common example would be a muxponder that would take twenty 10 Gbps signals and output a 200 Gbps signal.

4. Physical Trunking and Distribution

Figure 3 shows a practical approach for the evolution to DAA, particularly from what are typical starting points. In the scenario labeled “Analog Hub”, where the end to end signaling terminates at a hub, there are four downstream wavelengths and two upstream wavelengths used on a single fiber between the hub and legacy analog node. This is a byproduct of having an internally segmented node. In this type of deployment, the fiber is generally “point to point” from the hub to a physical legacy node and includes several spare fibers within the same sheath. In transition from analog to digital, one of the spare fibers can be used to provision 20 to 40 DWDM wavelengths to facilitate point to point DAA.

In the scenario labeled “Analog Secondary Hub”, two fibers carrying 16 analog wavelengths each from the Primary Headend are de-multiplexed to four groups of four wavelengths each. Each group of four wavelengths is used for downstream and upstream transmission between the Secondary Hub and the legacy analog node.

From Figure 3, we see that the transition to DAA can have several impacts to the connectivity. First there can be a replacement of the cores to a more central location where the connectivity from the new core position, like a primary headend or a data center, has to be accounted for. Because of the scale of the signal it is conceivable that the best solution for this link is in terms of TDM, OTN framing, allowing for multi-hundred Gig rate signals. Also, this implies that the hub location is a networking point. Which is the launch point towards a final aggregation point the FAR, facilitating the nature of replicated and concurrent signaling typical in cable plants.

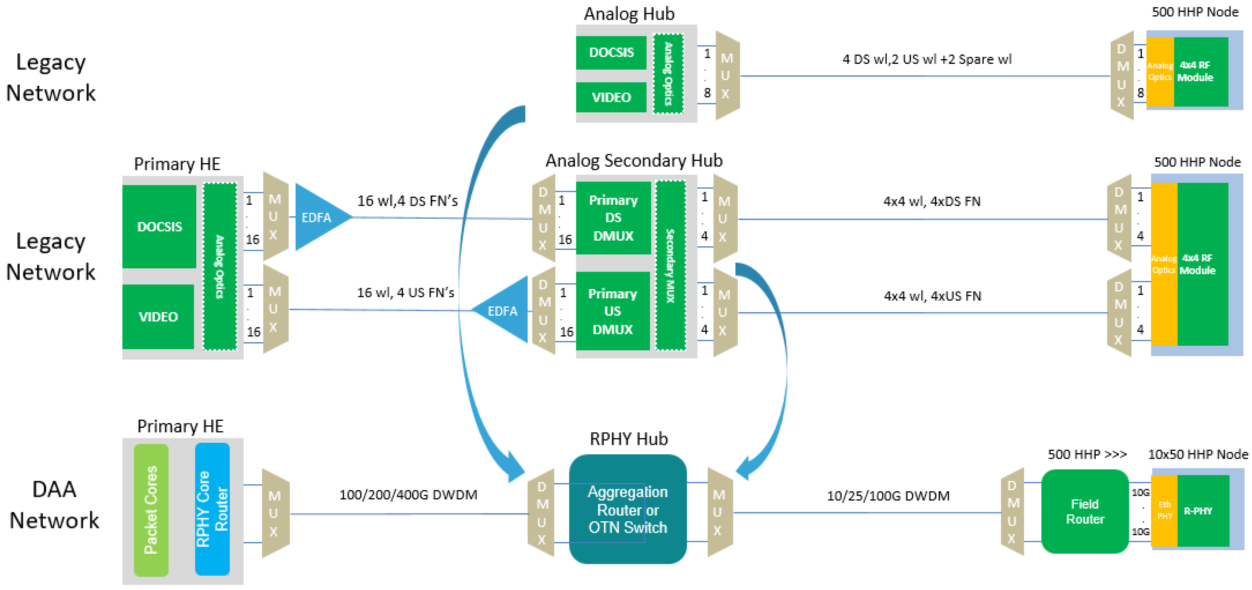

Figure 3 - Transition from Analog Transmission to Remote PHY

5. Componentry Comparison

Coming from the cable access world, we typically do not have a native understanding of what goes into these new remote digital technologies. Figure 4 represents the high-level componentry that make up the digital parts of both a router and a muxponder. Within the diagram we can see that there are some high-speed uplink optical inputs and accompanying PHY, lower speed optical downlinks and their accompanying PHY, and most importantly a function specific ASIC or FPGA along with a robust processor. The size of the silicon, the number of available gates drives its substantive differences in power consumption and functions. Interestingly, in the case of routers there is now a vast set of robust, third party off-the-shelf options to choose from that the industry refers to as merchant silicon. The silicon within the muxponder with functions such as an OTN framer and mapper can also function as an ONT switch, with increased gates and power. Finally, accompanying the ASIC or FPGA is a CPU processor complex that functions with the control plane and runs the operating system of the product.

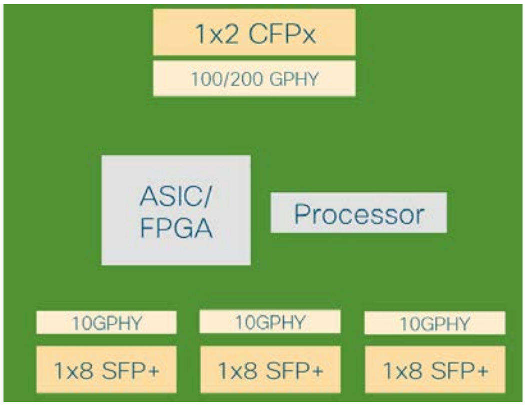

Figure 4 - Sample bill of Material for Field Router/Muxponder

6. Architectural Comparison

Figure 5 below, shows a comparison of features (in green) and challenges (in red) for each of the presented architectures. While we have already mentioned some items on the lists, it merits to compare them together by topic. Note that there is no one solution that is perfect for all situations but knowing how to evaluate them is a useful tool as these options make their way to market.

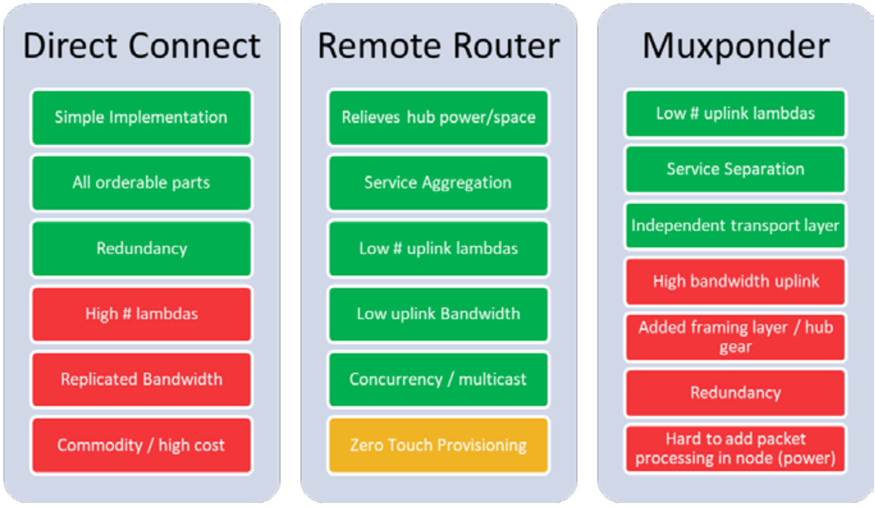

Figure 5 - Architecture Options Comparison Pictograph

6.1. Configuration

With regards to implementation, there is no doubt that the most straightforward option is direct connect. While adding another active device in the field adds complexity, it should be no more effort than configuring a remote PHY device. Since connection to packet cores already necessitates zero touch provisioning at scale, there is no new technology in packet processing being implemented here. Ultimately, any solution must be plug and play, with no settings intervention needed on site. Further, the target for actives in the field should require no manual interaction at all, even at a central location. The target should therefore be toward full automation of the service with general profile guidelines set by the network operator.

6.2. Usage of optical signals

It is worth considering the number of trunk lambdas used by each option. In the case of direct connect, there are as many lambdas needed as there are DAA endpoints. In the case of a FAR or muxponder there can be one lambda per trunk fiber servicing between 12 and 24 DAA endpoints. This allows the trunk fiber in a WDM environment to be used for other services. Note that this does put an added challenge on the description for optical uplinks. They must be able to overcome the passive losses of added WDM equipment while also being able to operate in a wavelength specific environment. These two items are not a given for high throughput long distance optics and are discussed further in section 7.

On the downlink side, there is an opportunity to use lower cost 10 Gbps optics that can easily adapt to the environment. Based on the lower cost of short reach optics, the total cost for a solution that introduces remote aggregation should be much less than direct connect. This is an achievable goal, as we will show in section 8.3.

6.3. Implementation

The physical aspect of implementation is also worth considering. Note from Figure 5 that for all cases there is a transition from a large uplink signal to a breakout distribution for DAA endpoints. In the direct connect solution, this transition happens once at the hub as signal leaves to the access plant. In the case of the muxponder, this transition happens twice where the infrastructure for 10 Gbps connectivity happens both at the hub and at the node. In the case of FAR there is an option to do one of two things: In a pay as you grow scenario, as we will show in section 8.3, the uplink to the FAR can use signaling in terms of 10 Gbps, as needed, allowing the use or reuse of the 10 Gbps layer at the hub. Note that options for dense routing gear can include options of 10 or 25 Gbps, which in cases where the optics allow, the uplink can be in terms of 25 Gbps making the time for the reuse longer. On the other hand, the uplink connectivity to the FAR can be in terms of 100 Gbps or more. This application then skips over the extra 10 Gbps aggregation layer in the hub, going directly from the cores router to the FAR. This is a savings in physical and carbon footprint. Effectively, the aggregation layer in the hub is moved directly to the node.

As we saw previously, the OTN layer is a whole separate logical function that is done independently, and addition to the routing layer that will also be in use. There are products that aim to combine these two functions so that from the outside it looks like one “box” is doing both. This approach will benefit connectivity, but there is no way of getting around the fact that muxponding and routing are two distinct functions that will evolve and be implemented separately. Combining them in one box negates the benefit of treating these networking functions in their own time and availability of scale.

6.4. Uplink Bandwidth

The bandwidth aligned with direct connect and muxponding solutions is equivalent to the physical connectivity of the remote PHY system, as seen in Figure 6 below. Without needing to know the functionality of the packet cores, one can deduce the bandwidth sizing of the backhaul link.

Note both the direct connect solution and the muxponder solution have to transmit as much (or more) bandwidth as is determined by the number of connections to the DAA endpoints. In the case of direct connect the necessary throughput is 10 Gbps times the number of DAA endpoints. In the case of the muxponder, it is 10 Gbps times the number of endpoints rounded to the nearest multiple of framing speed. For example, suppose the framing speed is 100 Gbps, then for 12 DAA endpoints at 10 Gbps, the throughput would have to be two connections at 100 Gbps or one connection at 200 Gbps composite. If the number of DAA endpoints is 20, the throughput would be the same 200 Gbps. At 24 endpoints, the throughput would have to be 300 or 400 Gbps, depending if the framing speed is at 100 or 200 Gbps multiples. As we will see in section 8, there can be a large difference, for a long period of time, between the connection speed at the DAA endpoints and how much bandwidth is being used in the uplink. Thus, part of looking for the right solution is the ability to predict the amount of bandwidth that will be used for the backhaul over time.

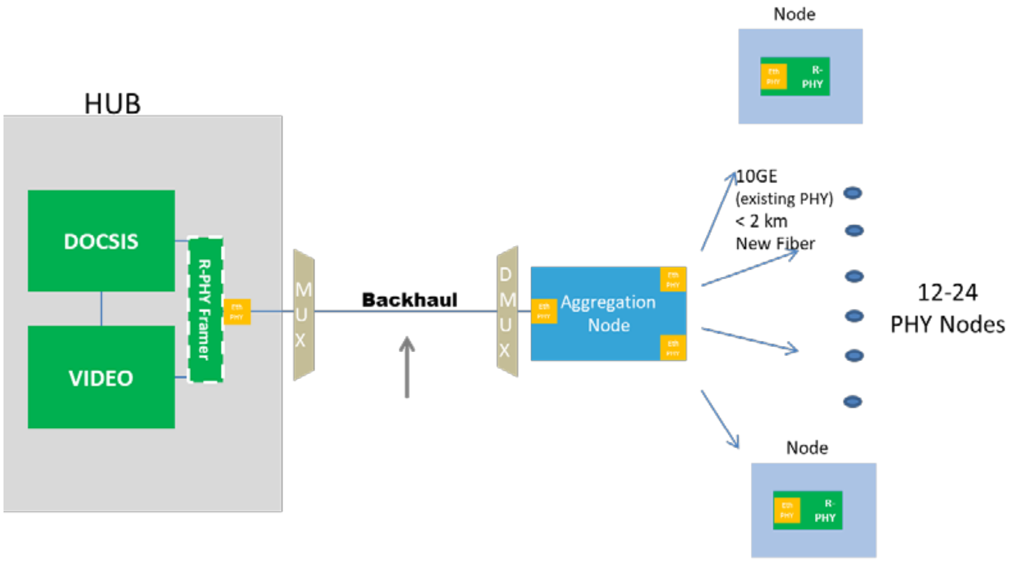

Figure 6 - Physical Connectivity of Remote PHY System

6.5. Multicast and Unicast Bandwidth

An additional consideration is how each solution handles multicast and unicast bandwidth to a group of RPD's. One notable feature of the packet cores used by MSOs is their ability to address their subscribers with sets of common bandwidth. Figure 7represents the logical connectivity of a DOCSIS packet core, which has at its disposal several unique service groups, and the relationship those service groups have with RPD's. It is not necessarily a one to one correspondence, and when it is 1xN each logical DOCSIS service group can service multiple RPD's. In practice we have seen that the number of RPD's per SG can be as high as eight, though it is more common that we see four RPD's or less sharing unicast DOCSIS bandwidth. A similar relationship exists when considering how unicast video content is shared across multiple RPD's. Although the operator could maintain a one to one relationship between video and DOCSIS service groups, we typically see a 1xN relationship whereas there are typically twice as many RPD's per video service group. This practice is generally in an effort to share unicast video content across a larger base of homes, thus minimizing the changes to the existing back office infrastructure. This is not the case for multicast video content. In theory, one set of multicast video channels could serve an entire headend or hub service area. From a more practical perspective, the number of ad insertion zones, thus unique copies of multicast content dictates how multicast video content is shared across multiple service groups. Since a single FAR would rarely, if ever, span multiple ad zones, it is fair to assume a single set of multicast video channels would typically serve all the RPD's connected to a single FAR.

In the case of the muxponder, there is no method to differentiate replicated versus unique bandwidth. The field aggregation router on the other hand by its very nature has the ability to differentiate packet relationship between DAA endpoints and the packet cores which can be multicast or unicast. This has the effect of significant savings in the overall backhaul bandwidth capacity needed and opens considerable options for managing the use of optics over time, as we will see in section 8.3.

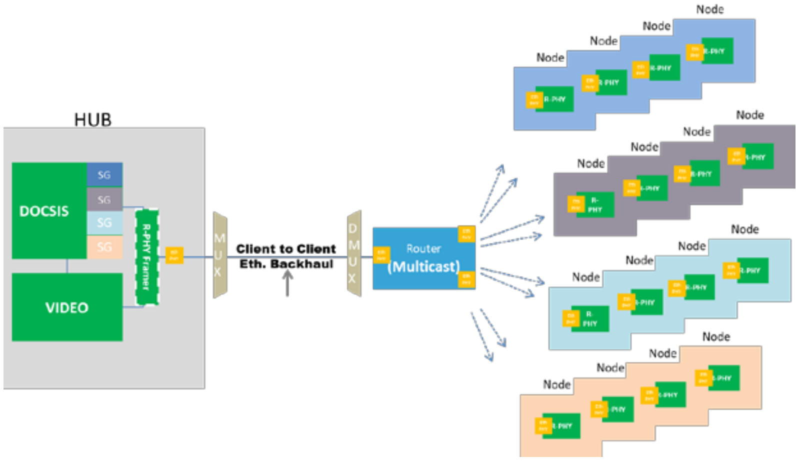

Figure 7 - DOCSIS Service Groups and Multicast Bandwidth

6.6. Converged Access

It is also worth considering that these solutions, in the context of an access network, could converge multiple services beyond what is typically used in Cable, for example mobile or PON. All the solutions here can address this need but there is a difference depending on whether the services are to be aggregated and transmitted (like in the FAR) or transmitted on the same medium but kept separately (like in direct connect or the Muxponder.) In this consideration there is no one answer, but there is a matter of preference for network engineers. This certainly means that there is space for a muxponder solutions that should be investigated, particularly for services that do not replicate bandwidth and run at line rate by contract.

7. Coherent Optics

For reasons of enabling an aggregation element in the field, the MSO community has put considerable effort into making the transmission of 100 Gbps and beyond accessible for the cable access plant. Just recently, CableLabs released a specification for a 100 Gbps ZR solution that is of coherent technology and capable of DWDM channel specificity. This technology also has the ability to cover distances up to 80km and facilitates the use of a FAR, (CableLabs, 2018). There is now an effort to extend a definition to 200 Gbps that also facilitates the natural use of muxponders.

The connectivity of high throughput optics to the FAR falls under Ethernet client to client connections and has been recognized within a greater market opportunity for similar Ethernet signals. This need is also being addressed at the IEEE 802.3 “Beyond 10km” group (802.3, 2018). This effort is tracked very closely in aim to leverage the eventual evolution of ZR 10 Gbps links towards their next transition at 100 Gbps. This will also have the effect of drastically reducing the cost of these optics, ideally in the time frame that will be needed by MSOs.

On the transport side, with distances that can span hundreds of kilometers, there are also efforts that overlap the work being done at IEEE and CableLabs, because the speeds that overlap at 100 and 200 Gbps. For more information see the work being done at OIF, (Forum, 2018), and at OpenRoadm (OpenRoadm, 2018).

8. Traffic Engineering and Backhaul Capacity

8.1. Estimating Uplink Bandwidth Capacity

Understanding and quantifying the difference between provisioned bandwidth and utilized bandwidth is a useful tool when estimating backhaul capacity. For the purposes of illustration in this section, we quantify downstream "provisioned bandwidth" as the DOCSIS bandwidth provisioned per service group. In comparison, we quantify downstream "utilized bandwidth" as the peak bandwidth per service group measured at the WAN port of the CCAP chassis.

In Figure 8 , we see a collection of data that represents several hundred CCAP chassis across a major North American market. The graph illustrates that 78% of the chassis "utilize" between 5 and 15 Gbps of bandwidth as measured at the WAN interface of the CCAP chassis. In an effort to err on the high side, we've elected to use 20 Gbps per CCAP chassis for the comparison to "provisioned bandwidth" per CCAP chassis. Using 20 Gbps per chassis as our representative estimate therefore ensures that we've captured data from 99% of the chassis within the sample network.

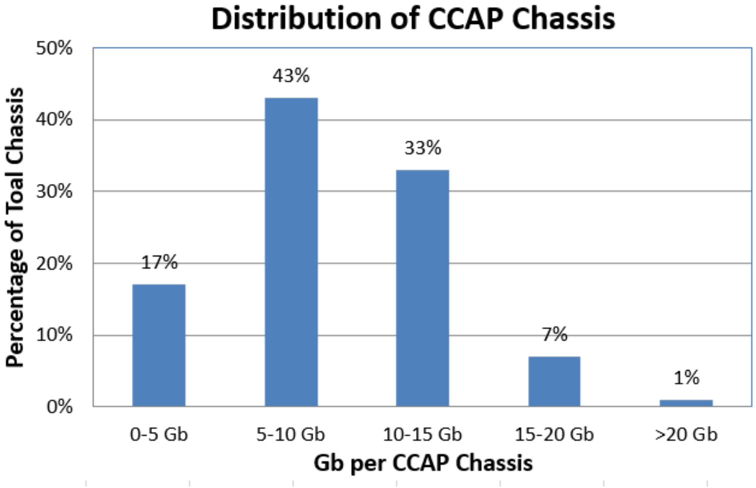

Figure 8 - Utilized Bandwidth per CCAP Chassis

In order to compare the utilized bandwidth to provisioned bandwidth per chassis, one simply needs to take the provisioned bandwidth per service group multiplied by the number of service groups per chassis. Using 32 DOCSIS 3.0 channels plus a 192 MHz of DOCSIS 3.1 OFDM, we could estimate the provisioned bandwidth at approximately 3 Gbps per service group. Using 100 service groups per chassis for our comparison, we can also estimate about a 15:1 ratio between provisioned and utilized bandwidth as represented in Figure 9 below.

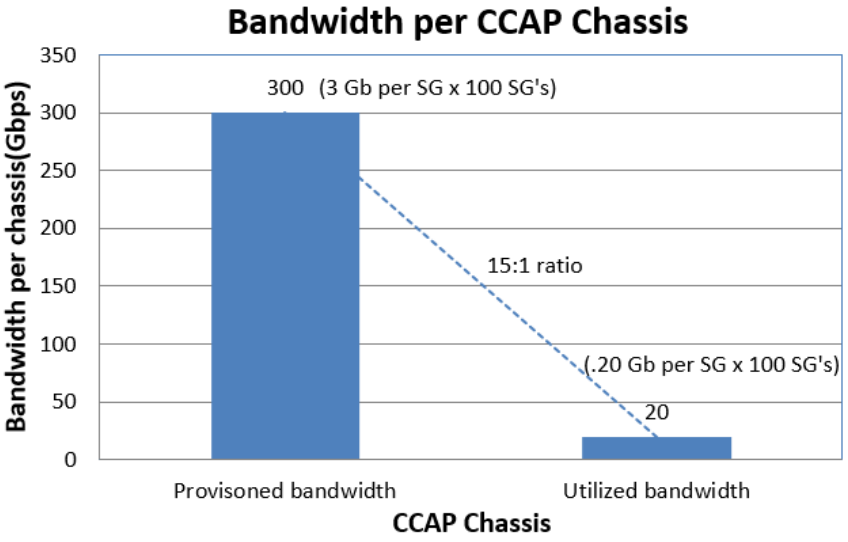

Figure 9 - Provisioned versus Utilized Bandwidth per CCAP Chassis

8.2. Engineering Backhaul Capacity

In the same manner as done in the section 8.1, we use a bandwidth consumption approach to project the backhaul capacity needed to a FAR over time. There are three type of data points to consider when projecting the backhaul capacity. One is the projection of actual subscriber usage over time, another is the aggregate usage of the subscribers being addressed by a packet core service group, and finally the composite usage for the subscribers being addressed by the FAR.

Based on empirical field data, 200 Mbps per Service Group is referenced as the “utilized bandwidth” in Figure 9 above. However, within the backhaul capacity modeling we use 400 Mbps per Service Group as an initial value, to err on the high side. The utilized bandwidth per Service Group is expressed on a per subscriber basis in Figure 10 below and is the initial value used to calculate backhaul capacity. We have also found that a compound annual growth rate (CAGR) of 40% is a representative value for consumption growth, for most MSOs in all parts of the world. The growth from 4 Mbps per subscriber in 2018 at 40% CAGR is therefore shown in Figure 10.

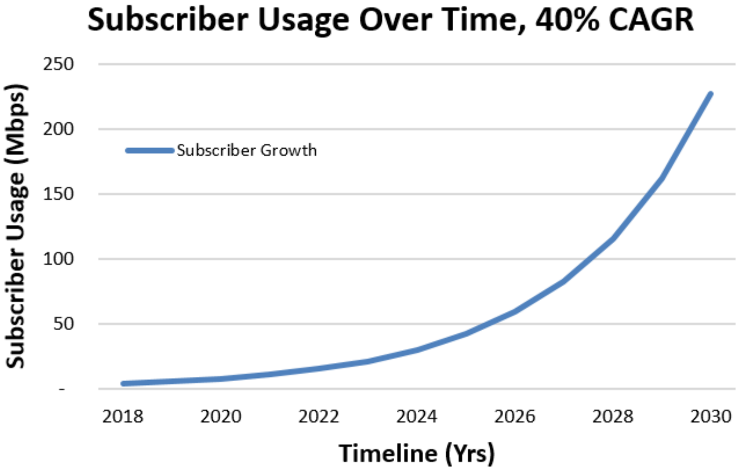

Figure 10 - Subscriber Usage Over Time, 40% CAGR

The usage per subscriber is meaningful in context of its grouping within the servicing packet core. In this example, we use DOCSIS as the packet core driving bandwidth usage to its service groups. In cable engineering circles we refer to these groupings in relation to the topology they are serving, the capability of subscribers addressed by a DAA endpoint, households passed (HPP), and the related take rate, which is the percentage of subscribers actually using the service. We note that in fiber deep applications we see topology arrangements such that each remote PHY is set up for 50HHP. In addition, we note that the typical take rate of service is about 50% so only 25 of the 50 possible customers is using the service. Nevertheless, we use 100% penetration in our backhaul estimates in order to err on the high side.

The subscriber size for the logical service groups of DOCSIS varies. This is driven by many factors, but overall it can range from 25 subs, creating a one to one correspondence between service group and remote PHY device, to several hundred subscribers having a one to many correspondence between DOCSIS service group and RPD’s. In practice, we see up to eight RPD’s per DOCSIS service group, with a common number being four RPD’s per DOCSIS service group. The linear addition of bandwidth from its subtending subscribers gives the usage for the DOCSIS service group. Figure 11 shows the aggregate DOCSIS service group capacity for various subscriber densities, in terms of RPD’s and homes passed, noting that 100% penetration HHP is in fact the number of subscribers.

One interesting note here is that the physical interface of DAA endpoints, including RPD’s of current generation is 10 Gbps, which from Figure 11 shows the time that capacity runs out according to how many subscribers are being serviced. In the case of 400 subscribers that capacity runs out early in the 2023 time-frame, while if you have only 50 subscribers that capacity runs out much later, nearing 2030. This relationship between service group granularity and physical line side capacity of RPD’s is important to understand and can help drive decisions for how the DOCSIS groupings and physical data rates of DAA endpoints will evolve. For example, it might be advantageous to have plans in the next decade to reduce number of subscribers towards one to one correspondence between DOCSIS service groups and RPD’s. It is also possible that a 25 Gbps RPD will be feasible before the end of the next decade and thus not necessarily force a one to one correspondence between DOCSIS service groups and RPD’s. These type of calculations facilitates those decisions.

Figure 11 also allows us to understand the capacity of peak bandwidth per customer. While the subscriber usage is a factual usage value assigned for traffic engineering, there is also the provisioned bandwidth that a Service Provider allocates on a per Service Group basis. The provisioned bandwidth that is shared across the DOCSIS service group, along with the statistical nature of subscriber usage, is what allow customers to peak beyond their individual usage allowance. Note for example the 400 HHP Service Group, where on day one the provisioned bandwidth for the whole service group is 3 Gbps. These 3 Gbps at any point in time, as given by the statistical nature of usage can be available anywhere within that group of subscribers. Naturally, this is the job of DOCSIS, to facilitate this statistical nature of usage. It would be a mistake to put this job on the network itself. In fact, if the network capacity were calculated in this form, with the given CAGR the throughput capacity would grow to unrealistic values very quickly, making a necessary change in physical interfaces beyond 10Gbps very soon for all cases.

Note: for completeness Figure 11 also includes the framing overhead for Ethernet, IP, remote PHY pseudowire and DOCSIS concatenation. The overhead for a PSP frame with 5 DOCSIS segments and MTU size target at 2000 bytes is calculated to be 3.7%.

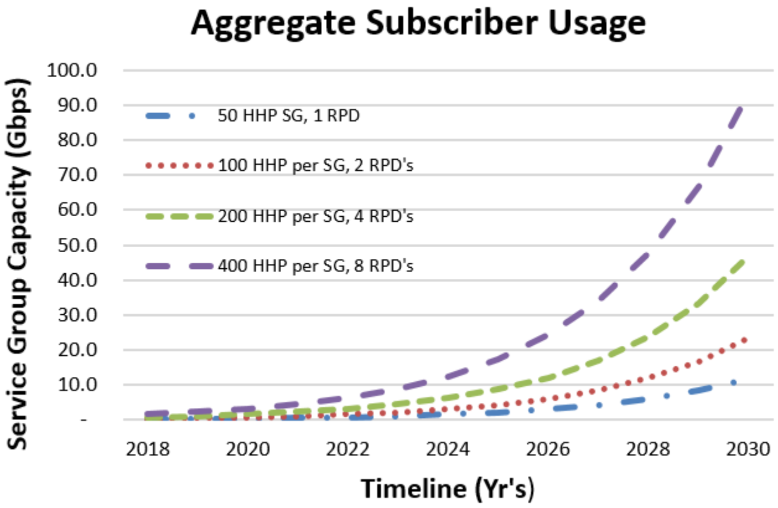

Figure 11 - Aggregate Derived DOCSIS Service Group Capacity Over Time

The FAR serves a collection of endpoints which themselves service a collection of subscribers within a number of DOCSIS service groups. In essence, the data in Figure 11 is linearly added according to the number of service groups and thus subscribers passed through the FAR. Figure 12 then shows the backhaul capacity needed on the uplink of the FAR over time. In its lifetime the FAR might service a number of RPD's, Figure 12 also shows curves for 12 and 24 RPD's in service. For completeness, Figure 12 includes a 1.25 Gbps addition to all users for a broadcast video tier, from the video packet core.

Of greatest interest are the actual values of the capacity needed over time as they drive both the connectivity design of the FAR and the optical solutions. With that in mind if we estimate the lifetime of the FAR between nine and eleven years then we see that capacity solutions needed will be within 100 Gbps, making that the obvious long-term bandwidth target for first product implementations. We also expect that towards the end of the decade there would be use for solutions that expand beyond 100 Gbps. The other item to note is that for the early years of the FAR the uplink capacity can be fully addressed without 100 Gbps solutions. We discuss this management of solutions in the next section.

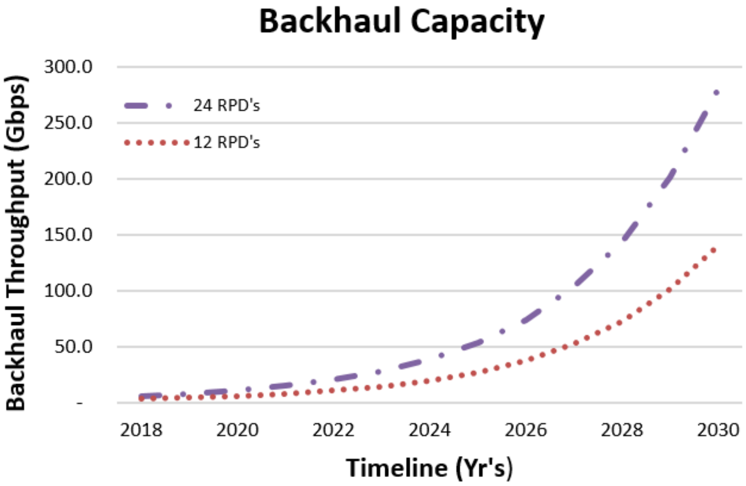

Figure 12 - Backhaul Capacity Needed for Field Aggregation Router, 12 and 24 RPD's Subtended

Note that Figure 12 includes in it the premise of high estimates of very generous 4.0 Mbps usage per subscriber and 100% penetration per DOCSIS service group. The data in Figure 12 therefore gives targets that are shortened in time and larger in capacity than what is more likely to happen, and thus it is an error on the high side as previously indicated. Figure 13 however, is plotted with a starting point of 3.0 Mbps of usage per subscriber and a penetration rate of 75%. Figure 13 is a more likely scenario of what backhaul capacity may be expected. Note that the need for throughput beyond 100 Gbps is further delayed, making the case that maximal capacity at 100 Gbps for the FAR is more than sufficient.

Also, note that backhaul capacity calculations like the ones presented below allow us to calculate the forwarding capacity of the routing fabric. For example, if the maximal capacity needed on the uplink is 100 Gbps, then the downlink capacity needed will be equal, so a 200 Gbps forwarding capacity for a routing chip is sufficient for non-blocking operation.

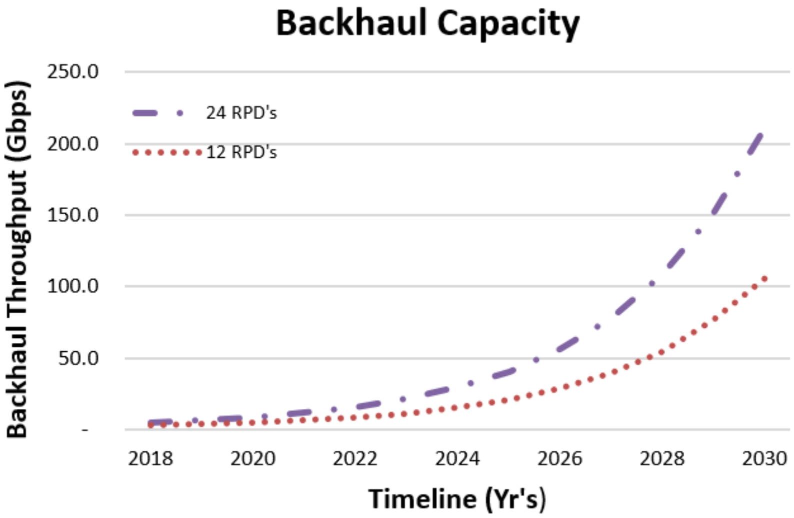

Figure 13 - Most Likely Scenario for Needed Backhaul Capacity, 12 and 24 RPD's Subtended

8.3. Connecting the Uplink

Note that the data we have shown gives an opportunity for a transitionary approach to the uplink. A transitionary approach might be beneficial in leveraging better cost or availability for the optics, or architectural flexibility as the network moves from classic analog to a fully formed digital plant. To that end we show the Figures below.

Figure 14 shows a possible initial uplink configuration. As we see from the modeled data, the first few years of the FAR the uplink can suffice with a few, in this case two, 10 Gbps connections. Leveraging for instance some of the open ports on the router if they are not already used for downlinks.

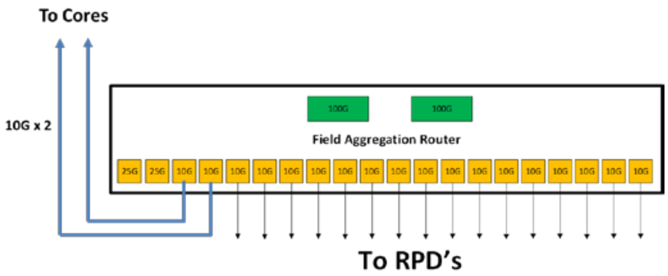

Figure 14 - Initial Uplink Configuration

Figure 15 , shows another possible configuration, where one 25 Gbps link can address the necessary uplink capacity for several years. It is important to note that the optical solutions for 25 Gbps are limited by distance, thereby this solution might not be available to all deployments, but there is a broad footprint within the MSO space where links from hubs to RPD's, and thereby the FAR are within 20 km range.

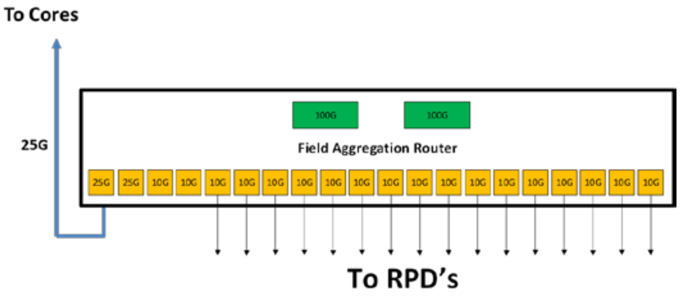

Figure 15 - Secondary Uplink Configuration

Figure 16 , shows the final transition leveraging the large bandwidth available at 100 Gbps. Note that in the router representation there are two 100 Gbps connections which allows for redundancy as needed.

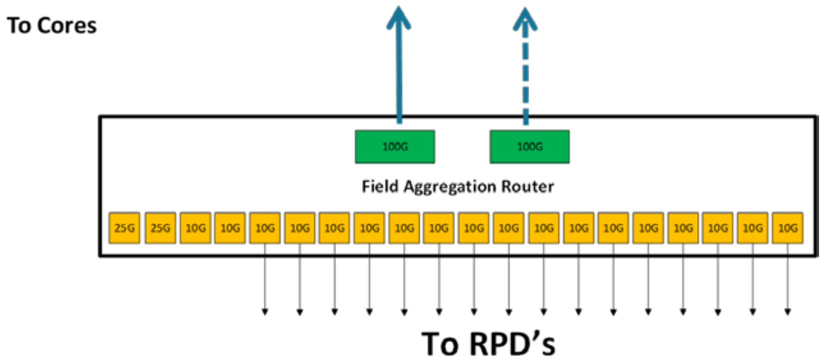

Figure 16 - Third Uplink Configuration

Note also that the Figures above give a method for a more granular approach to building the backhaul. For example, if in year 8-10 of the product an MSO has modeled usage to be below 120 Gbps, but generally over 100 Gbps. That connectivity then is straight forward by using a 100 Gbps connection and a handful of 10 or 25 Gbps connections. This could significantly minimize the cost of the optical connections while not taking away from the functionality of the FAR.

With regards to cost specifically, one of the main challenges to remote aggregation devices, like the FAR or the muxponder, is the cost of the uplink optics, particularly at the 100+ Gbps rates. The challenges of having optics that can generally work beyond 40 km and of wavelength specificity make the technology choices for high rate optics very limited, basically only to the use of coherent optical solutions. The drawback of coherent optical solutions is cost however. One way to limit cost then is to purchase these units as volumes and scales make them available. Another way to limit the cost of the coherent optics bought is to have them align with what will be realistically used, in this case not much more than 100 Gbps. In other words, if someone upsells 400 Gbps capability, there is very little value in that proposition as it might never be used in the timeframe of that product.

Conclusion

Distributed Cable Access Architectures have created a new form of backhaul market that did not exist just a few years. While the DAA backhaul has similarities to what is done for other non-cable access services (like Mobile backhaul, or PON backhaul) the nature of signal distribution for DOCSIS, video and other supporting packet cores allow for unique solutions. As the DAA backhaul is built, it would be best to do so in a cost-effective manner with a general toolset that already exists within the networking world.

This white paper has covered cost and technology comparisons for architectural options for the DAA backhaul in context of bandwidth growth over time. We have also detailed the optical and networking implementations needed to address DAA backhaul. Some of the specific topics that were covered included switching and routing in hubs and HFC nodes, distinctions in DOCSIS and video network transmission, and the applicability of subscriber usage when engineering the backhaul capacity. We also explored the effect of unicast and multicast content on provisioned bandwidth, applicability of OTN framing, implementation of direct detect and coherent optics, and relation to ITU and IEEE standard bodies.

After reading this white paper, the cable operator should have obtained the practical perspective necessary to compare the various DAA architecture options with the view towards deciding for deployment and long-term evolution.

Abbreviations

ASIC	Application Specific Integrated Circuit
BNG	Broadband Network Gateway
CAGR	Compounded Annual Growth Rate
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CPU	Central Processing Unit
DAA	Distributed Access Architectures
DOCSIS	Data Over Cable Service Interface Specification
DWDM	Dense Wavelength Division Multiplexing
FAR	Field Aggregated Router
FPGA	Field Programmable Gate Array
Gbps	Gigabit per second
HHP	Households Passed
IP	Internet Protocol
L2TPv3	Layer two Tunneling Protocol version 3
MAC	Media Access Control
Mbps	Megabit per second
MSO	Multiple System Operator
MTU	Maximum Transmission Unit
OFDM	Orthogonal Frequency-Division Multiplexing
OIF	Optical Internet Forum
OTN	Optical Transport Network
PON	Passive Optical Network
RMAC-PHY	Remote Mac and PHY
RPHY	Remote - PHY
TDM	Time Domain Multiplexing
WDM	Wavelength Division Multiplexing
ZR	Long Range Optic, 80km.

References

802.3, I. (2018, May). *Beyond 10km Adopted Objectives, Study Group*. Retrieved from http://ieee802.org/3/B10K/project_docs/objectives_180521.pdf

CableLabs. (2018, June 29). *P2P Coherent Optics Physical Layer 1.0 Specification*. Retrieved from https://apps.cablelabs.com/specification/P2PCO-SP-PHYv1.0

Forum, O. I. (2018). *Current Work done at OIF*. Retrieved from http://www.oiforum.com/technical-work/current-oif-work/

G.694.1, I. (n.d.). *Spectral grids for WDM applications: DWDM frequency grid*. Retrieved from https://www.itu.int/rec/T-REC-G.694.1-201202-I/en

Microsemi. (2017, March). *Microsemi Enables Terabit OTN Switching Cards for Flexible Optical Networks*. Retrieved from https://www.prnewswire.com/news-releases/microsemi-enables-terabit-otn-switching-cards-for-flexible-optical-networks-300608509.html

OpenRoadm. (2018). *Open Roadm MSA*.