Moving Towards the Light:

Migrating MSO FTTP Networks to a Distributed Access Architecture

A Technical Paper prepared for SCTE/ISBE by

Phil Miguelez
Executive Director, Network Architecture
Comcast
Comcast Center
Philadelphia, PA
215-286-1126
phil_miguelez@cable.comcast.com

Table of Contents

List of Figures

Introduction

Cable coaxial systems and more recent HFC (Hybrid Fiber-Coax) networks have been expanding across the North American continent for nearly seventy years. MSO coax access links pass more than 85% of the single family and multi-dwelling properties in the US. Cable systems have evolved over the decades from basic video carriage to multi-services transport including voice, high speed internet, and data services for both residential and commercial customers.

In 1982 the landmark United States v AT&T anti-trust settlement broke up the Bell System, which led to the creation of seven independent Regional Bell Operating Companies (RBOCs) just two years later. The Regional Bells no longer had the monopoly protection that the telephone network enjoyed in the past. Competition accelerated in the mid 1990's with the growth of the internet. Cable systems launched DOCSIS® which allowed both digital voice and data services over the MSO coax networks. The twisted pair copper wires of the phone companies reached almost every home in the US, but this basic network of copper lines had been in existence for a hundred years and had limited high speed data capacity. In order to survive, the Regional Bells started to merge and consolidate their coverage areas. They also began the process of migrating their networks from "Plain Old Telephone Service" (POTS) over twisted pair copper lines to fiber optics, in order to compete with the cable and satellite competitors that were rapidly growing by offering expanding video on demand content and faster data.

Verizon initiated FiOS (Fiber Optic Service) in 2005 and built out FTTH networks covering 18 million households, primarily in the northeastern portion of the US. AT&T developed U-verse, a hybrid network consisting of a fiber to the curb transport layer converting to a DSL over twisted pair access link to the home. The cable operator alternative to FTTH PON was RFoG (RF over Glass). RFoG involved very little new technology. The analog modulated laser that had been used to transport the RF channel load to the HFC node was now optically split to feed 32 mini nodes that were located at the subscriber homes. RFoG provided no additional bandwidth or capacity compared to HFC, but it did satisfy the demands of new home developers that insisted on a future proof "fiber to the premise" solution instead of coax. RFoG struggled to gain traction then (as now) due to a number of weaknesses with this technology – limited bandwidth (BW), limited capacity, higher construction costs associated with fiber, and the nagging concerns about optical beat interference (OBI) that was always hard to detect, initially, and had no known cure. The only major benefit of RFoG was that long fiber drops provided a much lower cost solution than HFC in rural serving areas with low homes per mile. The paucity of homes in these locations also provided lower upstream traffic congestion on the network, so the statistical chances of generating OBI was significantly reduced. Many smaller market MSOs serving these communities in the mid-west and southern areas of the country have never had a serious problem with OBI, and swear by RFoG.

The housing boom that generated the increasing interest in fiber to the home ended suddenly as a result of the 2008 Great Recession. New single-family home construction came to a standstill and financial tightening halted all but "business as usual" network maintenance projects.

Post-recession, as the financial and home construction markets recovered, a shift to multi-dwelling unit (MDU) new build projects (instead of single home communities) was clearly evident. But the most significant change in broadband access was the announcement in 2010 that Google Fiber would soon bring Gigabit fiber connectivity to a number of communities across America, starting with Kansas City, KS. Google's entry into the broadband delivery arena was the catalyst for a number of competitive changes. Google Fiber threatened both telcos and cable operators. Gigabit service was, at that time, out of reach for telco DSL or MSO DOCSIS networks. Cable franchise agreements were now being re-examined as cities lined up to be the next potential Google “Fiberhood.” Emboldened by Google’s negotiating power to extract concessions from municipalities -- such as access to utility poles, and expedited construction permitting -- the lowered barriers to entry encouraged a new wave of fiber overbuilders to enter the market. These new competitors began to target the incumbent operator’s MDU footprint, using the Google template.

“Gigafy America” became the rallying cry from these new competitors.

Cable MSOs now have to plan a fiber strategy -- or watch these competitors take away new greenfield opportunities, as well as existing MDU and bulk contract, gated community subscribers.

Migrating MSO FTTP Networks to DAA

Competitive Pressure

As many as 25% to 70% of US properties passed by MSOs can be classified as MDUs. The broad definition of MDU includes traditional apartment buildings and multi-use buildings containing commercial businesses, plus residential units and a wide range of gated properties -- such as “over 55,” golf course, and waterfront communities. What makes these MDU properties unique is that, for the most part, the developer of the building or community negotiates a long term, bulk contract with a service provider that typically guarantees 100% subscriber penetration. Service provider decision power for these situations is usually concentrated with the property developer or HOA (Home Owners Association). Many of these owner associations quickly learned to hire technology consultants to evaluate the competitive solutions being proposed, so as to force improvements in data delivery rates and other concessions. Gigabit download speeds are an attractive incentive to high-end consumers, who already use lots of internet-ready devices, in almost every room. Plus, “fiber-ready” buildings allow developers to raise the selling price of each condominium as “future proofed” homes. As a result, most developers began to demand fiber-only construction for their new greenfield properties.

New, startup fiber service providers seem to come into existence every week, competing for both new building developments and existing MDU contract renewals. And, for the first time, these upstart overbuilders are not competing with increased video content offerings or HD picture quality, but with offers of Gigabit-speed data -- at extremely aggressive pricing. Figure 1, below, shows the different competitors that have expanded into new cities and in some cases challenged existing Comcast properties.

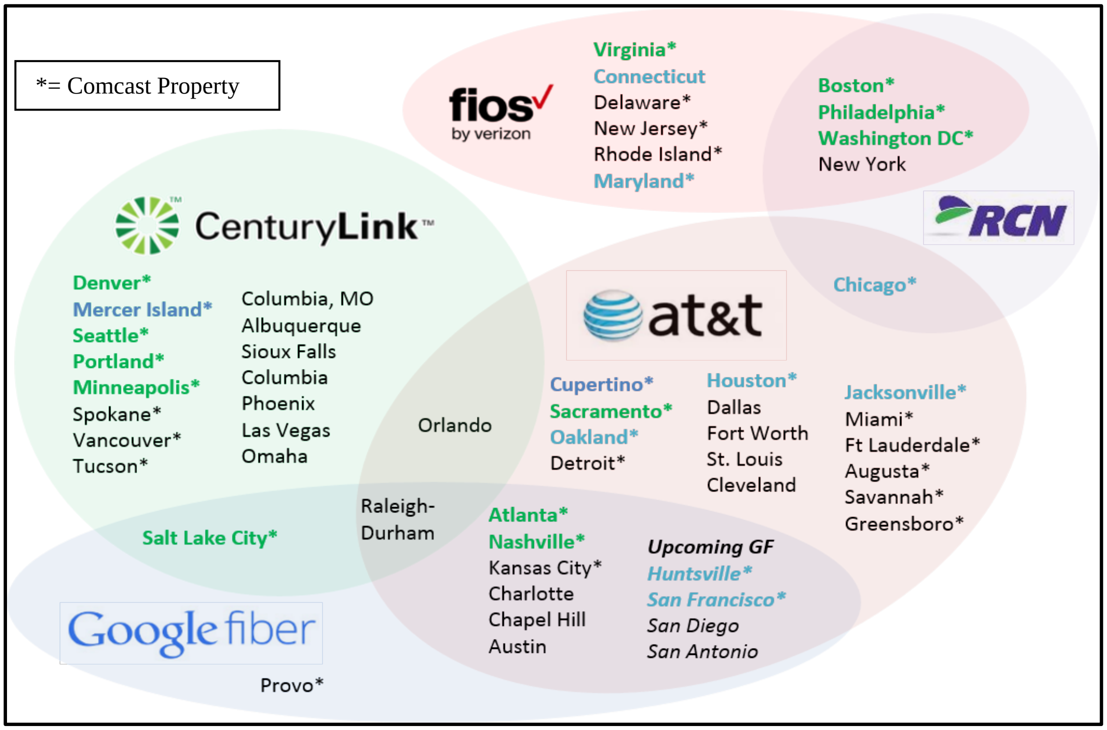

Figure 1 – Fiber Competitors Are Now Challenging Every Greenfield Opportunity

Which PON to Choose?

To meet the demands of greenfield developers for FTTP broadband solutions, Cable MSOs first need to select a PON (Passive Optical Network) operating system. RFoG was created to be the fiber transport equivalent of HFC. RFoG provides the same architecture as PON, while remaining completely transparent to back office systems and subscriber provisioning. The major drawback is that RF over Glass is limited to the same RF capacity as HFC. As a result, RFoG has the same asymmetric DS/US data capacity as HFC and cannot match the symmetric Gigabit speeds that fiber competitors are now offering.

Option 1: Gigabit Passive Optical Network / GPON

Google Fiber and all of the smaller overbuilders selected GPON as their initial architecture. GPON is a mature solution, deployed in volume worldwide, and adopted by all of the major North American telcos, including Verizon and AT&T.

The GPON specification was created by a working group of major telecommunications service providers and system vendors organized as the Full Service Access Network working group (FSAN). The spec was then issued by the International Telecommunications Union (ITU) as the ITU-T G.984 standard. The GPON standard allowed for larger, variable length packets than previous FSAN specifications, to provide 2.488 Gbit/s of downstream bandwidth and 1.244 Gbit/s of upstream bandwidth.This available data capacity is minimally sufficient to support a Gigabit symmetric HSD tier rate, depending on the PON architecture split ratio and the level of congestion on the network.

However, GPON has other limitations that impact MSO adoption. High among these is the lack of a DPoE mediation layer, which is critical for maintaining a common provisioning protocol across all supported services. While it is possible to create an equivalent DOCSIS mediation layer (DML) for GPON, the silicon and software development effort is estimated to be 12 to 18 months. Without support from the majority of MSOs, the cost and risks of investing in and waiting for DPoG were deemed to be too high. Another weakness is the limited interoperability between vendors. Although GPON is a common standard, many vendors customize the OLT and ONU's to achieve optimized performance as a competitive sales tool. An inability to mix and match vendors within the network results in a substantial barrier to lower costs.

Another hurdle in adopting GPON is the makeup of the FSAN / ITU-T standards body that manages the specification. FSAN is a very telco-centric organization. Many MSOs are members of FSAN, but the major contributors are telcos and the vendors that support them. It would be difficult to get substantial traction within FSAN without significant participation and consensus from a block of cable MSOs.

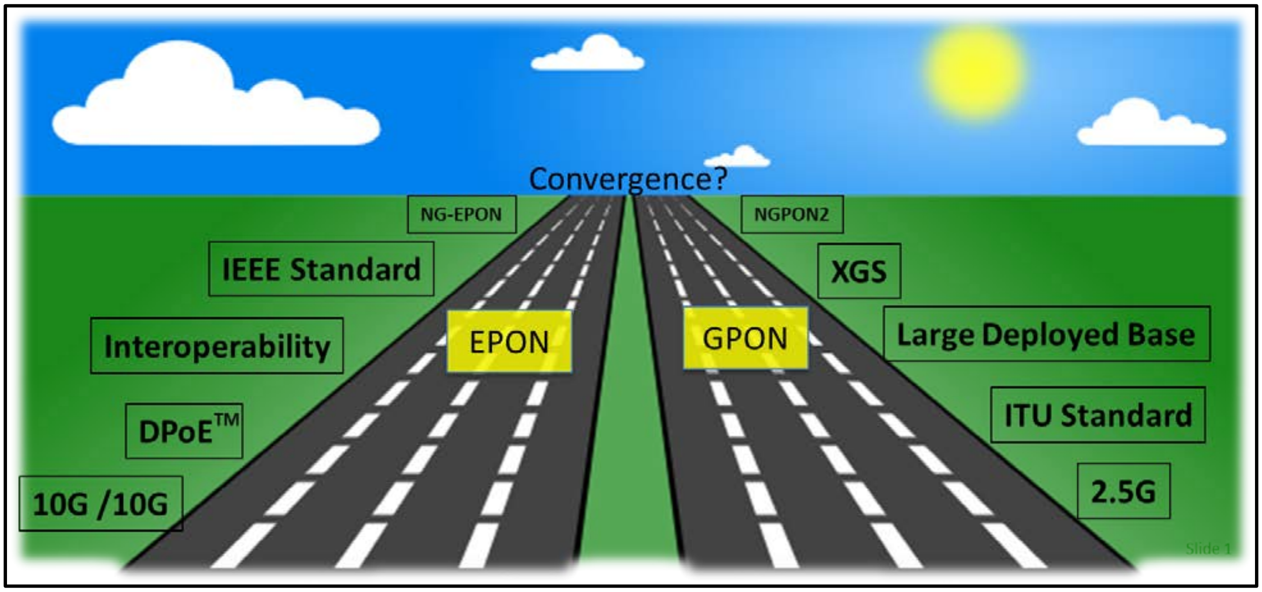

Figure 2 – EPON / GPON Comparisons (Courtesy of Curtis Knittle, CableLabs)

Option 2: Ethernet Passive Optical Network / EPON

Another option is Ethernet PON, developed by the Institute of Electrical and Electronic Engineers (IEEE) -- a leading standards organization that has been very supportive of MSO requirements. The IEEE 802.3ah and 802.3av EPON standards provide 1 Gb symmetric or 10 Gb symmetric and asymmetric options. EPON has also been widely deployed, DOCSIS provisioning is supported via DPoE ™, and there is proven interoperability between vendor equipment. The drawbacks regarding the currently available EPON standards are that GEPON (Gigabit Ethernet Passive Optical Network) is not capable of multi-subscriber gigabit symmetric service,10G EPON is not widely deployed in North America, and is considered too expensive, primarily due to the higher 10G optics cost.

Selecting which PON technology to deploy is always a balance between cost, operational complexity, and its ability to meet current and future capacity requirements. The rapid progression of high-speed data consumption, tracked by Nielsen and others for the past 30 years, continues to increase at a compound annual growth rate of close to 50% year over year. At current growth rates, the projected data capacity needed will approach 10 Gb by the early 2020's.

GEPON and GPON are the lowest cost PON solutions, because of lower data rates and uncooled optics used in optical network units (ONUs). Plus, scale economics play a role, because of annual, worldwide market volumes for ONUs. GPON's downstream data capacity is 2x higher than 1G EPON, giving it a clear advantage. But GPON provides no competitive advantages for MSOs, and based on the previously discussed projections regarding HSD CAGR, the relatively low, 2.5 Gb downstream data capacity will have a limited service lifetime in all but lower tier rate, best effort applications. The GPON standard also specifies a lower, asymmetric data rate for upstream traffic. This significantly reduces the usefulness of GPON in commercial services applications, which usually require symmetric data transport. NGPON2 is a more recent approved ITU standard that is capable of providing 10 Gb symmetric transport, but its use of ONU tunable wavelength filters makes it very expensive, particularly for residential subscriber deployments. As a result, NGPON2 has not been deployed in any operator systems to date.

10G EPON is the most recently deployed IEEE PON technology and has useful capacity to support multi-Gigabit tier rates for both downstream and upstream traffic. 10G is a significant technology leap over GPON, and provides a competitive edge against network overbuilders. It should have a serviceable life of several years based on the current CAGR models.

The current North American market volume for 10G PON lasers is still relatively low. This is one of the reasons for the higher cost of these devices. The announcements that Comcast and other major MSOs are planning to adopt 10G EPON has already produced improved pricing, as optical vendors position their products for the emerging cable operator FTTP business. 10G PON optical transceivers are also mainly available today only as pluggable XFP or SFP+ packaged devices. Conversion to a BOSA-style package, typically used in most 1G EPON and GPON ONU deployments, will help to further lower the cost curve of these transceivers.

Choosing which PON technology Comcast would pursue was a long and arduous process. As detailed above, GPON offered a low cost, mature technology that could immediately defend against the aggressive competitors that were eroding our MDU footprint. 10G EPON had not yet been deployed in a cable MSO network, and other than China, had only been selected for a few commercial installations. On the plus side, 10G EPON's DPoE capability would provide operational similarity with the existing DOCSIS back office, and 4X or higher data capacity than GPON.

In the end, after several months of vigorous internal debate and analysis 10/10 EPON was selected, based on two primary factors: First, symmetrical 10 Gb allowed the capacity to offer multi-Gigabit HSD for both commercial and residential subscribers, which was not possible with GPON. Second, and even more convincing, was an analysis that took into consideration the growth rate of data in the network. Deploying GPON would result in lower capital costs, but the serviceable lifetime of GPON was limited, due to its lower data rates. The estimated network upgrade costs to eventually replace GPON with a 10 Gb solution was significantly higher than going all-in with 10G EPON on day one.

The Role of PON in Cable MSO Networks

The initial move into FTTP by cable operators was driven by the threat of competition. While some markets will require a fiber to the home solution, the majority of the MSO residential network footprint will be adequately served with DOCSIS and HFC for many years to come. On the business services front, FTTP PON has been discussed for years as the eventual solution, when the available cable plant dark fiber is exhausted.

The introduction of D3.1 has raised new questions regarding the need for FTTP beyond new greenfield and competitive threat applications. MSOs have traditionally met the need for increased capacity by a combination of RF bandwidth expansions and node splitting.
Eliminating analog video channels, digital video compression technology, and higher order QAM modulation formats supported by DOCSIS has allowed operators to keep pace with the BW requirements of HD channel growth and the ever increasing array of new channel content. The arrival of D3.1 and OFDM has the potential of increasing the capacity of the HFC access network to 10 Gb or higher with a corresponding expansion in the RF system bandwidth. Migrating the network to an N+0 architecture paves the way to implementing Full Duplex DOCSIS and the possibility of gigabit symmetric service when it becomes available in a few years. This very real and near term technology will make HFC DOCSIS cable systems the equivalent of a 10G FTTP fiber network.

How long will Gigabit-per-subscriber data rates keep pace with the continuing growth rate of data? Maintaining the viability and lifetime of the HFC network is always the main priority. At what point does the crossover occur, between Fiber Deep and FTTP?

If the growth of data continues at its historical rate, an N+0 DOCSIS 3.1 network, supplemented with FTTP to support high end residential and SMB / home office commercial users, may be sustainable for several years. At some point, though, within the next 10 years, HSD growth will begin to require per-subscriber speeds that exceed the bandwidth limitations of balanced, well designed, N+0 fiber-deep D3.1 HFC networks. Fiber to the tap and similar proposed solutions result in an order of magnitude higher number of active devices -- which will be complicated, power hungry, and need to be provisioned and managed. A Fiber Deep N+0 architecture provides the natural demarcation point for the HFC-to-FTTP transition when needed. Changing the node to a modular virtualized OLT location and adding fiber drops, or possibly wireless access transceivers in place of RG6, could likely be the next evolutionary phase for MSO networks.

The next leap in PON technology is currently being debated at both IEEE and ITU FSAN working group meetings. Most likely the optical device advances that are now driving data center transport rates will be adapted for longer link networks. If so, then data rates or 25, 50, and 100 Gb are not very far off from being a reality.

10G EPON Network Architecture Challenges For MSOs

MSOs face a number of design and operational hurdles on the path to creating a workable FTTP playbook. The list below identifies the most significant issues. Many of these will be expanded on in later sections of this paper.

10G EPON FTTP Architecture Challenges:

- Available dark fiber rapidly decreasing
- Link reach >> 20 km
- Silicon dependency (one vendor provides 90% of all OLT chip sets)
- High cost of fiber construction and home wiring
- New Customer Care, Billing, Tools, Operations, Processes
- Fiber Handling (Field Tech training, tools, diagnostic equipment)
- Homeowner and MDU Consultants
- "Gigabit" speed, speed tests, servers, dependencies
- HSD tier offering in mixed DOCSIS and EPON footprints
- External positioning for FTTx (D3.1 versus fiber comparisons)
- Voice and Video integration (Migration to all IP)

FTTP in Cable Plant Realities

HFC networks are extremely efficient. By relying on active elements in the network, the number of customers that can be reached and supported with triple play services (voice, video, and HSD) in the access last mile link provides the lowest total cost-per-subscriber compared to any competitive solution. As a result, cable networks have been built with link reaches ranging from 20 km to over a 100 km, with the lowest amount of available fiber depending on the homes passed density of the geographic area, and individual operator design practices. Access network fiber availability for MSOs is further challenged by the success of business services. Available dark fiber has been used for adjacent market high revenue applications such as cell tower backhaul and commercial metro Ethernet links, reducing the amount of fiber accessible for a FTTP migration.

Classic PON architectures, whether centralized or distributed, require an extensive amount of fiber. The reference optical distribution network (ODN) design for FTTH is a 20 km passive fiber link to reach 32 subscribers. Higher split ratios -- to increase the number of subscribers per OLT port, connector attenuation loss, and other design margin factors -- reduce the maximum reach of the PON link. Figure 3 illustrates the link reach possible with a completely passive optical network.

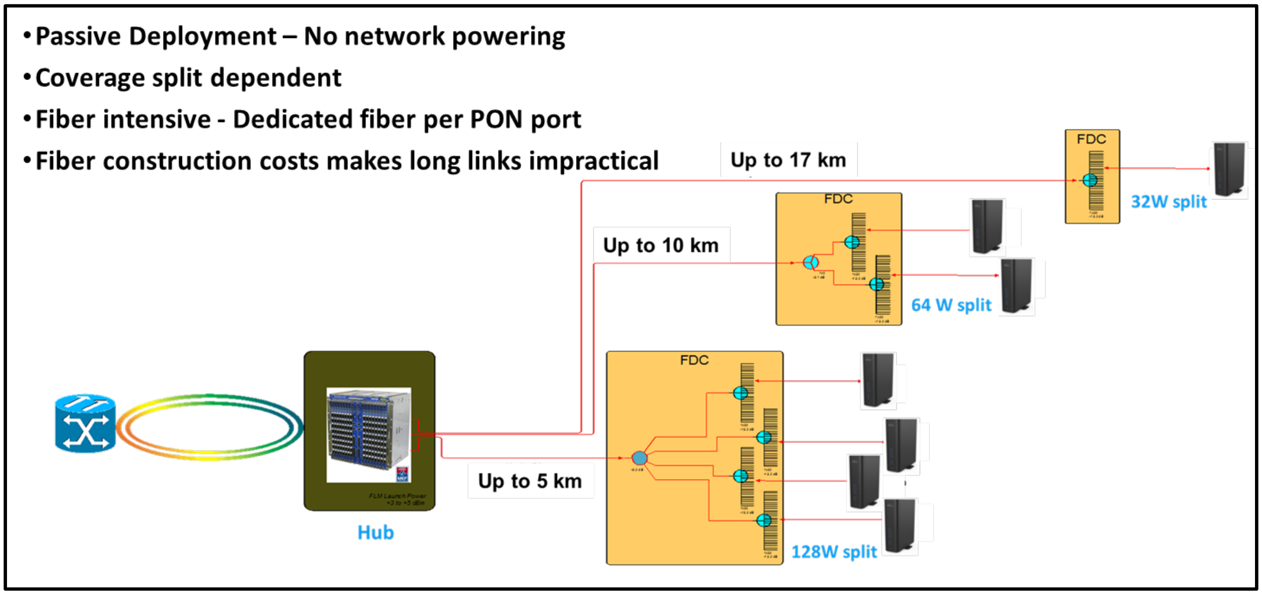

Figure 3 – FTTP Deployment Direct Feed Coverage

Using traditional PON ODN construction, most customers can only be reached by reducing the splitter ratio. This means few subscribers per port, and therefore much higher OLT port counts and cost per connected subscriber. Fiber construction represents one of the highest network deployment costs, particularly in existing legacy brownfield applications. Therefore most cable operators target FTTP primarily for greenfield applications, with the exception of a legacy serving area that is under competitive threat.

While the material and installation costs of fiber and coax are almost identical in new, greenfield construction cases, there are still major differences. Fiber splicing and connectorization is the most significant. Fiber cable is available with several standard fiber counts, from 12 fibers to 288 fibers. Splicing each of these individual fibers requires special equipment and training. Each fusion splice can cost approximately $25, depending on local construction rates in the build area. In the case of a 288-count fiber bundle, fusion splicing each strand can take several days to complete and verify. Splicing connector pigtails onto a fiber bundle at a distribution cabinet or field splitter is an equally expensive and time-consuming endeavor. Ribbon fiber and multi-fiber MPO connectors are being considered as a means to reduce the time and labor needed by these large-scale optical connections. The drawback is the service impact implications of repairing or replacing a multi-fiber connector.

HFC plant design has typically been 60% aerial construction and 40% underground. In new greenfield properties, 100% of the construction is underground. Underground construction costs can be 2X to 3X higher. This includes not only trenching the fiber, but also installing buried chambers for fiber field connections and pedestal fiber management.

Fiber drop cables to the home represent another significant cost premium to HFC deployments. HFC drop cables are typically limited to 150 feet maximum, due to the high RF attenuation of RG6 coax. As a result, HFC taps are distributed along the access link and sized to provide connections for 2, 4 or 8 homes depending on the home density in the particular serving area. When a new customer requests service, the drop cable usually only has to be trenched across one, or at most 2, building lots. Fiber drop cables, on the other hand have extremely low optical attenuation, and can span a kilometer or more depending on the location of the field splitter or fiber distribution cabinet (FDC). In this case, the drop cable must be trenched across several lots, which increases the drop's construction cost, not to mention the public relations issues caused by construction crews digging up people's yards.

Home wiring is the next major challenge for FTTP construction. In many new greenfield construction sites, and particularly MDUs, a media panel is usually built-in at a central location within the home. It aggregates all of the fiber and RF distribution connections, and provides a storage location for the ONU / Gateway. Unfortunately, the range of variations from one development to another is wide. There are also significant differences in opinions regarding which party owns or is responsible for installing the media panel and fiber wiring.

The RFoG Overlay

Today, FTTP EPON is primarily used only to deliver high-speed data. In order to provide the same triple play experience over FTTP that subscribers receive with the RF coaxial system, an HFC or RFoG overlay is required. HFC overlays are most likely in cases were a legacy, bulk MDU property is being converted to FTTP, triggered by a competitive contract renewal situation. In the majority of greenfield cases, a fiber-based transport solution is mandated, which means an analog / QAM RFoG overlay must be deployed.

An RFoG overlay is primarily needed to support analog and QAM DOCSIS video transport. The increased complexity and cost of multiple CPE devices needed for each fiber-connected home is a significant deterrent to the growth of FTTP in MSO networks. As shown in Figures 4 and 5, an RFoG overlay implementation for FTTP EPON requires additional fibers and OSP equipment. OBI mitigation is also needed in most cases.

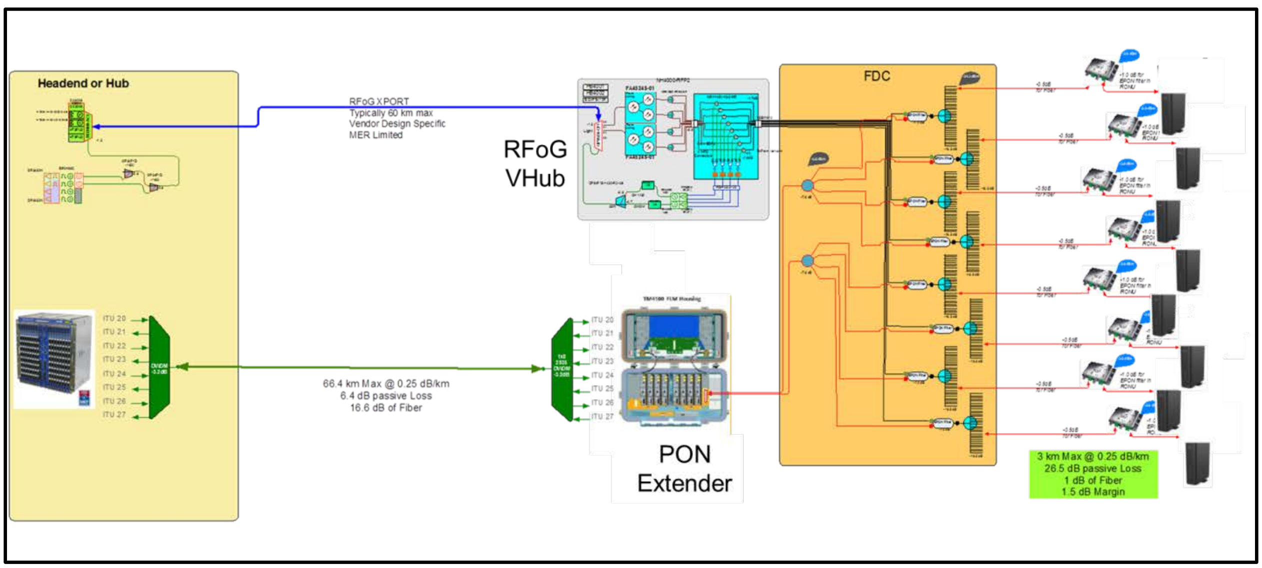

Figure 4 – RFoG / EPON Overlay Design Configuration

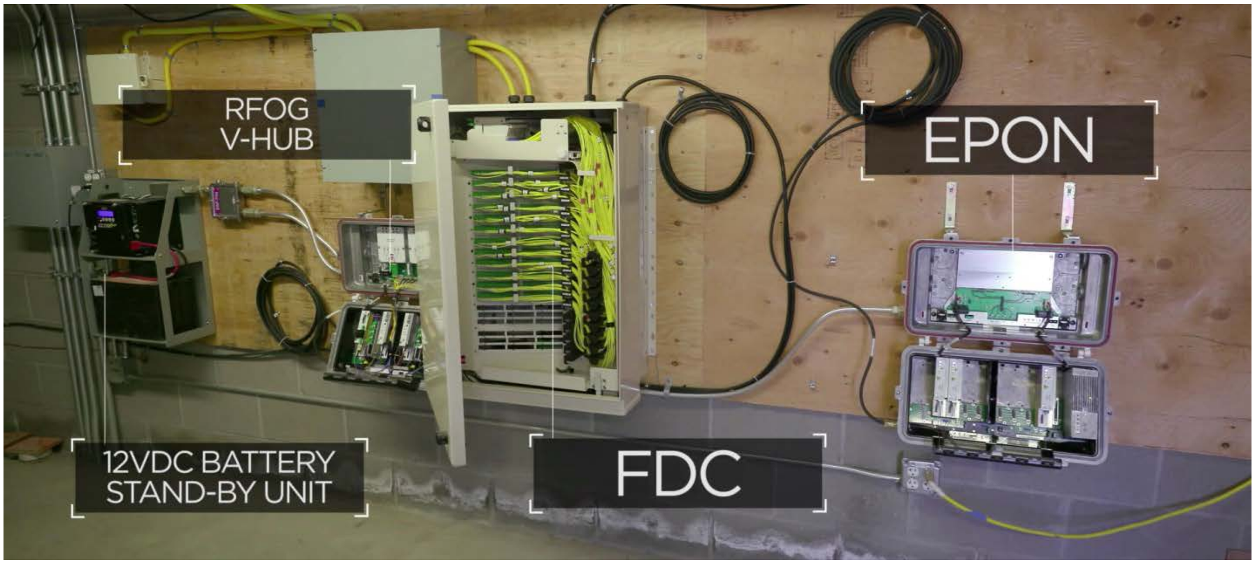

Figure 5 – RFoG / EPON Overlay Actual MDU Deployment

An RFoG Vhub is used to amplify the DS wavelengths, in order to compensate for the PON splitting losses. The Vhub also provides for aggregation and OEO analog to digital Ethernet conversion of the RFoG ONU US signals.

In an RFoG overlay, the RFoG and EPON wavelengths are muxed together on the access drop fiber at the fiber distribution cabinet. The RFoG ONU must support an added pass through optical coupler to allow the EPON wavelengths to be coupled from the RFoG ONU to the EPON gateway.

EPON is ultimately capable of duplicating every service flow provided by an all-QAM, DOCSIS provisioned system. The issue is updating, testing, and validating the OLT / ONU software to provide not only the IP packets containing voice, video, and data information, but also a transparent back office operation that is compatible with the existing CMTS. DpoE is only one condition of this compatibility. Other features and service flows within the DOCSIS-provisioned network also need to be replicated for the PON environment. At this time, 10G EPON equipment providers and MSO software developers are working to complete the qualification and field trials that will make a deployable, all IP solution possible by the end of 2017.

Triple play IP transport will not only allow for the elimination of analog RFoG overlays, but is also a necessary step in eventually migrating to a SDN/NFV network for both D3.1 HFC and EPON FTTP.

Cable Plant Designs Are Not PON Friendly

HFC networks consist of a headend, supported by several hub locations, arranged in a star or ring configuration, that serve the local population centers of subscribers. These primary hubs contain the CMTS, local channel content, ad insertion and optical / RF access node links for serving areas ranging from 60K to over 100K homes passed. As new communities developed or expanded, they sometimes exceeded the capacity of the existing primary hub. Smaller, secondary hubs have been added, by leasing space, or building a small facility sized for the new community being served. These secondary hubs mainly support the access edge for 15K to 30K HP. This system of headends and hubs results in over 85% of all node and subscriber locations being within 20 km of the nearest hub.

While this proximity to a local hub seems to be ideal for PON architectures, the complication is that many secondary hubs are too small to support the rack space needed for a mainframe OLT. Also, because the secondary hub is sized to serve a smaller number of homes, the power and HVAC capabilities of the hub are not adequate to support the added load of the OLT.

Secondary hubs are linked to the main hub or headend via a primary and secondary fiber trunk connection. The problem is that one or both of these link paths, when added to the subscriber access fiber link, may exceed the PON ODN optical budget. This is illustrated in Figure 6, which represents an actual FTTP trial deployment case.

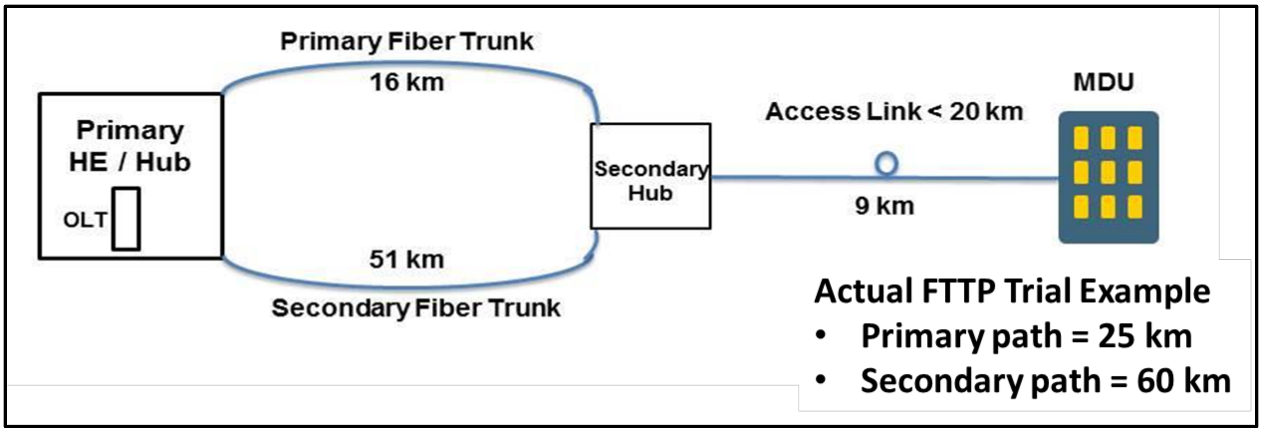

Figure 6 – Secondary Hub Limitations Can Create Extended Link Reach Situations

Although the MDU example in Figure 6 is only 9 km from the local hub, the OLT was located at the primary hub, with access to a larger number of potential FTTP subscriber sites. The additional path distance of the primary and secondary fiber trunks exceeded the PON optical power budget. Additionally, in this RFoG / EPON overlay case, the number of fibers needed for RFoG DS/US links, plus the EPON links for two buildings, exceeded the available dark fiber. This trial proved the need for a distributed access solution, and accelerated the qualification of a 10G capable PON Extender.

PON Extenders – An Interim Distributed Access Solution

A PON Extender is essentially an OEO repeater. Extenders are typically designed as strand-mount, clamshell node housings that can be cable plant-powered. Most extender designs usually support up to 8 OLT output ports. The traditional mainframe OLT PON optics are replaced with ITU grid DWDM optics, which can be muxed onto a single fiber and transported much further than the 20 km limitation of PON optics. DWDM optical transceivers are widely available, with reach capability extending up to 80 km. At the PON Extender side of this link, the DWDM wavelengths are de-muxed, received and then converted to PON wavelengths. The output of the PON Extender is the same as the output of the OLT.

PON Extenders do not manipulate the data or provide any management functions. They can provide basic diagnostic information on the status of the unit, and the optics in particular. The number of ports used can be sized to the property being served.

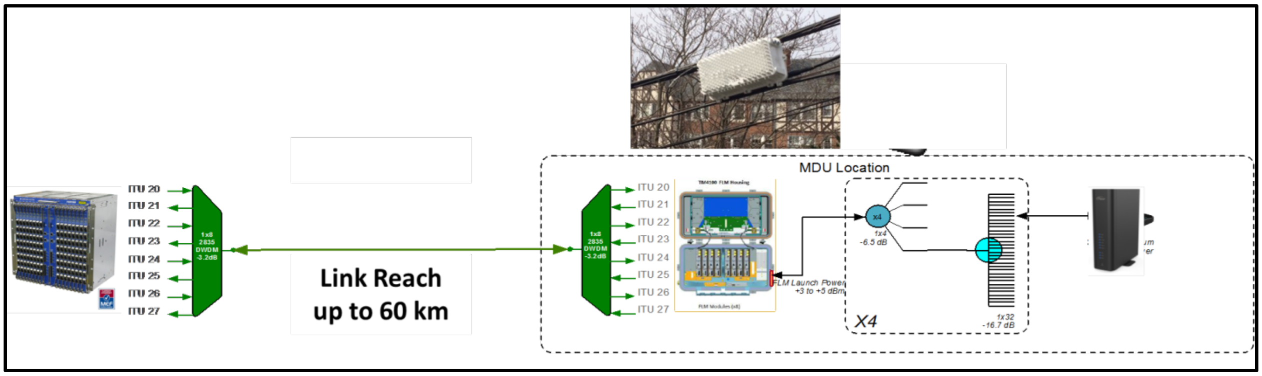

Figure 7 – 10G / 10G EPON Extender Network Configuration

The major advantages of the PON Extender are significantly increased fiber utilization -- up to 8 OLT port links over a single fiber, plus extended reach. The lack of MAC functionality substantially reduces the total power consumption of the unit. The node housing and basic OEO physical layer function make PON Extenders operationally familiar for cable systems engineers and field support technicians.

The main disadvantage of PON Extenders is that they do not eliminate or reduce the OLT equipment requirements in the Hub. The extender is basically the cost equivalent of a second line card needed to reach the same number of subscribers.

Evolution of Distributed Access for FTTP EPON

The driver for distributed access in FTTP EPON networks is identical to the driver for distributed access in HFC networks – improved scalability. To noticeably increase the data capacity per subscriber in HFC networks, the RF bandwidth available per subscriber must be increased. There are a number of ways to accomplish this: Increasing the raw RF bandwidth, node splitting, or reduced amplifier cascades (i.e. redesigning the network from N+5 to N+0). The consequence of increasing the number of nodes, and thereby reducing the number of subscribers sharing the node bandwidth, is the impact to the hubs supporting those nodes. Each new node requires additional CMTS ports, optical transmission lasers and receivers, plus rack space for signal distribution equipment and connecting cables. Additionally, floor space, power consumption, backup power generators, and HVAC capacity quickly come into play. The alternative -- building new hubs to support the racks of new equipment -- is a long and very expensive undertaking.

The situation with FTTP PON is only slightly different. Mainframe OLT port counts are usually 8 per line card, for a total of between 80 and 112, depending on the vendor. The power consumption of a fully loaded OLT is several kilowatts, and, similar to the HFC scenario previously described, each OLT needs backup power, fiber management and a controlled environment. At a split ratio of 128 subs per port, a

112-port OLT can serve over 14K customers; at 32 subs per port, the number of customers served drops to 3.5K. The impact to the hub facilities and the OSP fiber plant is equally dramatic.

10G EPON has enough data capacity to support multi-Gigabit service to a maximum of 128 subscribers per port, assuming normal contention factors for residential and SMB users. The problem is reaching those subscribers. The majority of PON applications over the next several years will serve non-contiguously located greenfield properties with a typical HHP size of 300 units. Almost all of these applications will be greater than the 5 km link reach limit for a 128-way PON split ratio. Most links will reach beyond 20 km from the primary hub, and almost all will have limited dark fiber available. PON Extenders, as discussed earlier, can bridge the gap -- but at a higher incremental cost per subscriber. Distributing the OLT edge network closer to the subscribers being served solves many of the issues associated with traditional FTTP networks.

The Benefits of FTTP Distributed Access

A Distributed Access Architecture (DAA) for FTTP has many of the same goals and follows the same trajectory as HFC Remote PHY or Remote MAC-PHY. DAA accomplishes the following:

- Dis-aggregation of the data, management, and control planes within the OLT platform
- Reduction of hub rack space and powering requirements
- Distribution of the PHY edge closer to the subscriber
- Provides for flexible deployment sizes
- Increased link reach and fiber utilization
- Interoperability between vendor solutions
- Provides a path to network function virtualization (NFV)

One significant distinction for FTTP distributed access is that it is impractical to completely separate the MAC and PHY functionality of the current OLT silicon. Therefore, a remote OLT design will be similar to MAC-PHY system architectures.

Simply scaling down the number of line cards and rack unit size of a mainframe OLT platform and moving it out into the field is not an efficient or network-safe means of creating a distributed access network. The backplane switching fabric of a mainframe OLT is designed for high port counts, and is not readily scalable for low homes-passed applications, such as serving an MDU with 250 apartments/condos. The existing Ethernet switch silicon is also one of the highest power consumption components within an OLT, so designs based on repackaging an existing line card will typically generate powering requirements of 200 watts AC or more. This is much higher than the cable plant network power grid can support without adding more power supplies to the network, and is definitely beyond the thermal capacity of the largest node housings that can be accommodated on the strand today. Rack-mountable outside cabinets with internal power supplies are a possible option, but with the higher power consumption of these designs, some level of environmental conditioning is typically needed.

The bigger issue with a downsized mainframe OLT platform for remote applications is that all of the provisioning, security, and control functions still reside on board, which dramatically increases operational risks. Maintaining this type of remote design will also require a much higher skill set level and diagnostic equipment tools than the typical field technician has today.

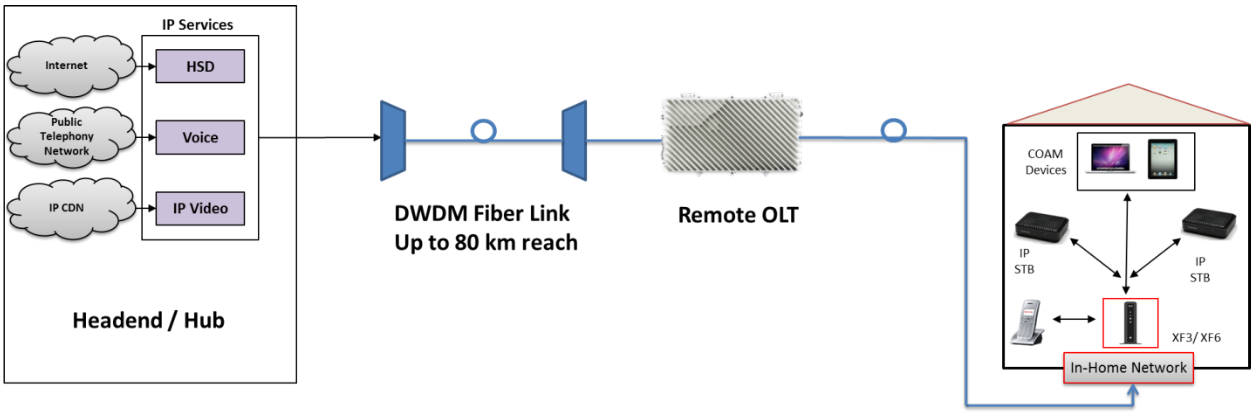

Figure 8 – 10G EPON Remote OLT All IP Triple Play

The target design for a remote OLT is a strand-mountable, node-housed device. Node housings are self-contained environments: Weather proofed, RFI shielded, and capable of being located in the building, in a pedestal cabinet, or on a pole (as originally intended.) R-OLT nodes are designed to operate on cable plant power over an extended ambient temperature range of -40C to +60C.

Creating a node based OLT that is compatible with existing cable plant guidelines and capable of operating in outdoor, uncontrolled environments is a unique challenge for legacy PON equipment suppliers with limited or no HFC plant experience. Newer switch silicon provides lower power consumption and the ability to power down unused ports. The restrictions on total power consumption, combined with the need to meet safe operating temperature limits for 10G optics and PON MAC silicon chip sets, limits the number of 10/10 EPON ports that can be supported to only four, in most cases. While this may appear to be to constraining, the ability to place a node-based remote OLT close to the property being served allows 128 subscribers to be connected, per port, or 512 subscribers per node. This is more than enough to cover 90% of typical MDU situations. Larger properties, such as new, single family home developments, would need multiple nodes -- but proximity to the subscribers would reduce the amount of drop fiber construction needed.

By disaggregating the data, management, and control planes of the OLT, these functions can remain at the headend / hub with DpoE / vCM clients at the R-OLT. The R-OLT can then be designed primarily as a layer 2 device, with limited layer 3 functions as required.

Virtualizing the FTTP Network

The ultimate goal for distributed architectures is network function virtualization (NFV) and software defined networking (SDN). Virtualizing the network will allow the use of common, interoperable hardware, with open software, to promote agile development and faster innovation. Under this virtual operating system, all subscribers can receive the same quality of experience -- whether they are connecting through an R-PHY HFC node, or a Remote OLT.

SDN separates data, control and management planes to enable an open architecture with standard interfaces and protocols. NFV decouples network functions from proprietary, built-in hardware, to provide distributed software functionality that runs over common, multi-sourced hardware. Cloud architectures shift network and operations intelligence into a centralized platform, providing cloud-based applications and service models.

SDN's open architecture provides an end-to-end orchestration and controller platform of both virtual and physical network components, that can be integrated dynamically and programmed automatically, based on service and network demands. Automation combined with programmability enable service agility. In addition, new services and deployments may be introduced with faster time, less investment and reduced cost.

An efficient NFV design for access networks requires breaking software into modular components that correspond to workflows -- that can be distributed at the edge, core or cloud, and coordinated by a centralized SDN orchestration platform. This can be achieved by using containerized software to implement micro-services, as opposed to traditional and monolithic systems. This approach provides smaller and simpler SW subcomponents, which can be created, added and modified independently, using agile development and continuous delivery techniques. These software subcomponents can then be implemented, tested and deployed to introduce new services and network functionality faster, with the help of efficient cross-functional DevOps teamwork.

Cloud-based platforms with distributed functionality provide agility and flexibility of both service and network operations, as wider communication pipes are controlled by a centralized intelligence architecture.

A centralized controller manages both modular and distributed virtual components as well as physical components serving smaller subscriber serving groups. This removes single point of failure issues that can disaffect a large domain, in traditional systems. It reduces CAPEX as the deployment can be scaled over time based on real-time and on-demand requirements. New features are introduced by deploying new software components on the same hardware platform without depending on vendor-specific releases with long lead times. Virtual components can be moved or scaled based on performance, network, and high availability constraints.

This centralized approach, along with powerful data mining techniques, provides end-to-end visibility, better search, and improved correlation of data. Furthermore, virtual probing, together with probing of physical components, may be implemented for troubleshooting, new service integration, and customer QoE assessment. The deployment of a virtualized EPON architecture is flexible and scalable, compared to static, rigid and mass rollouts. New workflows could be introduced as NFV micro-services during trials and deployments to introduce new services, and modify or scale existing services.

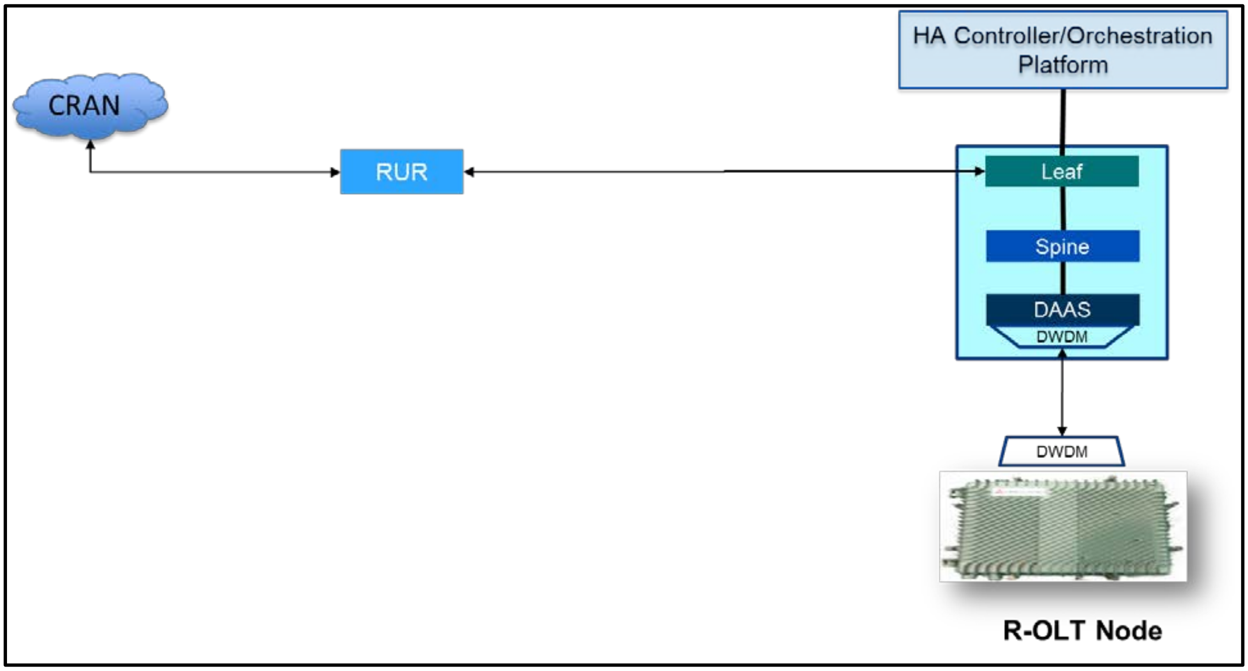

Figure 9 – Virtual OLT Configuration

Figure 9 above illustrates the implementation of a remote OLT in a virtualized network. The remote OLT PON ports are transported over 10 Gb optical links, muxed onto a single fiber and combined at the hub aggregation switch. The controller / orchestration platform directs the flow of data between the regional access router and the R-OLTs connected to the hub, while monitoring performance and policing security, provisioning, and operational policies.

Conclusion

Increased market competition and the continuing growth of data consumption is driving MSOs to deploy networks capable of Gigabit-per-subscriber throughput -- necessary to defend their subscriber footprint against fiber-based overbuilders. Greenfield housing developments, and, in particular, MDUs are being aggressively targeted by PON competitors, prompting cable operators to respond with a fiber to the premise solution.

The applications for FTTP in cable footprints for the next few years will be limited to new greenfield opportunities, entertainment venues, and commercial entities, like small to medium businesses. DOCSIS 3.1, enhanced by FDX, will remain the preferred architecture for legacy, residential, brownfield networks. In order to lower the cost of FTTP construction to a level comparable with HFC, an all-IP, triple-play-capable solution must be deployed to replace the current RFoG overlays and to migrate to a more distributed architecture.

FTTP distributed access architectures provide the same network advantages as remote PHY, in traditional coaxial HFC systems. The current development and eventual deployment of R-OLTs will distribute the

PHY edge closer to the subscriber, enabling increased link reach, higher fiber utilization, and flexible serving group sizes.

Distributed access technology also enables the transition to network function virtualization and a more cloud-based network. Virtualizing the network breaks the long-standing model of bookended, proprietary link equipment, permitting the use of interoperable hardware and open software.

A flexible network that can transparently integrate both HFC and PON subscribers -- while maintaining the exceptional quality of experience of DOCSIS -- may be the best defense against the new generation of fiber competitors.

Abbreviations

BOSA	bidirectional optical sub-assembly
BW	bandwidth
CAGR	compound annual growth rate
CMTS	cable modem termination system
CPE	consumer premise equipment
CRAN	Comcast Regional Access Network
DAA	distributed access architecture
DAAS	distributed architecture aggregation switch
DML	DOCSIS mediation layer
DOCSIS	data over cable system interface specification
DPOE™	DOCSIS provisioning of Ethernet
DPOG	DOCSIS provisioning of GPON
DS	downstream
DSL	digital subscriber line
DWDM	dense wavelength division multiplex
EPON	Ethernet passive optical network
FDC	fiber distribution cabinet
FDX	full duplex DOCSIS
FiOS	Fiber Optic Service
FSAN	Full Service Access Network (working group forum)
FTTH	fiber to the home
FTTP	fiber to the premise
Gb	gigabit
GEPON	gigabit EPON
GPON	gigabit passive optical network
HA	high availability
HD	high definition
HE	headend

HFC	hybrid fiber-coax
HOA	home owners association
HSD	high speed data
HVAC	heating, ventilation and air conditioning
IEEE	Institute of Electrical and Electronic Engineers
IP	internet protocol
ITU	International Telecommunications Union
ITU-T	ITU - Telecommunications Standardization Sector
km	kilometer
MAC	media access control
MDU	multi-dwelling unit
MPO	multiple fiber push on/pull off (optical connector)
MSO	multiple system operator
NFV	network function virtualization
NGPON2	Next generation passive optical network 2
OBI	optical beat interference
ODN	optical distribution network
OEO	optical-electrical-optical
OFDM	orthogonal frequency division multiplexing
OLT	optical line termination
ONU	optical network unit
OSP	outside plant
PON	passive optical network
POTS	plain old telephone system
QAM	quadrature amplitude modulation
QoE	Quality of Experience
RBOC	regional bell operating companies
RFI	radio frequency interference
RFOG	RF over glass
RUR	Regional Universal Router
SDN	software defined network
SFP	small form factor pluggable
SMB	small - medium business
SW	software
US	upstream
vCM	virtual cable modem
Vhub	Virtual hub
XFP	10 gigabit small form factor pluggable

Bibliography & References

Nielsen's Law of Internet Bandwidth; https://www.nngroup.com/articles/law-of-bandwidth/

2016 Sping Technical Forum; Predictions On The Evolution Of Access Networks To The Year 2030 & Beyond, Tom Cloonan, http://www.nctatechnicalpapers.com/Paper/2016/2016-using-docsis-to-meet-the-larger-bandwidth-demand-of-the

ITU G.984-1 Gigabit-capable passive optical networks; International Telecommunications Union, https://www.itu.int/rec/T-REC-G.984.1/en

IEEE P802.3ca 100G – EPON Task Force; *Institute of Electrical and Electronic Engineers, www.ieee802.org/3/ca*

