Leveraging DAA Investments to Accelerate FTTP

Considerations for FTTP that are Specific to Cable Operators

A Technical Paper prepared for SCTE by

Brian Soloducha, P.Eng
Principal Network Architect
Rogers Communications Inc
2400 32 Ave NE Calgary AB Canada
403-716-6041
brian.soloducha@sjrb.ca

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

Cable Operators know how to build and operate high-quality Hybrid Fiber-Coax (HFC) networks for Residential and Business customers. Recent developments in the last few years have focused energies on new Passive Optical Networks (PON) for serving new customers in Greenfield builds with Fiber-to-the-Premise (FTTP) architectures and technologies. With network expansion initiatives in both the United States and Canada, Cable Operators need to be flexible with strategies for growing their networks and increasing revenues from new Commercial and Residential customers.

This Technical Paper discusses a physical architecture for FTTP that will look familiar to those Cable Operators that have already adopted Distributed Access Architecture (DAA) for their HFC deployments. Starting with a reminder of legacy HFC architecture and critical Hubsite components, the paper then discusses legacy PON architecture and Hubsite components, focusing on key limitations that have previously prevented widespread FTTP deployments at Cable Operators. Topics under discussion will progress to recent developments that overcome these limitations. The paper will conclude by looking into the future at problems that are already foreseeable and that will need to be addressed in due time.

This Technical Paper is targeted for an audience curious about starting (or restarting) an FTTP program, and for Cable Operators who have deployed DAA for HFC and are looking to leverage their investment towards PON architectures in Greenfield builds. For those Operators who have already deployed Remote OLTs, the later sections that are forward-looking may be of interest.

2. Background – Legacy FTTP deployments at a Cable Operator

Back in the early 2010s in Western Canada, the regional Cable Operator trialed FTTP builds for greenfield deployments to learn both the strengths and weaknesses of the architecture, technology, and operationalization thereof, that a strong and ambitious competitor Telco was deploying. Residential households-passed and Business drops exceeded 10,000 as an experiment to determine if and how FTTP could be successfully adopted at scale by a Cable Operator. Two major limitations of these legacy PON deployments were directly experienced in Western Canada that are specific to Cable Operators. These limitations were not relevant to competitors that did not operate legacy HFC networks. Other Cable Operators may have experienced other challenges above & beyond those observed in Western Canada.

The initial years of this experimental deployment program featured a 1G EPON protocol providing data / Internet services, with Radio-Frequency-over-Glass (RFoG) technology overlays being used to provide Video and Voice services. This allowed for the use of a DOCSIS Provisioning of EPON (DPoE) back-office provisioning engine. In later years, the EPON technology was slowly deprecated and migrated over to GPON protocols in some instances.

The physical architecture of this broad experiment utilized Centralized Optical Line Terminals (OLT) with the OLT located in Hubsites that fed the Optical Distribution Network (ODN). A split ratio of 1:32 with a distance limit of 20.0 km was achieved with Central Splitting at a single location per ODN. This location was called a Fiber Transition Cabinet (FTC), although other literature for the industry may call this a “Fiber Distribution Hub” (FDH) or “Local Convergence Point” (LCP). The idea was to use existing Backbone Fiber from a Hubsite to a Node Service Area and place the FTC/FDH for each ODN inside the Node Service Area.

2.1. Limitations of Backbone Fiber Infrastructure

The program for FTTP was designed to use existing Backbone Fiber cables from Hubsites out to new Node Service Areas over the Access Network Infrastructure that was in-place at that time.

The existing legacy HFC network was designed with the intended following parameters:

- Node Service Areas would typically be sized at approximately 500 Households-passed (HHP)
- One analog RF node per Node Service Area
- 2 fibers (one forward and one return) per analog RF node
- Backbone Fiber Cabling sized at 288 fibers per cable
    - non-ribbonized to reduce Mean-Time-to-Repair (MTTR) in the event of fiber cuts
- Backbone Fiber allocated as follows:
    - 48f reserved for Hubsite to Hubsite traffic such as Broadcast Video and Metro Dense-Wavelength-Division-Multiplexing (DWDM)
    - 240f reserved for Node Service Areas
- Fiber for Node Service Areas were allocated as follows:
    - 2f used by the current analog RF node
    - 2f reserved for future node segmentation activities
    - 2f reserved for Business services

Not all deployments meet these intended parameters; recent deployments have higher rates of compliance.

Given that 6f are allocated per Node Service Area out of a total of 240f reserved for Node Service Areas on the Backbone Fiber, this allows for 40 Node Service Areas to be served (unprotected, single-ended) out of a Hubsite per Backbone Fiber cable. In Metro areas, the Backbone Fiber cables are typically built from one Hubsite to another Hubsite, so another 40 Node Service Areas are reverse-fed from the far-end Hubsite that the Backbone Fiber cable connects to.

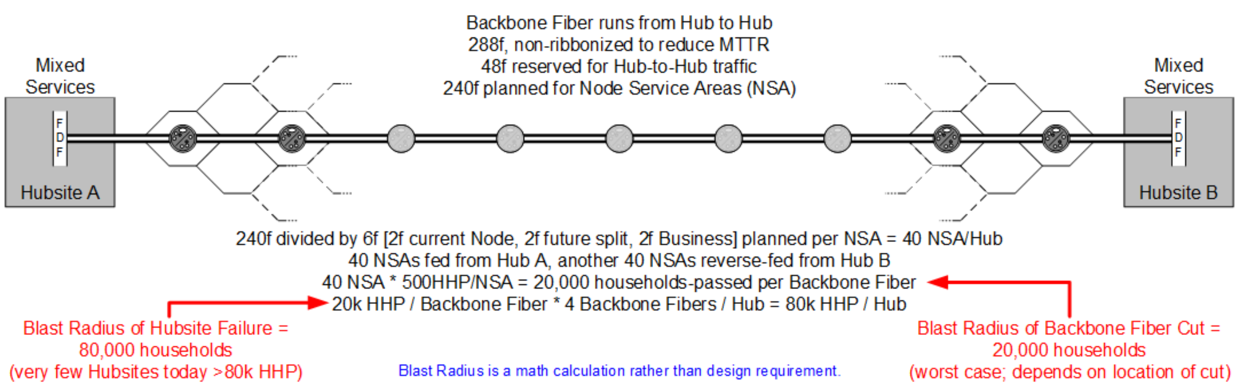

Figure 1 - Backbone Fiber Design Parameters

Western Canada typically has smaller Hubsites than what might be common in Central Canada or the United States, with typical Hubsites serving approximately 50,000 HHP and connecting 3 or 4 Backbone Fiber cables. Vaults for separate entrances for each Backbone Fiber cable would be ideal.

The math calculations for Backbone Fiber then result in a ‘blast radius’ of up to 20,000 HHP for a Backbone Fiber cable cut, or up to 80,000 HHP in the event of a complete failure at a larger Hubsite (power outage & generator failure, forest fire, catastrophic vehicle crash, etc). This concept of a blast radius is not a design requirement; rather, it is just a resultant of the sizing parameters used elsewhere in the design. The (in)frequency of these failure events has not yet warranted a program to build protection into the Node backhaul connectivity. While some Cable Operators have built protection paths to Nodes out of the same Hubsite for Layer 1 protection; and other Cable Operators are looking at Layer 2 or Layer 3 protection out of the far-end Hubsites (when Backbone Fiber is built from one Hubsite to another), protection paths for Node Service Area connectivity has not been prioritized in Western Canada.

This architecture for the HFC network was deployed in field trials in the late 1980s. Design parameters and fiber cable sizing (288f total, 240f reserved for Node Service Areas) were finalized in the late 1990s.

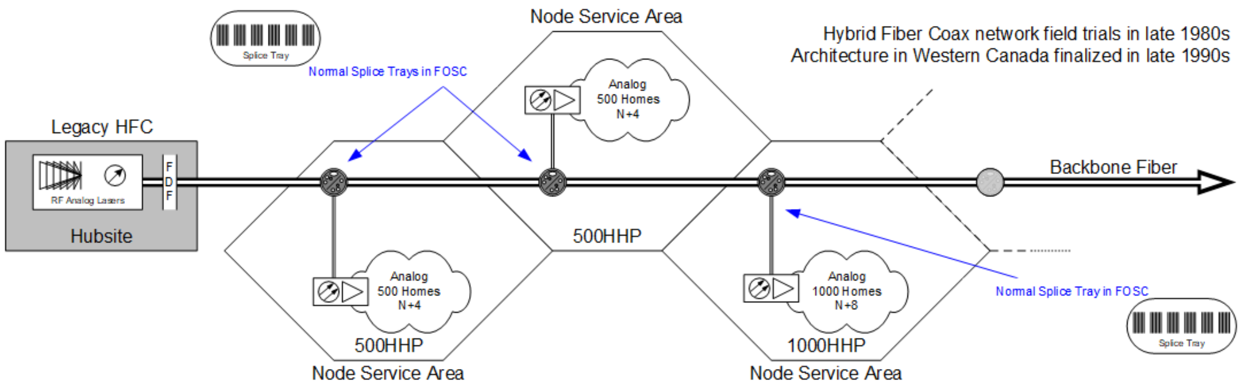

Figure 2 - HFC Node Service Areas

As the experimental FTTP deployment program kicked-off in approximately 2010, it became apparent that the consumption of Backbone Fiber for this FTTP architecture vastly exceeded the planned parameters for fibers assigned per Node Service Area. If all customer drops were connected to 1:32 splitters at the FDH, then a maximum of only 7680 drops could be served off of that Backbone Fiber cable, instead of the planned 20,000 HHP. Further, there would be no room for growth, no back-up fibers for new technologies, and no fibers available for dedicated Business services.

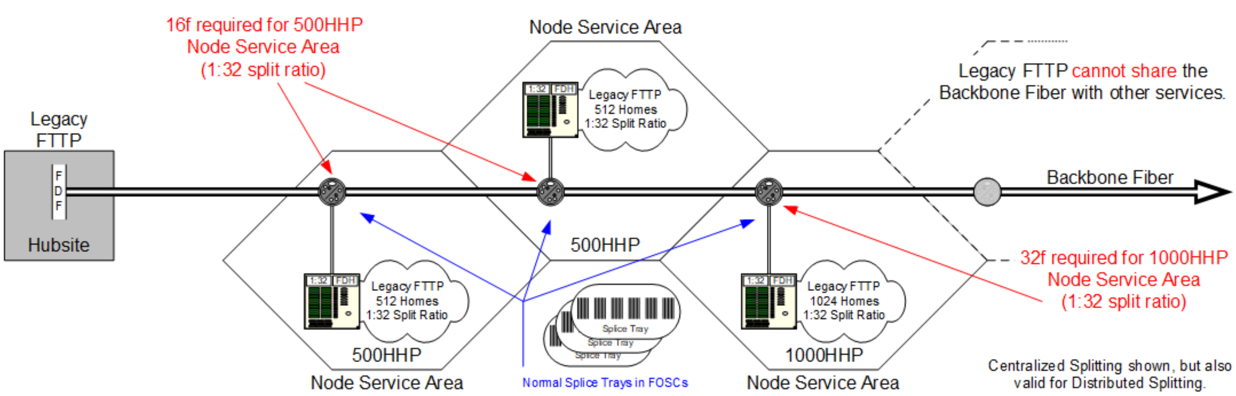

Figure 3 - FTTP Node Service Areas

An attempt was made to minimize unnecessary consumption by using an "Active Drop" configuration at the FDH; rather than connecting every drop to a 1:32 splitter, only active customers would be connected, thereby saving Backbone Fiber usage back to the Hubsite. This configuration also helps to limit the number of PON ports required at the Centralized OLT.

This proved out to be an effective way to minimize the Backbone Fiber usage, but had three horrendous side effects:

1. A truck roll became mandatory for every single customer add AND disconnect.
2. The documentation became difficult to keep up-to-date.
3. The FDH became difficult to properly maintain and route fiber connections over time, making subsequent modifications (customer additions and removals) progressively more difficult.

With the Active Drop configuration, some FDHs became so unwieldly as to be considered "rats' nests".

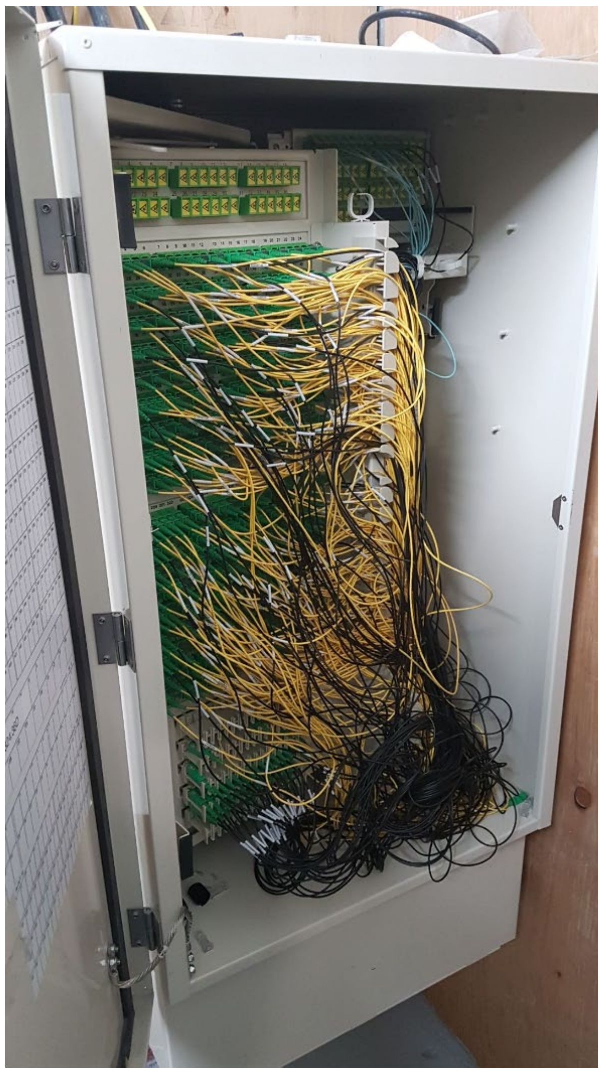

Figure 4 - FTC with Active Drop Configuration

In addition to Backbone Fiber usage, another problem is starting to crop up in older FTTP builds in various parts of both the United States and Canada, namely OLT PON port availability. Some Hubsites and Central Offices do not have rackspace or power to add additional OLT PON ports, and depending on planned penetration rates, there may not be OLT PON ports or Backbone Fiber available to light up new 1:32 power splitters regardless of whether the customer location has a fiber drop. In these cases, the additional revenue from a single new customer would not justify the costs for either a Backbone Fiber build or a new OLT PON port activation, and hence that customer must wait until another customer is disconnected. This does not lead to happy customers.

Backbone Fiber networks used for deploying HFC in Western Canada were largely sized and built decades ago. It is on this infrastructure that the widescale experimental FTTP deployment program was attempted in Western Canada, and where the severity of the Backbone Fiber usage requirement became apparent, despite trying to mitigate the challenge through the use of an Active Drop configuration. This limitation was one of the two key reasons why the experimental FTTP deployment program was shuttered for a decade.

2.2. Limitations of Video Service Delivery

There is a multitude of equipment required in Hubsites for Cable Operators to serve Voice, Video, and Internet product solutions to customers over legacy HFC and legacy FTTP architectures. Internet service is easy over the data path provided by 1G EPON and GPON signals, and Voice services can be added with Session Initiation Protocol (SIP). Unfortunately, Video services requires either an IP Video solution or an RFoG overlay solution.

Back when the experimental FTTP deployment program kicked off, IP Video solutions were complex, costly, and seen as unnecessary by most Cable Operators who already had working Video solutions for their HFC networks. As such, an RFoG overlay solution was chosen for this experimental FTTP program.

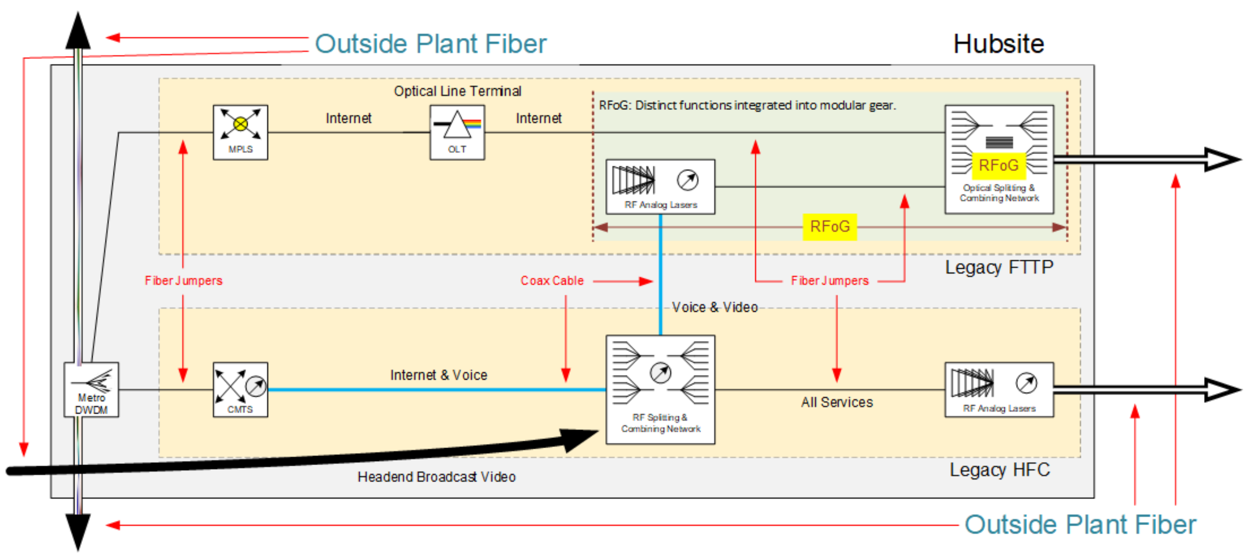

Figure 5 - Hubsite Configuration for Legacy FTTP and Legacy HFC Architectures

Service Delivery for Legacy FTTP was accomplished by combining Voice & Video signals from the Cable Modem Termination Station (CMTS) and Broadcast Video feeds, together with the Data / Internet feed from the OLT. This combining for legacy FTTP occurred at RFoG equipment in the Hubsite, and additional Customer Premise Equipment (CPE) was required to separate out the signals at the customer location. Additional challenges with RFoG, such as Optical Beat Interference, became apparent over time.

As more Cable Operators became aware of the requirements and operational issues of RFoG, demand tapered off while costs for RFoG equipment remained high with a limited number of suppliers in the market.

This requirement for an RFoG overlay for a Video product proved to be the other key limitation as to why the experimental program for FTTP deployment was shuttered.

3. Evolution of the Access Network and Hubsite Infrastructure

There have been several changes in the past decade worthy of discussion that alleviate some of the limitations for FTTP deployments that are relevant to Cable Operators.

Many Cable Operators have now pivoted, or are in the process of pivoting, to an IP Video solution as part of larger operational considerations. A discussion of this general trend is beyond the scope of this paper. This change negates the need for an RFoG overlay solution for Video services and therefore removes entirely the need for costly RFoG equipment in both the Hubsite and the customer premise.

Backbone Fiber utilization trends and the implications of changing to DAA for HFC are worthy of a more detailed discussion.

3.1. Sharing of Backbone Fiber within an Analog Node Service Area

Fiber congestion mitigation projects are still undertaken in small batches on a regular basis, but large scale Backbone Fiber build programs spanning many months or even years for permitting and construction are to be avoided or deferred as much as possible. As traffic grew on legacy analog RF nodes, and as new neighborhoods were built and new Business services deployed, alternative methods were investigated for growing capacity without requiring large-scale Backbone Fiber overbuilds.

In the mid-2010s, passive DWDM filters and lasers for the analog RF nodes were introduced into the network to allow for segmentation of serving groups within the same Node Enclosure. This allowed for additional capacity within the Node Service Area while still using the same amount of Backbone Fiber for that pocket. This passive DWDM architecture was internally called MWL (Multi-Wavelength) and deployed as-needed in a Pay-As-You-Grow fashion in both 4-channel and 8-channel configurations. This first generation of passive DWDM was only intended and used for connecting analog RF nodes within a common node housing. [For Business customers and possibly a few rural applications, Coarse-Wavelength-Division-Multiplexing (CWDM) has alternatively been used.]

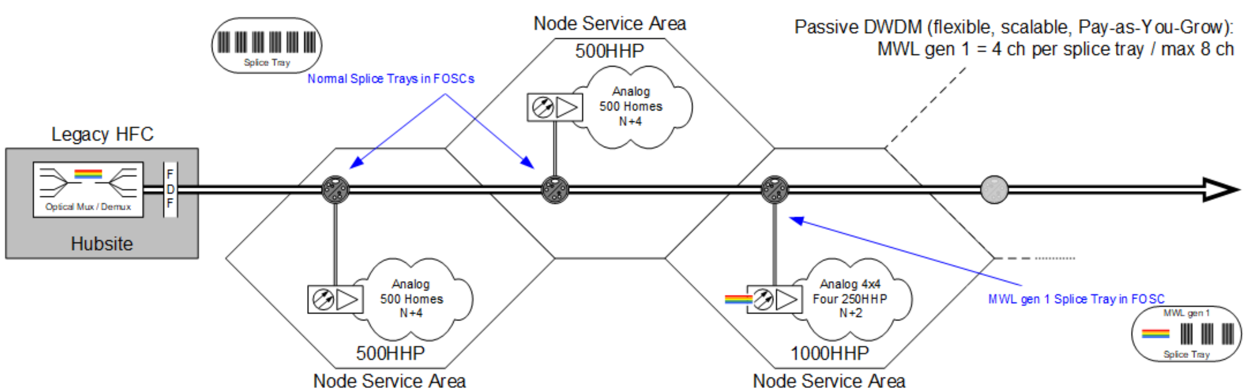

Figure 6 - Deploying Passive DWDM (MWL gen 1) into Node Service Areas

Of all metro Node Service Areas in Western Canada, approximately 20% have 4-channel passive DWDM configurations, and 5% have 8-channel passive DWDM configurations. Typically, a 4-channel passive

DWDM mux was placed into a splice tray in the FOSC; however, a 4-channel mux could also be placed inside the Node Enclosure (rare).

The deployment of passive DWDM at existing Node Service Areas allowed for relatively easy segmentation without the need for re-spacing of amplifiers in the coax portion of the network. This work separated one Node in one Node Enclosure into multiple nodes in the same Node Enclosure, thereby lowering the number of customers on each coax feed from the Node Enclosure and hence lowering total utilization and congestion by reducing the denominator (total coax capacity / HHP).

This approach allowed the same pair of fibers from the Backbone Fiber cable into a Node Service Area to be shared amongst multiple Nodes for that Node Service Area. This was an important technology breakthrough and operational mode for staying ahead of congestion on the HFC network in Western Canada.

Another important technology breakthrough and operational mode vital to staying ahead of congestion on the HFC networks in Western Canada was the conversion of the entire network to an 85 MHz return. A discussion of that mid-split upgrade program is beyond the scope of this paper.

3.2. Deploying Distributed Access Architecture

Approximately 5 years ago, a program was kicked-off in Western Canada to evaluate and deploy Distributed Access Architecture (DAA) for HFC. There have been many Technical Papers in recent years relating to DAA, so a full summary of the benefits and risks of DAA is unwarranted, but a short summary for both the Outside Plant and for the Hubsite as DAA relates to FTTP for Cable Operators is as follows:

In the Outside Plant, DAA allows for:

1. 10G Ethernet backhaul, allowing for
2. I-Temp DWDM full-band tunable pluggable optics, allowing for
3. Full-band passive DWDM filters and muxes

In the Hubsite, DAA allows for:

1. Converged Interconnect Network (CIN) routers, allowing for
2. Simplified Hubsite Configurations

These changes are key to a new physical architecture that will allow Cable Operators to deploy FTTP in greenfield Node Service Areas, or for in-fill demands from new or unexpected Multiple Dwelling Unit (MDU) buildings such as apartment buildings, condo highrise towers, row townhouses, etc.

3.2.1. Upgraded 40ch passive DWDM in the Access Network

The move to DAA has led to a flourishing of 10G DWDM optical links powered by Full-band Tunable DWDM pluggable optics. These optics are hosted by Remote PHY Devices (RPDs) in the Access Network and by CIN routers in Hubsites. Tunable DWDM pluggable optics in the SFP+ form-factor have now been used in Western Canada since late 2019. Some Cable Operators may choose to use fixed-wavelength DWDM optics; that conversation is beyond the scope of this paper.

The first generation of passive DWDM filters and muxes deployed in Western Canada were 4-channel and 8-channel configurations. These were based on the fixed-wavelength DWDM transceivers suggested by the analog RF node vendors and based on supply availability.

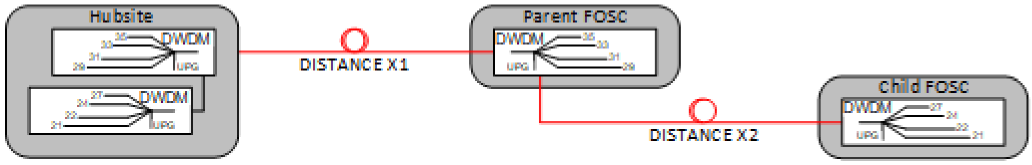

Figure 7 - 4ch/8ch Passive DWDM (MWL gen 1)

With the adoption of DAA and tunable optics, a new generation of passive DWDM filters and muxes were warranted. While the tunable optics were capable of 96-channels across a 50GHz-spaced grid, this was deemed excessive for foreseeable demands.

A next-generation passive DWDM system was developed:

- 100 GHz-spaced, 40-channel total capacity
- 10-channel initial capacity in Hubsites, 5-channel initial capacity in FOSCs in the Outside Plant
- Flexible, Scalable, Pay-as-You-Grow
- Allows for 8 total FOSCs in a Node Service Area
- Allows for 16 channels reserved for HFC (allowing for N+0 in 500 HHP Node Service Areas)
- Allows for 24 channels for other services (Business waves, Small Cells / Wireless, Remote OLTs, etc), and more if N+0 is not a required HFC configuration
- Adds channel groups in parallel

A key feature of this next-generation system is that channel groups are added in parallel rather than cascaded in series; this design minimizes total insertion loss at higher channel counts when compared to a cascaded design.

This next-generation system is called MWL gen 2 to maintain internal consistency for users in Western Canada, and features a few mandatory components in the Hubsite and one mandatory component in the Outside Plant. Surprisingly, due to the optical transmit and receive specifications of the tunable optics and the low total insertion losses achievable per channel, attenuators (Global Pads and Group Pads in the diagram below) are usually required for these wavelengths. [Short range tunable optics with looser transmit and receive specifications did not make sense from a supply chain perspective.]

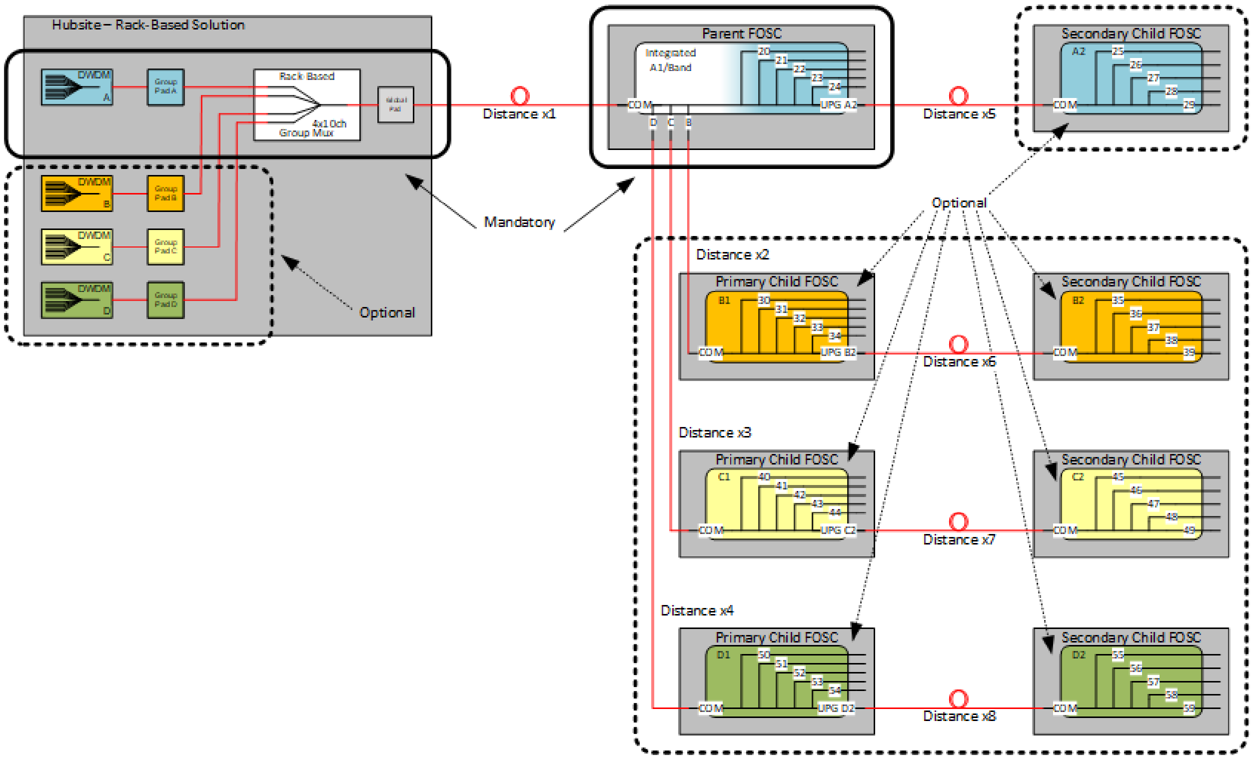

Figure 8 - 5ch/40ch Passive DWDM (MWL gen 2)

The equipment qualified and selected for the MWL gen 2 passive DWDM was tested with 10G wavelengths at all channels, and with higher bitrate signals of 100G, 200G, and 400G. There was some clipping observed on an Optical Spectrum Analyzer (OSA) for the 400G wave, but the Forward Error Correction (FEC) was sufficient for a clean Bit Error Rate (BER) Test of 72 hours over 50 km (lab spools of fiber). Therefore, if utilization of the MWL gen 2 passive DWDM systems start to congest with a plethora of 10G waves to a common Node Service Area, it will be possible to upgrade to higher bitrate wavelengths by using Point-to-Point Coherent waves with a Coherent Termination Device in the field to aggregate those smaller 10G demands. This is expected to future-proof the passive DWDM architecture.

With the introduction of this MWL gen 2 passive DWDM system, there is now a clear delineation for how Node Service Areas are to be served:

- Older analog RF nodes that have not yet been segmented will remain on non-shared fiber pairs.
- Analog RF nodes deployed after ~2015 will remain on MWL gen 1 passive DWDM systems.
- All Remote PHY Devices for DAA, all Remote OLTs for FTTP, all waves for Business services, all Small Cell & Wireless Backhaul waves, and all other wavelengths are to use the MWL gen 2 passive DWDM systems to share the Backbone Fiber.

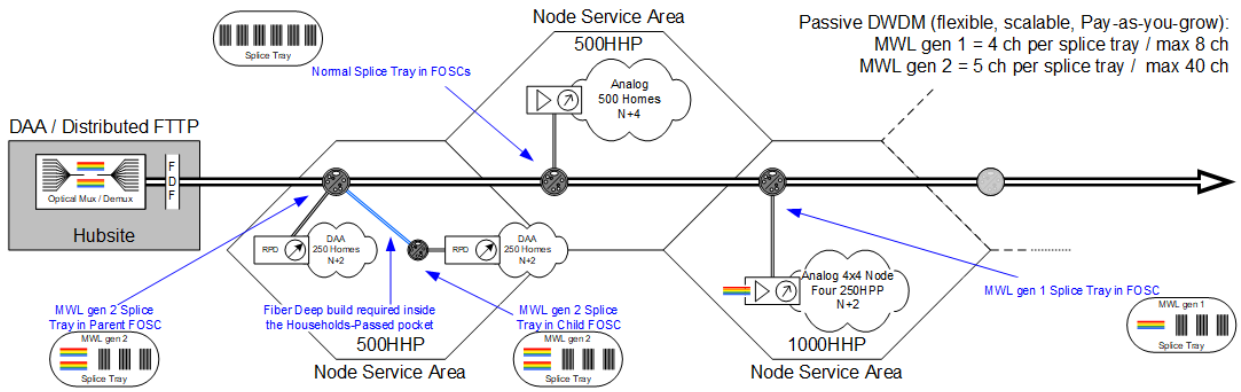

Figure 9 - Deploying Passive DWDM (MWL gen 2) into Node Service Areas

In this manner, DAA nodes can be deployed in Node Service Areas using the existing Backbone Fiber cables. Fiber can be built deeper into Node Service Areas and additional RPDs can be placed to allow for segmenting down from N+4 configurations (or higher) to N+2 configurations (or lower, if desired).

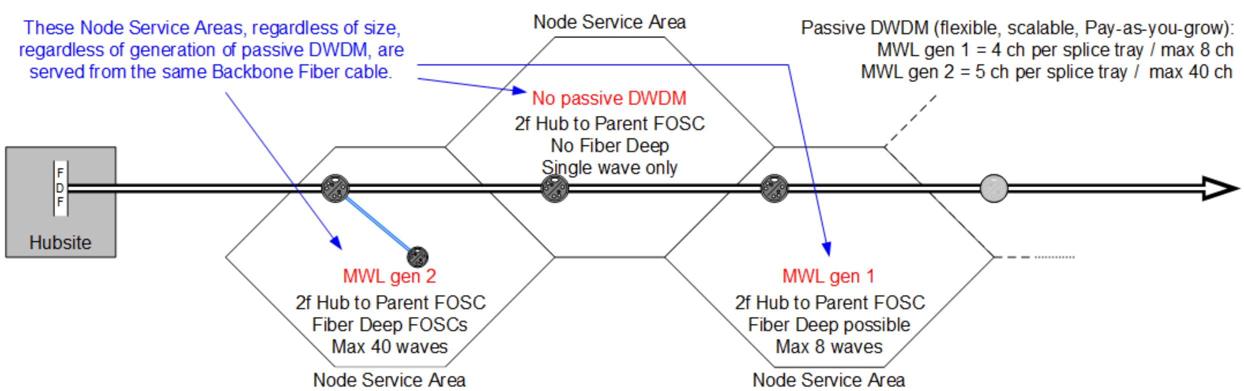

Figure 10 - Backbone Fiber Cable with Passive DWDM at different Node Service Areas

The important aspect under consideration is that the limitation of Backbone Fiber cable utilization remains; these are 288f cables and only 240f can be used for Node Service Areas. Adding passive DWDM filters and muxes allows for more efficient utilization of the fiber already in the ground.

There is no current plan to insert passive DWDM onto existing paths that do not have it already, however, all migrations from analog RF nodes to RPDs are required to include passive DWDM on the Backbone Fiber as part of that conversion process.

By pivoting from analog nodes with very specific fixed DWDM wavelengths over to RPDs with tunable optics, the move to DAA allows Cable Operators to save time, effort, and money. While narrowband DWDM optics and the additional filters add costs when compared to wideband 'grey' optics, this approach allows Cable Operators to defer, or maybe avoid altogether, costly Backbone Fiber overbuild programs for the Access Network.

There are many operational benefits for tunable DWDM optics that outweigh the higher cost when compared to fixed-wavelength DWDM optics, but that discussion is beyond the scope of this paper.

3.2.1. Simplified Hubsite Configurations

The move to DAA will eventually lead to simplified Hubsite configurations.

One of the many promises of DAA is a significant reduction in Hubsite rackspace and power required if all service delivery is moved to a DAA configuration. This is both very true and highly unrealistic in the next 5 to 10 years. Changing Hubsite configurations (and analog nodes for RPDs in the Outside Plant) for existing customers is a lot of work that does not produce new revenue.

Unless there are pressing reasons for this DAA conversion for existing customers, such as a need to vacate the Hubsite due to a sale of the facility or being unable to renew a lease, or a need to upgrade to DOCSIS 4.0 for congestion or competitive reasons, or a need to remove End-of-Life / End-of-Support equipment, then the most likely operational mode for Cable Operators will be to Cap & Grow. Existing service delivery would be capped on existing platforms (CMTS, Centralized OLTs), and then new service delivery would be brought up on infrastructure used for DAA. This would include new RPDs (DOCSIS 4.0 upgrades) and Remote OLTs (new Distributed FTTP deployments in Greenfield SFUs, MDUs, and commercial / business parks), as well as Business and Small Cell / Wireless wavelengths (depending on the functionality of the CIN routers).

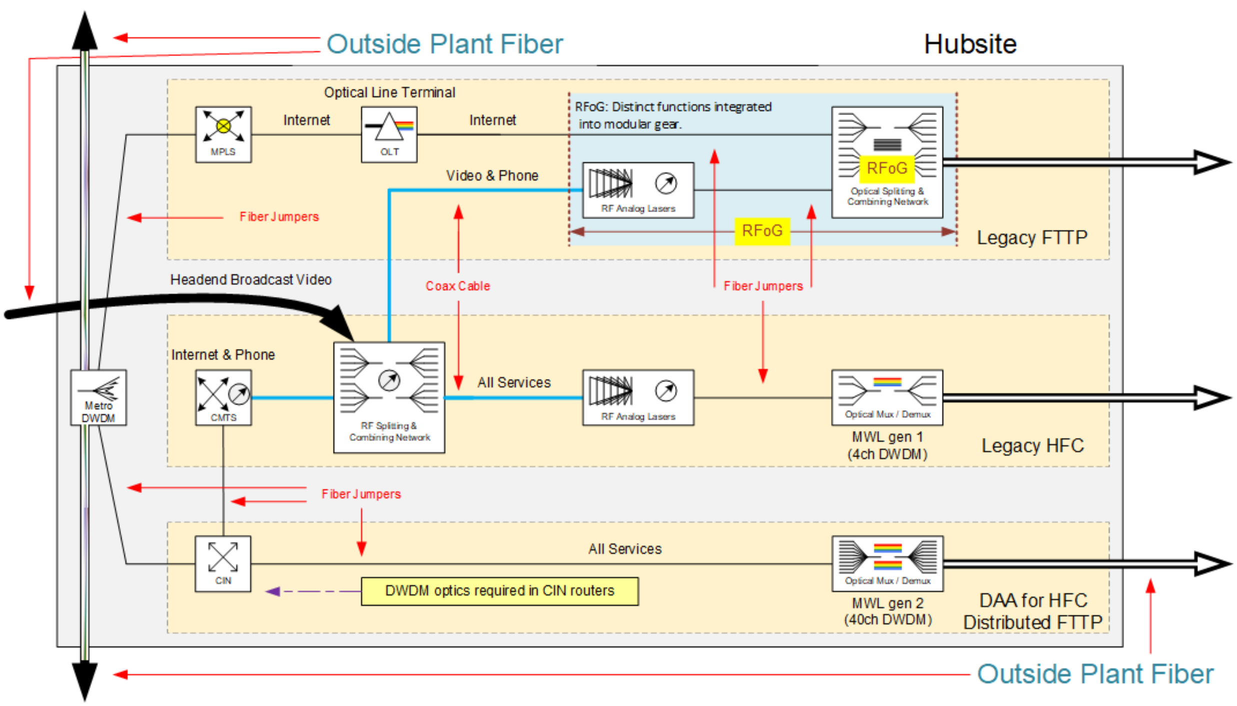

Figure 11 - Hubsite Configuration for currently-supported architectures

Hubsites will still be pressed for rackspace and power as this Cap & Grow model is implemented. CMTS links through Metro DWDM transponders back to Core Routers can be moved over to the CIN routers to alleviate some transponder costs; this might also be possible for Business Services routers (MPLS in the diagram above) depending on functionality.

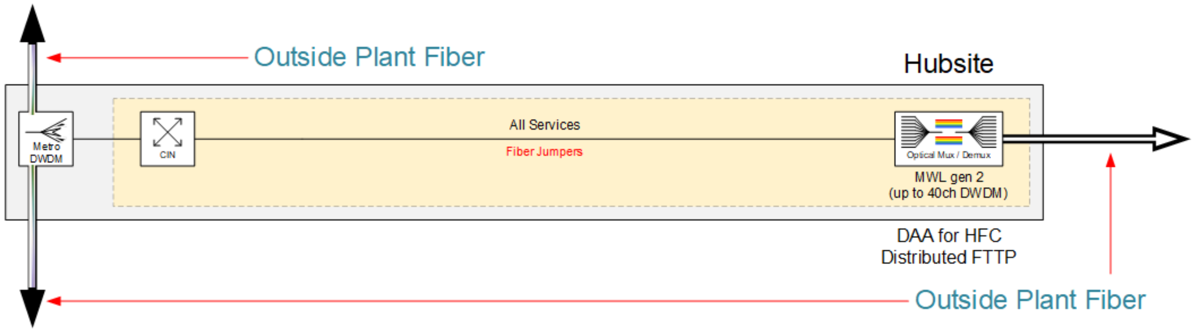

Figure 12 - Hubsite Configuration for DAA and Distributed FTTP

The end goal once all legacy HFC has been switched to DAA, and if Central OLTs can be replaced with Remote OLTs, is a simplified Hubsite configuration. This will take many years for most Hubsites. Future PON protocols may warrant new Centralized OLT equipment; consolidation & DAA conversion activities to make rackspace and power available at existing Hubsites would be required in Western Canada and likely for many Cable Operators.

4. Remote OLTs and Distributed FTTP

The terms Remote OLT and Distributed FTTP have been lightly-used previously in this Technical Paper. This discussion will now define these terms and proceed with discussing their merits and drawbacks.

4.1. Defining Remote OLTs

Remote OLTs are the functional equivalent for FTTP as a Remote PHY Device is for DAA. These Remote OLT devices have a small number of PON ports that terminate the ODNs for FTTP networks, have small number of Ethernet backhaul ports for connectivity back to a CIN router, and are environmentally-sealed and temperature-hardened with no moving parts (no fans). Similar to an RPD, a Remote OLT will have fins for heat dissipation, a coax seizure mechanism for power from an RF-isolated coax plant, and a fiber entrance port for a Node Service Cable. A Remote OLT will look like an RPD, but the inside guts will be different.

Figure 13 - Remote OLT Node Enclosure

Perhaps controversially, there is a need to differentiate Remote OLTs from remotely-placed OLTs.

Table 1 – Defining Remote OLT Characteristics

Remote OLT	Remotely-Placed OLT
• Sealed Node Enclosure (sometimes referred to as 'clamshell') • 90 Vac from RF-isolated coax • No moving parts • Small number of PON and backhaul ports • Temperature hardened	• Rackmount in controlled environment (Hubsite, Environment Cabinet, etc) • -48 Vdc or 120 Vac power • Modular Chassis with different cards • # PON ports based on modules in slots • Requires fans

A rack-mount chassis-based OLT where the number of PON ports is expandable by adding modules to the chassis, where -48Vdc or 120 Vac powering is required, and where the temperature is both controlled in an environment cabinet and moderated by fans is considered to be a remotely-placed OLT.

Over the past 2 years, a total of four different Remote OLTs were put through evaluation and qualification activities in Western Canada, including 10G EPON and XGS-PON variants. On the physical specification requirements side, testing was done to ensure compliance with temperature conditions (-40 C to +60 C), powering needs (compatibility with 90 Vac coax power supplies), and Access Network compatibility (must utilize tunable DWDM optics to share the Backbone Fiber over MWL gen 2 passive DWDM).

Similar to RPDs that were evaluated, the Remote OLTs evaluated required pluggable optics to be rated for Industrial Temperature use (I-Temp, -40C to +85 C), and suffered from reboots during minor power bumps (such as may be experienced when a Plant Technician switches from utility-feed to battery at the coax power supply during routine maintenance activities).These reboots for Remote OLTs, similar to RPD behavior and performance, result in a loss of customer service for a duration measured in minutes as opposed to durations measured in seconds for power bumps on analog RF nodes.

Cable Operators looking at Remote OLTs will also need to consider Node Service Cable requirements.

Some of the Remote OLTs tested required a laptop for initial System Line-up and Test (SLAT). Unfortunately, this appears to be a holdover for systems designed for Telcos and then ported over to a form-factor suitable for Cable Operators. An automated, zero-touch provisioning system for initial deployments is preferable, and vendors for these Remote OLTs should be encouraged to adapt their products accordingly.

All four Remote OLTs that were tested used less power than typical RPDs; Cable Operators may want to look at smaller power supplies for deployments where only a single Remote OLT is required.

4.2. Defining Distributed FTTP

For this discussion, the physical architecture when discussing Distributed Access Architecture consists of:

- CIN routers in Hubsites
- Passive DWDM and tunable optics
- Remote PHY Devices terminating the coax portion of HFC service delivery

Distributed FTTP is then defined for this discussion as the physical architecture consisting of:

- CIN routers in Hubsites
- Passive DWDM and tunable optics
- Remote OLTs terminating the ODN portion of PON service delivery

When talking about Distributed FTTP, this explicitly means a physical architecture that involves a Remote OLT as defined above.

4.3. Benefits

Moving the PHY layer closer to the customer has benefits with both DAA and Distributed FTTP.

For DAA, RPDs close to customers results in high modulation rates and hence more capacity on the plant for those customers.

For Distributed FTTP, Remote OLTs close to customers results in lower optical attenuation due to fiber distance, and as a result there is more optical power available; this allows for higher split ratios.

Legacy FTTP with Centralized OLTs standardized on 1:32 split ratios due to distances from Hubsites & Central Offices (COs) being relatively far. This was not too problematic for operators running 1G EPON or GPON because there was a limited amount of bandwidth to be shared between customers. With newer protocols like 10G EPON and XGS-PON, the bandwidth ceased to be a limiting factor, and hence higher split ratios become desirable as this is a more efficient use of available PON ports and fiber required from the PON port into the ODN. For Centralized OLTs, because of the distance to the customers, higher split ratios are not always possible, and therefore the Telco or Cable Operator is stuck with inefficient use of PON ports, and more importantly, inefficient use of the Backbone Fiber used for that PON port.

Remote OLTs have two strong benefits; by being close to the customers, higher split ratios result in a lower number of PON ports required for the same number of Households-passed. More importantly, by being able to utilize tunable DWDM optics for the backhaul, a Remote OLT is able to share fiber into a Node Service Area from the Backbone Fiber cable, and therefore this allows the operator to avoid or defer costly and intensive Backbone Fiber overbuild construction activities.

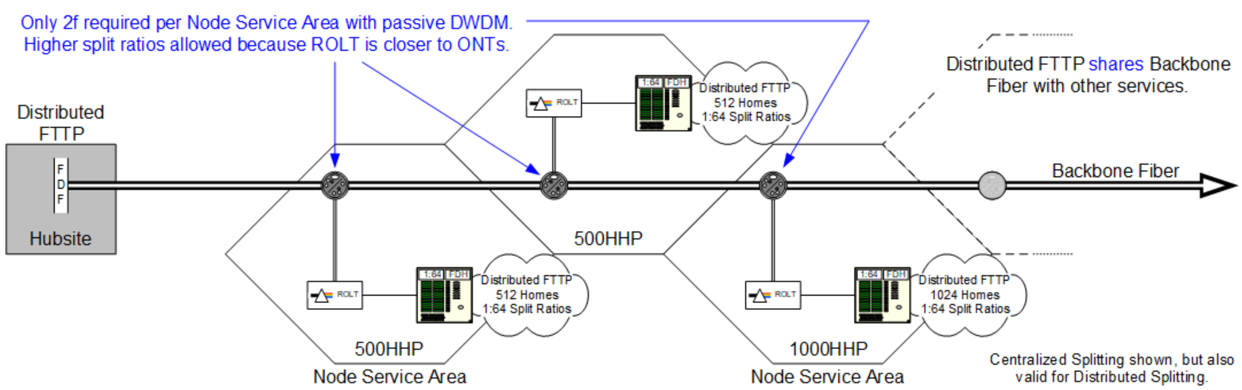

Figure 14 - Remote OLTs sharing Backbone Fiber through Passive DWDM

This also allows for unexpected growth within a Node Service Area without breaking the design parameters for the number of Backbone Fibers assigned to that Node Service Area, such as when a Node Service Area with RPDs is suddenly targeted by a developer for a new in-fill MDU.

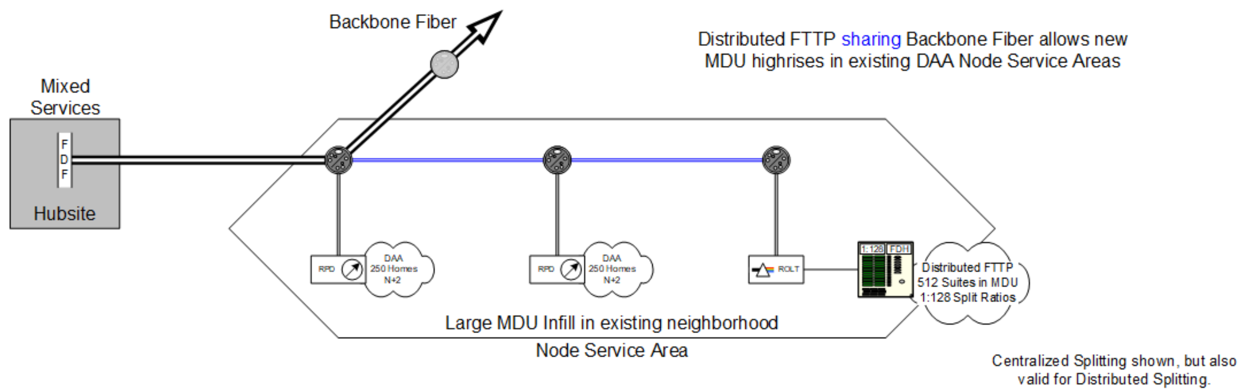

Figure 15 - New RPDs and Remote OLTs into existing Node Service Areas

This idea of adding services to an existing Node Service Area also applies to Rural Network Expansion opportunities where the expansion for that rural network would be single-ended. If the Node Service Area has passive DWDM but few active wavelengths, it would be possible to serve a low number of new Remote OLTs from that Node Service Area directly without consuming new fibers on the existing Backbone Fiber cable.

Distributed FTTP with Remote OLTs also creates an interesting dynamic where the linkage between Backbone Fiber utilization and PON port activation is severed. With Centralized OLTs, there is a 1:1 linkage between Backbone Fibers used per PON port activated. With Remote OLTs, there is no correlation between Backbone Fiber utilization and PON port activation; with a 40-channel passive DWDM system, the Backbone Fiber utilization now becomes dependent on how many channels an operator is assigning to a Remote OLT (for both capacity and protection), and how many PON ports need to be activated inside that Node Service Area, and further by the number of PON ports per Remote OLT. [All Remote OLTs tested in Western Canada had at least four PON ports per device, and some had more.]

With this linkage intentionally severed, additional PON ports can be lit without concern over Backbone Fiber usage. This allows operators to be intentionally inefficient in their PON port utilization. While this would require additional capital costs, it can be of great operational benefit. With Centralized OLTs, operators use Active Drop configurations to squeeze every available port out of a lit 1:32 splitter. As discussed above, this now requires a truck roll for every customer add or disconnect. By being intentionally inefficient on PON port utilizations, operators can instead adopt a "Hot Drop" configuration, whereby all drops are built, tested, and connected during the construction phase of the ODN's lifecycle.

This greatly simplifies operations and maintenance at a central splitting location, reducing the need to visit the location and move fiber jumper connections for every single customer addition or disconnect.

Figure 16 - FTC with Hot Drop Configuration

Hot Drop configurations also allows an operator to plan for extra drops to be available and spliced during the initial construction and build, useful for common scenarios experienced in Western Canada:

1. Basement suites rented to tenants in SFU houses & neighborhoods.
2. Professional Work-from-Home accounts requiring billing that is separate from a family account at the same address.
3. Convenience stores requiring Point-of-Sale accounts while co-located bank Automatic Teller Machines require separate accounts.

If limits are applied to how split ratios and distances are managed, these scenarios could also be accommodated with a 1:2 splitter at the side of the building. An operator would have to measure optical receive power levels prior to intentionally inserting additional optical loss into the path of an existing customer to add a 1:2 splitter; for this reason, having margin for additional drops pre-spliced and tested back to the central splitting location is ideal.

4.4. Drawbacks

One of the drawbacks for Remote OLTs is that it is an active device requiring power and maintenance in the Outside Plant. For a Cable Operator with decades of experience placing actives in the field, this should be possible. For Telcos that have decommissioned & removed DSLAMs out in the field, this might be an untenable position. This creates uncertainty for vendors whereby Centralized OLTs are a known market (Telcos), but Remote OLTs have unknown potentials because they are an entirely new market for both products and customers (Cable Operators).

Another drawback of Remote OLTs is that the product cycle with the latest generation of protocols has lagged development of Centralized OLT products with the same protocols by several years. This is a function of temperature-hardening, form-factor, and the effort required to produce suitable prototypes, amongst other reasons. It is hoped that CableLabs' efforts on a next-generation Coherent PON (CPON) protocol will shorten or remove this gap that exists between Telcos able to deploy FTTP with Centralized OLTs and Cable Operators who require Backbone Fiber to be shared amongst services.

4.5. Considerations for Distributed FTTP specific to Cable Operators

Having briefly touched on the difference between Telcos and Cable Operators above, a focused discussion is warranted. The DOCSIS protocols have been wildly successful for Cable Operators providing Internet services to customers over the past two decades. Telcos with legacy twisted-pair networks were faced with a horrible reality and eventually acknowledged that they either needed to build expensive FTTP networks or they would be left out of (or even have to abandon) the Residential market for Internet, and potentially the market for Residential Voice, while leaving the potential market for Residential Video untapped.

In the United States, operators started ambitious FTTP programs like AT&T U-Verse, Verizon Fios, and Google Fiber, but these programs delivered a product roughly comparable to what was available with DOCSIS at Cable Operators at the time, and these FTTP programs scaled back as costs escalated.

In Canada, Telcos also started ambitious FTTP programs, but continued building as the costs per household-passed increased, partially-enabled by a falling interest rate environment that allowed for a debt-fueled spending binge.

With the protocols available at the time, companies that chose to build FTTP networks to salvage Residential markets knew that Centralized OLTs were required, and probably planned for 1:32 split ratios. They therefore knew that large Backbone Fiber constructions programs would be required everywhere in their serving areas.

In Western Canada, the regional Cable Operator's experimental FTTP deployment program was halted in 2011 when the two monumental limitations, both the lack of an IP Video solution and FTTP-related Backbone Fiber consumption, proved to be insurmountable with the technology solutions available at that time. There was still sunshine among the clouds; DOCSIS was a great service delivery method with a known architecture supported by sufficient infrastructure resources in both the Hubsite and the Outside Plant, and the protocols being used by the competitor Telcos did not provide a superior downstream for their Internet product.

In the years since the suspension of the experimental FTTP deployment program, success-based growth in network traffic required upgrades and augments to the infrastructure in Hubsites, where rackspace and power usage has now approached upper limits of what the Hubsites can support. The pivot to DAA has allowed for a cap & grow operational mode that has kept these constraints manageable so far, but these

Hubsite constraints are still large enough to prevent a widescale deployment of Centralized OLTs across the network.

In the Hubsite, DAA requires approximately 1 port in a CIN router per 250 HHP, and Distributed FTTP appears to have similar requirements for unprotected links for Remote OLT backhaul connectivity.

Growth in upstream traffic has been accommodated by a 3-year mid-split (85 MHz) return program that was fortuitously completed just prior to the COVID-19 pandemic lockdowns.

Considering Brownfield overbuilds, recent deployments of DOCSIS 3.1 and the mid-split program, together with DOCSIS 4.0 coming soon, make it hard to justify overbuilding existing HFC plants with FTTP. Additionally, existing HFC costs points are currently much lower than FTTP on a per HHP basis. HFC is still a great architecture that has a long life ahead of it and that will still meet the needs of most customers for many years with good economics.

Considering fresh Greenfield builds, today's FTTP protocols of 10G EPON and XGS-PON provide adequate capacities, in both the downstream and upstream, while being capable of serving a higher number of customers per PON port than what was possible in the past. Developers are starting to prefer FTTP over HFC for new builds, and customers are generating more upstream traffic, enough to consider symmetrical services in their decision-making process. [Price still appears to be the primary consideration where customers have a choice between ISPs that use different technologies for their service delivery.]

With these thoughts in mind, the experimental FTTP deployment program was restarted in Western Canada a few years ago. The two prohibitive limitations had been resolved:

- An IP Video solution is now in place that allows for a Residential triple-play Voice, Video and Internet product offering.
- Remote OLT products are now available that allow for the use of passive DWDM and tunable optics so as to allow the Backbone Fiber to be shared amongst a variety of users.

There are now revenue-generating customers on multiple Distributed FTTP networks in Western Canada.

FTTP deployments are no longer generally-problematic for Cable Operators, but constraints in the existing Backbone Fiber might require an architecture like Distributed FTTP at other operators like it is required in Western Canada. Ultimately, the Backbone Fiber in Western Canada must be shared amongst many services (RPDs, Remote OLTs, Business services, Small Cell / Wireless services, etc), and it is anticipated that other Cable Operators face similar constraints.

To summarize, for the existing Access Network Infrastructure in Western Canada, Remote OLTs are absolutely required to enable a standardized architecture that simplifies planning, construction, deployment, acceptance testing, and service delivery for new FTTP customers in Greenfield SFU, MDU, Commercial Business Park, and Rural Network Expansion builds. Centralized OLTs consume too much Backbone Fiber, and there are not enough resources available in Hubsites, whether considering rackspace or power, to assume that new equipment can be added at existing Hubsite facilities.

4.6. Future PON Protocols

There are many developments with regards to future PON protocols, but the following discussion is meant to narrow the focus to those details specifically relevant to Distributed FTTP with Remote OLTs.

4.6.1. 25G PON (25G EPON and 25GS-PON MSA)

The 25GS-PON Multi-Source Agreement (MSA) is an initiative spearheaded by Nokia and AT&T to create a market ecosystem for a PON protocol with a higher capacity than today's XGS-PON protocol. To summarize, the PHY layer essentially copies what the IEEE has specified for 25G EPON, while extending the Transmission Convergence (TC) layer used for XGS-PON.

As of July 2023, there were 54 member companies officially participating in the initiative, but notably absent were large suppliers Huawei and ZTE, and the large operator Verizon. Speculation as to why they have not signed on as of the writing of this Technical Paper is beyond the scope of this paper.

There are two specific benefits of 25GS-PON worth discussing in the context of Distributed FTTP when compared to 10G EPON or XGS-PON protocols:

1. 25GS-PON provides a higher aggregate data capacity over an ODN.
2. 25GS-PON provides a higher data capacity to a single subscriber.

For the first specific benefit of 25GS-PON to provide higher aggregate data capacities, the protocol should provide approximately 21 Gbps over the shared ODN as compared to a data rate of approximately 8.8 Gbps over an ODN running 10G EPON or XGS-PON protocols.

This first specific benefit of 25GS-PON is necessary for legacy architectures with Centralized OLTs because each PON port consumes a single strand of Backbone Fiber, and any constraints on Hubsite rackspace or power, or Backbone Fiber availability, will shackle the ability of the operator to add new revenue-generating customers if there are constraints or congestion anywhere in that network.

This first specific benefit of 25GS-PON is completely negated for Distributed FTTP architectures since lighting up new PON ports on Remote OLTs does not require new assignments of Backbone Fiber. In a Distributed FTTP network, a congested ODN with a 1:128 split ratio can be easily segmented into two ODNs with 1:64 split ratios, or four ODNs with 1:32 split ratios. Another alternative would be to 'cherry-pick' high-traffic users off the congested ODN and onto a new ODN with a smaller split ratio.

Operators with Distributed FTTP architectures do not require 25GS-PON to provide higher aggregate data capacities to customers.

In Western Canada, the splitters at the central splitting location are deliberately planned with 1:2 and 1:4 splitters cascading into 1:32 splitters so that segmentation activities can be easily accomplished in future years. This adds a small amount of additional insertion loss from the additional connectors, but that loss is easily overcome because the Remote OLTs are so close to the customers.

For the second specific benefit of 25GS-PON to provide higher data capacity to a single subscriber, the customer premise equipment would need to support higher flows for 25G Ethernet, probably through a pluggable SFP28 module. If customers have high traffic loads that consistently exceed the capacity of a 10G EPON or XGS-PON network, then a dedicated wavelength service would be more appropriate for them over a shared-media PON service.

For operators with Centralized OLTs, or for operators with distributed splitting, it is difficult to provide a dedicated wavelength service as separate fibers and splicing activities are required. For operators running a Distributed FTTP architecture with central splitting and with passive DWDM between the Hubsite and the Node Service Area, adding a dedicated wavelength bypass service would still require splicing at the central splitting location and at the FOSC(s) with the passive DWDM breakout, but no new Backbone

Fibers would be required. A router port with DWDM optics in the Hubsite would be required to terminate the wavelength.

For those customers with a high peak burst traffic approaching 20 Gbps, but with much lower day-to-day traffic, a shared-media PON service running 25G EPON or 25GS-PON protocols could be a workable solution. Operators with Distributed FTTP architectures should evaluate Remote OLTs capable of these 25G PON protocols for these customers with specialized needs of this burst traffic that exceeds the capabilities of today's 10G EPON and XGS-PON protocols. Some of the participants in the 25GS-PON MSA are working on 25GS-PON Remote OLTs, and it is believed that there are some 25G EPON Remote OLT products also available or in development.

If operators do not have customers with high peak burst traffic patterns, AND these operators have Distributed FTTP architecture, AND future generations of PON protocols develop as expected (CableLabs Coherent PON Working Group, for example), then the decision as to whether to deploy 25G PON protocols (25G EPON or 25GS-PON) would come down to economics and cost points.

Another thing to consider for 25G PON specific to Remote OLTs is the Ethernet backhaul component. Development is ongoing for pluggable 25G optics that meet the requirements for devices in the field:

- Industrial Temperature (-40 C to +85 C)
- Distances preferably up to 80 km (optical power and dispersion tolerance)
- DWDM full-band tunable to share Backbone Fiber

Products that meet these requirements hopefully will be generally available in 2024.

If two 10G Ethernet pluggable optics are used, then the full capacity of the 25G protocol would not be available back to the router, and only a single 25G PON port would be available at the Remote OLT.

For operators needing multiple 25G PON ports activated at single Remote OLT, then 100G DWDM wavelengths could be used, like those envisioned and standardized by the CableLabs Coherent Optics P2P Working Group, published specification P2PCO-SP-PHYv1.0-I03.

An interesting dynamic for Cable Operators thinking about deploying Remote OLTs with 25G PON is that Remote PHY Devices for DOCSIS 4.0 might have similar backhaul requirements, if multiple serving groups are fed out of the same Node Enclosure, depending on the achievable throughput.

Table 2 – Expected DOCSIS 4.0 Modulation Throughput

Modulation Rate (QAM)	Theoretical Throughput (bits/s/Hz)	Achievable Throughput (bits/s/Hz)
256	8	6
1024	10	7.5
2048	11	
4096	12	9
8192		
16384		10.5

The DOCSIS 4.0 spectrum available for downstream (DS) is dependent on where the split for upstream (US) return occurs.

Table 3 – Expected DOCSIS 4.0 Downstream Spectrum Usage

DS Lower Band Edge	54 MHz	87.5 MHz	108 MHz	258 MHz	372 MHz	492 MHz	606 MHz	834 MHz
192 MHz DS blocks	9	9	9	8	8	7	6	5
Total DS Spectrum (MHz)	1728	1728	1728	1536	1536	1344	1152	960

Noting that achievable throughput varies and is lower than the theoretical throughput, an estimate for the downstream traffic possible on a DOCSIS 4.0 network can be made based on the number of 192 MHz downstream blocks are possible (which is a function of the where the upstream split is, as per the DOCSIS 4.0 specifications).

Table 4 – Expected DOCSIS 4.0 RPD Backhaul Requirements

192 MHz Blocks	Achievable QAM	Achievable Throughput (bits/s/Hz)	Usable Downstream Spectrum (MHz)	Capacity Calculation (Gbps), per Serving Group
1	4096	9	192	1.728
2	4096	9	384	3.456
3	4096	9	576	5.184
4	4096	9	768	6.912
5	4096	9	960	8.64
6	4096	9	1152	10.368
7	4096	9	1344	12.096
8	4096	9	1536	13.824
9	4096	9	1728	15.552

It is expected that initial iterations of DOCSIS 4.0 RPDs will only be capable of 1 serving group, and therefore only two 10G Ethernet backhaul connections combined in a Link Aggregation Group (LAG) would be required. Even at 12 bits/s/Hz and 1536 MHz of downstream spectrum (upstream 5 MHz to 300 MHz, downstream 372 MHz to 1794 MHz), a single serving group would require 18.432 Gbps and thus be satisfied by a 20 Gbps LAG for backhaul.

Regardless of whether an operator is looking at 25G PON in Remote OLT, or DOCSIS 4.0 in an RPD, once the number of PON ports or serving groups grows beyond one per Node Enclosure, the backhaul will start to get non-trivial and will require some planning and forethought. This is not expected for a number of years, but will need to be a consideration for how Remote OLTs and Remote PHY Devices are designed in the future.

4.6.2. 50G PON (50G EPON and an ITU-T variant)

For Cable Operators with Distributed FTTP, 50G PON variants are not interesting. These 50G PON variants address the same challenges as 25G PON, namely higher aggregate data and higher data to a singular customer. For higher singular data flow requirements, Distributed FTTP with central splitting allows problematic traffic flows, whether 25G or 50G, to be moved to separate ODNs or dedicated wavelength services without additional Backbone Fiber, as discussed above. For aggregate data flow requirements, Distributed FTTP with Remote OLTs allows for easy segmentation of congested ODNs.

4.6.3. 100G PON (CableLabs Coherent PON Working Group)

In Western Canada, any PON protocol that cannot easily share Backbone Fiber must be relegated to Remote OLTs; building more Backbone Fiber is not an acceptable pre-requisite for connecting new FTTP customers.

If a PON protocol was made available that could easily share the Backbone Fiber by being tunable across a wide range of the DWDM spectrum currently used to share this Backbone Fiber, and if the optical link budget was sufficient enough to overcome distances and insertion losses typical for these passive DWDM deployments, while still providing a high enough split ratio to be considered efficient, then a Centralized OLT solution would be acceptable.

Remote OLTs are not necessarily desirable because they become another active in the field that requires planning, permitting, construction, deployment, and on-going maintenance. Remote OLTs are not being deployed as a function of choice; rather, Remote OLTs are being deployed because Centralized OLTs do not work with the 10G PON protocols today for the requirements that Cable Operators face, namely using existing Backbone Fiber sized and built decades ago.

This presents a gap in PON protocol developments. The market adoption of XGS-PON vs NG-PON2 at Telcos leads to the assumption that future development at the ITU-T will not prioritize sharing of Backbone Fibers. CableLabs' Coherent PON Working Group is working on a specification for a PON protocol for a next-generation 100G coherent PON (CPON). For deployment in Western Canada, CPON will need to be able to share the Backbone Fiber with other DWDM wavelength services.

When considering the differences between Centralized OLTs and Remote OLTs, the backhaul portion for Remote OLTs at 10 Gbps is similar to RPDs running DOCSIS 3.1, and is a known quantity supportable by the CIN routers used for the initial DAA deployments. However, at 100 Gbps DWDM backhaul, the costs are much higher for that 100 Gbps backhaul link with Remote OLTs. This is because 100G router ports and 100G tunable DWDM optics will be required, which are currently generally more expensive for Cable Operators than 10G ports.

This 100G cost premium becomes a Day 1 cost incurred at the initial deployment of the Remote OLT regardless of the number of planned customers. If the tunability (required for sharing the Backbone Fiber) and the optical link budget (required for high split ratios and long distances) can be obtained through design of the OLT and ONTs with sufficient specification and forethought in the PON protocol, then the economics place the cost burden on a success-based ONT deployment model rather than a Day 1 backhaul requirement equipment investment burden. This allows for a Pay-as-You-Grow model to be applied for deploying 100G PON overlays on top of existing 10G PON deployments. This is desirable for upgrading individual existing customers on existing ODNs, and will extend the product lifecycle for the 10G FTTP networks.

The scenario for minimizing 100G PON ONT costs with limited tunability and optical link budget specifications will tend to force deployment models at Cable Operators towards Distributed FTTP with Remote OLTs and therefore lead to higher Day 1 deployment costs. This might be acceptable in larger Greenfield builds but will be problematic when trying to proceed with success-based upgrades of existing 10G PON deployments.

For the upgrade scenario where a CPON overlay is being used to cherry-pick problematic traffic flows off existing 10G PON ODNs, Centralized CPON OLT with tunability and high optical link budget specifications split at a Hubsite feeding multiple Node Service Areas should be dramatically simpler, and hopefully cheaper, than building new Remote OLTs at each of the Node Service Areas where this traffic grooming effort is required. An operator could even proactively deploy CPON signals into Node Service Areas with existing 10G FTTP ODNs and migrate the highest 2 or 3 customers per ODN from 10G PON to CPON and thus alleviate congestion before other customers even start to notice.

This ability to cherry-pick customers does require a central splitting architecture; a complete comparison of different splitting architectures is beyond the scope of this paper, but central splitting is used in Western Canada.

As previously discussed in this paper, an experimental FTTP deployment program was halted a decade ago because of unacceptable Backbone Fiber overbuild requirements; the CableLabs CPON protocol must not shackle Cable Operators in the future with a similar requirement. A low cost ONT for Cable Operators that forces Backbone Fiber overbuilds will result in a much-higher total overall deployment cost that will ultimately slow, or even halt, 100G PON deployments at Cable Operators in the future.

It is therefore recommended that the CPON protocol development efforts focus on PHY-layer specifications that meet the tunability and optical link budget specifications that Cable Operators need rather than having a narrow focus on ONT material costs for a total overall lowest-cost ONT.

As a cautionary note, deployments of equipment in Hubsites always require rackspace and power; as such, while this is still many years away, Cable Operators should proceed with DAA conversions so that rackspace and power can be reclaimed in sufficient quantities and timeframes so as to allow for new Centralized OLT deployments once future generations of PON protocols are available (such as CPON). Remote OLTs can be used for 10G PON deployments, and probably 25G PON deployments, but it is unclear whether this will be desirable, let alone possible, for 100G PON protocols and beyond.

5. Conclusion

It is possible for operators to deploy FTTP architectures for Greenfield MDU, SFU, Commercial Business Park, and Rural Network Expansion builds. For Cable Operators who have migrated to an IP Video solution and who have Backbone Fiber cables that were sized & built decades ago, existing DAA equipment can be leveraged to accelerate the deployment of FTTP.

Distributed FTTP architecture utilizes new Remote OLTs for PON port activations, and together with existing DAA equipment in both the Hubsite (CIN routers) and Outside Plant (passive DWDM), allows Cable Operators to light up new ODNs while decoupling the linkage between PON port utilization and Backbone Fiber utilization. Cable Operators can now deploy FTTP for customers without worrying about exhausting scarce Backbone Fiber resources, which was an insurmountable constraint for those operators trying to deploy legacy FTTP architectures.

Care should be taken that future PON protocols allow Cable Operators to revert to FTTP architectures with Centralized OLTs, rather than forcing Cable Operators to adopt Distributed FTTP architectures with Remote OLTs. This will allow Cable Operators to choose architectures and parameters best suited to their needs, rather than be forced into a specific architecture with limited product availability. This will require future PON protocols to be tunable across multiple wavelengths and to have sufficient optical link budget specifications to overcome passive DWDM insertion losses, fiber attenuation loss from distance, and optical power loss from power splitters inside the ODN.

Abbreviations

BER	bit error rate
bps	bits per second
CableLabs	Cable Television Laboratories
CIN	Converged Interconnect Network
CO	central office
CPON	coherent passive optical network
DAA	Distributed Access Architecture
DOCSIS	Data-Over-Cable Service Interface Specifications
DPoE	DOCSIS provisioning of EPON
DS	downstream
DWDM	dense wavelength division multiplexing
EPON	Ethernet passive optical network
FDH	fiber distribution hub
FEC	forward error correction
FOSC	fiber optic splice cabinet
FTC	fiber transition cabinet
FTTP	fiber-to-the-premises
GPON	gigabit passive optical network
HFC	hybrid fiber-coaxial
HHP	households-passed
Hz	hertz
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
ITU-T	Telecommunication Standardization Sector of International Telecommunication Union
LAG	link aggregation group
LCP	local convergence point
MDU	multiple dwelling unit
MPLS	multiple protocol label switching
MSA	multi-source agreement
MTTR	mean-time-to-repair
MWL	multiple wavelength
ODN	optical distribution network
OLT	optical line terminal
ONT	optical network terminal
OSA	optical spectrum analyzer
PHY	physical layer
PON	passive optical network
RFoG	Radio-Frequency-over-Glass
RPD	remote PHY device
SFP	small formfactor pluggable
SFU	single-family unit
SIP	Session Initiation Protocol
US	upstream
XGS-PON	10-Gigabit-capable symmetric passive optical network

Bibliography & References

CableLabs CM-SP-PHYv4.0-I06-221019: *Data-Over-Cable Service Interface Specifications DOCSIS 4.0 – Physical Layer Specification*

CableLabs CPON-SP-ARCH-I01-230503: *Coherent Passive Optical Networks 100 Gbps Single-Wavelength PON – Coherent PON Architecture Specification*

CableLabs P2PCO-SP-PHYv1.0-I03-200501: *Point-to-Point Coherent Optics – P2P Coherent Optics Physical Layer 1.0 Specification*

ITU-T G.9807.1 Amendment 1: *10-Gigabit-capable symmetric passive optical network*

