Learnings From a DAA/DOCSIS 3.1 Early Adopter

Launching and Maintaining a Next-Gen HFC Plant

An Operational Practice prepared for SCTE•ISBE by

Jim Walsh
Solutions Marketing Manager
VIAVI Solutions

David Hering
Senior Product Line Manager
VIAVI Solutions

David Judge
Cable Broadband Engineer
Vodafone New Zealand

Table of Contents

List of Figures

Introduction

The Hybrid Fiber Coaxial (HFC) plant is rapidly evolving to economically address subscriber bandwidth and service quality demands, and technologies like DOCSIS 3.1 and distributed access architectures (DAA) are key to this evolution. While these breakthrough technologies will enable HFC to remain the dominant broadband service medium for years to come, they also present challenges to the folks who must maintain the plant. This paper will touch on the drivers behind adoption of these technologies and follow the early planning, preparation, and live-plant rollout of DAA and DOCSIS 3.1 from an early-adopter MSO and test vendor including early successes and lessons learned. Examples will be provided demonstrating how industry-wide collaboration and standardization have been and will remain keys to success.

Distributed Access Architecture Drivers

As an industry, we have grown accustomed to our subscribers demanding more bandwidth each year, usually at the same price as last year's service. Each time we think that bandwidth demand models must be wrong, and that there is no way subscribers can possibly consume data at levels predicted for just a couple of years out, a new killer app emerges to prove us wrong. Over the top video (OTT) is a prime example. Even just a few years ago not many could have imagined our customers' binge-watching habits or that three kids could be simultaneously talking with their friends via FaceTime. In the past cable operators have used the same playbook to address the need for increased bandwidth:

- Add more carriers
- Get more bits/Hz of spectrum
- If all else fails – node splits

Unfortunately, the traditional tools in MSO toolboxes alone are not enough anymore

- **Add More Carriers** – Analog reclamation has freed up space, and optional 1.2GHz/204MHz DOCSIS 3.1 extensions help but are not always practical/affordable to implement
- **Get more bits/Hz of spectrum** – DOCSIS 3.1 helps here with the potential to move to modulations as high as 4096QAM, but moving above 1024QAM often requires SNR improvements that are challenging with current system limitations
- **Node Splits** – Splitting nodes enables smaller service groups, but increases the total number of nodes. Each added node brings with it more gear in the hubs that require CAPEX, rack space, and power/cooling. Rack space and power/cooling budgets alone take this option off the table in many cases.

The Solution: Distributed Access Architectures

Fortunately, the best and brightest minds in cable predicted this crunch and developed/specified a solution to enable the required node split rate while also enabling more bits to be squeezed from each Hz of plant. Hub constraints can be eased by distributing network functions out into the field – functions that previously resided in equipment that occupied space and consumed power in the hubs. The fidgety analog optical link can be converted to a robust, commodity 10G Ethernet optical link with higher SNR's, enabling higher modulation orders (and therefore more bits/Hz).

1. Distributed Access Architecture Variants

Presently the two dominant DAA variants are Remote PHY (R-PHY) and Remote CCAP (R-CCAP), although R-MACPHY solutions under development are gaining mindshare among operators. Within R-PHY there are several sub-variants. We will not cover the technical differences or pros/cons of the variants in detail, as there have been many great papers on this topic in SCTE Cable Tech Expo proceedings from previous years. We will cover just enough to enable an understanding of their impacts on test and measurement.

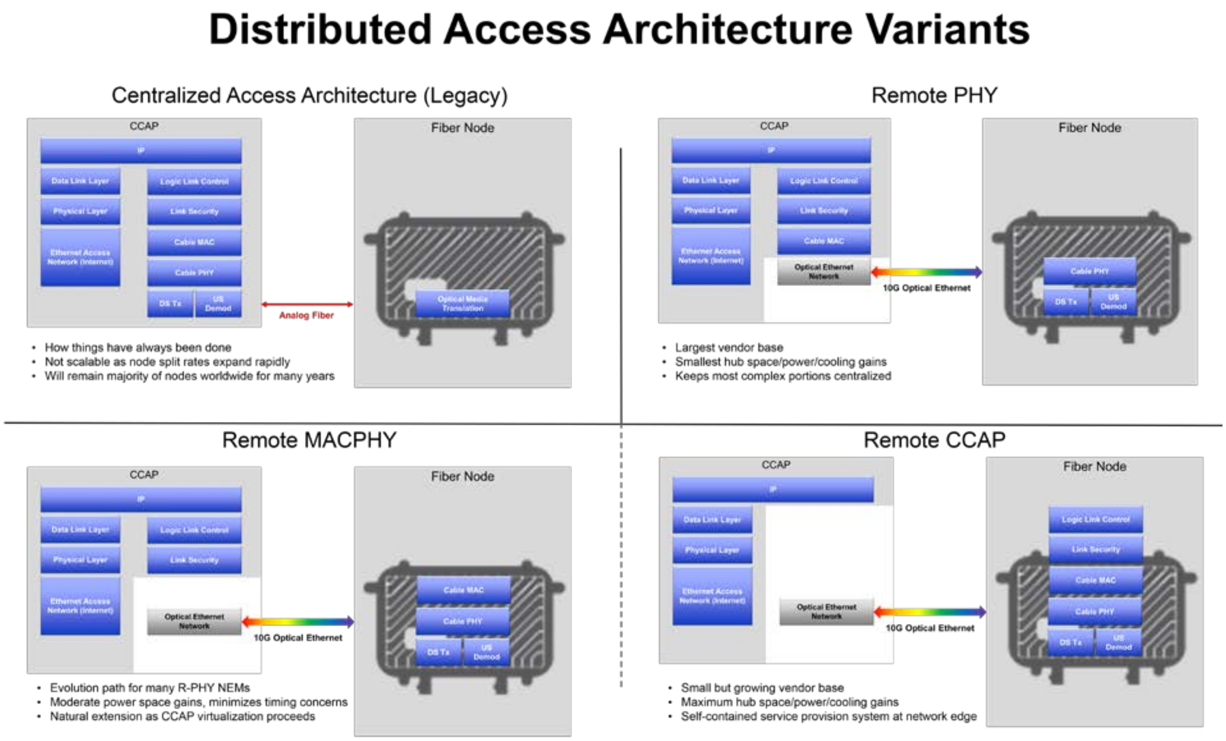

Figure 1 - Distributed Access Architecture Variants

As its name implies, R-PHY moves just the PHY layer out into the node while the MAC layer remains centralized. This MAC/PHY separation creates challenges related to Ethernet timing. This will be discussed in detail in later sections. R-PHY has been the most popular option from established CCAP vendors although they have also been following up with R-MACPHY offerings as well. R-PHY tends to deliver the smallest rack space/power/cooling savings but is viewed by some as the simplest from an overall implementation and management standpoint as the MAC layer management remains centralized.

At the opposite end of the spectrum is R-CCAP, moving nearly all CCAP functionality out into the field. This option took off, as one vendor was very early to market with a solution enabling simultaneous transition to DAA and DOCSIS 3.1. Maximum rack space/power/cooling benefits are generally obtained with this option but concerns about single-sourcing and decentralized management have limited adoption among major operators.

Commonalities for all DAA variants:

- Redistribution of hub/headend-based capabilities into the outside plant

- Replacement of analog optical link with commodity Ethernet optical link
- Removal of RF combining network and related test points from hub/headend

The first point above is perhaps key to most of the challenges that operators will face with DAA deployments. By splitting up previously co-located layers, a brand new optical media translation interface has been created. This is disruptive to traditional demarcation points for existing MSO groups and processes. It is not clear nor obvious where the split between headend or field groups lies with this messy separation. This creates an environment rife with opportunities for responsibilities to fall through the cracks and finger-pointing during troubleshooting. Figuring out and clearly defining ownership of all aspects of this new interface is critical to long-term DAA success.

2. Other DAA Benefits

The obvious benefit of DAA is the enablement of rapidly-accelerating node split rates without expanding current hub/headend real estate and power-cooling footprints, but there are several other compelling benefits.

- Higher SNR digital optical link enables higher modulation order attainment in DOCSIS 3.1 (More bits/Hz!)
- More robust optical link – higher plant reliability and reduced maintenance costs due to replacement of fidgety analog optical link
- Increased service offering flexibility is enabled as Ethernet pushes deeper into the plant. Serve high-usage customers or FTTH clusters via EPON, provide Ethernet for business services, or 5G backhaul from the mux deployed deep into the network instead of running dedicated fiber for each.

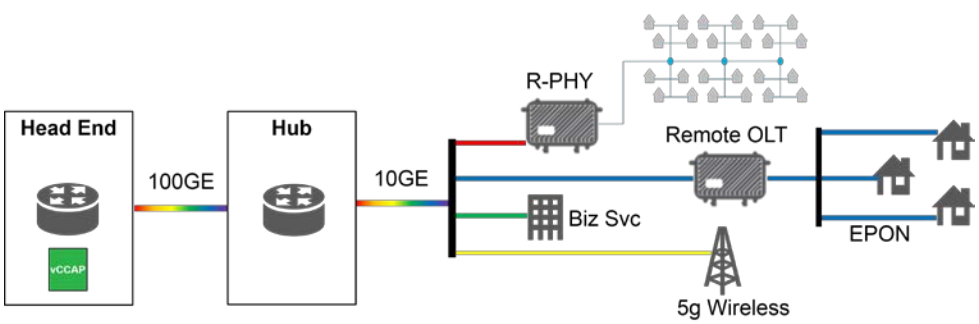

Figure 2 - Benefits of Deeper Ethernet

Expected Challenges for Operators

While the benefits of DAA are significant, they come with challenges. In this section we will focus on the challenges faced by those who must turn-up, monitor, and maintain the HFC plant. The challenges generally fall into several general categories:

- Visibility & injection limitations via removal of RF from hubs
- Ethernet timing due to MAC-PHY separation
- Expanded use of previously-specialized technologies
- Increased field complexity amid architectures/vendor proliferation

1. Removal of RF from hubs

Upstream ingress has been the bane of cable operators since the earliest days of 2-way HFC communications, and there are no indications that this will change even in fiber deep architectures. Historically 75%+ of upstream ingress enters the plant through sources originating as home or drop issues. This is not expected to change as plants transition to fiber deep. Customers will continue messing with their home wiring and drops will continue to degrade. Hub-based return path monitoring systems have been a critical tool used in conjunction with field meters to efficiently address upstream noise, but these hub-based systems rely on a return RF feed to operate. All DAA variants have the common trait of removing the RF combining network from hubs disallowing the use of traditional rack-mounted gear and forcing a change in methodologies. Return sweep capabilities face a similar predicament, but with the added challenge related to the inability to combine telemetry signaling into downstream RF in the hub. On the downstream side, forward sweep and leakage taggers also rely on RF combining access which will no longer be present in hubs in a DAA environment.

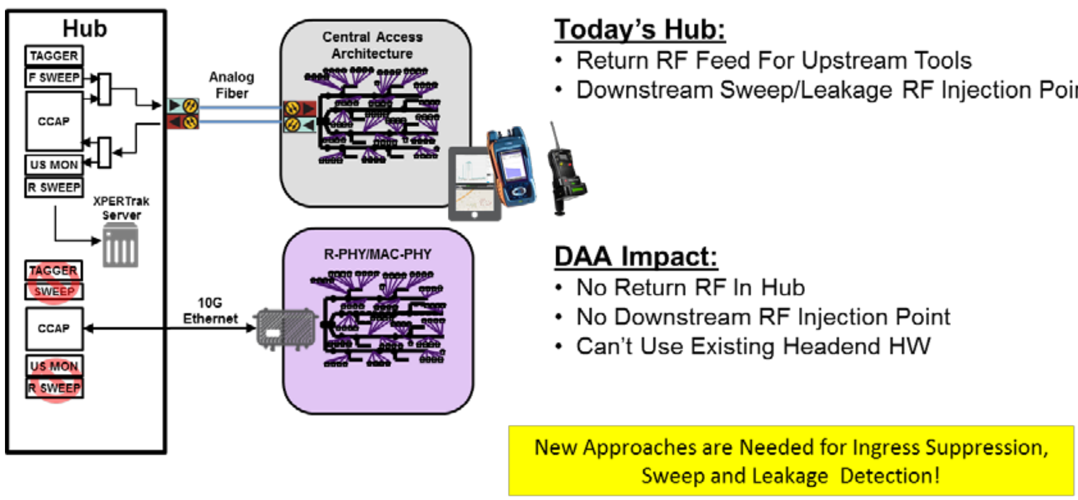

Figure 3 - Challenges Created By Removal of RF From Hubs

One solution to the upstream ingress and return sweep visibility challenge is to virtualize the upstream spectrum analysis and return sweep receiver functionalities into the Remote PHY Device (RPD). This virtualization, combined with a software agent (See RCI in figure below) orchestrating communications between the field meter, central server, CCAP, and RPD enables technicians to use the same process and field meter for sweeping the upstream of both legacy centralized access architecture (CAA) nodes and DAA nodes. The RCI would typically run as an instance on a virtual machine co-located with CCAP or could run on standalone commercial off the shelf (COTS) x86 server.

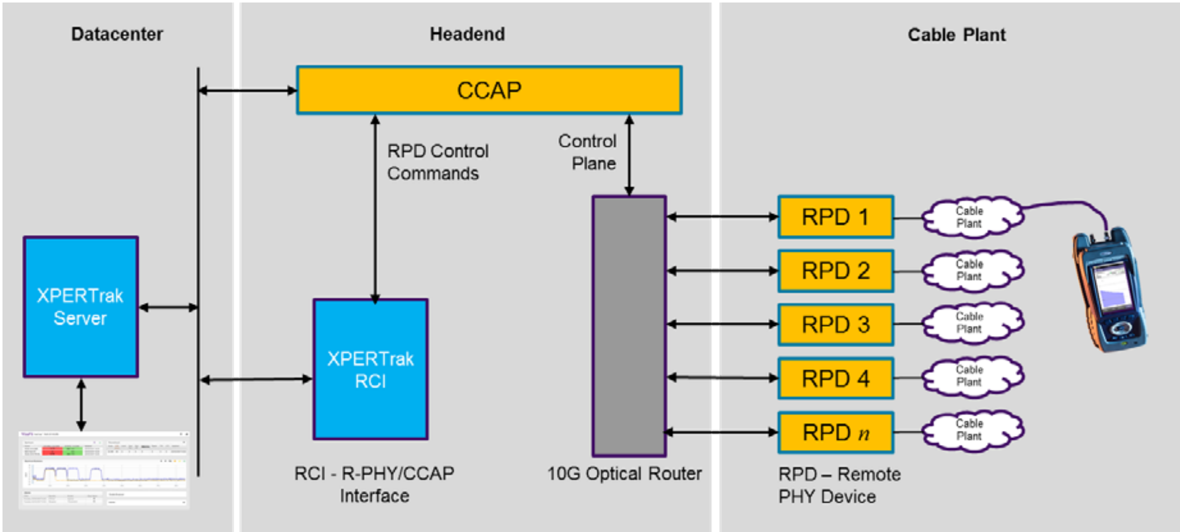

Figure 4 - Virtualized Upstream Ingress and Sweep System

Sweepless sweep has generally been accepted as a substitute for meter-based sweep for the downstream, and leakage tagger capabilities are currently implemented in several DAA vendors systems, with more in development. Through these virtualized solutions, critical maintenance capabilities can be retained throughout the DAA transition while minimizing the complexity faced by technicians in the field.

2. Ethernet Timing Concerns Due To Separation of MAC and PHY

DOCSIS upstream communications (DOCSIS 3.0 and 3.1) at their very core are reliant upon the maintenance of precision timing within MAC domain groups. Modems are assigned precise time slots in which to transmit packets across a shared upstream path, and if a modem transmits out of their assigned slot the potential exists for collisions with packets from other modems resulting in BER or dropped packets. To achieve this synchronization, time information is communicated to the CCAP and RPD's by the Precision Timing Protocol (PTP) Master, and from there each RPD communicates it to CPE via DOCSIS Timing Protocol (DTP). The illustration below shows the typical PTP process.

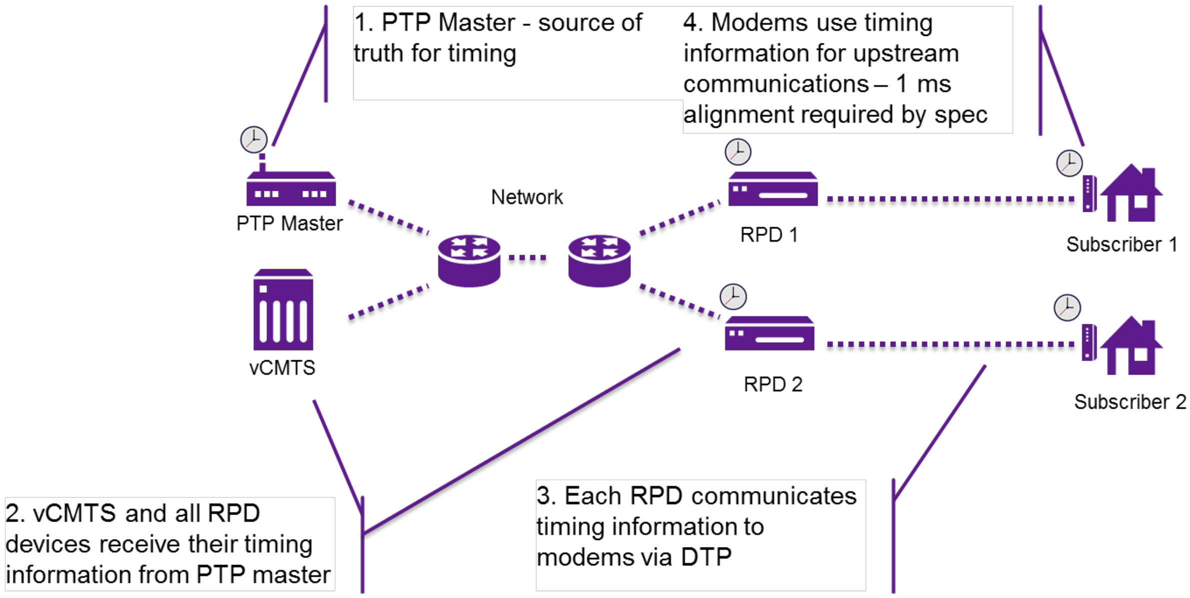

Figure 5 - PTP Use in Remote PHY

The inherent problem with PTP when the MAC and PHY are split is the assumption of symmetry in delay from RPD and PTP Master. The PTP Master communicates total bi-directional delay time for the round trip from itself to each RPD and back. The RPD's assume that this delay is symmetrical and simply divides the delay by two to calculate the delay in each direction. If something impacts only one leg of this journey creating an asymmetrical delay, an error in the delay reported vs actual for the delayed leg of one half of the asymmetry will occur. See below for a numerical example of how delayed PTP messages could result in improper clock synchronization.

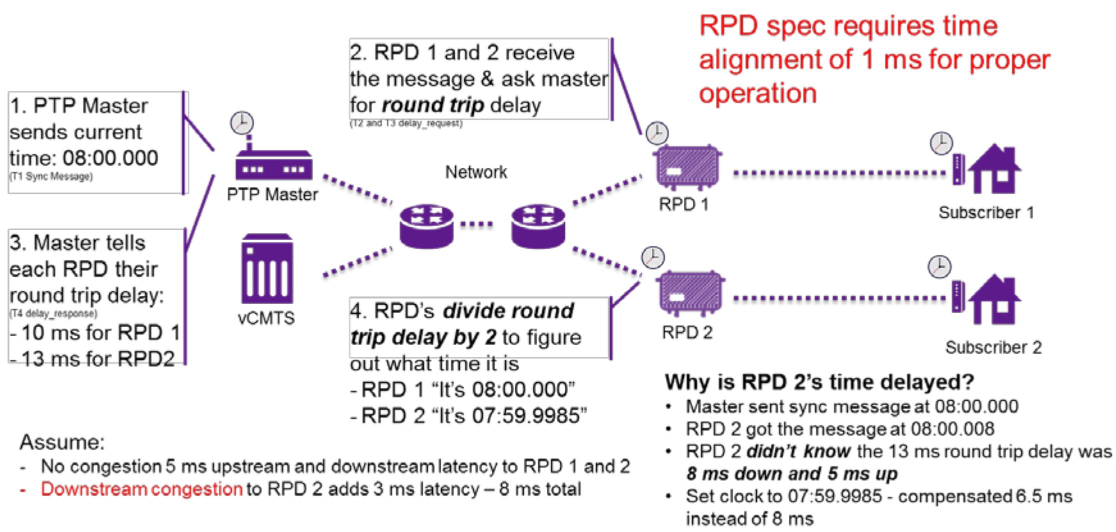

Figure 6 - Example of How PTP Error Can Manifest

3. Expanded Use of Previously-Specialized Technologies

While fiber and Ethernet technologies and testing are not new to the HFC, historically they have been delegated to a specially trained and equipped subset of technicians. The rollout of DAA will push both technologies deeper into the network and force operators to train and equip a much larger portion of their workforce to properly install, maintain, and troubleshoot issues with these technologies. The criticality of clean fiber connections cannot be understated with up to 80% of issues in fiber networks resulting from dirty or damaged connections. Technicians will need training on proper fiber handling and inspection/cleaning and access to the tools to do this efficiently. Dense Wavelength Division Multiplexing (DWDM) has found limited applications in the HFC to date but will play a much larger role in future DAA-enabled networks. The same goes for Ethernet testing and troubleshooting. This can no longer be the exclusive realm of business service techs. Many more maintenance techs will now need to test Ethernet services as part of RPD turn-up and troubleshooting.

4. Increased Complexity Amid Architecture/Vendor Proliferation

The transition to DAA will be a multi-year process for most operators, not a one-time event. It is expected that the majority of most operators' nodes will still be CAA for at least the next 5-10 years or more. It is also expected that large operators (at least) will implement multiple DAA architectures as specific situations dictate, resulting in a multi-year period where any given region may consist of a mix of legacy CAA nodes and DAA nodes using various architectures from multiple vendors.

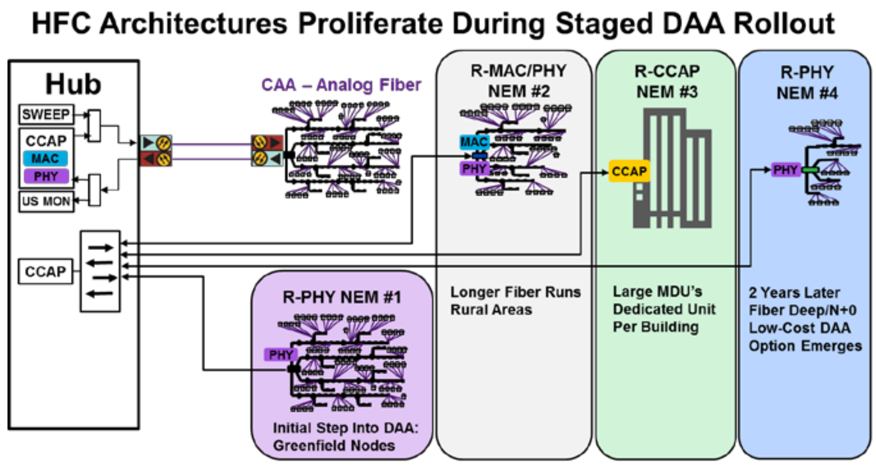

Figure 7 - Complexity Introduced By DAA Variant/Vendor Proliferation

Unless standard maintenance processes that apply across the heterogeneous mix of nodes can be implemented, life for technicians will become quite complex. Imagine the confusion if a tech must use one set of processes and tools for a CAA node they are working on in the morning, a different set for the R-PHY node assigned after lunch, and yet another for the R-CCAP node serving the MDU they are assigned to troubleshoot in the afternoon. Not exactly a recipe for success or operational efficiency! Care must be taken to ensure that maintenance needs are considered during network design and tool selection to ensure that saving pennies on the front end doesn't result in extra dollars of cost to monitor and maintain them.

One Early-Adopter Company's DAA Voyage & How They Overcame Challenges

In this section we will chronical the voyage undertaken by Vodafone New Zealand (NZ), an early adopter for both DOCSIS 3.1 and DAA technologies. Vodafone is a leading broadband provider in New Zealand and has a sizeable video and broadband customer base in their HFC footprint. Vodafone NZ is historically an early adopter and has not shied away from being on the bleeding edge of new technologies to stay ahead of the competition and provide best-in-class services to their subscribers. Like many cable operators globally, the entry of FTTH competitors was a major driver behind initiatives enabling faster broadband offerings and improved service quality.

1. Architecture Details

Both Christchurch and Wellington employed traditional HFC architecture. The fiber infrastructure was in good shape and could therefore be reused. The new architecture incorporated WDM for the forward path, combining the GPoN traffic, digital TV and telemetry onto a single fiber. The other fiber is used purely to transport return RF back to the hub for use by existing return path monitoring and sweep gear. This solution required no new fiber therefore existing test and measurement gear could be reused.

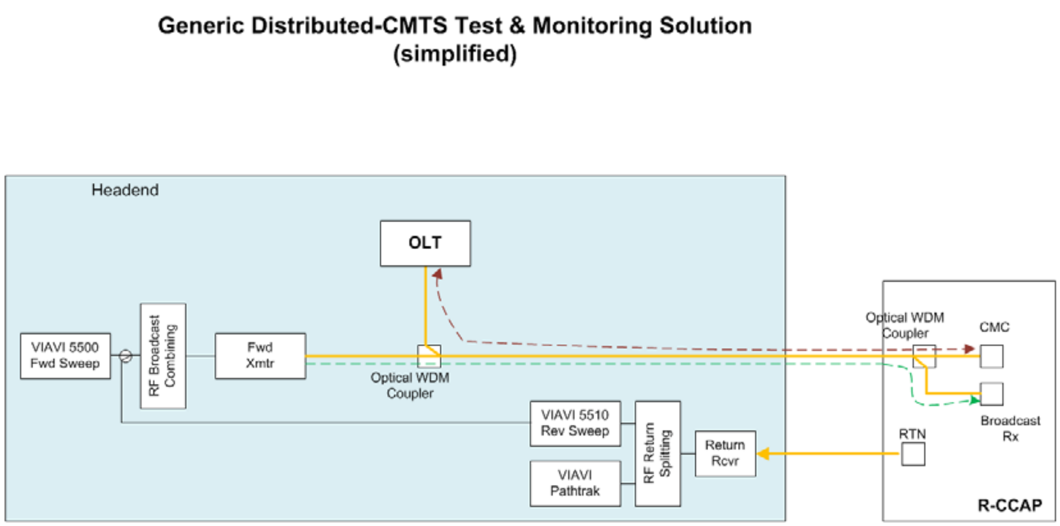

Figure 8 - Interim DAA Maintenance Architecture

Nodes with lower subscriber numbers use a single DAA CMTS to provide desired capacity. For larger nodes, multiple DAA CMTS's are co-located at the optical node. VFNZ has up to two DAA CMTS co-located in an optical node. The forward fiber is split across the two DAA CMTS, resulting in a 3-fiber solution (one shared forward, one each reverse). This solution required both existing fibers to operate, so only one new fiber was required to provide back-feed of RF to the hub for monitoring and sweep use. This was recognized as an interim solution until the D-CMTS vendor implemented the CableLabs specified Narrowband Digital Forward and Return (NDF/NDR) capabilities to support a virtualized test solution.

2. The Deployment Process:

The deployment process itself occurred very quickly and went better than most could have reasonably expected. The entire footprint was cut over to R-CCAP and DOCSIS 3.1 in 6 months, and customer disruption was minimal with outages generally limited to ~60 minutes for single D-CMTS nodes and only during maintenance windows. Sounds easy – right? The secret to the success achieved by Vodafone NZ was careful planning.

Discussions with NEMS and test partners began as far back as 2015 to plot the course for a successful deployment. Attention to detail in system architecture/design assured that maximum capacity could be achieved at minimum cost while still retaining test capabilities critical to maintaining the plant once turned up. After architecture details were nailed down focus shifted to deployment planning.

Configuration and test of all gear before sending to the field for deployment paid great dividends. Centralizing the configuration function helped ensure consistency of configurations across the footprint, enabled use of lower-skilled personnel for field deployment, and minimized opportunities for error.

All units were tested before being sent to the field for deployment. These tests included basic unit functionality, as well as service testing, using racks of CPE in the lab. Pre-deployment testing minimized time spent troubleshooting issues in the field, enabling further standardization of the deployment process. It is much quicker to diagnose defective gear in a controlled environment with skilled personnel available vs in the rain at 2AM with minimal support available. Swapping-in a properly configured replacement unit is also much easier as part of a controlled work procedure vs configuring a replacement in the field.

The field deployment process was also carefully planned and refined over time before large-scale rollout occurred. Contractors performed much of the field work although some Vodafone NZ personnel were available to support as-needed. Typically, two people worked in parallel:

- One tech would splice in new fiber tails and wire in the WDM
- The second tech would swap out the OTN

Once both techs were finished, they would use field meters to perform basic checks (Level/MER), and multiple homes off each node would be polled every 5 min by Vodafone NZ systems to ensure satisfactory performance.

3. The Results

The result of this careful planning and disciplined execution is that everyone won!

Vodafone NZ Subscribers

- Gained access to gigabit service options
- Receive even better quality of service

Vodafone NZ

- Achieved a 300% gain in plant capacity reusing much of their infrastructure
- Implemented future-proofed architecture enabling scalability for future growth
- Retained key maintenance capabilities
- Gained technical knowledge/expertise by being early adopter

4. Lessons Learned

Many things went right in the Vodafone NZ R-CCAP/DOCSIS 3.1 rollout, but there were still some lessons learned along the way. The risks of being an early adopter were known up front, so it was no surprise when these risks materialized into challenges. Being one of the first large-scale DOCSIS 3.1 deployments globally, we surfaced issues related to the immaturity of support for the new protocol and firmware on both the CCAP and CPE. Interoperability issues frequently emerged requiring close cooperation with vendors to address. While these troubleshooting fire drills are disruptive and time-consuming, they are also invaluable learning opportunities for the engineers involved. Working side-by-side with the vendors' subject matter experts allowed Vodafone NZ personnel to gain a deeper understanding of the service provision equipment and DOCSIS protocols than would have otherwise been possible through routine system operation. Anyone who has been through the early adopter experience and the invaluable learning that it provides knows the truth and wisdom behind the quote "that which doesn't kill you only makes you stronger."

Outside of this specific vendor example there are other common themes encountered in working with operators globally on DAA planning and deployments. One is that planned/specified test capabilities often lag behind the launch of new network gear. DOCSIS 3.1 has been deployed in the field for quite some time, yet we are just now starting to see NEMS implement some of the exciting PNM capabilities documented in Section 9 of the DOCSIS 3.1 PHY Layer Specification. Understandably, functionality to ensure the successful passing of packets receives top priority, but as a result, test capabilities often don't release until well after CCAPs and CPE are fielded. The same has held true for the NDF/NDR capabilities documented in the R-PHY spec – few vendors have released support to date although most have a committed roadmap to do so.

General Conclusions

Through work with cable operators globally in developing best practices the following framework has been established for recommended testing for Remote PHY rollouts. This list is generally applicable to all DAA rollouts with the exception of PTP timing tests that apply specifically to Remote PHY.

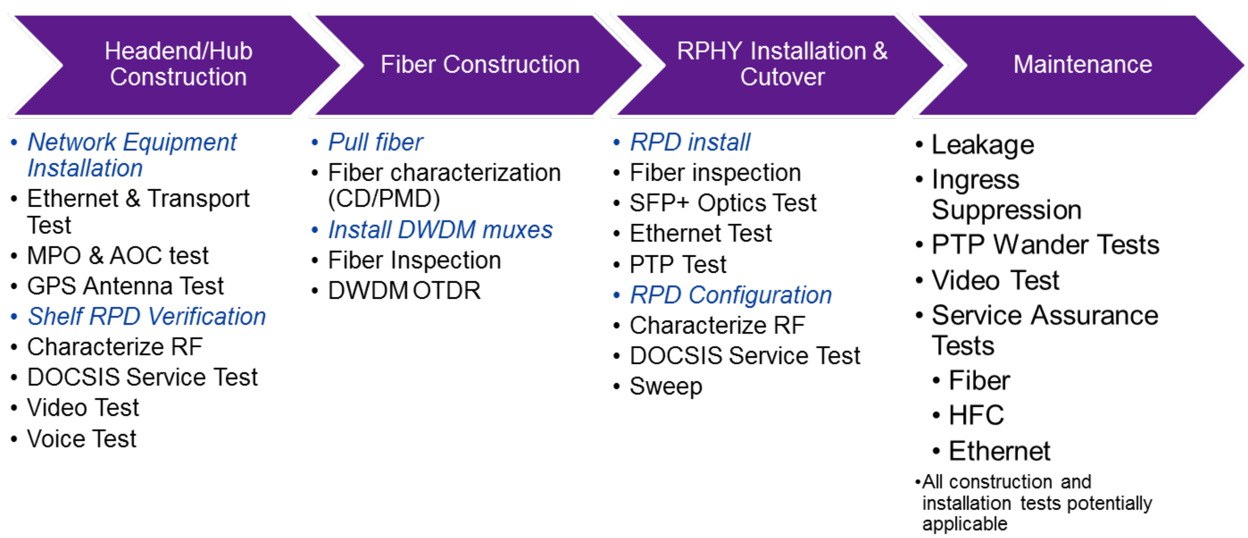

Figure 9 - Framework for DAA Deployment and Maintenance

Again – detailed discussion of each point in this framework is beyond the scope of this document but a brief description is provided below with general test content and rationale.

1. Headend/Hub Construction

The goal of this section is to ensure that the infrastructure needed to feed an R-PHY deployment in the field is functional and correct before any work is done to validate the field aspects. Ethernet and transport gear can be tested including optical cabling interconnections (MPO & AOC test) as well as GPS antennas required for PTP functionality. Shelf RPD verification can be used to pre-test the hub/headend components and services delivered before field R-PHY gear is deployed.

2. Fiber Construction

This section focuses on validating that fiber has been successfully deployed without any damage or excessive bending. Like most aspects of new technologies, the secret to success often lies in disciplined execution of the fundamentals. Techs need to remember “inspect before you connect” – dirty or damaged fiber interfaces are still a leading cause of fiber network performance issues. The DWDM muxes will often be a new ingredient for many HFC plants and will require specialized equipment to properly validate.

3. R-PHY Installation and Cutover

This is the most exciting part of the process, where end-to-end performance can finally be tested for the first time. As part of the RPD install process the SFP and related optics must be tested after ensuring that the fiber is clean and undamaged. Once operational, Ethernet testing is required, as well as verifying adequate throughput and that timing is working properly via PTP testing. As a matter of good practice this is an appropriate time for sweep to characterize the RF performance and to ensure DOCSIS services are operational and working at full performance as well.

4. Maintenance

Most of the plant issues that techs spent their time addressing in legacy plant will still be present in DAA nodes (minus the dreaded analog optical link setup!). If the R-PHY planning process, including system/tool selection has been thoughtfully carried out, the tools and processes that techs use will be very similar, other than the changes noted earlier in this paper. Upstream ingress will still be the #1 issue and having capable tools to address this is still critical. Leakage is still a valuable maintenance capability, even in shorter cascades that often accompany DAA implementations, for localizing potential upstream ingress sources along with many other plant issues. Most RPD vendors have integrated leakage tagging capabilities to enable continued use of deployed field gear for drive-out and walk-out use cases. One new test for techs working on R-PHY nodes will be the PTP wander test as a validation and troubleshooting tool. Service assurance is also still important. Be sure to consider how this will be handled for HFC, fiber and Ethernet aspects of the network. It’s also important to note that all construction and installation tests are also potentially applicable in the maintenance category.

Conclusion

Hopefully after reading this paper you have a better understanding of how DAA is enabling the explosive node split rate required to meet subscribers’ ever-increasing broadband appetite. Redistributing certain network functions from overcrowded/overheated hubs out into the field eliminates many barriers to gigabit service delivery but creates challenges. At the heart of most challenges is the new optical media translation interface that is created by the redistribution. Understanding the technical and organization challenges will determine whether a project succeeds or fails. Specific challenges including the elimination of RF in the hub for test gear has been addressed by building virtualization capabilities into the RPD's themselves and orchestration software to tie it all together. All this theoretical guidance is great, but the Vodafone New Zealand example demonstrates that DAA can indeed resolve critical business problems for operators and continue to be maintained by the same workforce with mostly same field tools and processes. Keys to success include careful planning of the architecture, and the deployment/validation process, and plant maintenance during what is for most a multi-year transition period. Finally, a framework was provided to guide thought regarding test considerations for the entire lifecycle of a typical Remote PHY deployment.

Abbreviations

AOC	active optical cable
CAA	centralized access architecture
CCAP	converged cable access platform
CMTS	cable modem termination system
COTS	commercial off the shelf
CPE	customer Premises Equipment
CWDM	coarse Wavelength Division Multiplexing
DAA	distributed Access Architecture
DTP	DOCSIS timing protocol
DWDM	dense wavelength division multiplexing
HERD	headend rearchitected as datacenter
HFC	hybrid fiber-coax
MPO	multi-fiber push on connector
NDF	narrowband digital Forward
NDR	narrowband digital Return
OTT	over the top (video)
PTP	precision Timing Protocol
R-CCAP	remote ccap
R-PHY	remote phy
RCI	remote phy / ccap interface
RPD	remote phy device
SCTE	Society of Cable Telecommunications Engineers
SFP	small form-factor pluggable
SNR	signal to noise ratio

