SCTE·ISBE CABLE-TEC EXPO 2019 | NEW ORLEANS, LA SEPT. 30-OCT. 3

2019 Fall Technical Forum
SCTE·ISBE • NCTA • CABLELABS

Practical Lessons of a DAA Deployment with a Virtualized CMTS

A Technical Paper prepared for SCTE•ISBE by

Asaf Matatyaou
Vice President, Solutions and Product Management, Cable Access Business
Harmonic, Inc.
4300 North First Street, San Jose, CA 95134
1-408-490-6834
Asaf.Matatyaou@harmonicinc.com

Brian Bendt
Engineering Director, P.E.
Comporium
471 Lakeshore Parkway, Rock Hill, SC 29730
1-803-326-7173
Brian.Bendt@Comporium.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

The promise and potential of virtualizing a cable hub has been discussed over the past few years. There are many opportunities when discussing virtualization. In 2018, Harmonic presented one possible starting point, by virtualizing a CMTS in a centralized deployment model in a headend or hub. The focus of this paper will be specific to a Distributed Access Architecture (DAA) deployment, leveraging the same approach of virtualizing the CMTS.

While the benefits and general considerations of a virtualized CMTS (vCMTS) have been described in past technical papers, this paper will focus on real-world experiences and lessons learned deploying a vCMTS in a DAA. The combination of virtualization and the use of Remote PHY as a protocol between a vCMTS Core in a centralized facility and distributed Remote PHY Devices (RPDs), when migrating from a hardware-based CMTS deployed on a traditional HFC infrastructure, will be described. Important areas of consideration when deploying DAA will include usage of IEEE-1588 over aware and unaware networks, as well as the Converged Interconnect Network (CIN) connecting the vCMTS with the RPDs. Supporting legacy services, such as video and out-of-band, and operational aspects necessary to field-deploy DAA, such as leakage and upstream spectrum, will be described. The paper will also examine the differences between DAA and a traditional HFC deployment, where the RF is generated in the headend or hub in comparison to the Remote PHY node. Lastly, the usage of streaming telemetry will be explored in comparison to legacy monitoring techniques, paying particular attention to important metrics that are gathered by the RPDs.

Architecture Description

The architecture described in this paper is made possible by advancements in standards and technology, specifically the Remote PHY specifications published by CableLabs and the virtualization of the CMTS into commercial off-the-shelf (COTS) x86-based servers. CableLabs issued the first set of Remote PHY specifications in June 2015 and most recently updated the specifications in March 2019 with the twelfth revision of the specifications. [8]

"In a Remote PHY Architecture, the classic integrated CCAP (I-CCAP) is separated into two distinct components. The first component is the CCAP Core and the second component is the Remote PHY Device (RPD)."[1] The RPD consists of the physical layer functionality defined for an I-CCAP, with the remainder of the I-CCAP functionality residing in the CCAP Core. The CCAP Core is logically the combination of a CMTS Core and EQAM Core, and is connected to the RPD via IP transported over digital fiber. In this paper, the CCAP Core is implemented as a vCMTS with a separate legacy EQAM pre-existing in the operator's network and will be referred to as "vCMTS."

Since 2015, SCTE technical papers, such as "Transforming the HFC Access Network with a Software-Based CCAP" [4] and "Real-World Deployment of a Virtual Cable Hub" [3], have defined virtualization and the benefits of a vCMTS. The benefits and general considerations for vCMTS are beyond the scope of this paper.

"Remote PHY" is an implementation of a DAA, but not necessarily restricted to DAA deployments. In fact, the term "remote" doesn't restrict the RPD from physically being co-located with the vCMTS, and the CableLabs specifications describe examples where the RPD and RF may be located in the network, in

[1] Remote PHY Specifications, CM-SP-R-PHY-I10-180509, pg. 10

the headend/hub or in an optical node. The "remote" aspect is that the physical layer or PHY is separate or remote from the CMTS Core.

Figure 1 shows an I-CCAP deployment architecture, as well as Remote PHY-based centralized and DAA deployment architectures, both of which use the Remote PHY standard signaling to communicate between the vCMTS and the RPD (existing in the PHY shelf and Remote PHY Node (RPN)). Remote PHY signaling includes the Downstream External PHY Interface (DEPI), Upstream External PHY Interface (UEPI) and the Generic Control Plane (GCP).

The benefits of Remote PHY and a detailed description of the specifications, including the signaling, are beyond the scope of this paper and are well documented in the industry over the past few years.

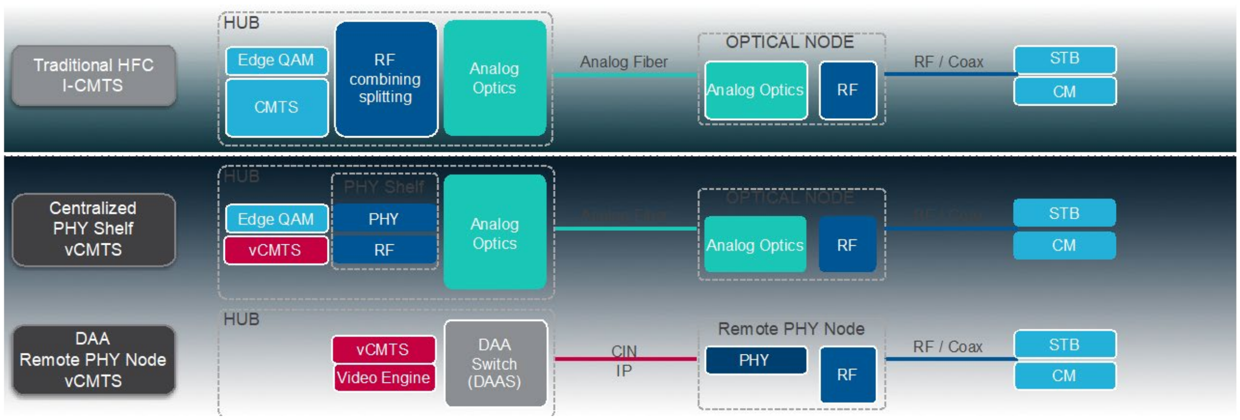

Figure 1 - Traditional HFC, Centralized and DAA Deployment Comparison

This paper focuses on a DAA deployment, whereby the RPDs are deployed in outdoor Remote PHY Nodes. The terms "Distributed" or "DAA" are used to describe this deployment architecture as the RPDs are deployed in a distributed fashion, deeper in the HFC infrastructure. Per the CableLabs DAA overview, this involves "moving functions into the network reduces the amount of hardware the headend (hub) needs to house, thus creating efficiencies in speed, reliability, latency and security in support of 10G."[2] A traditional deployment with I-CMTS is based on chassis deployed in each hub location within the operator's footprint. This is the common architecture today. In Figure 1, this means that the top drawing is replicated over and over again. This was a necessary approach at the dawn of DOCSIS services, but technology has evolved with DAA and vCMTS, enabling much more efficient implementations, while simultaneously being more scalable and capable.

This traditional I-CMTS deployment was the starting point before the transformation to deploy vCMTS and Remote PHY Nodes. Cable operators continue to increase their bandwidth and service offerings to keep up with subscriber usage demands and telco competition. Continuous bandwidth growth has historically created challenges when deploying in a legacy I-CMTS deployment, such as:

- Equipment requirements driven by repeated node splits, including RF splitters and combiners, and optical transmitters and receivers
- Substantial equipment space and power requirements leading to facility/building expansions to accommodate more legacy optical nodes and systems
- Substantial headend cabling required for each of the nodes, impacting space and airflow
- Financial and space requirements needed to distribute equipment to additional locations

[2] https://www.cablelabs.com/technologies/distributed-access-architecture, July 27, 2019.

- Financial requirements to provide fiber path diversity from headend or hub to legacy optical nodes, if required

On the other hand, benefits driving the transformation to a vCMTS and DAA deployment include[3]:

- Reduced space, power, wiring and cooling
- Speed to complete node splits
- More frequent and shorter development cycles
- Sustainable capacity growth, elastic scalability and increased flexibility
- Improved total cost of ownership, including reduced operational (opex) and capital expenditure (capex)

In this deployment example, the vCMTS running on COTS x86-based 1-RU servers replaced legacy I-CMTS chassis in the same locations. The Remote PHY Nodes were deployed in place of legacy analog optical nodes and connected back to the vCMTS Cores over the Converged Interconnect Network or CIN. This approach allows operators to deploy the Remote PHY Nodes in existing locations or while performing node segmentation, leveraging the existing HFC infrastructure (i.e. leveraging existing amplifiers downstream of the node).

Deployment Details

The deployment details describe the ending point after replacing the legacy I-CMTS chassis with the vCMTS Core servers, looking at a specific hub location and taking into account all data, video, out-of-band and proactive network maintenance services required for a production deployment.

Let's describe each device in this deployment type and where the device is located.

1. vCMTS: the CMTS Core functionality is implemented on a set of COTS x86-based servers. Each vCMTS Core replaced legacy I-CMTSes at each location and are connected to the Converged Interconnect Network (CIN).
2. Video Engine: functionally similar to an edge QAM, supporting broadcast video and video-on-demand. However, a Video Engine outputs video services over an IP-based L2TPv3 (i.e. DEPI) tunnel, over the CIN, to each RPD. In this deployment, there is no additional statmux, rate shaping or video processing performed by the Video Engine.
3. NDF/R Engine: narrowband digital fiber and return "digitizes a small portion of the spectrum, and sends the digital samples as payload within packets that traverse between the [vCMTS] and the RPD. This approach works with any type of OOB signal as long as the signal can be contained within the defined pass bands."[4]
4. QPSK Modulator/Demodulator: legacy SCTE 55-2 equipment that is used unchanged from legacy deployments. In this deployment, the existing SCTE 55-2 equipment is used in combination with the NDF/R Engine, which digitizes the RF to IP packets traversing over the CIN.
5. Core Routing Engines (CRE): standard COTS networking equipment are the switch fabric connecting the vCMTS Core servers to the rest of the CIN. The CREs are co-located in the same location as the vCMTS Core servers and are part of the Layer 2 CIN. They also perform the function of aggregation, sometimes referred to as the Distributed Access Architecture Switch (DAAS).

---

[3] Real-World Deployment of a Virtual Cable Hub, pg. 5
[4] Remote Out-of-Band Specification, CM-SP-R-OOB-I09-180509, pg. 14

6. Edge routers: large-scale edge routers connecting the access network with the core backbone network, which are typically previously established.
7. Remote PHY Nodes and RPDs: each RPN outdoor enclosure can house one or two RPDs. Each RPD connects to the vCMTS Core servers over the CIN, and outputs RF over the existing HFC infrastructure. A test RPN (without production subscribers) is deployed at each hub.
8. IEEE-1588 PTP Grandmasters: "Remote DTI provides timing synchronization between CCAP Cores and RPDs based on the IEEE 1588v2 standard. The protocol supports the basic synchronization between the CCAP Core and Remote PHY Device for DOCSIS/video/OOB services."[5]

Figure 2 shows the devices and connectivity between the vCMTS Core, Video Engine and NDF/R Engine over the CIN, which is the network between the CCAP Core and the RPD. In this deployment type, the CIN traverses between the headend, each hub and the RPDs over a Layer 2 network.

One of the benefits in this deployment type is that the CIN infrastructure is transported over digital fiber which has the benefit of longer distances, more wavelengths and increased performance compared with legacy HFC transport. While the RPDs deployed in this deployment leverage the existing HFC infrastructure south of the Remote PHY Node (i.e. existing external amplifiers), the same CIN can be used to deploy fiber deep Remote PHY Nodes (i.e. high-power amplifiers are built into the node and no additional external amplifiers are needed downstream).

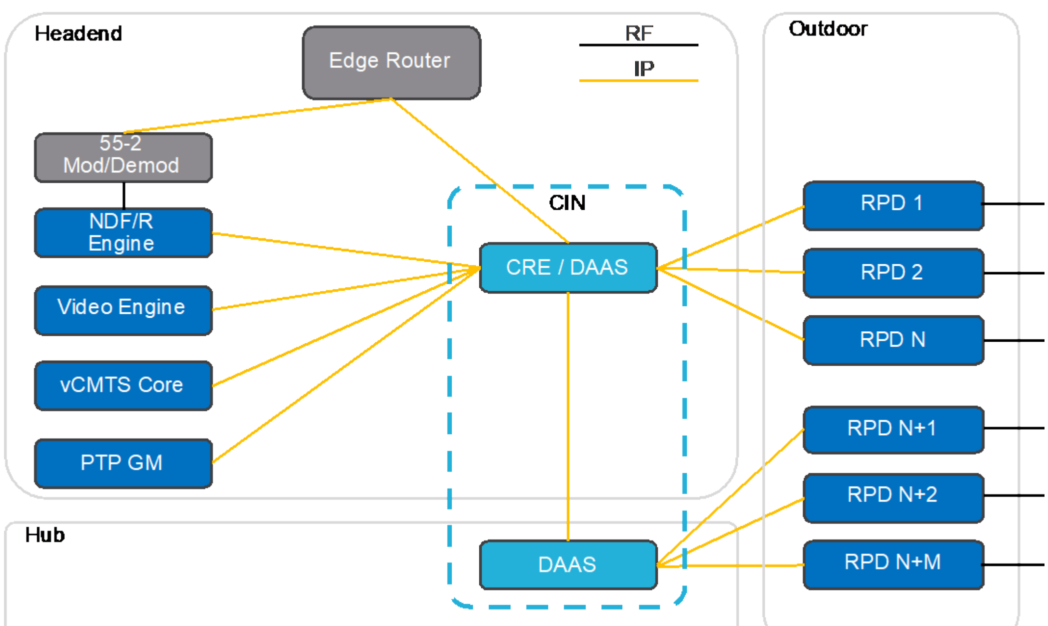

Figure 2 - vCMTS Deployed in a Distributed Access Architecture

[5] Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.5.2, pg. 30

Real-World Considerations

While there are many benefits and opportunities with this new Remote PHY-enabled vCMTS in a DAA deployment, there are real-world considerations that should be considered. In particular, the following topics are described: services, compute location, networking, timing, network maintenance and operations.

1. Services

Cable operators deliver multiple services to subscribers, notably broadband, video and voice services. Delivery of the different services to a subscriber's cable modem, set-top box and other IP-based devices is generated and processed by different headend equipment, transported over the HFC network as RF signals and use different technologies and protocols.

In a traditional deployment, each service is processed and generated at the headend with RF output over a non-conflicting frequency span (with other services). For example:

- Broadband and voice: an I-CMTS transmits and receives RF for DOCSIS-based data for broadband and voice services.
- Broadcast video and video-on-demand: an edge QAM transmits RF for broadcast video and video-on-demand.
- Set-top box control: Out-of-band systems (e.g. SCTE 55-1 and SCTE 55-2) transmit and receive RF for controlling set-top boxes.
- Network maintenance and monitoring: Test equipment transmit and receive RF for network maintenance and monitoring purposes.

The RF for each of the services is combined at the headend, or at each hub, to have all services available when received at each subscriber's devices. After the RF is combined, it is then split to different service groups and analog optical nodes. Additionally, transmitters in the headend send the RF signals to the analog optical nodes as well as amplified over the traditional HFC infrastructure. In the reverse direction, receivers in the headend acquire the upstream RF signals.

This is a simplified description of how, for decades, subscribers receive all services over a single coaxial cable in the home, as well as what equipment is used by the cable operator to deliver these services. This description is important to understand, specifically for the sake of contrasting service delivery and transport over a traditional HFC infrastructure and DAA.

A key difference between a traditional HFC deployment and DAA:

- A traditional HFC deployment transmits and receives all services from the headend over RF, as shown in Figure 3
- A DAA deployment transmits and receives all services from the headend via IP over Ethernet

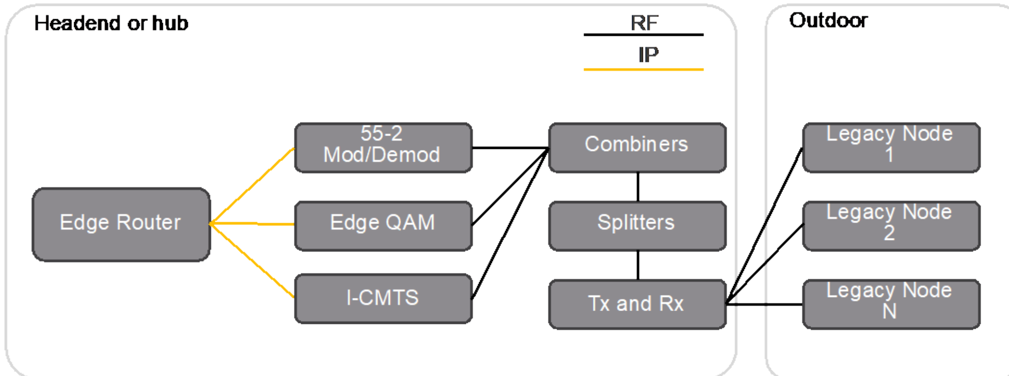

Figure 3 - A Traditional HFC Deployment Transmits and Receives RF to Analog Optical Nodes

With DAA, and unlike a traditional deployment, all the services and signals are transported over an Ethernet-based IP network, the CIN. Transporting over the CIN means that what was once transmitted over RF from the headend or hub, is now transported over IP, eliminating the splitters, combiners, transmitters and receivers at the headend or hub. The facility savings (i.e., equipment, space, power and cooling savings) is substantial, as well as the savings due to reduction in ongoing cabling, maintenance and operational expenses associated with continuous bandwidth growth and annual expansion.

Let's briefly revisit each service described in a traditional deployment and the changes required to deliver each with DAA, over the CIN:

- Broadband and voice: A CMTS Core (in this case, a vCMTS) transmits and receives IP for DOCSIS and Remote PHY-based data for broadband and voice services.
- Broadcast video and video-on-demand: A Video Engine transmits Remote PHY-based IP for broadcast video and video-on-demand, and is independent of the vCMTS Core and Out-of-band engines.
- Set-top box control: Out-of-band systems (e.g. SCTE 55-1 and SCTE 55-2) transmit and receive Remote PHY-based IP for controlling set-top boxes. Alternatively, out-of-band systems can transmit and receive RF, similar to traditional deployments, with an NDF/R Engine performing the conversion from RF to Remote PHY-based IP and vice versa.
- Network maintenance and monitoring: Test equipment transmit and receive Remote-PHY IP for network maintenance and monitoring purposes. Additionally, some RF signals such as pilots and alignment carriers can be generated at the Remote PHY Device.

Remote PHY-based IP refers to the different Remote PHY protocols specified by CableLabs, specifically, Downstream External PHY Interface (DEPI), Upstream External PHY Interface (UEPI) and the Generic Control Plane (GCP).

2. Compute Location Options

One of the key opportunities when transitioning to a DAA with a vCMTS Core is where the CMTS Core is located, enabled by two technologies: Remote PHY and virtualization. The Remote PHY specifications enable separating the CMTS Core from the RF and provide the operator an opportunity to decide where the CMTS Core should be located. I-CMTS equipment perform DOCSIS processing, transmit and receive RF in the same chassis. There is no choice in location selection as I-CMTSs are deployed at each hub location connected to the access network, regardless of the size and scale required to support the nearby footprint of subscribers. The Remote PHY specifications enabled separating the CMTS Core from the RF and provide the operator an opportunity to decide where the CMTS Core should be located.

The second technology enabler for flexibility in CMTS Core location is virtualization. A vCMTS Core solution provides the operator an opportunity to determine where the vCMTS Cores should be located, with options such as installing the vCMTS Cores at each hub location with the RF (similar to I-CCAP) or consolidating the vCMTS Cores at a few hubs or even a single centralized hub location. Having at least two vCMTS Core server locations may also provide geographical redundancy.

In this real-world deployment example, the vCMTS Core servers were deployed in the same locations as the legacy I-CMTS chassis. This approach allowed the operator to reduce space, power and cooling requirements at existing facilities while expanding their capability to deploy faster broadband services. Additionally, it was the simplest transition strategy from I-CMTS to vCMTS and DAA, as it required the least amount of change. The same replacement approach, installing Remote PHY Nodes at the same points in the network and HFC infrastructure as legacy analog optical nodes was used when transitioning from analog optical nodes to Remote PHY Nodes.

Expandability with more Remote PHY Nodes can be accomplished by adding each node to an existing vCMTS Core server with spare capacity, via the DAAS or by adding extra vCMTS Core servers for more capacity.

"More flexibility in deployment location and scalability requires increased attention on the scalability of each element in the end-to-end network. While expandability is easier and provides quicker time-to-market to add capacity, operators must pay attention to the scale limits of the compute and networking resources separately, as each type of resource type may require additional devices when those limits are reached. On the other hand, when scale limits are reached for a particular resource, those can be expanded in a focused method and the operator won't have to scale everything at once. For example, when I-CCAP chassis reaches any of its scale limitations, another I-CCAP needs to be installed. In comparison, when a DAAS reaches a scale limitation, such as port count, another DAAS switch can be added without adding additional vCMTS Core servers. Nevertheless, an operator must pay attention to each element's specifications and plan a network for existing and additional subscribers."[6]

3. Converged Interconnect Network

It's worth repeating, that a key difference between a traditional HFC deployment and DAA:

- A traditional HFC deployment transmits and receives all services from the headend over RF
- A DAA deployment transmits and receives all services from the headend using IP over Ethernet

"The network between the CCAP Core and the RPD is known as the Converged Interconnect Network (CIN). The CIN encompasses either or both the hub access network and the optical access network. The CIN can contain both Layer 2 switches and Layer 3 routers."[7]

While the benefits of an Ethernet-based IP network (digital fiber) have been noted, there are a few CIN related decisions that need to be made:

1. Location and number of CIN switches between the vCMTS Core and the RPD
2. Layer 2 or Layer 3 configuration

[6] Practical Lessons of a Centralized Virtualized CMTS, pg. 8.
[7] Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.2.5, pg. 26

3. Traffic prioritization
4. Redundancy: Network redundancy is another deployment decision with many options available, including, link redundancy, chassis redundancy, and line card redundancy within a chassis. This topic is beyond the scope of this paper.

3.1. CIN Switch Location Options

The simplest way to describe the CIN is as a logical network switch that connects the vCMTS Core, Video Engine and OOB Engine with the RPDs. In fact, the CIN can be deployed as a single switch co-located with a few ports connected to the vCMTS Core, Video Engine and OOB Engine, and many ports connected to RPDs. In the case of a vCMTS deployment, there are three types of traffic sources/destinations that are supported via the CIN:

- The vCMTS Core, Video Engine and OOB Engine
- The RPDs
- The Edge Router – the same one connected to legacy I-CMTSes

The CIN serves as the fabric that connects the northbound content sources, the devices performing the data and video processing (i.e. vCMTS Core, Video Engine) and the southbound RPDs. Consequently, the CIN can also be deployed as a collection of switches, creating this fabric.

The decision of whether to directly connect the Edge Router, vCMTS Core and RPDs to the same CIN switches is determined on a few factors:

1. An existing network infrastructure: it's possible to leverage an existing network to transport the CIN traffic.
2. Facility location of Edge Router, vCMTS Core, Video Engine and OOB Engine: there is typically a switch co-located with the vCMTS Core, Video Engine and OOB Engine.
3. Distances between devices: switches can be used to extend the distances between the vCMTS Core and the RPDs, if necessary, as well as provide fiber route redundancy.
4. Number of devices needing to connect to each switch: the main factor of switch ports is the number of RPDs, factoring in the initial set of installed RPDs and potential future RPD growth.

In this deployment, there are two locations for the CIN switches: centrally located at the headend and at remote hubs. The CIN switch located at the headend is connected to the Edge Router, Video Engine and OOB Engine. The CIN switch located at the remote hub is connected to the headend CIN switch, the vCMTS Core and the RPDs. In this case, the headend CIN switch partially performs the function of the CRE (the networking fabric between the Edge Router and the multiple vCMTS Core servers), while the remote hub switch partially performs the function of the CRE and is the DAAS (physical connectivity to the RPDs).

In this deployment, layer 2 networking was selected, as it was the quickest and simplest to deploy with CIN switches deployed at only two locations: at the headend and hub.

3.2. Traffic Prioritization and Capacity Management

Remote PHY and DOCSIS demand that latency sensitive traffic receives the correct priority and handling in the CIN. Specifically, IEEE 1588 PTP packets used for timing synchronization between the vCMTS Core and the RPDs need to meet latency and jitter requirements based on how PTP is configured and deployed between the Timing Server Grandmaster and each PTP slave device (e.g,. vCMTS Core, Video

Engine, RPD). DOCSIS MAC management messages such as the MAP also have real-time timing information and require proper traffic priority.

In an isolated CIN, where there is no other unrelated network traffic handled by the CIN switches, the design considerations are simpler, as congestion can be managed with good network design and capacity management practice, especially when planning expanding CIN network resources, before the congestion actually occurs in deployment. However, if the CIN can't be guaranteed to be congestion free, capacity management isn't sufficient. Congestion management is required.

If congestion does occur, congestion management is performed by traffic prioritization. Traffic prioritization effectively allows higher priority control and user packets to survive the congestion. There are multiple ways DOCSIS and IP-based traffic prioritization can be used to maintain the Quality of Service (QoS) when manageable congestion occurs in the CIN. However, if congestion reaches the point that it is impactful to critical control packets, such as DOCSIS MAP MAC Management Messages (MMMs), no amount of traffic prioritization will resolve the impact on stable network operations.

The Remote PHY specification also guides: "to prevent the MAP from being slowed down by other traffic in the CIN, the DOCSIS traffic (or a subset containing the MAP messages) may be sent in an independent L2TPv3 flow that can have a unique DSCP. The value of the marked DSCP value should be consistent with a configured "per hop behavior (PHB)" that will provide MAP messages with the highest priority and lowest latency across the CIN to the RPD."[8]

4. Timing[9]

With all Remote PHY deployments, timing specifications such as R-DTI need to be adhered to, regardless of the location of the RPD. "The MHAv2 version of DTI (i.e., R-DTI) defines how to distribute phase and frequency information from the CCAP Core device to remote PHY devices within the HFC network."

"For Ethernet based networks, IEEE 1588 allows both phase and frequency information to be transferred between nodes across an existing packet network with switches or routers, thus making it ideal for R-DTI.

In order to reduce any phase offset introduced by latencies through the network, IEEE 1588 defines a protocol for calculating the latency across sections of the network, and then compensating for those latencies. The latency calculations assume that the link is symmetric, and therefore the protocol works well for traditional full duplex Ethernet networks. IEEE 1588 also defines a protocol for determining the latency through any intervening switches or routers within the network, but the device is to be IEEE 1588 capable [referred to as PTP aware]. If the devices are not IEEE 1588 capable, the phase offsets and convergence times within the network will be greater [referred to as PTP unaware]."[10]

Devices in the network which are not IEEE 1588 capable are considered as non-participating. For the IEEE 1588 capable devices, they can operate as either a transparent clock or boundary clock. A transparent clock device modifies the PTP message, accounting for the processing time of the PTP message within the device, while a boundary clock receives the PTP message as a slave device and re-generates the PTP message as a master.

This real-world deployment experience is across a PTP unaware network, as not all existing network devices were IEEE 1588 capable. While the convergence times within the network are greater, as

[8] Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.6, pg. 30
[9] Practical Lessons of a Centralized Virtualized CMTS, section 3.1, pg. 9.
[10] Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, pg. 6, 16

expected, they have not been operationally significant to justify immediate replacement of all network devices to be IEEE 1588 capable. However, it is critical to evaluate and consider the jitter and latency conditions of the CIN regardless of selecting an PTP aware or unaware mode. Both PTP deployment modes demand meeting jitter and latency requirements, which can be impacted by the CIN network device capabilities, and the number of network hops and congestion conditions.

As timing is critical for Remote PHY operation, there are a couple of other options to consider when deploying IEEE 1588 PTP grandmaster(s) in the network, which transmit the synchronization information to the other clocks in the same network. The first option is the reliability of the PTP grandmaster, as there exist a range of products which are small (SFP form factor) and less reliable as compared to full carrier-grade products which have redundant input/output clock (IOC) cards. The second option to consider is whether to use the best master clock algorithm (BMCA), which determines the highest quality or "best" clock within the network, in case the grandmaster clock quality is compromised or fails.

In this deployment, two carrier-grade grandmasters (each independently deployed at different headends) with high availability features were deployed. High availability features include redundant inputs, outputs, clock and power supplies.

5. Access Network Maintenance and Monitoring

Access network maintenance and monitoring is vital to effectively install, maintain and monitor the end-to-end system. Performing this has historically been performed by injecting or inspecting RF signals with test or monitoring equipment in the headend or in the field. With DAA, RF inspection is no longer available at the headend, as there is simply no RF to inspect with all services running over an IP-based CIN.

Typical proactive network maintenance (PNM) includes upstream spectrum analysis, return sweep, ingress detection and common path distortion. Typical test signals that need to be injected include leakage detection, AGC pilots and alignment carriers. With DAA, these test signals can be generated at the RPD. However, other test signals can also be converted from RF to IP and vice versa with an NDF/R Engine so that they can be transported over the CIN.

In a DAA deployment, PNM examples of the upstream spectrum analysis and return sweep include the following devices:

- vCMTS Core: configures the NDF/R functionality for FFT data to be transported from the RPD to the PNM server.
- RPD: performs FFT processing and transmits FFT data over the CIN to the PNM server.
- CIN: the IP-based network connecting the different devices
- PNM Server: processes the FFT data to display upstream spectrum and sweep results
- Handheld device: telemetry to/from via NDF/R

Similar types of processing can be used to perform other types of PNM, such as CPD and ingress detection. Note that, as an outside plant device supported by maintenance technicians with deep backgrounds focused on RF technology, there is an operational transition period. Some fundamental expectations of prior generation devices should be anticipated, such as RF test points and modular design, in order to ease this transition.

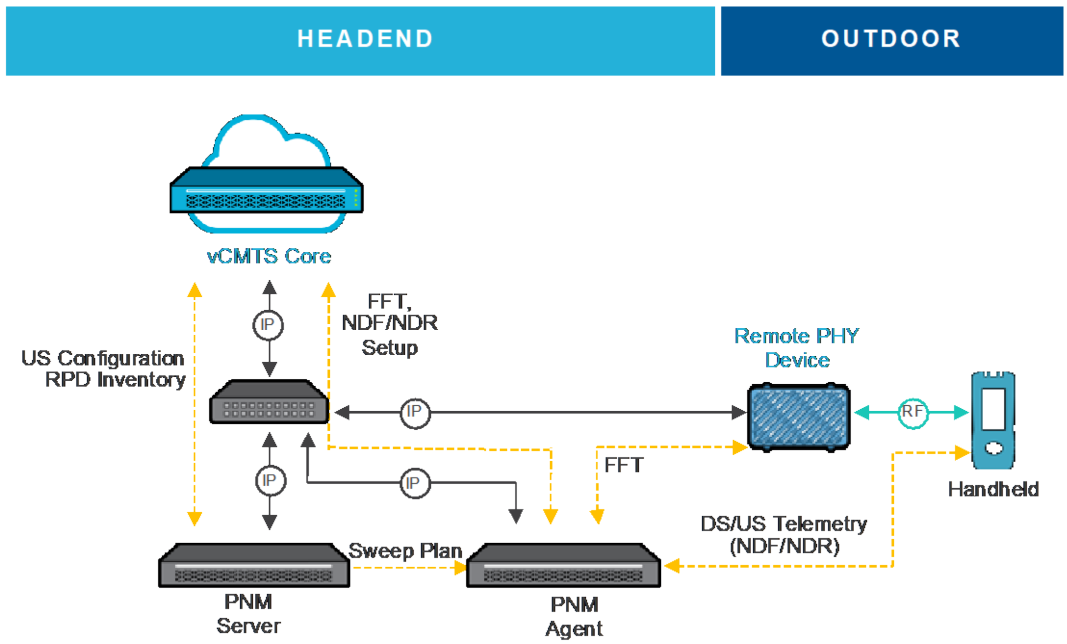

Figure 4 - Upstream Spectrum Analysis and Sweep with DAA

6. Operations

New technologies such as Remote PHY and virtualization bring many benefits, but they must be supported by real-world operational practices that enable deployment. In the experience of deploying a vCMTS in a DAA, there is an immediate opportunity and need to start by consolidating monitoring of the vCMTS Cores, the CIN and RPDs in a single consolidated operational view. With the CIN as the fabric connecting all the devices the access network, the data can be gathered and visualized with other modern technologies.

Streaming telemetry data coupled with Grafana, as an open visualization tool for analytics and monitoring, is used to give operational organizations a continuous stream of health and performance for the system. The vCMTS streams its data, as well as the RPD data it gathers via the Remote PHY GCP. Other vital system information is gathered and displayed via customizable dashboards. With devices distributed in the network, having a single tool that visualizes a continuous stream of data across topic-based dashboards allows DAA to be operationalized in scale. This is a departure from the tried-and-true command-line-interface CLI-based “show” approach, which is limited to the text being displayed on a screen, a human being inspecting the output at a single instance in time. CLI still has its place for troubleshooting very specific issues but is also isolated to connecting to individual devices and piecing together the story to get a complete understanding of the end-to-end system health and performance.

With a software-based virtualized CMTS Core, there are many points of inspection in the software that can be made visible based on field experience and extended over time, based on new findings. Continuous monitoring improvements will also be augmented over time with configuration, deployment and automation capabilities. In addition, future operations, services, or CMTS upgrades – channel re-mapping, adding DOCSIS carriers, PMA optimization, Full Duplex DOCSIS (FDX) – all can take advantage of the significantly enhanced platform agility to increase the velocity of introduction of such changes, features, and services.

Conclusion

7. Lessons Learned

Looking back at the field experiences from deploying vCMTS in a DAA, there are some key lessons learned. vCMTS has been deployed in different environments over the past few years, and consequently most of the lessons learned recently are related to scaling DAA deployments and specific operational aspects related to moving from a lab to the field. In particular:

1. New technologies require substantial lab testing.
2. DAA technologies may cross typical operational group responsibility boundaries, including installation and support. Operational responsibilities require planning and discussion amongst different organization before deployment.
3. Out-of-band designs, such as SCTE 55-2, need to be planned. Specifically, there are options for centralizing or distributing the out-of-band systems across different headend and hub sites.
4. Proactive network monitoring integration should be planned to gain operational familiarly with new DAA methods of performing sweep, upstream spectrum analysis and reverse ingress monitoring.
5. Disaggregation of the end-to-end solution components, as well as the separation of software from customized hardware results in separate hardware and software roadmaps and multi-dimensional scalability. With disaggregated swim lanes, the end-to-end solution may be multi-partner and multi-dimensional, as well as some swim lanes managed by the operator.
6. Tools, tools, tools: visibility into all aspects of software is a priority at the outset and is essential to effectively debug and develop utilizing Agile-based iteration management.

8. Summary

Cable operators have an existing footprint that they continue to grow and improve upon but it can't be overhauled overnight. Technologies such as Remote PHY and virtualization continue to extend the tool set that cable operators can use to stay competitive with capacity growth demands and challenging market environments. This paper has described a transition from I-CMTS in a traditional HFC deployment to a vCMTS in a Remote PHY-based DAA deployment.

DAA addresses historical challenges by eliminating some headend equipment altogether (e.g. splitters, combiners, transmitters, receivers) and by moving RF processing from the headend to the field, further reducing power, space, cooling and cabling requirements at the headend.

Additionally, when vCMTS is coupled with DAA, cable operators immediately benefit from additional savings of space, power, cooling and cabling, while unlocking the path to sustainably growing capacity, adapting quickly to customer demands, and a solution that is flexible and elastic enough to dynamically augment and shift resources to the most in-demand applications.

In summary, vCMTS deployed in a DAA is operationally and financially compelling, addressing today's access network demands as well as looking into the future, with no end in sight to continuous broadband and service demands.

Abbreviations

CapEx	Capital Expenditure
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CLI	Command Line Interface
CMTS	Cable Modem Termination System
COTS	Commercial Off-The-Shelf
CPE	Customer Premise Equipment
CPU	Central Processing Unit
DAA	Distributed Access Architecture
DEPI	Downstream External-PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
Gbps	Gigabits Per Second
GCP	Generic Control Plane
HFC	Hybrid Fiber-Coaxial
HW	Hardware
I/O	Input/output
MAC	Media Access Control
NFV	Network Function Virtualization
NIC	Network Interface Controller
NOC	Network Operations Center
OOB	Out-of-band
OpEx	Operating Expenditure
OS	Operating System
PHY	Physical
PNM	Proactive Network Maintenance
RF	Radio Frequency
RPD	Remote PHY Device
RPN	Remote PHY Node
RU	Rack Unit
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networking
SW	Software
TCO	Total Cost of Ownership
TTM	Time to Market
UEPI	Upstream External-PHY Interface
vCMTS	Virtual CMTS
vCPE	Virtual CPE
VOD	Video on Demand

Bibliography & References

[1] DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-118-, April 22, 2019, Cable Television Laboratories, Inc.

[2] Matatyaou, Asaf. Practical Lessons of a Centralized Virtualized CMTS. San Jose: Harmonic, 2018. Web.

[3] Matatyaou, Asaf. Real-World Deployment of a Virtual Cable Hub. Web. San Jose: Harmonic, 2017. Web.

[4] Matatyaou, Asaf. Transforming the HFC Access Network with a Software-Based CCAP. San Jose: Harmonic, 2015. Web.

[5] Modular Headend Architecture v2 Technical Report, CM-TR-MHAv2-V01-150615, June 15, 2015, Cable Television Laboratories, Inc.

[6] Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, May 9, 2018, Cable Television Laboratories, Inc.

[7] Remote Out-of-Band Specification, CM-SP-R-OOB-I09-180509, May 9, 2018, Cable Television Laboratories, Inc.

[8] Remote PHY Specification, CM-SP-R-PHY-I12-190307, March 7, 2019, Cable Television Laboratories, Inc.

