Lessons from Operating Tens of Thousands of Remote PHY Devices

A Technical Paper prepared for SCTE by

Jorge Salinger
VP, Access Architecture
Comcast Cable Communications
1701 JFK Blvd – Philadelphia, PA 19103
+1 (215) 439-1721
jorge_salinger@cable.comcast.com

Steve Sigman
Director, Access Architecture
Comcast Cable Communications
1701 JFK Blvd – Philadelphia, PA 19103
+1 (609) 685-3480
steve_sigman@cable.comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

Cable operators have been actively converging video and data services into a common converged cable access platform (CCAP) for many years now. This trend, which requires an evolution towards newer, more modern, and denser equipment, was intended to free up space in the headend.

However, as the success of high-speed data and on-demand services continues, the evolution of the access network progresses towards expanded capacity and/or ever-smaller service groups. For the former, the spectrum allocated to narrowcast services increases, driving operators to deploy Data-Over-Cable Service Interface Specifications (DOCSIS©) services including more SC-QAM (single carrier quadrature amplitude modulation) channels, more capacity for more or wider OFDM (orthogonal frequency division multiplexing) channels, as well as more narrowcast video services. For the latter, more CCAP ports are needed, which drives the deployment of more line cards and eventually more chassis. These expansion trends result in a continuous growth of headend equipment, which is already starting to exceed the capacity that headend facilities can support.

The above trends are now intractably linked to additional evolutions: distribution of components of the access network, implementing a distributed access architecture (DAA), and virtualization of the core network functions. By now, deployment of DAA devices, such as remote PHY (R-PHY) nodes, is accelerating to tens of thousands of remote PHY devices (RPDs), hosting millions of cable modems.

As a result of the experience, an agile approach to network and device management is now required to ensure that we can effectively monitor our networks, these distributed devices, and customer service status. Traditional proactive network maintenance (PNM) and operations support system (OSS) toolsets provide views into current system health, but unexpected behavior on the RPDs and the systems that work with them provide clear opportunities for new approaches to network and system monitoring and management.

This paper explores lessons learned the hard way, gaps discovered and filled, as well as processes developed to improve issue detection and lower mean time to repair (MTTR). The paper focuses upon factors to keep both maintenance technicians and customers happy, including:

- Whole system monitoring: Real-time telemetry, visual dashboarding tools, and error condition detection and alarming for not just RPDs, but the entire architecture, including R-PHY cores, engines, switches, and timing servers;
- Outside plant considerations, including plant powering;
- Provisioning management, including addressing challenges with having all required systems configured at time of node cutover;
- Hardware and software management, including how lab testing can lead to a better customer experience.

This paper begins by outlining the evolution of service provider networks, and then describes why and how the migration to a distributed access architecture is necessary and beneficial. The paper then expands into features that can be implemented with DAAs and discusses the topic of virtualization. Finally, the paper explores how the implementation of DOCSIS 4.0 could be implemented in DAA networks.

2. Cable Network Evolution

2.1. Typical HFC Networks Today

Most hybrid fiber/coax (HFC) networks have been designed with an upper frequency boundary of 750 MHz or 860 MHz, while some are designed to support 1 GHz and other newer networks are designed to support 1.2 GHz. For the more abundant 750 MHz or 860 MHz networks, if not already fully utilized, it is expected that use of their capacity will be increased to the point of exhaustion. This will happen as a result of 1) increased DOCSIS usage for even faster high-speed data (HSD) service tiers; 2) additional high-definition (HD) programs (for broadcast [BC] and especially narrowcast [NC] services, such as video on demand [VOD] and switched digital video [SDV]); and 3) new service additions such as Internet Protocol (IP) video and cloud-based digital video recorder (cDVR.)

In recent years the growth in, and demand for, HD programming has resulted in the need for allocation of large numbers of CTA (Consumer Technology Association) channels[1] for HD services, both for BC and NC, which has filled every available portion of the spectrum. This is especially true for BC, where large numbers of programs are offered in HD format, while simultaneously the need for distributing the standard definition (SD) version has persisted. This has resulted in the need for use of 3x to 5x the number of CTA channels than previously required. For example, a typical digital multiplex including 10 to 15 programs would require an additional three to five CTA channels for the HD-equivalent streams, even assuming the newer, more sophisticated multiplexing schemes available in the market. Of course not every program is available, or still sought by subscribers, in HD format. But very large numbers of them are, including 100 to 150 BC programs.

The above is also applicable to a great extent in systems utilizing SDV technology for content distribution. The difference is that the HD and SD versions of the program are not distributed unless a subscriber requests them, which reduces the marginal increase in capacity. Assuming that all programs are distributed in only one format, which is certainly a valid expectation for programs of low viewership, then the increase in capacity for a conversion from SD to HD would just be the increase in capacity required for the transmission of the HD program without requiring the simultaneous use of bandwidth for both formats.

Additionally, considerable spectrum is needed to deploy high-capacity narrowcast legacy video services, especially cDVR, and thousands of HD video-on-demand titles. For the former, initial observations suggest that network requirements for cDVR may be as high as 4x to 5x that of VOD, and that peak utilization overlaps, at least partially, with that of peak use for other narrowcast services, such as HSD.

Finally, the growth in HSD services continues. All network operators have offered increased service tiers and observed an increased use of broadband capacity for well over a decade now, as shown in Figure 1, which amounts to a constant year-over-year compounded growth. The applications have changed throughout this time, and the demand has continued to increase at the same relentless rate.

[1] 6 MHz-wide channel allocations used in North American and some other cable networks are defined in *CTA Standard "Cable Television Channel Identification Plan" CTA-542-D R-2018 (Formerly CTA-542-D), updated February 2019.*

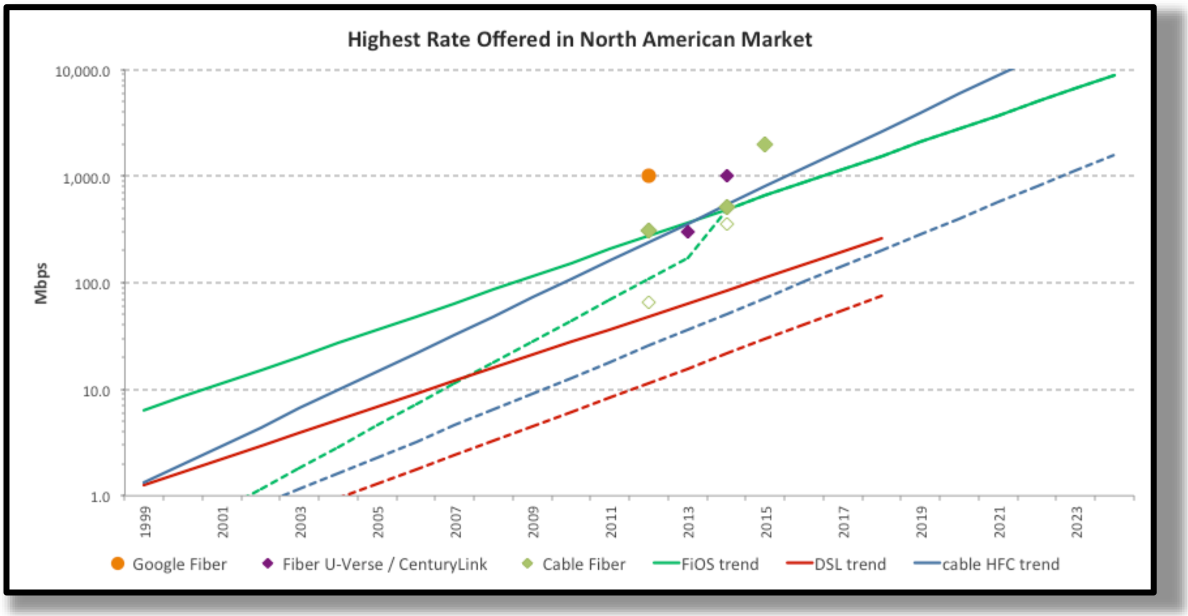

Figure 1 - Examples of HSD service tier capacity increase over time

How does this compare to other operators' data services and a longer period? Projecting an operator's HSD service growth back in time to when Internet services started 25+ years ago, services should have been about 100 bps. This coincides with the history of telephone modems from 110 and 300 baud modems from the mid-1980s, to 56 kbps/V.42, into ISDN (integrated services digital network) services.

This demonstrates that the growth seen in cable industry HSD services is typically over a much longer period of time, rather than an exception observed by operators in recent years.

2.2. Growth Projections

From all of the above, it follows that, should the usage growth pattern continue at the same rate as in the past, networks will be required to provide >1 Gbps HSD services within the next few years. This growth, coupled with the surge in HD video formats (8K TV sets are increasingly available, even as compatible content catches up), and more personalized narrowcast services, will result in a significant growth in narrowcast capacity, as shown in Figure 2.

To support this growth, cable operators have deployed bandwidth reclamation tools such as SDV for digital broadcast, digital transport adapters (DTAs) for analog service reclamation and the shift to all-IP, increased spectrum (1 GHz and above), or a combination of all. These tools have been extremely valuable to operators, and their operational complexity and cost remains well justified.

In the case of SDV, early predictions from industry analysts projected that the efficiency of SDV would reach 40% (e.g., programs requiring 10 CTA channels could be carried in 6). This has proven to be understated, since it was based on the use of SDV to reduce bandwidth required for existing services. As SDV's role in the network grew, the efficiencies have been even greater, especially as SDV expanded in scope to support niche service introductions with low initial viewership that would otherwise be difficult to deploy.

The benefit of DTAs has been just as, or perhaps even more striking. Cable operators deploying DTA devices are able to eliminate the need to distribute the analog channels in the network. Even if DTAs are distributed to top tier analog customers, such as only the traditional "expanded basic" subscribers, the move would reduce a channel line up from perhaps 50 6 MHz channels, or 300 MHz of capacity, dedicated to 50 analog programs to perhaps as little as four such 6 MHz channels (24 MHz) dedicated to transport the 50 programs in their digital-equivalent transport. Using the same comparison method as the above SDV case, this is a >90% efficiency. If extended to the entire analog tier, the efficiency gains are very significant.

Despite the availability of these tools, they are not universally applicable. With respect to SDV, in general it is not likely that all broadcast programs will be switched since experience shows that many broadcast programs are constantly viewed by someone in the service group during peak hours, which will leave a large portion of the spectrum still used for broadcast. Similarly, not all analog channels can be removed in the short term due to operational and/or cost constraints.

Additionally, while many cable operators will use some or all of the capacity tools, in general they won't be used by every operator for all applications because of a range of variables that are out of scope for this discussion.

Finally, there are also significant potential gains to be achieved from the use of advanced video compression standards (VCS), and variable bit rate (VBR) multiplexing. In the case of VCS, coding efficiencies of approximately 50%, depending on implementation and content type, can be obtained with H.264/MPEG-4 Part 10. Furthermore, with the release of the H.265 standard in April of 2013, it is possible to achieve a 50% improvement over H.264. The use of VBR promises to result in a capacity efficiency gain of as much as 70% versus constant bit rate (CBR) . The combined gains from using the above approaches for multiple services are even more significant.

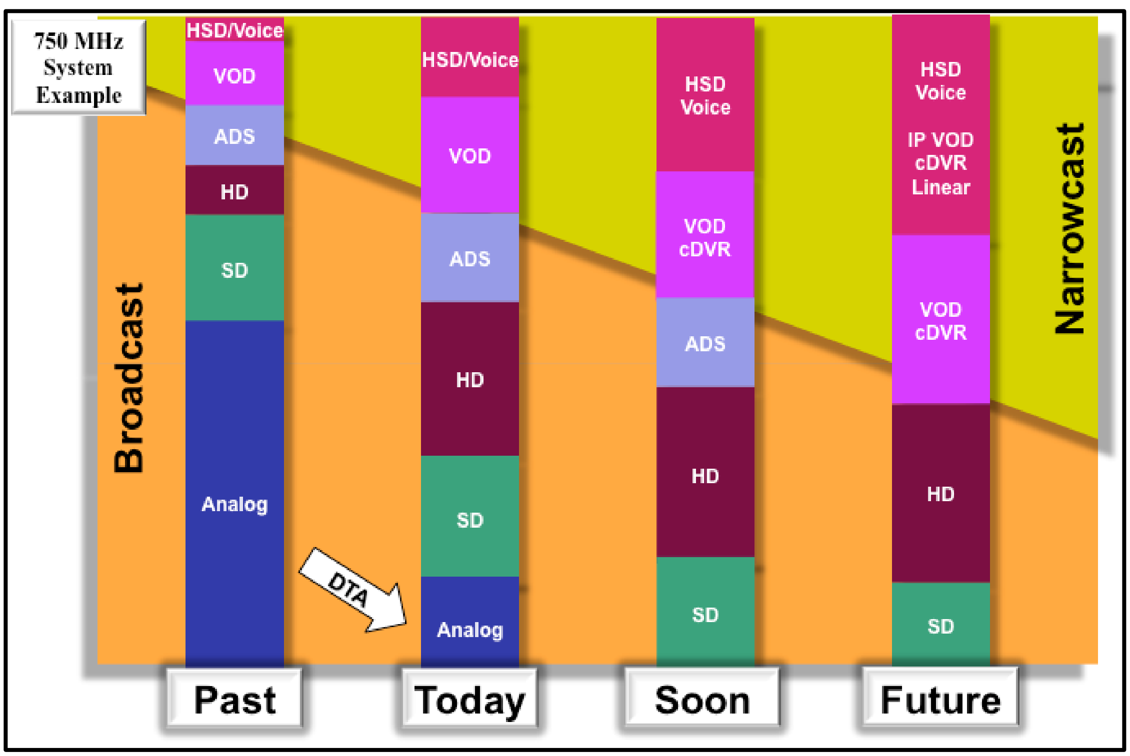

Figure 2 - Example of narrowcast service growth over time

However, these are difficult tools to take advantage of from a network perspective, because on a proportional basis, relatively few legacy set-tops will support all these technical advances, especially H.265. These tools are more likely to find significant support in equipment designed to handle newer, IP-video based services.

And, this approach will nonetheless require additional capacity from the network. This is especially true when considering that the deployment of these advanced video services will result in yet another simulcast of video programs, at least initially. This is because, realistically, advanced services will not, at least initially, replace the currently deployed service formats.

Furthermore, ubiquitous support for such devices would require considerable spectrum if the legacy services are maintained for an extended period – which is expected, given that legacy devices will continue to be deployed, for some amount of time. Moreover, this increase in simultaneous use of the more advanced IP video services while maintaining legacy services will be especially impactful over time, as the number of IP video services increases.

All of the above, coupled with the success experienced by operators in recent years with business services, homes security, etc., will likely require the deployment of IP capacity beyond what can be supported today. As well, it will require the development of tools for increased spectral efficiency and/or unleashing additional spectrum in the HFC network. The following sections of this paper enumerate some ways in which this can be achieved.

2.3. The Advent of DOCSIS 3.1

As it has been pretty well covered in the trade media, DOCSIS 3.1 deployment has been quite extensive. Most cable operators have deployed DOCSIS 3.1 across their markets, and several have even deployed DOCSIS 3.1 throughout their entire footprint.

The key motivation for the 3.1 version of DOCSIS technology is, in a nutshell, to scale DOCSIS more efficiently, both from cost and operations perspectives.

For the first 10 years of DOCSIS deployments, it was possible to offer Internet services and support its growth with just one downstream 6 MHz DOCSIS channel. Over the last five to 10 years, service speed demands increased to speeds of hundreds of Mbps, which require bonding of many 6 MHz channels. Therefore, the industry deployed multiple DOCSIS channels using DOCSIS 3.0, sometimes using 32 or more channels, and even requiring capacity beyond that supported by DOCSIS 3.0.

To that end, the three key goals and features of DOCSIS 3.1 were:

1. A much more efficient use of spectrum, with up to a 50% improvement in bandwidth efficiency (or bps/Hz). This is a result of more efficient forward error correction (i.e., replacing the older and less efficient Reed-Solomon approach with the far more efficient low density parity check [LDPC]), and the addition of the higher-order modulations1024- and 4096-QAM downstream, and 256- and 1024-QAM in the upstream.)

   These new modulation schemes provide two and four bits per second per hertz of improvement in both the upstream and downstream signal directions, while the use of the new forward error correction approach provides approximately 5 dB better RF performance. The end result is that networks are able to transport 1 Gbps of DOCSIS capacity in about 120 MHz of spectrum. For context, doing the same with DOCSIS 3.0, using SC-QAM requires about 180 MHz of spectrum.

2. Cost reduction, mainly by leveraging technologies commonly used in other transmission media. Specifically, the inclusion of OFDM, which is used extensively in wireless and wireline transmission media. The addition of OFDM for the downstream and OFDMA (orthogonal frequency division multiple access) for the upstream should enable operators to reduce costs by "packing" more bits in the HFC network more efficiently. As a result, these technologies will likely attract a larger supplier ecosystem, increasing innovation and fueling competition.

3. To enable a simple and orderly transition strategy. This applies doubly, in terms of compatibility with the previous generation of cable modem termination system (CMTS) and cable modem (CM) equipment, and simultaneously supporting an expanded spectrum capacity in the HFC network.

Specifically, DOCSIS 3.1 cable modems operate with DOCSIS 2.0 and 3.0 CMTS/CCAP equipment, enabling deployment of DOCSIS 3.1 CPE (customer premises equipment) as soon as it became available. Similarly, DOCSIS 3.1 CCAPs support DOCSIS 2.0 and 3.0 CPE, allowing operators to upgrade headend equipment without having to change any of the existing CPE. And, DOCSIS 3.1-based CM and CMTS equipment both support the currently required upstream and downstream spectrum, with expandability of the upstream to 85 MHz and beyond, and of the downstream up to 1.2 GHz.

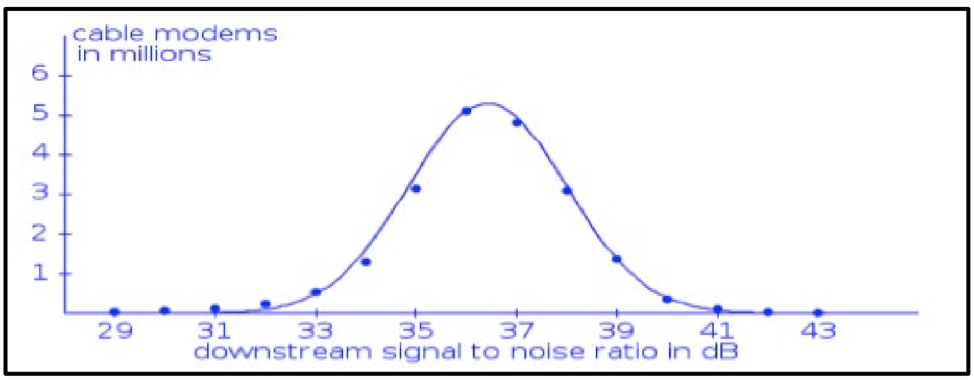

Figure 3 - Example of downstream SNR for a large population of cable modems

Figure 3 depicts the downstream signal-to-noise ratio (SNR) as reported by a very large population of cable modems. This data shows that many cable modems will be able to support the high-order modulation profiles included in DOCSIS 3.1.

Table 1 - SNR required for DOCSIS 3.1

Modulation	Signal-to-Noise Ratio
512-QAM	27 dB
1024-QAM	30 dB
2048-QAM	33 dB
4096-QAM	36 dB
8196-QAM[2]	39 dB
16384-QAM	42 dB

[2] 8196-QAM and 16384-QAM are included for future consideration in the DOCSIS 3.1 specifications.

Assuming an 8/9 LDPC coding ratio, Table 1 shows the required SNR for the modulation rates included in DOCSIS 3.1.

Applying the SNR requirements from Table 1 to the population of modems shown in Figure 3, we can see that a large population of cable modems would not achieve sufficient SNR to operate at 4096-QAM. Furthermore, if sufficient headroom is allowed to account for environmental fluctuations, the population of cable modems that would not receive signals with sufficient SNR to operate at 4096-QAM would be significant.

2.4. The Analog Intensity Modulated Forward Link in HFC Networks

As the name indicates, HFC networks use a fiber transport between the headend and the coaxial plant. This fiber link, intended to reduce the size of amplifier cascades, which improves performance, was originally developed with analog intensity modulated[3] lasers and compatible receivers in both signal directions, upstream and downstream.

Over time, the performance of the upstream link was improved by replacing the analog modulation with a digital transport. This improved performance significantly, and allowed for longer distances between the headend and the node. Different vendors implemented their own methods and technical capabilities to implement a digital transport, which resulted in incompatible systems and required the use of the same vendor's components for both the node and the headend.

However, the downstream link remained almost unchanged over time, with the only enhancements focused on improving distance and RF spectrum capacity. Performance has not really been an issue like it was in the upstream.

But more importantly, while the digital capacity of the upstream was limited to a few megabits per second, well under a gigabit of digital capacity which could easily be digitized and carried with Ethernet optics, the downstream digital capacity needed to transport the downstream spectrum has been considerably higher, reaching and even exceeding 10 gigabits per second.

Because of the above, analog forward links continue to be used. Even though headend equipment is currently capable of launching signals with >47 dB modulation error ratio (MER) performance, which is sufficient to generate and transport 16,384-QAM signals, analog lasers are limited to about 35 to 38 dB of MER performance, which would limit end-of-line performance to barely enough for 2,048-QAM or 4,096-QAM in short cascades in the best of the cases. In short, QAM density from headend equipment development outpaced the capabilities of the analog lasers we use.

2.5. Description of Options for Digital Forward Link

As time has gone by, technology evolution and certain developments as described below have enabled options for implementing a digital forward link. These include:

1. Evolution of edge-QAM modulators, which have gone from single and/or a few modulators to supporting 32, 64 or even more modulators,

2. Development of the CCAP, combining the functions of the video QAM modulator and DOCSIS into a single platform, and

---

[3] The majority of optical links deployed in HFC networks use linear fiber optic signal transmission, based upon analog intensity modulation. Many in the field use the term "AM fiber link" or similar.

3. Migration to digital video, either partially or completely.

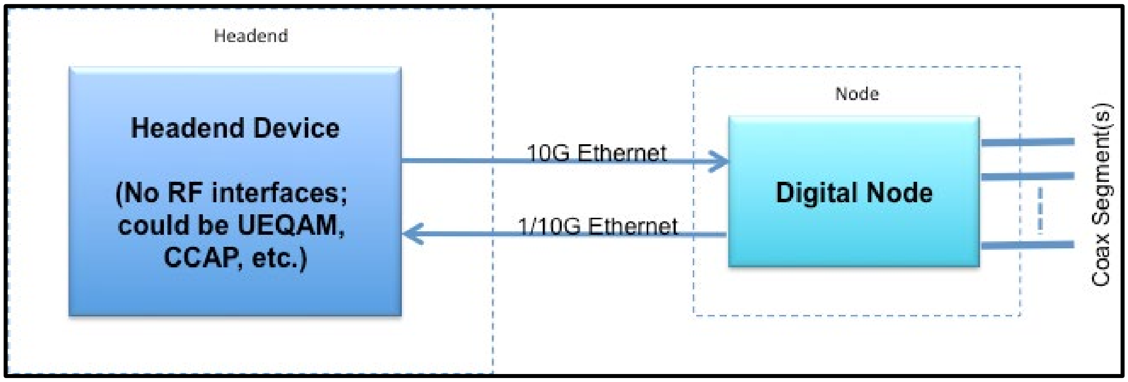

Figure 4 - Digital forward high-level architecture

With this technological evolution, it is conceivable to remove the RF combining network, and instead implement it digitally in the edge device, such as the CCAP. This evolution of the edge headend devices makes it possible to envision several options for digitizing the forward link.

Fundamentally, the migration to a digital forward includes the components included in Figure 4, and described as follows:

- The headend device, such as a CCAP, a high-density edge-QAM modulator supporting QAM for the entire spectrum,
- The node, which contains components normally implemented in the edge-QAM modulator or CCAP, to generate the RF signals,
- The link between the headend device and the node, comprising a digital interface, such as an Ethernet link.

There are the various approaches for how a digital forward link can be implemented to replace the currently used analog link. These approaches can be categorized into four groups, plus one option that would still leave RF generation at the headend device, as outlined in Table 2:

Table 2 - Categories of options for implementing a digital forward link

Option	Description and Approach
1. Maintain RF output in the headend	1.a Headend equipment remains unchanged 1.b Headend RF output is digitized, transported digitally, and RF is regenerated in the node
2. Remote the DAC from the PHY	2.a The DAC is removed from the headend 2.b Digital samples are transported digitally to the node where the DAC generates the RF signals
3. Partition the PHY and remote the lower portion of the PHY	3.a The PHY is split between the headend and the node 3.b The digital bit stream between upper and lower PHY is transported from headend to node
4. Remote the entire PHY	4.a The entire PHY modulation is moved to the node 4.b The MAC remains in the headend, and MAC frames are transmitted from the headend to modulator that resides in the node
5. Remote the entire PHY and MAC	5.a The entire PHY and MAC are removed from the headend device and placed in the node 5.b IP frames are transported from the headend to the node.

2.6. Comparison of Options for Digital Forward Link

There are pros and cons for each of the options. The following sections outline these tradeoffs.

2.6.1. *Option 1: RF remains in the headend*

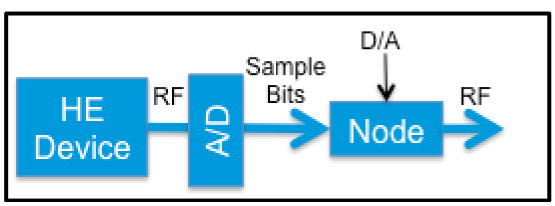

Figure 5 - Block diagram for Option 1

Equivalent to digital return, the RF output from the headend device is digitized, transported digitally, and converted back to RF in the node. This maintains HFC transparency. This option results in the highest bit rate over fiber; the capacity for multiple nodes would not fit into the available capacity of one 10 Gbps fiber. There is a loss of MER in the double conversion, so this option provides the least performance improvement. This option results in the least intelligence placed in the node, but an additional conversion stage is added in the headend.

2.6.2. Option 2: Digital-to-analog conversion is moved to the node

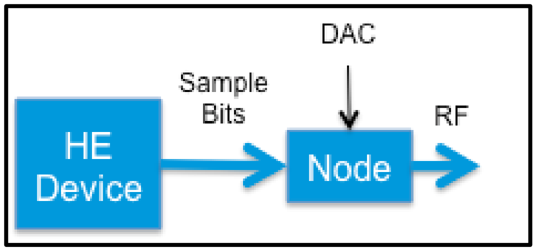

Figure 6 - Block diagram for Option 2

This option requires separation of the digital-to-analog conversion from the modulator. Together with Option 1, it results in the least intelligence in a node. It has a similarly high bit rate over fiber as Option 1; capacity for multiple nodes would not fit into the available capacity of one 10 Gbps fiber.

2.6.3. Option 3: Lower PHY is moved to the node

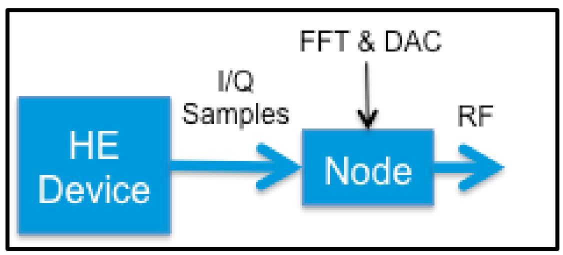

Figure 7 - Block diagram for Option 3

In this option, the PHY layer needs to be split into two components: upper and lower. This implements more node intelligence than in either of the previous options. Although it offers a lower bit rate over fiber than the previous options, it is still reasonably high. However, this option would require an industry proprietary point-to-point link between the headend port and the node, to transport the I and Q samples. Also, implementing this option would require the definition of interfaces which have never been defined in previous versions of the DOCSIS specifications, which in turn would result in modification of the silicon used and/or planned to date, and therefore results in the highest implementation complexity.

2.6.4. Option 4: Entire PHY is moved to the node

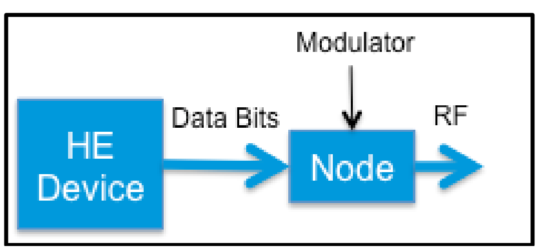

Figure 8 - Block diagram for Option 4

Moving the entire PHY layer into the node places more intelligence in it than with all previous options. This option results in the lowest bit rate over fiber; multiple nodes fit into the capacity of a 10 Gbps fiber. It enables a packet-based link between the headend and node, which results in significant benefits (which are outlined later in this paper.) It could use existing/planned silicon devices, and thus may be the easiest and quickest to implement. Lastly, this option offers the best MER performance improvement over analog.

2.6.5. Option 5: Move PHY and MAC to the node

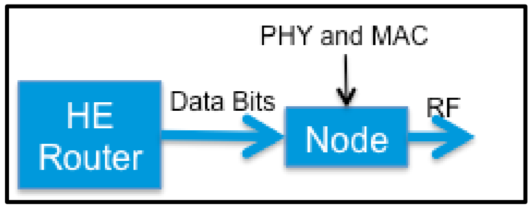

Figure 9 - Block diagram for Option 5

This option puts the most intelligence in the node. The data rate between the headend and the node is equivalent to the actual data transmitted, except for the addition of ancillary network data. It offers the same packet-based network benefits, and the same (highest) MER performance, as Option 4.

2.6.6. Comparison of Options and Implications

Any of the five options described above accomplish a migration away from the analog intensity modulated forward link and into a new era, where the link between the headend and the node becomes a digital link. And, while any of the above approaches accomplish a migration to a digital link, over the years, options 4 and 5 have received the most attention because of their relative implementation simplicity versus options 1, 2 and 3. We now call these options remote PHY and remote MAC-PHY, and we call the devices that implement them RPDs and RMDs (remote MAC-PHY devices), respectively.

As we migrated towards the implementation of a digital link, and separated either the physical layer in a remote PHY implementation, or also migrated the MAC in a remote MAC-PHY implementation, we stepped into the era of DAA. In these distributed access architectures, the remainder of the CCAP in the headend no longer needs to be implemented in an application-specific hardware design. Instead, the remainder of the CCAP in the headend can be implemented entirely as software running in general purpose compute platforms, which we now call a virtualized headend platform.

3. Benefits of Distributed Access Architectures

We will now focus on the benefits of a distributed access architecture, discuss some of the features of DAA, and outline network evolution strategies.

There are many benefits from the implementation of DAA. The following sections of this paper describe them.

3.1. Improved performance

Improvements on performance are achieved in multiple ways, including:

- Improved SNR (and MER) characteristics
- Longer link distances
- Higher reliability
- Better use of capacity

As described in the above sections of this paper, one key benefit of digital forward link is the improved performance resulting from the migration from an analog to a digital link. This gain varies depending on the characteristics of the analog link being replaced, but can be generalized as an improvement of 5 dB in SNR at the end of the line. This gain will result in higher capacity/Hz as it will be possible to run the higher order modulations as shown in Figure 3 and Table 1 for more of the cable modems in the network. This will enable significantly higher transport capacity for customers in the HFC network.

In addition, the digital forward link will enable longer distances between the headend and the node. This is because digital interfaces, such as an Ethernet link, are designed to operate over much longer distances while carrying the designated capacity. Extending the distance between the CCAP and the digital node would enable operators to move their CCAP devices back in the network to more centralized facilities, leaving the hub or OTN free of CCAP equipment. The benefit of such change could be very big for some operators, especially as segmentation of the network continues towards smaller service groups, for which additional CCAP equipment needs to be deployed.

A third benefit from the digital forward link is improved reliability of the optical link. It is well known that analog links require periodic maintenance and are subject to the effects of environmental changes. By contrast, Ethernet optical links are far more stable across a wider range of environmental conditions, and require little to no maintenance. The impact of this benefit could be very significant to operators.

Finally, the data transmitted through the link can be used more efficiently. One key example of such efficiency is the case where one link is used for multiple remote devices. As shown in Figure 10, one link from the headend CCAP device can be used to transport broadcast services once for multiple remote devices. This is achieved by using multicast addressing, whereby each of the remote devices uses the same lineup for each of the respective service groups. In doing so, a single link from the headend CCAP can be used for all the remote devices without exhausting the transport link capacity.

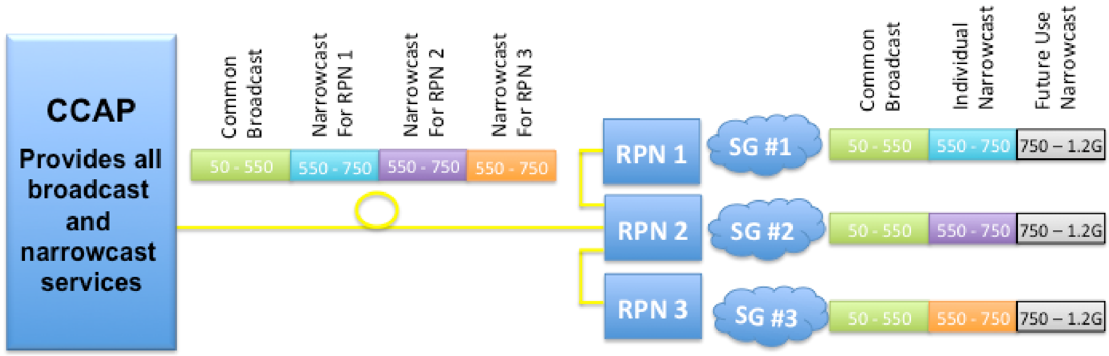

Figure 10 - Reuse of broadcast capacity across multiple RPNs

3.2. Increased Headend Equipment Density

The implementation of DAA makes it possible to improve the density of CCAP devices in several ways.

First, while CCAP devices are normally implemented via separate upstream and downstream line cards, a DAA line card could implement both upstream and downstream. This, in effect, doubles the capacity of a CCAP chassis.

In addition, a typical CCAP downstream line card will house eight or perhaps 12 RF ports, as defined by the printed circuit board space consumed by the components required for RF modulation, plus the sheer connector spacing required. However, Ethernet connectors can be placed considerably closer to one another, allowing a similar line card to easily house 16 to 24 ports. This additional density gain once again doubles the capacity of a CCAP chassis.

Finally, it is possible to consider "daisy chaining" remote PHY nodes (RPNs) off of a single CCAP Ethernet port. This is because, on the one hand, the capacity of a 10 Gbps Ethernet link would support the capacity needed for a single RPN. Plus, it is possible to generate an RPN "channel lineup" by transmitting the broadcast content once to multiple RPNs. As depicted in Figure 10, the data stream transmitted from the CCAP could contain a single "copy" of the broadcast line-up content, plus individual versions of the narrowcast content for each of the RPNs. The RPNs would then reuse the broadcast lineup content to recreate the individual RPN channel lineup. In this way, each service group served by the CCAP port would contain the same broadcast lineup while allowing for its own unique narrowcast line-up.

Then, as the narrowcast lineup capacity grows over time, CCAP ports would be segmented to support fewer RPNs, akin to the way service groups are split today to provide more narrowcast capacity as it is required.

As summarized in Table 3, the combined effect of the three factors described above is very significant, ranging from 8x to 18x of headend capacity gain. From a space and power perspective, this can facilitate huge savings.

Table 3 - DAA headend density gain

Density Factor	Density Gain
Combined US/DS line card	2x
Greater number of ports per line card	2x to 3x
Multiple RPNs per CCAP port	2x to 3x
Combined capacity gain	8x to 18x

But, just how meaningful is this headend density gain?

Consider: A migration from an HFC architecture with an average of N+5 (meaning an optical-to-RF node followed by five cascaded amplifiers) to N+0 would require about 10x the number of nodes, and the headend density benefits resulting from the DAA would neutralize the potential increase in CCAP equipment.

It is then quite clear that from a space and power savings, distributed access architectures take the benefit of CCAP to a whole new level.

3.3. Integration of HFC and Fiber Services

One of the largest areas of growth for operators is business services. Cable operators have deployed business services via both cable modems and fiber-based infrastructure. The fiber-based services are either point-to-point, using Ethernet and wavelength division multiplexing (WDM), or point-to-multipoint, using PON (passive optical network) technologies (either EPON [Ethernet passive optical network] or GPON [gigabit passive optical network]).

This duality results in the existence of two parallel networks. One of them, the HFC infrastructure, uses fiber from the headend to the node via an analog intensity modulated link for the forward direction and either analog or proprietary digital return, followed by coax infrastructure from the node to the home. The other consists of digital fiber from the headend to the customer, which is often used for commercial services.

Given the use of a digital fiber in both the forward and the return for the RPN, and especially because this digital fiber is based on Ethernet technology, it is possible to collapse both of these networks into a single infrastructure. Even without fully collapsing the Ethernet network better utilization of physical fiber can be enabled by the move to common DWDM wavelengths and multiplexers.

Therefore, the implementation of RPNs with an Ethernet interface between the CCAP and the RPN would make it possible to implement a PON interface at the RPN.

The benefits from this integration include:

- Reduction of the optical link for PON to the distance between the node to the customer premise
  - The typical distance from a node to a customer premise in an N+0 architecture is 1 to 2 kilometers. This would virtually eliminate any distance limitations for PON, making it possible to implement the largest possible densities.
  - In addition, this shortened distance would enable the use of lower power optics, which can translate into significant savings – especially for 10 Gbps optics, and for the upstream, which results in significant savings in the ONU.

- Leverage a single network for multiple services, which will reduce maintenance and increase operational efficiencies.

3.4. Migration Strategy

One of the more concerning issues to operators is the migration strategy when going to DAA.

Any migration that requires synchronized cutovers, or which requires changes in multiple locations to execute, is problematic, and usually results in a barrier to adoption. Therefore, it is very important that the migration to DAA allows for unsynchronized changes.

Ideally, the migration to DAA allows for opportunistic changes in the network. For example, one such change would be to migrate a single node, such as would be the case in a multiple dwelling unit (MDU) to increase capacity.

As it turns out, DAA enables such gradual, unsynchronized and opportunistic changes in the network. What follows is an overview of the steps and components involved in the migration to DAA.

Starting with the components of the network on both sides of the DAA, neither the back-office nor the various components in the customer premise need to be modified in any way. All back-office components are unaffected by the migration to DAA, and any additional MIBs (management information bases) for management and/or commands for configuration can be added well before the first CCAP line card or node is deployed. Customer premise devices are not affected in any way when deploying DAA, and any enhancements that are made possible through the introduction of DAA would be implemented in CPE that can be introduced before or after the migration to DAA.

The critical portion of the network where changes need to be made are in the headend and the plant.

To begin with, the changes required in the headend are primarily in the CCAP platform. The CCAP architecture was specifically designed to support multiple technologies simultaneously, which makes it possible to install regular RF upstream and downstream line cards and DAA line cards in the same chassis. While some operators may choose to deploy a separate CCAP platform for DAA, it is certainly possible to support both types of line cards in the same chassis. Of course, these DAA line cards can be installed at any time prior to beginning the migration in the plant, and any removal of RF upstream or downstream line cards can follow the deployment of any number of DAA line cards or nodes.

Turning our attention to the plant, it is similarly possible to migrate regular nodes to DAA nodes in any sequence. As an example, what follows is a sequence of steps where a single node is gradually converted from standard HFC to DAA.

Figure 11 depicts a single HFC node connected to a CCAP device.

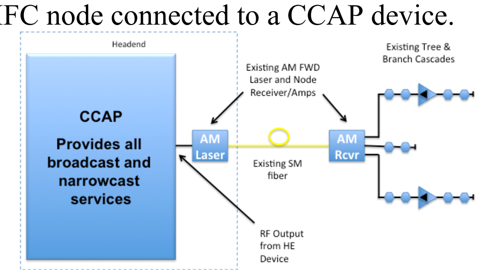

Figure 11 - Single traditional HFC node

Figure 12 shows how the HFC node would be converted to RPN while the rest of the HFC network remains unchanged. The DAA line card in the CCAP would have been deployed in the headend in a prior activity, and even the RPN could have been deployed before the day of the cutover. Then on the day of the change the fiber cable could be swung in the headend from one AM laser to the CCAP DAA card, and in the field from the HFC node to the RPN. Of course it is not necessary to perform the migration in such a fashion, but it would be possible if desired.

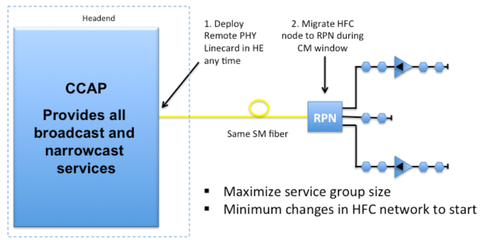

Figure 12 - RPN deployment step 1

Figure 13 depicts a possible step 2 in the process, whereby additional RPNs are installed to segment the original service group further. These additional RPNs could be daisy chained from the original RPN by taking advantage of the broadcast reuse feature, minimizing complexity in the deployment process.

NOTE: The example depicted is one in which fiber is run to every amplifier station. However, a more efficient segmentation scheme would include optimal placement of RPNs in an N+0 HFC architecture with some turnaround of passive components.

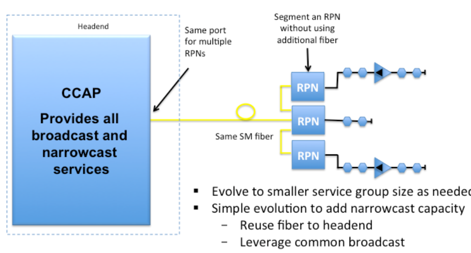

Figure 13 - RPN deployment step 2

Figure 14 shows how further segmentation could take place by replacing the remaining amplifiers in the network with RPNs.

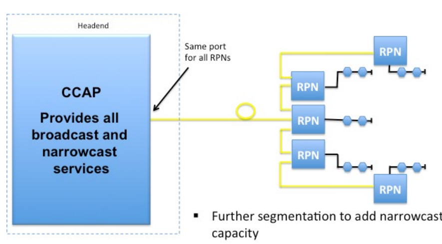

Figure 14 - RPN deployment step 3

Figure 15 shows that the RPN service group depicted above is segmented as additional narrowcast capacity is required. In this example, two of the RPNs from the DAA service group shown in Figure 14 are split into separate service groups using separate CCAP ports.

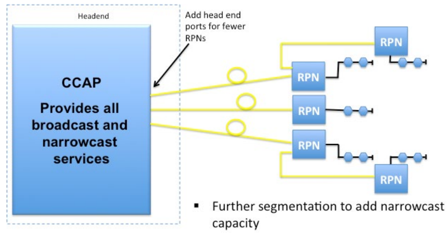

Figure 15 - RPN deployment step 4

Eventually each of the RPNs could be connected to an individual CCAP RPN port as shown in Figure 16. This would provide up to 10 Gbps of capacity to each RPN. This could, for example, be desirable to provide both RF and PON services from the RPN.

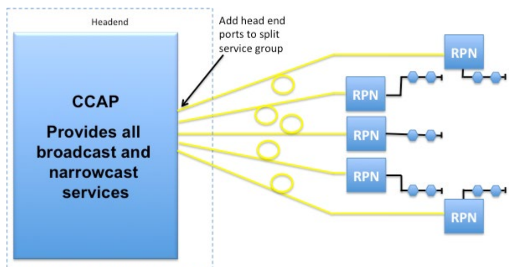

Figure 16 - RPN deployment step 5

Similarly, the DAA line card in the CCAP could be upgraded to support even more capacity as such capacity is needed and becomes cost effective. For example, the Ethernet link from the CCAP to the RPN could eventually be upgraded to 40 or 100 Gbps, both of which are already commercially available.

3.5. From Today to Virtual CCAP

As the network has to continue in operation through the transition, virtualizing the CCAP requires careful planning and a sensibly staged process. As with roads, where cars must be kept moving during any lengthy highway reconstruction, in the network customer traffic must continue flowing day after day. In a sense, while road work is visible to car drivers, in a network the modifications remain invisible to the end user.

One way to do so is to migrate individual functions, one at a time. So, one must develop a list of the functions that would be virtualized, and this list would be prioritized, such as on the basis of complexity of implementation and benefit. Those features with the lowest implementation complexity and the highest benefit would be prioritized higher in the list, and consequently implemented first.

In DOCSIS 3.1, one of the functions that would rise to the top of any such list is modulation profile management (MPM). This is because MPM will take time to be implemented by vendors in a CCAP chassis, but implementing externally via virtualization could be quite simple. In the process, its benefit to operators is quite significant since it would enable better efficiencies from DOCSIS 3.1.

Over time, implementing virtualization of the various functions of the CCAP would lead to a completely virtualized CCAP platform. Such a platform would be more easily scalable than CCAP platforms are today, where segmentation of service groups requires the addition of more chassis in a linear relationship fashion.

In addition, and perhaps more importantly, virtualizing the CCAP will enable the development of additional functionality, and improvements to such functionality, to occur much more rapidly than it is possible to do today.

4. DAA Components, Use Cases and Generations

As operators move forward with the implementation of DAA, the evolution of DAA components, implications of the various use cases, and the generational aspects of DAA have become important to understand and track.

4.1. DAA Components

As the link from the headend to the node is converted from analog intensity modulated forward to digital, using Ethernet as the transport, several approaches can be taken for the implementation of the remaining headend components. In this section we examine one approach in some detail, for which a key goal is to convert all required components into functionally individual software pieces implemented independently.

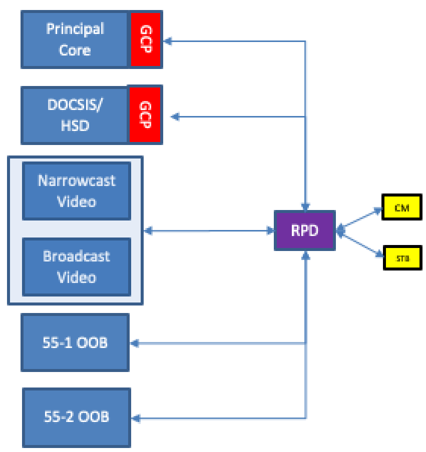

Figure 17 - DAA implementation components

4.1.1. Advantages and disadvantages of discrete components

As shown in Figure 17, an implementation approach for DAA is to develop discrete software (SW) components for each of the various DAA components. Some of the advantages for doing so include:

- The implementation can consist of a multi-vendor platform, where each component can be developed by a different party.
- By having smaller functional components, their implementation tends to be simpler.
- Time to market also tends to be reduced.

However, implementation of smaller discrete components has its downsides, such as:

- There is an implied requirement to more tightly specify the behavior of each component to ensure that the overall system will operate as intended.
- Interface specifications between the various components is required.
- Management of the various components, including their configuration and upgrade is generally more complicated, and requires more elaborate orchestration.

4.1.2. Key DAA discrete components

The key components depicted in Figure 17 include:

- The principal core, which is the first component that the RPD will contact after receiving an IP address.
  - The principal core is implemented such that it will configure all the RPD functions except DOCSIS channels and behavior, which will be implemented by the DOCSIS CMTS.
  - As depicted in Figure 17, the principal core communicates with the RPD using the GCP protocol, for which it is known as the Generic Control Protocol Principal, or GCPP for short.
  - Included in the GCP principal core are all the non-DOCSIS command and control functions for the RPD, including configuration, management and reporting.
- The DOCSIS core, which is the second component that the RPD will contact in the network.
  - The DOCSIS core also communicates with the RPD using GCP.
  - The DOCSIS core provides all configuration, command and control for DOCSIS channels, both downstream and upstream.
- Narrowcast and broadcast video engines
  - Implemented as separate components, the narrowcast and broadcast video engines provide all the content services for the various RPDs in the network.
  - Neither the narrowcast nor the broadcast video engines communicate with, nor have knowledge of, the RPDs.
  - Services are configured statically in the narrowcast and broadcast video engines upon their bring-up, and are multicasted to all RPDs, which listen for these services as configured by the GCPP.
  - The narrowcast and broadcast video engines could be implemented separately, but they could be operated together as a single functional system.
- Out-of-Band engines or cores
  - The OOB (out of band) functions are implemented separately from the video engines.

  - Given that video systems are implemented using a single encryption and command/control technology, only one (i.e., either SCTE 55-1 or SCTE 55-2) of them is deployed in any one system.
  - The OOB function may or may not implement GCP for communicating with the RPD. When GCP is implemented the OOB server is a core, and it will configure the OOB downstream and upstream OOB channels in the RPD. However, when GCP is not implemented the OOB server is an engine, and the GCPP will configure the downstream and upstream OOB channels.

- Finally, not depicted in Figure 17 is a very important component: the timing server.
  - The timing server, also known as the grandmaster, provides the critical timing synchronization for all the DAA components.
  - Each of the DAA components will include a timing client, which will communicate with the timing server to maintain timing synchronization.
  - While timing synchronization is not absolutely critical for video services, it is imperative for DOCSIS service to operate. Therefore, video services may be initiated before timing synchronization is achieved, but DOCSIS services will not.

4.2. Key aspects of DAA interoperability

The base implementation of a DAA system is generally simple. However, significant complexity is introduced when interoperability with different vendors' components is introduced.

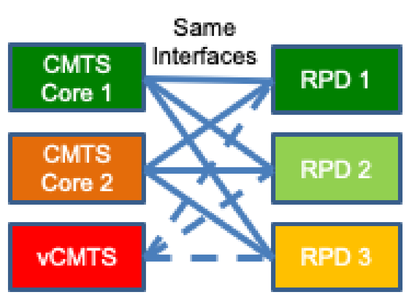

Figure 18 - Functional CMTS-RPD interoperability matrix

As depicted in Figure 18, the number of combinations of interoperable components increases geometrically as additional components are added on either side of the interoperability matrix. Having a single CMTS interoperate with multiple vendors' RPDs is complex and requires a lot of careful planning and implementation. If the number of CMTS implementations is increased to two or three, the interoperability complexity doubles and triples respectively.

When considering the overall CCAP system, the complexity to achieve multi-vendor interoperability is even larger. For example, if multiple GCPPs and/or multiple video engines and/or multiple OOB engines/cores are introduced into the mix, the amount of complexity and work required for testing and interoperability results in increases by orders of magnitude.

Therefore, a multi-vendor RPD deployment coupled with a single headend implementation is a sensible approach to an interoperable DAA ecosystem.

4.3. Use Cases

In the same way as there are different kinds of nodes for different HFC network applications, there are RPDs with different characteristics that are best suited for each of the specific HFC network use cases. Similarly, while there are use cases for RPDs in the outside plant, there are also applications for RPDs in headends, or the "inside plant" as it is frequently called, which will have different implementation characteristics. The following sections cover the key scenarios.

4.3.1. Outside plant

The environmental characteristics of RPDs developed for outside plant make the design of such devices very different than for RPDs developed for inside plant. The key characteristics for outside plant RPDs are as follows:

- Designs must conform to very tight space availability requirements inside of a node enclosure
- RPDs must support an environment where heat dissipation without the use of fans is critical
- Powered from quasi-square wave power supplies used in HFC networks
- Minimize power consumption to the extent possible given the limited amount of power during normal operation and especially during stand-by power mode

In addition, and perhaps more importantly, there are different kinds of nodes for different HFC network applications, which will drive varying designs for RPDs for outside plant, as follows:

- Traditional HFC networks include cascades of multiple amplifiers and cover a plant footprint of a few hundred homes. In such cases the network capacity offered by the RPD should be maximized, such as including multiple downstream and multiple upstream ports.
- Newer HFC networks are built with fewer, or even no amplifiers, and are targeted to cover smaller network footprints. In such cases it is not necessary for the RPD to support much more than a single downstream port, with either a single or dual upstream ports.
- Finally, given that scaling is needed as in any other network application, it should be possible to support greater capacity over time to the extent possible. For example, while initial deployment may only require a single downstream and/or upstream, over time service group segmentation may require additional downstream and/or upstream ports in a single node. For that purpose it is usually a design requirement that multiple individual RPDs fit into a single node enclosure.

Given the above, RPDs that are built with a single downstream and a single upstream (frequently called 1x1) or a single downstream with dual upstreams (1x2), such that operators can place a single one in the node or place additional units when capacity demands require it. Newer silicon designs include more capacity at lower power levels, making it possible to develop RPDs that contain multiple downstream ports and multiple upstream ports, such as 2x2 and 2x4 designs in a single RPD.

4.3.2. Inside plant

In contrast to the outside plant environmental characteristics, RPDs developed for inside plant have other constraints that make the design of such devices very different than for RPDs developed for outside plant. The key differences in the design of RPDs for inside plant are as follows:

- Rather than the physical volume of the allocated space, the layout is a primary concern so that dense set-ups in a rack are possible, including cabling distribution in the front and/or back of the rack.
- Designs must support an environment where forced air is used for heat dissipation, requiring airflow from front-to-back or back-to-front, sometimes allowing airflow from side-to-side and/or in a vertical direction.
- Powering frequently requires DC power supplies, but AC power supplies are used in other cases.

In the case of inside plant RPDs, these are frequently implemented in one of two different form factors:

- Modular, where RPDs are individually removable in a chassis-based design, or
- Fixed, where the entire set of RPDs are part of a monolithic device

The modular design is generally used in larger headends where the ability to replace a defective individual unit is a paramount concern. The fixed design is targeted for a smaller facility, or even a cabinet, where space is the primary concern.

4.4. Generational considerations

One additional consideration is the evolution of the DAA to support new generations of equipment.

The initial implementation of DAA components included support for DOCSIS 3.1. The main component that is specifically developed and implemented for DOCSIS 3.1 is the RPD, which incorporates application specific integrated circuit (ASIC) devices which support up to DOCSIS 3.1, but will not support DOCSIS 4.0 functionality.

As newer parts of the DOCSIS 4.0 specifications are implemented, such as full duplex (FDX) DOCSIS, the RPDs will have to be swapped out in order to expand their support for DOCSIS 4.0. This process is akin to what had to be done with CCAP line cards in the past, where either upstream and/or downstream line cards are swapped over time as new versions become available. And, as part of this upgrade, the older equipment is reused in other locations where the newer equipment is not yet needed.

However, for the remaining DAA components, if these are implemented in software on general purpose compute platforms, these should be upgradeable to support newer DOCSIS specifications such as DOCSIS 4.0 and/or other enhancements by simply expanding the functionality implemented in software and downloading it to the platforms in which they run.

In fact, the process for upgrading the DAA components to support changes in the functionality including enhancements to DOCSIS, becomes easier than ever before given the nature of the DAA platform, especially when the DAA implementation includes a minimalistic RPD at the edge of the HFC network and virtualized components for the remainder of the DAA components.

5. Deployment experiences and lessons learned

As operators move forward with broad deployment of DAA, including tens of thousands of RPDs in hundreds of locations, there are many useful lessons learned. The following sections of this paper describe findings in four key areas: monitoring system approaches, plant powering, provisioning processes and approaches, and hardware and software management techniques and experience.

5.1. Monitoring System Approaches

As the number of DAA sites and components increases rapidly, the use of tools such as Simple Network Management Protocol (SNMP) become exceedingly difficult to operate and scale. With just a few thousand RPDs, operators would have more end points from which to gather data than they have CMTS or CCAP systems in total. Having deployed multiple tens of thousands of RPDs, and projecting that deployments will pass the hundreds of thousands in just a few years, streaming telemetry approaches become much better suited for monitoring such vastly distributed systems, and perhaps the only viable approach.

Open source tools like Grafana can be very easily used to plot streaming data collected from network components. And, gathering data from various sources, such as those outlined in Section 4.1, allows for quicker correlation of issues in the network.

As an example, Figure 19 depicts the status of RPDs in a particular area of the network. The view in the figure shows the status of RPDs as seen by a vCMTS, including the number of active and inactive RPDs and the corresponding number of cable modems in those RPDs. The vCMTS is one of the network cores in the DAA, and similar views are available as seen by other network cores as outlined in Section 4.1.

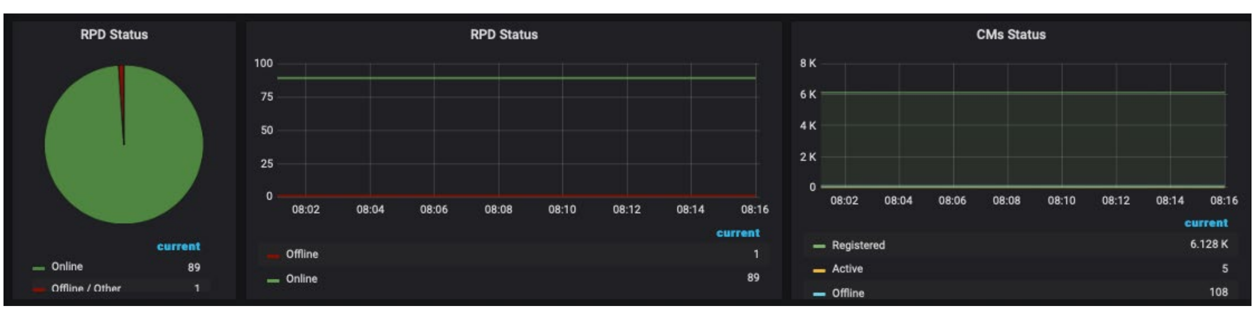

Figure 19 - RPDs and cable modems in a particular section of the network

Streaming telemetry can be easily obtained from modern compute resources, such as those used to virtualize the CMTS and other DAA components. However, streaming telemetry is not currently available from RPDs. Instead, an RPD supports polling telemetry through protocols such as GCPP, which is used in the RPD for other purposes, including configuration.

To further expand on the use of streaming telemetry to RPDs, work is currently underway in an effort led by CableLabs, and supported by cable operators and suppliers alike, to incorporate streaming telemetry in the remote PHY specifications. This specification development effort includes a prototype being developed between CableLabs, Comcast and the support from vendors.

5.2. Plant powering

Plant powering of RPDs has become surprisingly challenging, and as a consequence represents a rich set of "lessons learned." This section provides some relevant background, describes the challenges encountered, and the solutions applied to date.

5.2.1. HFC network topology and component powering

Network power supplies have always provided power to multiple HFC plant devices. Historically this included nodes, amplifiers and line extenders. In addition to HFC plant devices, network power supplies provided power for any other device connected to the plant, such as Wi-Fi access points, cellular microcells, etc. In the past, the plant powered some individual CPE devices used for specific applications, such as telephony (which are largely removed from cable networks now).

The footprint covered by an individual network power supply, which we might refer to as the power serving group, depended on the specific characteristics of the HFC network, especially related to the density of the area in question. In areas where traditional N+x HFC architecture is deployed, an individual power supply may provide power to a single node and a collection of amplifiers and line extenders, plus whichever other devices may be present in that power serving group. In the case of passive HFC plants, such as fiber deep and/or N+0 plants, a single network power supply might power a number of nodes and any other devices that may be present in the power serving groups.

5.2.2. Traditional network maintenance activities

Before the introduction of DAA, the HFC plant devices became operational almost instantly upon receiving power from a network supply. By contrast, other devices that may be present in the power serving group, such as a Wi-Fi access point or microcells, would take several minutes to become operational once network power is supplied.

Over the years, the normal process of a maintenance technician included some interruptions of network power for troubleshooting purposes. This might include procedures intended to identify sources of signal interference or impairments, or replacement of HFC components that might be causing signal interference or impairments, such as taps, splitters, power inserters, and any other passive component, all of which could be done very quickly. In doing so, power would be interrupted for a matter of seconds, or even fractions of seconds, which would cause a very brief loss of power for an amplifier or line extender, which CPE devices would not perceive, and therefore not cause an interruption of service beyond the brief moment of the power interruption itself.

5.2.3. Impact on DAA network devices

As DAA components have been deployed, those nearly imperceptible power interruptions described above became much more impactful.

For example, while a traditional node, amplifier or line extended would recover from a power interruption nearly instantly when power was restored, the same does not apply to an RPD. Instead, the RPD, which has a processor (or multiple, actually), requires considerable time to boot (in the minutes range), and then go through a power-on self-test (POST) process, which takes additional time to complete. This was followed by processes such as time-synchronization, which take even more time to complete.

Therefore, while a second or sub-second power interruption would have been imperceptible to most users in a traditional HFC network, it now causes multiple minutes of interruption until an RPD becomes operational, and even longer until the CPE in the home becomes operational once again.

To make matters worse, in cases where a serving group power supply provides power to multiple RPDs, such as is the case in passive networks, the service interruption would affect an even larger footprint.

5.2.4. Solutions

Several actions can be taken to overcome these problems, including training, shortening the power cycle process, and using power holding devices.

The most immediately available solution to minimize power interruptions is training, to prevent unnecessary power interruptions. As all cable operators' maintenance team leaders well know, this is much easier said than done. Training requires time, practice and repetition, all of which takes months or years to implement. Furthermore, there are many cases when interrupting power becomes impossible to avoid, such as when component replacement is necessary for impairment resolution.

Shortening the power cycle process requires, in most cases, substantial changes in the SW running in an RPD and the related systems, and in their integration. For example, it is relatively easy to simplify the POST process and/or to shorten the bring-up time. It requires more effort to reduce the configuration time, such as by keeping and reusing configurations whenever possible. Other aspects of the power-up process are more difficult to shorten, such as the time synchronization, which can only be shortened, and even eliminated, with a reboot from active operation that did not follow a power cycle.

An additional approach for eliminating short power loss events is to implement power holding devices within the RPD or in the node housing the RPD. Such devices, which are based on hardware, take even more effort to implement and have a linearly increasing cost (i.e., each unit has a cost, and the cost does not decrease substantially with the deployment volume). In addition, power holding devices can't hold power for considerably long periods of time, so they only prevent short power interruptions (i.e., power holding devices only prevent power interruptions lasting a few seconds).

The combination of all of the above approaches is probably the most effective combination.

5.3. Provisioning process and approaches

The installation, configuration, bring-up and provisioning process for an RPD requires information related to all the services provided by cable operators, plant characteristics and operational requirements. The process is complex given the various sources of information, and requires automation and verification. The following paragraphs provide an overview of those topics.

5.3.1. RPD Installation

There are two fundamental types of RPDs: those that are installed in nodes and those installed in headends (shelf-based). While the two types are functionally the same, the use case is different, and therefore some aspects of the installation and configuration are different.

For example, for node-based RPDs, there is an individual RPD to identify, configure and verify (unless two RPDs are placed inside of the same node). In the case of shelf-based RPDs, the identifier label used is typically for the entire shelf instead of being specific for an individual RPD, and identifying the specific RPD that is being provisioned requires an understanding of the shelf architecture.

While RPD identification may seem trivial, actually identifying an RPD that is part of a shelf presents several complications. Finding the identifier label may be difficult, especially if the shelves are installed immediately above or below another shelf (i.e., without any space between them in the rack). Additionally, the media access control (MAC) address of the RPD, which is a fundamental component of the configuration process, may be derived from a base MAC address for the shelf, instead of included on the label as it would be for an individual RPD in the node. While these may seem inconsequential aspects of the installation process, they become significant complexities operationally, especially when speed of deployment is a key goal.

An additional twist to the installation process is the replacement process. When an RPD needs to be replaced, the process has to involve some form of removal plus an addition. Given that the RPD configuration is specific to the RPD, the MAC address of the RPD being replaced needs to be removed, and the MAC address of the new RPD needs to be used in its place.

Finally, the installation process needs to take into account the isolation of known defective RPDs, a return material authorization (RMA) and repair, and a manufacturing, purchase and installation identification for tracking devices with problems and defects, all of which has to be taken into account as part of the process.

5.3.2. Information and Sources

The information that is used in the configuration and provisioning process includes plant characteristics such as power levels, tilt, leakage detection tones, pilot tones, etc. It also includes service information, such as broadcast and narrowcast video services' frequencies and channels, DOCSIS channel frequencies and modulation characteristics, and OOB signals, to name the most common.

The above information comes from different sources within the operator's organization, and therefore various individuals determine what's the correct information and use various types of source data. This can naturally lead to conflicting information and configuration.

In addition, there is added complexity in having multiple configuration sources for an RPD. For example, in the DAA described in this paper (see Section 4.1), the RPD receives configuration information from two or three different sources within the network, which include the GCPP, followed by the vCMTS auxiliary core, and the OOB engine auxiliary core for 55-2 encryption networks. Having multiple sources of configuration information can easily lead to conflicts that are hard to identify.

Finally, parameters change over time for a multitude of reasons on a regional or per-node basis. Some of the changes will apply to a number of nodes simultaneously, such as changes in channel line-up and/or addition of DOCSIS channels. Other changes will be unique to a node, such as changes in channel configuration to optimize capacity usage through profile management application (PMA) tools.

Given the multitude of information sources and configuration components, it is almost inevitable that errors will occur. One example involves the specific frequency assignments for either of the configuration components which might be incorrectly captured or communicated, including frequency overlaps. Given that the RPD does not have the inherent intelligence to understand which overlaps are intended and which are errors, it is possible to create a configuration that will not work.

5.3.3. Configuration and Verification Tools

Resolution of configuration errors, such as overlaps and conflicts, could be considered to be relatively straightforward. Configuration tools can certainly be developed to capture the information, and automation can be implemented to minimize errors, and even prevent them in some cases.

However, there are factors that make the detection of errors more challenging, such as when the information arrives at a central configuration tool from different data sources. In that case, an additional tool would have to consolidate all the configuration information before the final configuration is created, which is especially the case if the configuration is created "on the fly." In such cases, the change in the configuration could have been made a priori, perhaps hours or even days before the configuration is applied, and therefore the person or system that commanded the change would no longer be notified of the conflict.

Other information and/or changes are much harder to verify, such as configuration of pilot tones or leakage tones, either frequencies and/or signal magnitude. Only experts in the specific field might understand the information that is being entered, and therefore determine that a different value would have been appropriate.

Perhaps the ultimate approach to verify the configuration is to conduct a field configuration test. Tools have been developed for quite some time to measure RF characteristics of a channel line-up, including signal levels, channel content, etc. More recently, tools have been expanded to read an entire channel lineup, and compare it to a known, "good" lineup, to determine if there are any errors. Such tools are very useful because they verify the actual RF signals being generated by the RPD in their entirety. However, such testing does take resources, time and additional cost.

5.4. Software and hardware management

As with any device containing a processor, RPDs require software upgrades. While that is very much expected, how frequently changes are required may be surprising. In addition, hardware changes and use cases have been even more frequent than anticipated. The following paragraphs provide an overview of some of these situations found through field experience.

5.4.1. Software upgrades

Everyone involved in access networks has experienced their fast-paced evolution. Changes not only include the expansion of capacity, which is, of course, a very big challenge, but also of functionality, uses, and most recently of configuration, including dynamic changes.

The most typical of such changes involves incorporating new or modified functionality. For example, almost all operators deployed downstream channel bonding before upstream channel bonding, or downstream OFDM channels before upstream OFDMA channels. Similarly, lots of functionality included in the DOCSIS specifications is deployed over time. As a result, equipment manufacturers implement such functionality progressively, as it becomes necessary and useful to the operator. Therefore, software upgrades are required to make such additional functionality available.

Less predictable are new or different use cases. For example, while initially an operator may deploy RPDs in passive networks, they may over time expand to using RPDs in a more traditional N+x HFC network. Therefore, while the operator may not have needed to support pilot tones initially, such functionality may be required later on in the deployment lifecycle. As with any other functionality that is not needed initially, it may not have been incorporated into the initial software load and then deployed over time, requiring software upgrades.

An additional type of software change involves making configurable information dynamically possible. For example, initial deployments of DOCSIS channels may not require that modulation profile configuration changes be made without rebooting the RPD, such changes may need to be made dynamically later on, without requiring the reboot of the RPD, and even without requiring cable modem reset or even re-ranging. As with the above cases, such changes in the RPD operation require software upgrades.

Finally, as with any other software, bugs are almost inevitable. Therefore, when bugs are found they have to be identified and fixed, which of course requires software upgrades.

Upgrading RPDs that have been deployed requires the use of software deployment tools. This is important when the number of RPDs that have been deployed is in the thousands, and it becomes especially important when the number of RPDs increases by one, two and possibly three orders of magnitude.

In addition, it is inevitable that RPDs are manufactured over time, bought over even longer periods of time, and then placed in warehouses until these are deployed. Therefore, RPDs are almost always upgraded upon initial deployment, which generates the need to manage multiple software versions and upgrade processes.

Moreover, RPDs become even more challenging when their variety increases, for example when multiple vendors are deployed in a network, and especially when multiple models from each vendor are deployed. Performing field tests becomes necessary, requiring handling of exceptions, and gradual or partial deployments are inevitably necessary to support operational requirements. All this makes the need for a sophisticated software management and deployment tool.

5.4.2. Hardware upgrades

Just like having to support multiple software versions, over time multiple hardware versions are also inevitable. This can happen because of product evolution by manufacturers, new use cases by the operator, and industry technology evolution. Let's examine a few examples of each.

Product evolution is probably the least appreciated, but the most real form of hardware evolution. Manufacturers will almost always evolve their products to reduce costs, resolve component obsolescence, or achieve better performance. Therefore, multiple hardware versions are almost always necessary and inevitable. In some cases, it might be necessary to keep track of different hardware versions for purchase tracking requirements, especially when costs change. But keeping track of hardware versions is also important for other reasons, such as for certain troubleshooting activities, and even to understand and support certain software differences, such as when certain newer versions support features not available to older hardware versions.

Less predictable is the evolution of use cases. For example, an operator may initially deploy RPDs with fewer downstream and/or upstream ports than may be required over time. While it may be possible to double the use of some RPDs on certain nodes (i.e., start with one RPD in the node and later move to two RPDs in the node), in other cases it may be necessary to deploy a different type of RPD (i.e., start with one model of RPD in the node and over time move to use a different model of RPD in the node). In most cases the evolution would require one to keep track of the number or modem of RPDs in each node. Yet another example of change in use case would be the initial need for node-based RPDs, which eventually changes to incorporate the use of shelf-based RPDs, or vice versa.

Finally, the almost certain hardware evolution scenario is that the technology used initially evolves over time. For example, initially RPDs were developed to support DOCSIS 3.1, but more recently their design has been modified to support FDX. Cable operators have dealt with the evolution of access network technology for decades, migrating from pre-standards to DOCSIS 1.0, then to DOCSIS 2.0, followed by DOCSIS 3.0 and eventually DOCSIS 3.1. Equipment management, including shifting equipment between locations, is a process that operators have mastered and will continue to apply to support technology evolution.

6. Conclusions

Demand for more narrowcast service capacity has driven many changes, allowing operators to reclaim spectrum. Splitting nodes to reduce the number of homes passed in an HFC node's footprint is one key method used to provide more narrowcast capacity. Enabling narrowcast services on this reclaimed and newly added spectrum requires more equipment in the headend. Given the trajectory of growth, using traditional equipment technology creates an increased demand for more headend space. Deploying DAA enables the required growth without expanding facility footprints.

The deployment of the highest modulation orders available with DOCSIS 3.1 required higher MER out of the node. Transitioning to a digital forward link, which became possible with the newer DAA, is a key enabler to achieve this demand. Moving the entire PHY layer to the node has become the standard method to implement a digital forward link.

Deploying DAA in the outside plant allows for more efficient infrastructure utilization between services. DAA and commercial services can share physical fiber due to commonly used DWDM wavelengths and spacing. There is also the possibility of a converged Ethernet switching network to serve both DAA nodes and commercial customers. Having Ethernet inside of a DAA node allows for efficient PON deployments sourced from the same node.

Moving from specialized CCAP hardware to virtualized components leads to a more scalable system, where additional capacity can be added without the step function of adding new CCAP chassis. Furthermore, feature velocity can be improved on a virtualized system. DAA is implemented with discrete components that allow for a multi-vendor platform and includes smaller functional components.

On the flip side, implementing discrete components requires well behaved interfaces that adhere to a well-defined set of specifications. In the DAA implementation described in this paper, key functions include the GCPP core, the DOCSIS core, the narrowcast and broadcast video engine, the SCTE 55-1 OOB engine or SCTE 55-2 OOB core, PTP timing distribution, etc.

RPDs are available for both inside and outside plant applications, implemented as individual units or in shelves containing multiple RPDs, and including single or multiple services groups. Power level and segmentation capability/capacity are key attributes of outside plant RPDs. Density and serviceability are key attributes of inside plant RPDs.

Even though current generation RPDs lack support for FDX operation, other DAA components may have the ability to be upgraded for FDX operation using the current hardware. This enables an easier transition to FDX on an as-needed basis, by swapping hardware on desired nodes while utilizing the existing platform.

The experience acquired over the course of the last five years, especially the last two years, leads to several key lessons learned that are outlined in this paper. One such lesson is the need for flexible and scalable monitoring systems, using streaming telemetry and open source reporting tools. Another is the impact of plant powering considerations, which resulted in the development of software systems and hardware components to alleviate the effects of power interruptions during normal maintenance activities.

Having provisioned tens of thousands of RPDs already, a well thought out provisioning system process and approaches for configuration verification are very important to ensure smooth, rapid and correct deployment. It is especially important to consider the variety of information and the number of sources, and to create a process for gathering and applying the information that minimizes errors, some of which are easier to foresee than others.

Finally, given the need to deploy RPDs of multiple types and from multiple suppliers, software and hardware upgrade management becomes very important. Both foreseen and unforeseen changes require the use of flexible version management tools.

Abbreviations

A/D	analog-to-digital
AC	alternating current
AM	amplitude modulation
ASIC	application specific integrated circuit
BC	broadcast
CBR	constant bit rate
CCAP	converged cable access platform
cDVR	cloud digital video recorder

CM	cable modem
CMTS	cable modem termination system
CPE	customer premises equipment
CTA	Consumer Technology Association
DAA	distributed access architecture
DAC	digital-to-analog converter
dB	decibel
DC	direct current
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
DTA	digital transport adapter
DWDM	dense wavelength division multiplexing
EPON	Ethernet passive optical network
FDX	full duplex [DOCSIS]
Gbps	gigabits per second
GCP	generic control plane
GCPP	Generic Control Protocol Principal
GHz	gigahertz
GPON	gigabit passive optical network
HD	high definition
HFC	hybrid fiber/coax
HSD	high speed data
Hz	hertz
I	in-phase
IP	Internet Protocol
ISDN	integrated services digital network
LDPC	low density parity check
kbps	kilobits per second
MAC	media access control
MDU	multiple dwelling unit
MER	modulation error ratio
MHz	megahertz
MIB	management information base
MPEG	Moving Picture Experts Group
MPM	modulation profile management
MTTR	mean time to repair (sometimes mean time to restore)
NC	narrowcast
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
ONU	optical network unit
OOB	out-of-band
OSS	operations support system
OTN	optical transport network
PHY	physical layer
PMA	profile management application
PNM	proactive network maintenance
PON	passive optical network
POST	power-on self-test

Q	quadrature
QAM	quadrature amplitude modulation
RF	radio frequency
RMA	return material authorization
RMD	remote MAC-PHY device
RPD	remote PHY device
R-PHY	remote PHY
RPN	remote PHY node
SC-QAM	single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
SD	standard definition
SDV	switched digital video
SNMP	Simple Network Management Protocol
SNR	signal-to-noise ratio
SW	software
US	upstream
VBR	variable bit rate
vCMTS	virtualized cable modem termination system
VCS	video compression standards
VOD	video on demand
WDM	wavelength division multiplexing