



**VIRTUAL EXPERIENCE
OCTOBER 11-14**



Distributed Access Architecture Is Now Widely Distributed – And Delivering On It's Promise

A Technical Paper prepared for SCTE by

Dr. Robert Howald

Fellow
Comcast
1800 Arch Street, Philadelphia PA
robert_howald@comcast.com

Frank Eichenlaub

Executive Director, Access Network Technology
Comcast
frank_eichenlaub@comcast.com

Tobias Peck

Director of Product Management
EnerSys
tobias.peck@enersys.com

Adi Bonen

Executive Director
Comcast
frank_eichenlaub@comcast.com

Table of Contents

| Title | Page Number |
|--|-------------|
| 1. Introduction..... | 4 |
| 2. Why DAA?..... | 4 |
| 2.1. Key Benefits | 5 |
| 2.2. DAA as a Path to Virtualization | 7 |
| 3. DAA-Powered Next Generation Access | 8 |
| 3.1. Architecture Migration | 8 |
| 3.1. Spectrum Migration | 9 |
| 4. DAA Node Design Principles | 11 |
| 4.1. Leveraging a Connected Platform | 11 |
| 4.2. New Powering Challenges | 12 |
| 4.2.1. Node Power Design Considerations | 12 |
| 4.2.2. Node Power Efficiencies | 12 |
| 4.2.3. Practical Considerations and Solutions | 13 |
| 5. Network Powering Implications for DAA | 14 |
| 5.1. Example Power Interruptions Example Power Interruptions | 15 |
| 5.2. Design and Processes for DAA Power Resilience..... | 16 |
| 6. DAA Production Findings | 16 |
| 6.1. Fidelity Metrics | 17 |
| 6.2. Customer Experience Metrics | 19 |
| 6.3. Network Power Consumption Metrics..... | 20 |
| 7. Conclusion..... | 25 |
| Acknowledgments | 26 |
| Abbreviations | 26 |
| Bibliography & References..... | 26 |

List of Figures

| Title | Page Number |
|---|-------------|
| Figure 1 - Distributed Access Architecture Variants: RPHY, RMD, and R-CCAP | 5 |
| Figure 2 - DAA Enables Access-Agnostic IP Network Architecture Convergence | 7 |
| Figure 3 - DAA and Alignment to a Path Towards Virtualization and vCMTS | 8 |
| Figure 4 - Breakdown of Upgrade Approaches by Tier, Criteria, and Technology | 9 |
| Figure 5 - Migration Options for HFC Spectrum Allocation..... | 10 |
| Figure 6 - Temporary Power Packs for DAA Nodes Simplify Cutover..... | 13 |
| Figure 7 - Deployment Model of Maintenance Power Supply for DAA Node | 14 |
| Figure 8 - Use of Power Straps in HFC Networks | 14 |
| Figure 9 - AC Interruption to Node Power Supply | 15 |
| Figure 10 - AC Interruption to Node DC Power Supply Caused by Tap Plate Removal | 15 |
| Figure 11 - AC Interruption to Node DC Power Supply Caused by Short at Seizure Screw..... | 16 |
| Figure 12 - Downstream Rx MER Comparison: N+x HFC, N+0 Analog, and N+0 DAA | 17 |
| Figure 13 - Upstream Rx MER Comparison: N+x HFC (top) and N+0 DAA (bottom)..... | 18 |
| Figure 14 - Operational Metrics Proceed to Maturity as the DAA Platform Grows in Scale..... | 20 |
| Figure 15 - Pre/Post Input Power Comparison by Original Node Boundary..... | 22 |



**UNLEASH THE
POWER OF LIMITLESS
CONNECTIVITY**
VIRTUAL EXPERIENCE
OCTOBER 11-14



Figure 16 - Pre/Post Output Power Comparison by Original Node Boundary 22
 Figure 17 - Pre/Post Input Power Comparison by Power Supply 23
 Figure 18 - Pre/Post Output Power Comparison by Power Supply 23
 Figure 19 - Pre/Post Input Power Comparison Consolidated 24
 Figure 20 - Average Battery Backup Time Pre/Post Comparison 24

List of Tables

| Title | Page Number |
|---------------------------------|--------------------|
| Table 1 - Node Statistics | 21 |

1. Introduction

A large-scale production Distributed Access Architecture (DAA) footprint in terms of both homes connected and the number of digital nodes is being built by Comcast right now. Comcast began deploying DAA in 2018, focusing on the mature, multi-vendor Remote-PHY (RPHY) standard. The company continues to aggressively deploy DAA across all regions. A core premise of RPHY is that because it is a CableLabs standard, devices from different manufacturers of CMTS and Digital node are interoperable. Comcast has succeeded in delivering on this promise with Digital Nodes from three technology partners connected to its virtual CMTS (vCMTS) platform.

Many of today's DAA deployments are in areas with an N+0 (node plus zero amplifiers) architecture, which is sometimes referred to as "Fiber Deep." There is also an existing footprint of N+0 based on traditional analog fiber nodes. In both cases, there are no actives between the fiber node and customer homes. This creates the perfect opportunity to make A/B comparisons between the best-in-class analog HFC and DAA technology. Without RF amplifiers "watering down" any differences between the two networks, such as might exist in N+x deployments, we can attain a clear and unvarnished, side-by-side view of yesterday's and today's cable access technologies in production scale.

In this paper we will compare observations and numerical results obtained from operating these two N+0 variants. We will:

- 1) Describe lessons learned building both architectures, and adjustments made along the journey
- 2) Describe the transitional operational challenges and best practices
- 3) Quantify key performance metrics
 - a. EOL MER (End of Line Modulation Error Ratio)
 - b. Capacity
 - c. Network power consumption – a critical area of focus for 1.2GHz DAA over N+0 in particular

This paper will address the journey, results, and future expectations of the industry's largest known initiative to take DAA from concept to reality..

2. Why DAA?

A Distributed Access Architecture (DAA) breaks the CMTS function into two components and distributes one part of it into the HFC network, typically the fiber node. The functionality distributed has narrowed to two DAA options. These are Remote PHY (or RPHY) and Remote MAC-PHY (known interchangeably as "RMD" for Remote MAC Device, or FMA = Flexible MAC Architecture). Figure 1 depicts the variations. Also shown is the fully Remote CCAP (R-CCAP) approach, which has faded from consideration among most cable operators due to the complexity.

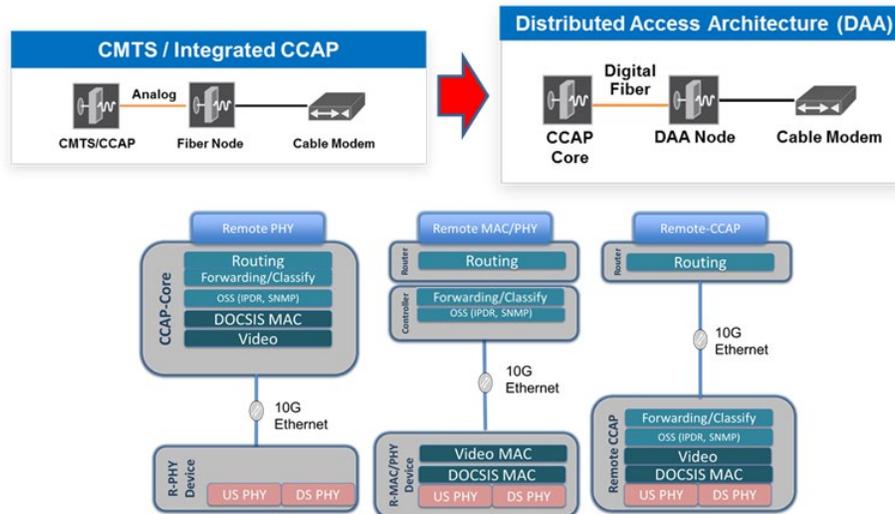


Figure 1 - Distributed Access Architecture Variants: RPHY, RMD, and R-CCAP

Remote PHY is currently the dominant deployed architecture, with tens of thousands of RPHY DAA nodes deployed in production networks in the US market alone, while only a relative handful of RMD devices have been installed. The FMA specification governing RMD interoperability got off to a later start, and this architecture takes on additional HW and SW complexity at the node. These two factors and forward-looking capacity advantages are contributing to the slower uptake of this technology. Whether this technology is brought to market in volumes will be determined over the next few years as more MSOs embrace DAA in their next generation architecture plans.

2.1. Key Benefits

Regardless of the amount of DOCSIS processing functionality distributed, DAA delivers powerful advantages associated with the standard digital Ethernet optics that are fundamental to the architecture. Until DAA came along, every node was implemented with Amplitude Modulated (AM) optics, created by, and unique to, the cable industry. Until that time, fiber optics had been based on binary digital transport only. A DAA node, fed by digital optics, offers significant network, infrastructure, and performance benefits over traditional AM-based optical links.

These chief advantages are:

1) *DWDM Wavelength Efficiency*

The use of digital optics means that there is a path to 80 DWDM wavelengths/nodes on a single fiber. Typical HFC AM optics are restricted to 16 wavelengths, depending on reach, to manage nonlinear effects that degrade MER. More than 16 are possible but they come with compromises in performance and distance. These nonlinear effects are much less significant in digital optics.

2) *Link Reach*

Using AM optics, as the link from Headend to node increases, fewer wavelengths can be used to meet a fixed end-of-line (EOL) MER requirement. HFC networks are constructed around these limitations. Digital optics eliminate this dependency for typical HFC reach requirements and enable link distances much longer than AM optics. The above aspects of reach and wavelength multiplication taken together simplify

and add significant flexibility to network architecture, promote cost savings due to the reduced fiber construction entailed, and enable the potential for consolidation of physical sites.

3) *EOL MER performance improvement*

DOCSIS 3.1 enables higher order modulation profiles, increasing the bandwidth efficiency up to 50% over DOCSIS 3.0 downstream and 67% in the upstream. Typically, EOL MER in HFC is dominated by the performance of the AM optics and degrades a little with each ensuing amplifier. In an Integrated CMTS, the fidelity requirement at the RF port feeding the optical transmitter is very high – 43 dB minimum for DOCSIS 3.0 and up to 48 dB for DOCSIS 3.1. The AM optical link then degrades the MER delivered to the node. It is a common node requirement to achieve a 38 dB minimum MER at the RF output port. This represents a significant fidelity loss of 5-10 dB. The MER degrades more as it is passed through HFC amplifiers.

Implementing DAA eliminates this AM Optical MER degradation situation completely. Instead, the CMTS fidelity requirement is met at the RPHY device (RPD) RF interface inside of the node. This increases the margin available to deliver DOCSIS 3.0 capacity, while increasing the capacity available for DOCSIS 3.1.

A similar situation exists for the upstream. Elimination of HFC upstream technology, whether an analog or digital return, occurs by placing the DOCSIS US Receiver in the node within the RPD, which recovers significant (67%) upstream MER.

4) *Space, power, cooling efficiencies in the Hubs and Headends*

Moving CMTS functionality to the node leaves less Headend components to power, cool, and consume space in a Hub site. Some of this power is distributed into the plant, which, as we shall see, quantifiably can provide a net positive return to the outside plant (OSP) due to the overall technology upgrade over the equipment it is replacing.

With respect to space, with the RF ports of the CMTS distributed, the connector density of the DOCSIS processing core can now be bounded by a higher density of optical connectors.

As traffic continues to grow, and new ports and computational power are necessary to support new demand, the incremental increase in footprint, power, and subsequent cooling needed to deliver the necessary processing power can be more granular, in particular as DAA leads to a virtualized processing core.

5) *Alignment with virtualization and convergence across last-mile access technologies*

As the HFC architecture evolves to DAA, nodes are more readily adaptable to other last mile access technologies that leverage Ethernet connectivity. Wireless, PON, and direct Ethernet services are all applicable, enabling convergence opportunity in the processing core.

Figure 2 shows the simple example of a Passive Optical Network (PON) Remote Optical Line Terminal (R-OLT) module. The concept is similar in principle to a plug-in RPD module, distributed to and backhauled from the same Ethernet-connected DAA architecture to deliver last mile residential and business services, but over an FTTH network.

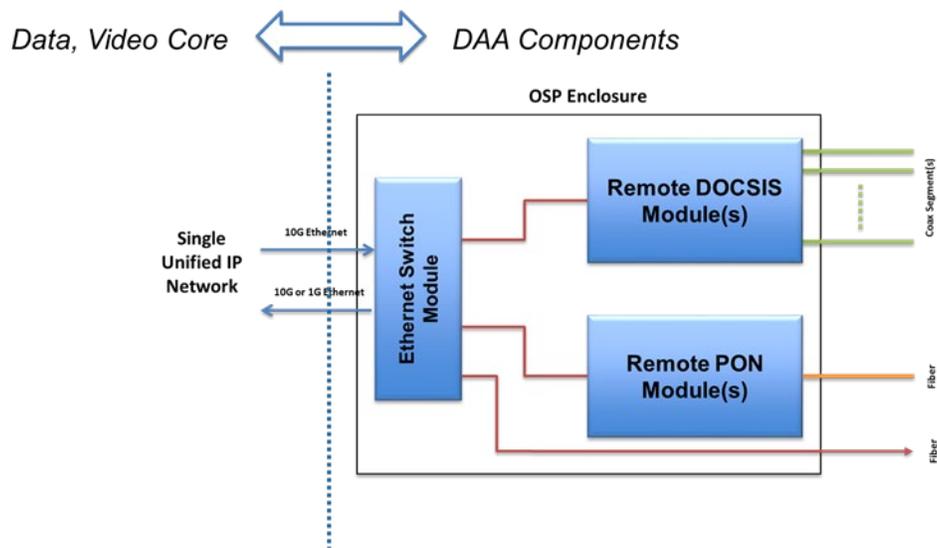


Figure 2 - DAA Enables Access-Agnostic IP Network Architecture Convergence

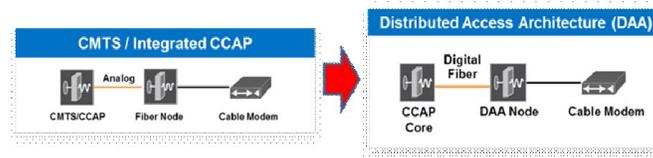
2.2. DAA as a Path to Virtualization

DAA can be considered the starting point of a larger shift in the HFC technology ecosystem away from purpose-built cable hardware where possible, and embracing the broader trends towards digital processing, software, and cloud-based services and applications. With the CMTS packet processing function in the headend separated from major DOCSIS-specific functions now implemented in the RPD, the CMTS can be revisited in the context of packet processing, switching, storage, and its DOCSIS scheduling role, all of which can be translated into virtualized real-time computational resources. With today's compute power, and the capability of software to deliver real-time services, a purpose-built DOCSIS machine is no longer required.

Instead, commercial off-the-shelf (COTS) servers and switches, combined with software running the required DOCSIS MULPI layer, can be used to implement the CMTS function. This leads to a simplified and more flexible virtualized CMTS platform, or vCMTS. Significant benefits can be leveraged over time through Moore's Law, delivering ever more compute power in increasingly dense physically dense footprints. Comcast introduced its first production DAA deployment hosted by a vCMTS in 2018, and since that time every DAA node has been connected to the Internet via a vCMTS, as the architecture has rapidly scaled.

Figure 3 conceptualizes DAA and its alignment to virtualization and vCMTS. As noted in the figure and pertaining to the idea of an access-agnostic last mile enabled by DAA, virtualization can apply to DOCSIS or PON / FTTH last mile, or other natural alternatives such as wireless. The virtual platform itself is not access-specific when it comes to the switching and routing of packets, simplifying the path to convergence of access technology at the packet processing layer when the access technology specific-SW is abstracted.

Distributed Access



Virtualized and Distributed Access - Disaggregation

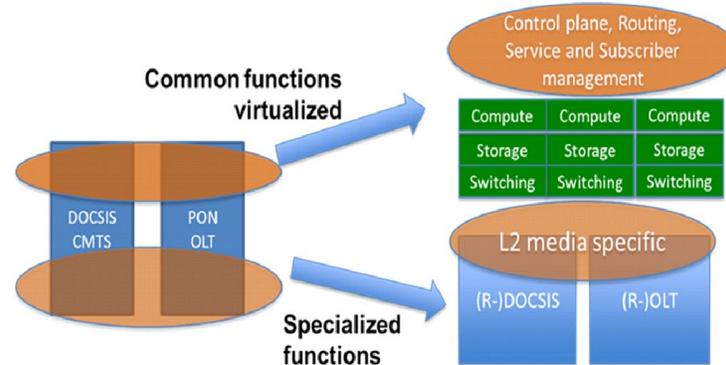


Figure 3 - DAA and Alignment to a Path Towards Virtualization and vCMTS

3. DAA-Powered Next Generation Access

3.1. Architecture Migration

Like most MSOs, Comcast is constantly assessing its access network upgrade strategy, adapting accordingly to changes in services, customer demands, traffic, and business priorities. The environment has never been more dynamic and perhaps – as the pandemic-related impacts blend into a new normal – never less certain. And, like all MSOs, while many network tasks are well-exercised muscles that are developed over years of experience, upgrades effecting the access network are generally not quick to implement nor simple tasks. Adding to the degree of difficulty is the need to make these upgrades and transitions as seamless as possible to customers.

Given that access network upgrades are multi-year initiatives, there is a need to prioritize accordingly. There is also a need to recognize that the pace of upgrades plays a role in applicable “tools in the toolbox.” Considering the relatively long-term nature of an upgrade cycle, it is important there are not “starving” areas of the network left unattended until the upgrade comes along. No part of the network can afford to stand still. While it varies across the footprint, the pace of traffic growth and service demand outpaces upgrade pace, forcing repeated touches. This limitation of physical hardware and construction is also one of the drivers for introducing more software into the network through DAA and virtualization, which we will discuss shortly. Lastly, it is important that there be simplicity and repeatability in the operation to make these upgrades seamless, efficient, and delivering the benefits expected.

With the above dynamics in mind, a multi-pronged strategy was developed to make sure network upgrades are in-tune with growth and service requirements, as they vary by region, and the projected timing and scale at which the network must be upgraded to do so. **Figure 4** outlines a strategy based on these objectives.

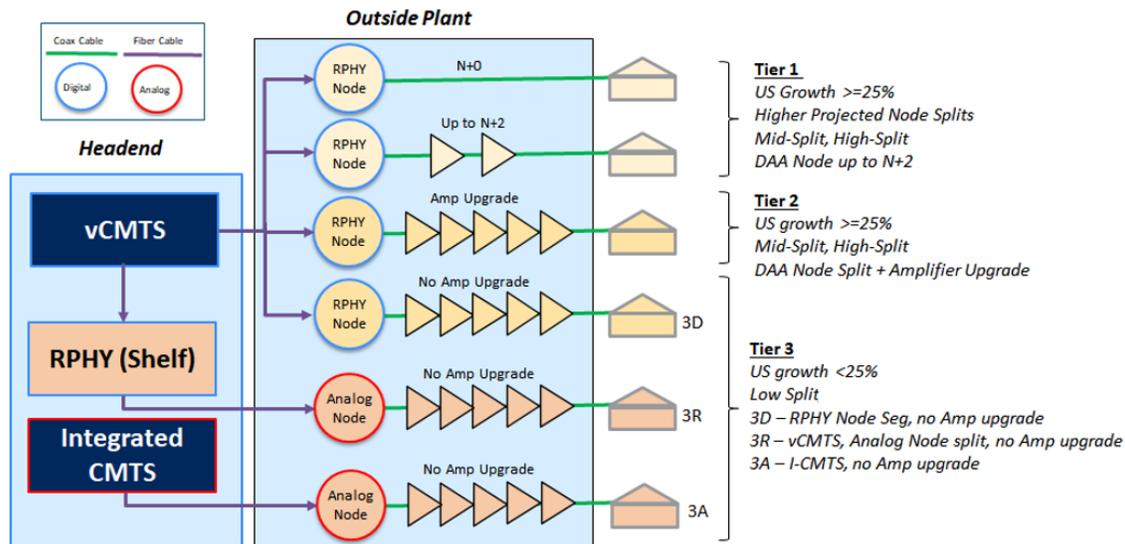


Figure 4 - Breakdown of Upgrade Approaches by Tier, Criteria, and Technology

In **Figure 4** we can observe that, across the roughly 60M homes and businesses passed by the Comcast network, the variations have been categorized into three Tiers of defined activity. Within the three Tiers, five migration options depend on various criteria and existing conditions. These are summarized as follows:

Tier 1 – The very highest HSD utilization and highest growth areas require the most aggressive change to network to stay ahead of demand – pulling new fiber deeper into the network, maximizing the downstream available bandwidth out to 1.2GHz, and decreasing the node size to a maximum allowable homes past while building more headroom into the capacity runway, deferring subsequent upgrades.

Tier 2 – High utilization in % growth, although not as high overall as the Tier 1 category. This approach defers subsequent upgrades for at least five years. For this category of network, this deferment can be accomplished with less aggressive means and can be done more quickly and more broadly by continuing to use node splits, but in doing so convert the network to vCMTS and DAA, while adding spectrum downstream and upstream via RF amplifier upgrades.

Tier 3 – Lower upstream growth and utilization than Tier 1 and Tier 2. Splitting of nodes “BAU” is sufficient to stay ahead of demand. However, while doing this work, begin introduction of key new enabling technologies of Tier 1 and Tier 2 (DAA and vCMTS) to place the foundation for a subsequent technology upgrade wave via Tier 1 or Tier 2 when necessary.

3.1. Spectrum Migration

As mentioned above, while migrating the access network to DAA, taking the opportunity to adjust the spectrum allocation is a practical and efficient way to increase long-term capacity. In addition, in order to provide higher upstream speeds and service tiers, a wider upstream path is required to deliver the instantaneous capacity necessary. With the commitment to install new equipment into the access network, it is the ideal opportunity to ensure that this equipment has a long runway of traffic capacity.

Figure 5 depicts stages or options available within the DOCSIS standards by which the upstream and downstream can be defined. It also indicates the HSD speeds available with these implementations for residential services under typical serving group sizes and capacity utilizations. Cable operators have been living in the “Low Split” (5 MHz – 42 MHz)_world for decades and have now begun moving to these higher upstream spectrum options.

While the Mid-Split architecture (upstream extended to 85 MHz) will cover an extended capacity runway for typical node sizes, its maximum speed, when carrying DOCSIS 3.0 traffic in the Low Split band, will be limited to about 300 Mbps. This increases to about 500 Mbps with an all-OFDMA upstream in the years ahead as DOCSIS 3.1 penetration continues to grow and eventually dominates the available bandwidth.

Both High Split (with an upper boundary of 204 MHz) and DOCSIS 4.0 Full Duplex (10G) provide the opportunity to offer, respectively, Gigabit and Multi-Gigabit symmetric services. Not shown below is a DOCSIS 4.0 Extended Spectrum DOCSIS (FDD) spectrum allocation, which, like FDX, can take the upstream as high as 684 MHz. However, as a Frequency Domain Duplex (FDD) architecture, it requires replacing all actives and passives in the plant, to extend the network to 1.8 GHz. That’s necessary to provide sufficient downstream spectrum to offset the loss due to the added upstream allocation, plus the increased diplex crossover guardband.

By contrast, FDX uses common spectrum for both downstream and upstream to maximize efficiency in this known, mature, well-behaved RF spectrum.

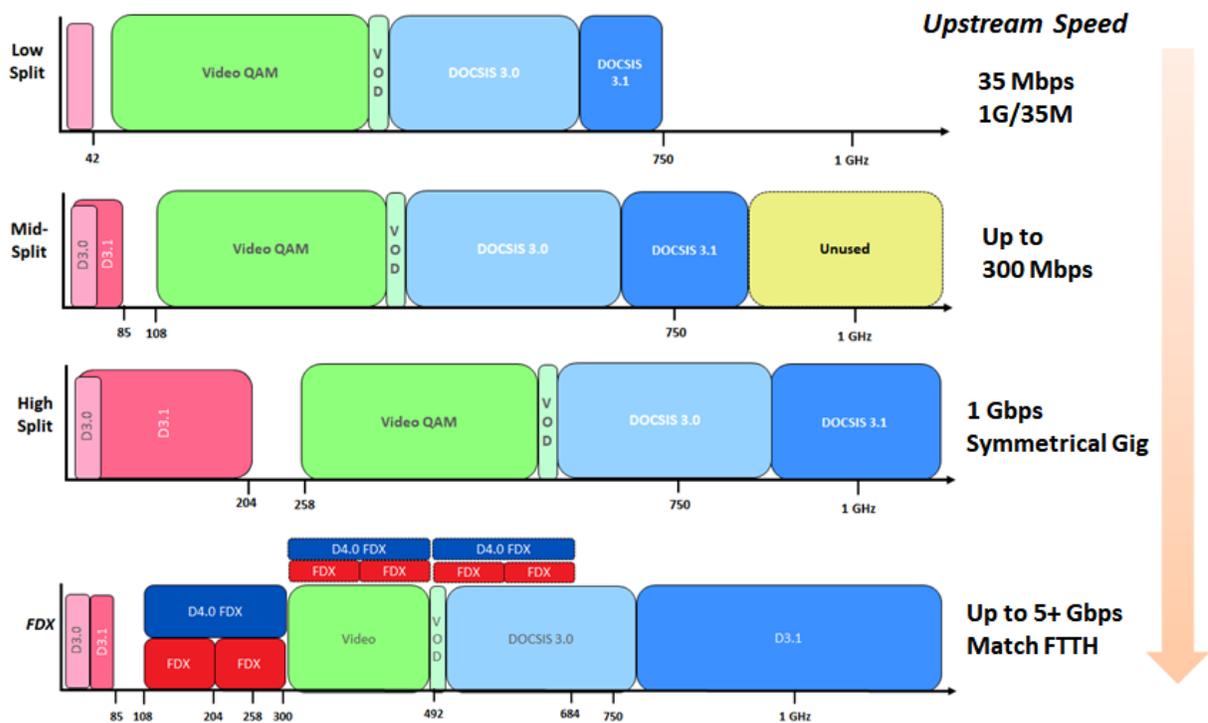


Figure 5 - Migration Options for HFC Spectrum Allocation

4. DAA Node Design Principles

4.1. Leveraging a Connected Platform

A DAA node has many of the expected essential features of an HFC node on the coaxial side of the device. It must place high fidelity RF signals onto the coax to reach the customers connected to its port, and it must receive RF upstream signals from those same customers. However, a “BAU” HFC node is a “dumb” transducer device, converting an analog optical waveform on fiber to an RF waveform on coax, and the opposite in the return path. Digital return upstream systems are a nuanced exception – they translate the RF waveform, but do so by digitizing it into a series of numbers and transmitting those numbers digitally before re-creating the waveform at the end of the optical link.

By contrast, a DAA node’s RPD is a smart and IP-connected device. The DAA node design can take advantage of this connectivity. To start with, common node functions can be accommodated without the RF plug-ins that are often used in traditional HFC nodes. These plug-ins can tempt technicians, resulting in excess tweaking flexibility that can be misused without careful guidance. This is particularly the case for attenuation pads and tilt equalizers. Furthermore, local control using plug-ins results in loss of traceability, and unknown equipment settings. Instead, a DAA node can implement remote configuration of gain and tilt, enabling centralized power profile management, including outage-free spectrum enhancement and lineup changes. This provides state visibility and tracking, and can be enabled (or not) to technicians at the expertise depth of the control desired.

In the RF chain leading to and from the node ports, a DAA node that provides per node-port US attenuation can be controlled to identify a port experiencing an abnormally large ingress, minimizing triage and recovery. Moreover, if the RF launch amplifier includes an US RF switch matrix, mitigation schemes can be used to minimize the group of subscribers experiencing service reduction. For example, in a 4-port node segmented to 2 US service groups, the node port suffering ingress can be segregated to be on its own US service group while the other 3 node ports are combined to share the other US service group without interference from the ingress. This preserves the highest possible level of service for the most customers until the issue can be resolved.

In terms of inventory management, a “smart” DAA node can be designed to enable remote module inventory by reporting the model and serial numbers of every pluggable module in the node as well as other factory parameters. In addition, status monitoring of various module parameters enables tracking the component and node health.

Lastly, and perhaps the most powerful benefit of a DAA architecture and as previously described, is introducing the ability to easily alternative last-mile access technologies, such as PON, Ethernet services, or wireless. DAA can go beyond just DOCSIS-based RPHY and RMD nodes. Devices implementing an outdoor switch (aka “Switch On A Pole” or SOAP) functionality can provide OSP IP-link aggregation for other OSP devices, as well as access points for Metro Ethernet subscriber lines. A SOAP device can share the same DAA node with an RPD or be independently housed in an RF-less node housing. In both cases, the enclosure for the SOAP device should have similar traits to a “regular” DAA node.

Moreover, as MSOs deploy more FTTH PON to their customers, a Remote-OLT device housed in an OSP DAA node is becoming many MSOs’ preferred approach. It provides both success-based deployment granularity, and serves as a more localized intermediate point from which to launch fiber-budget limited PON signals that may not be able to be passively delivered from the site where a typical chassis-based OLT would reside for a cable network. Similar to the Ethernet case, whether the R-OLT device shares the same DAA node with an RPD or is independently housed in an RF-less node housing, the DAA node should utilize the same key principles of resilient powering, remote monitoring, and control.

In short, the smart and connected nature of the DAA node, in particular the use of standard Ethernet optics for transport and virtualization, opens up a wealth of new architecture flexibility and operational benefits to DAA nodes that cannot be achieved in a typical HFC node.

4.2. New Powering Challenges

DAA nodes can drive an already existing N+x HFC system, be part of the more aggressive “Fiber Deep” migration that removes RF amplifiers (N+0) and extends downstream and upstream spectrum, and of course can also support anything between those two boundaries of plant state. The most challenging of these cases is the DAA node that powers the N+0 system, because these systems rely on higher RF output power to maximize reach, have the highest tilt to cover the longer length of coaxial cable that accompanies this reach, extends the downstream spectrum to 1.2 GHz, driving the Total Composite Power (TCP) higher than any previous node. It subsequently draws the most AC power to achieve these goals. Therefore, if a DAA node can support fiber deep, it will support all possible N+x variants it is likely to be deployed in with different RF configurations aligned to those architectures.

4.2.1. Node Power Design Considerations

The upgrade of a traditional N+X HFC system to a N+0 DAA requires paying even closer attention to plant powering. The DAA nodes that effectively replace (not 1:1) existing amplifiers are likely to consume significantly more power, and care should be taken not to surpass the line power supply’s rating or the hardline current carrying limit. DAA nodes can implement several key features to help mitigate such issues as part of an upgrade.

In a typical older N+X system, power is passed from a line power supply to a node, and through it to multiple amplifiers. A node is rarely fed power which first passes through another node. If, as a result of bringing down a node for maintenance that power to the amplifiers is also lost, no additional subscribers are affected. However, in N+0 systems with smaller node domains, a single line power supply often feeds multiple nodes, often passing through one node to reach another. In such cases, it is very desirable to maintain uninterrupted power passing through a node even when it is undergoing maintenance (such as a module replacement), such that a service interruption due to maintenance is not extended to other node domains that are not undergoing maintenance.

Note that while the customer-facing blast radius may be similar from a numerical perspective because of the smaller number of customers per N+0 node, DAA nodes are less tolerant to short outages, due to the possibility of a reboot being triggered, than their traditional HFC node or amplifier counterparts. In addition, node outages create new software challenges in the back office for state management and recovery that could be otherwise avoided with thoughtful OSP design practices. Accordingly, it is desirable for fiber deep nodes to support power passing even while their internal modules are swapped.

4.2.2. Node Power Efficiencies

In a node’s RF sections, most of the power is consumed by the output hybrid amplifiers. However, due to the insertion loss of a 4-way split as well as gain and tilt temperature compensation circuitry normally built into such an amplifier, it typically requires significant power to be allocated to drive these output hybrids. A DAA node with a smart RPD can significantly reduce this power requirement by implementing gain and tilt control in the RPD itself. Similarly, in the US, including configurable step attenuators prior to the return amplifiers enables the RPD to keep tight control on the signal level at these amplifiers, thus reducing the dynamic range requirement for them, allowing them consume significantly less power. In an N+0 node, about 15~20 watt can be saved with such schemes, which is significant savings for a device that in the 100-140W range, depending on configuration.

Power delivery of both line power supplies and hardline cable is limited by delivered current, not delivered power. Implementing power factor correction (PFC) in the DAA node, which aims to align the phase of the AC voltage and current waveforms, reduces the quasi-square-wave current at the given power it consumes. Therefore, a DAA node implementing power factor correction enables more power to be delivered by the existing system. Moreover, by implementing power factor correction to reduce the quasi-square-wave current, additional efficiency is achieved by reducing the power loss on the hardline itself due to its loop resistance.

4.2.3. Practical Considerations and Solutions

One of the lessons learned in early DAA deployments is the importance of immunity to short power interruptions, as mentioned previously. The HFC plant is prone to split-second power interruptions, with most happening as a result of tap faceplate removals. In a traditional node such an interruption likely results in a split-second signal loss to customers, an event which is hardly noticeable. However, in a DAA node this can cause the RPD to reboot, extending the signal loss to several minutes. A solution we Comcast utilizes a capacitor-based scheme which can maintain power to the RPD for a few seconds during power interruption. Unlike battery backup, capacitor-based backup is maintenance-free and its lifespan surpasses the node's.

Another lesson learned is the need for a temporary node powering option *during* installation, burn-in and provisioning. Upgrading a traditional HFC plant to a DAA system typically involves installing a new DAA node in parallel to an active old node, burning in the new DAA node, provisioning it, and only then cutting it in instead of the old node. Often, there is not enough available power in the active system to power the DAA node on top of the existing system, and thus typically a generator is used to power the DAA node during these stages. However, it is desirable that the node be independently ready for plant power in advance, without causing interruption when moving from temporary generator power to permanent plant power. A solution we used involves a special temporary-powering power pack (PP) that is mounted in the second PP position in the node. Examples of the temporary power packs are shown in **Figure 6**.

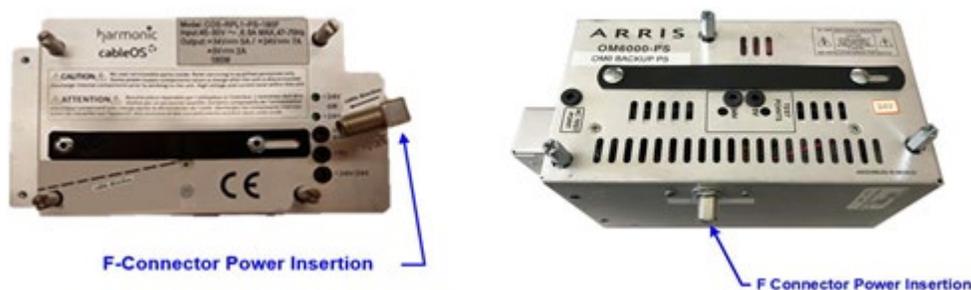


Figure 6 - Temporary Power Packs for DAA Nodes Simplify Cutover

The temporary PP drives the node's internal DC power rails as any PP, but gets its QSW power through a directly connected coax from an external generator, and not through the plant's and node's permanent power. During cutover, when the primary PP first gets its power from the plant, a node is operated briefly with redundant power (permanent and generator), and thereafter the temporary PP is removed (to be reused in the next node to be installed). **Figure 7** illustrates this temporary installation arrangement.

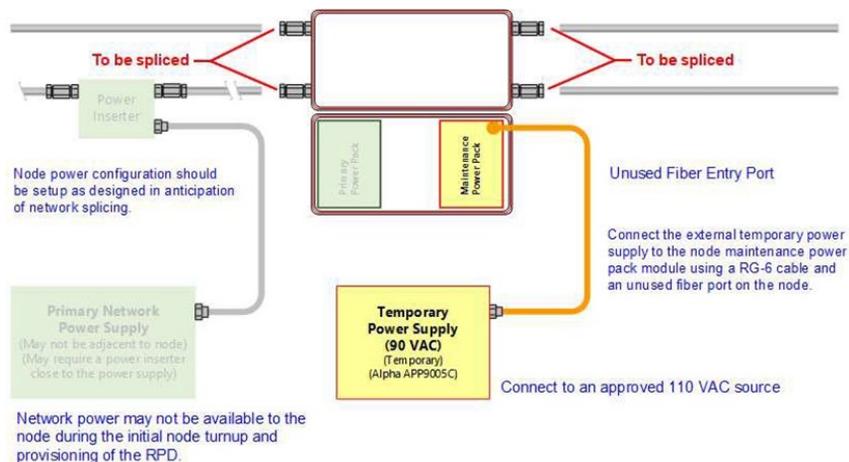


Figure 7 - Deployment Model of Maintenance Power Supply for DAA Node

5. Network Powering Implications for DAA

In a typical Comcast N+X deployment, it is common to have three power supplies providing AC to the network actives. Management of the distribution of power is provided by blocking power at strategic points to best utilize the available power. In an N+0 architecture there are three primary methods of distributing power to the nodes and other active devices (non-amplifiers) attached to the plant:

1. Centralized power using parallel feeds of 0.875” or 23-ohm power hardline cable to distribute power independently of the RF distribution.
2. Utilize power straps to connect node service areas to existing power supply locations.
3. Deploy independent, low wattage power supplies for each node service area.

There are advantages and disadvantages to each power architecture. To minimize the amount of parallel cable placement, permitting and construction activities, power straps were used to distribute power in the Comcast N+0 network. Simply put, a power strap is a power-only connection using 0.875” cable or 23-ohm power cable to route power from node area to node area. **Figure 8** is a simplified drawing of a power strap.

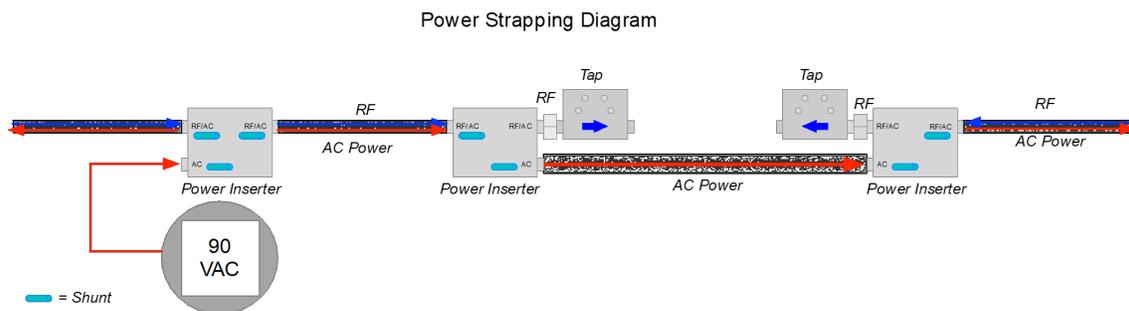


Figure 8 - Use of Power Straps in HFC Networks

Power straps provide a low-cost solution to manage power in a N+0 architecture but introduce a more complex network for maintenance teams to manage, troubleshoot and maintain. In addition to the potential for power glitches previously described, the common practices of pulling fuses or shunts to isolate plant issues now can shut down many nodes at once. Because of this, the power network design for DAA requires some additional levels of resilience to ensure a more robust experience.

5.1. Example Power Interruptions Example Power Interruptions

Typical holdup time for the AC power packs in DAA nodes extends from about 75ms to 250ms. The holdup time is defined as how long the power supply can maintain the DC voltage in the node after the input AC is interrupted. Loss of DC to the RPD initiates a reboot that can take many minutes to re-establish service to devices in the serving area. **Figure 9** is a capture of an oscilloscope displaying the impact of loss of input AC waveform and the ability of the node power supply to hold up the DC voltage.



Figure 9 - AC Interruption to Node Power Supply

Several common maintenance practices can also interrupt the AC waveform. Replacement of a tap faceplate or shorting a nut driver while tightening a seizure screw can occur during maintenance. **Figure 10** illustrates examples where a tap faceplate removal or installation can impact the AC waveform enough to cause an RPD to reboot.

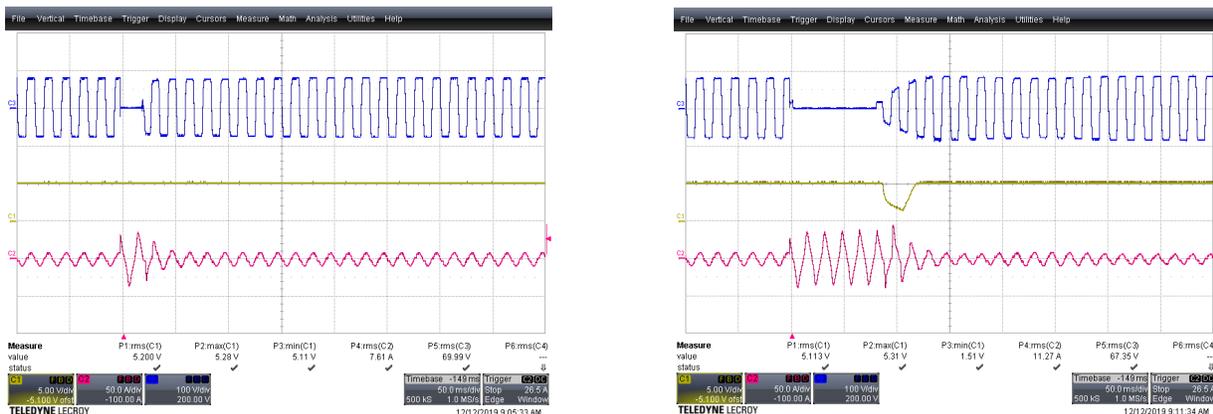


Figure 10 - AC Interruption to Node DC Power Supply Caused by Tap Plate Removal

In the example on the left, the faceplate was removed with little interruption to the AC waveform and no RPD reboot occurred. The example on the right demonstrates where a tap faceplate change took longer,

and the DC voltage was not able to be maintained to the RPD. In this case, an RPD reboot would be initiated.

AC interruptions impacted the nodes downstream of where the maintenance was being performed. Shorts, however, impact all nodes being fed from a common network power supply. **Figure 11** is an example of a technician accidentally shorting a nutdriver while tightening a seizure screw. As shown, there are several interruptions to the AC waveform eventually leading to an RPD reboot.

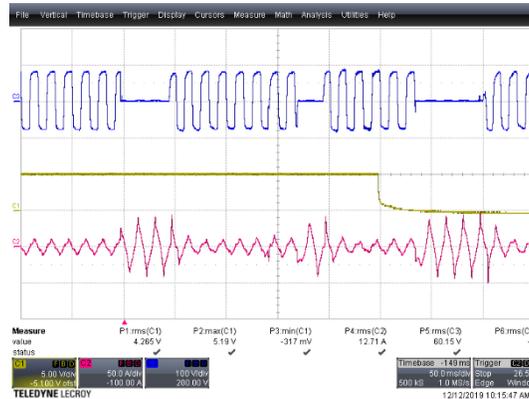


Figure 11 - AC Interruption to Node DC Power Supply Caused by Short at Seizure Screw

5.2. Design and Processes for DAA Power Resilience

To reduce the impact of normal maintenance procedures on the DAA nodes and improve the customer experience, several elements of the RPD deployments were focused on to improve the resiliency of the power network.

The impact on power from network passives and taps was reviewed. Currently available network passives are designed such that when the faceplate is removed, RF and Power continuity is interrupted. This is an undesirable characteristic in DAA deployments. Comcast is working with vendors to redesign the network passives to allow for cover removal and troubleshooting without impacting RF or Power to all ports.

In addition, taps were evaluated to measure their make-before-break characteristics. These characteristics vary by manufacturer and specific tap design. Processes were established to ensure that taps have no more than an 8 ms impact on the AC waveform when removing or installing a faceplate.

From a network powering point of view, some power supply models can be deployed with a fault-isolating dual output controller option that can limit the impact of a short. This controller allows the nodes to be split into two fault-isolated groups. When a short is detected on one leg, the device isolates that leg from the others in order to protect one group from the effects of the short. To fully utilize the feature some additional coax should be placed to maximize the coverage of the power supply.

6. DAA Production Findings

As of June 30th, Comcast has thousands of DAA RPHY nodes deployed, with the majority of these in an N+0 architecture, servicing over 1.5M homes passed across about 10k miles of plant. Beginning in 2021, DAA nodes are being deployed in traditional HFC architectures to accelerate capacity expansion and vCMTS footprint. These will number in the thousands this year and will consist of nodes with a single or two RPDs installed.

6.1. Fidelity Metrics

As noted in the overview describing the advantages of DAA, one of the key benefits is End-of-Line (EOL) MER. By eliminating the analog optical link, the MER loss that accompanies that fiber connection is eliminated. The MER generated by the RPD is essentially the same as the MER that would be delivered from a legacy I-CMTS port in a Hub site prior to traversing the AM Optical link into the field. It is therefore a high-fidelity signal placed onto the coax by the DAA node.

With the large scale of production of DAA nodes in the field there is now sufficient statistical sample size to do A:B MER comparisons among BAU HFC, N+0 with analog nodes, and DAA nodes.

Figure 12 shows distributions of Rx MERs of DOCSIS 3.1 OFDM for “BAU” HFC networks (top), analog N+0 networks, and DAA nodes (bottom).

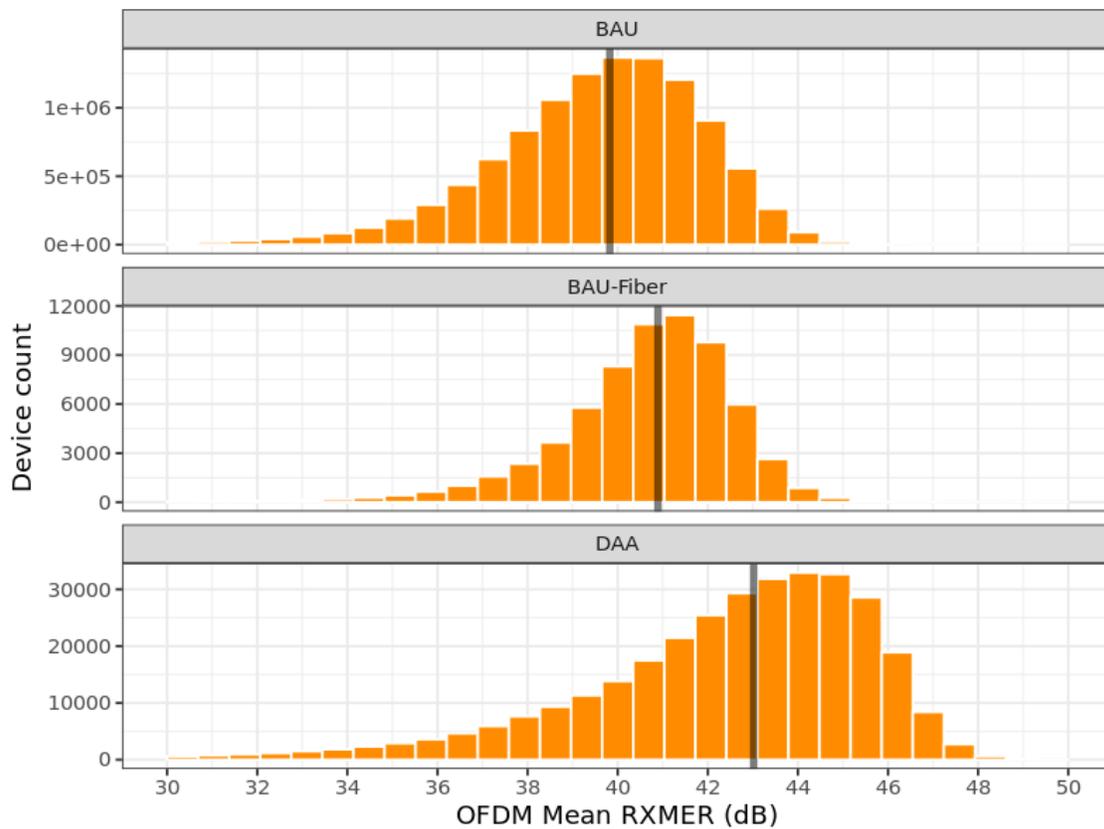


Figure 12 - Downstream Rx MER Comparison: N+x HFC, N+0 Analog, and N+0 DAA

Figure 12 trends exactly as would be expected. Analog N+0 has a better EOL MER than BAU HFC N+X MER, and DAA N+0 has a better EOL MER than Analog N+0 MER. There is a noticeable and quantifiable MER improvement in the bottom plot of **Figure 12**. The average MER increases with Analog N+0 over HFC N+X is a little over 1 dB. The MER gain from HFC N+X to DAA N+0 is over 3 dB. A 3 dB increase is equivalent to one modulation order of efficiency, such as the difference between the fidelity needed to support 512-QAM, for example, and 1024-QAM, which is a more efficient modulation in bps/Hz.

The numbers shown here are measured at the CM downstream receiver, and as such include all channel impairments up to that point, including the Noise Figure of the CM itself. The minimum DOCSIS 3.1 MER requirement for 4096-QAM is 41 dB at -6 dBmV/6 MHz. The minimum requirement for 2048-QAM is 37 dB @ -9 dBmv. Thus, these distributions would suggest that the DAA N+0 case will enable the 4096-QAM modulation profile *on average*, while for non-DAA there is a bandwidth efficiency penalty of about 9%. Whereas MER differences did not translate to anything except for performance margin in DOCSIS 3.0 systems, the use of DAA can add nearly 200 Mbps per OFDM block allocated in the downstream. With the use of multiple modulation profiles supported by the DOCSIS 3.1 Profile Management Application (PMA), this incremental new capacity is added to the network, optimizing coaxial bandwidth efficiency.

Lastly, note that the DOCSIS 3.1 downstream has **MUST** QAM formats up to 4096-QAM, but also offers optional profiles of 8192-QAM (8K-QAM) and 16,384-QAM (16k-QAM). The incremental capacity of increasing QAM formats is relatively modest (about 8%) for the increased MER needed to achieve them. However, with DAA, once impossible MERs are now within reach. While the full system specifications were not completed at the time of the DOCSIS 3.1 standards development for these QAM formats, based on the 4k-QAM requirements for the DOCSIS 3.1 downstream, the 8k-QAM MER requirement is likely to be in the range of 45 dB. From the distribution in **Figure 13**, a substantial number of CM MERs could reach this profile if it became available.

Now, pivoting to **Figure 13**, we see the side-by-side for the upstream, in this case only comparing BAU HFC vs DAA N+0.

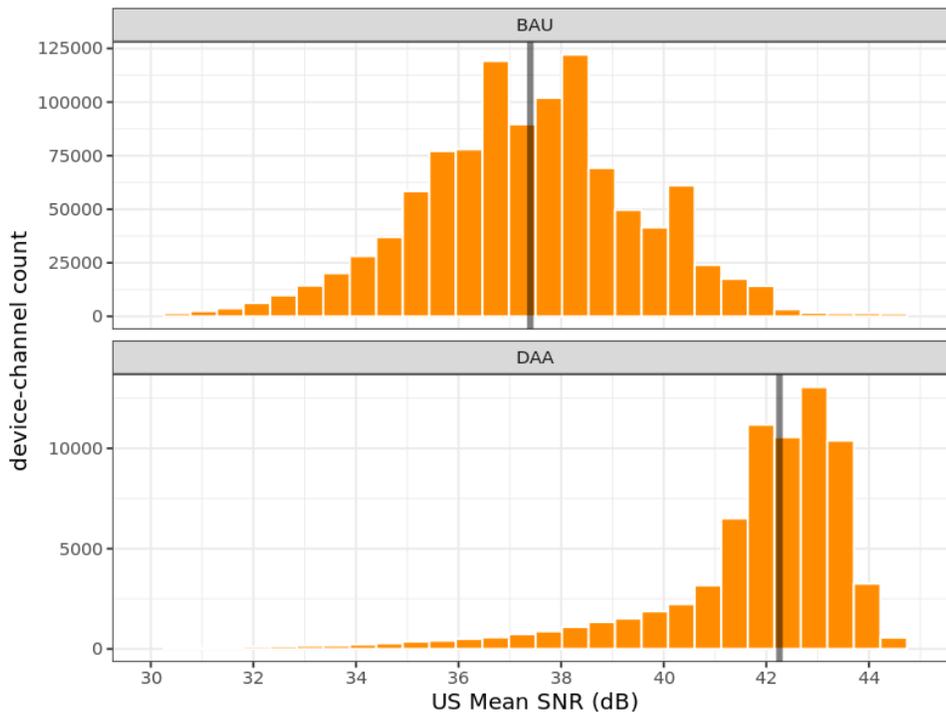


Figure 13 - Upstream Rx MER Comparison: N+x HFC (top) and N+0 DAA (bottom)

The large majority of node splits are driven by upstream capacity limitations, so we will consider the impacts of higher MER on capacity of the upstream.

The increase in bandwidth enabled by DAA in a production upstream is not determined completely by the capacity multiplier noted above for DAA and DOCSIS 3.1. There is an existing DOCSIS 3.0 traffic load in the Low Split band that will work to limit the percent capacity gain, as these carriers deliver a fixed DOCSIS 3.0 64-QAM maximum capacity regardless of network architecture. However, since we are considering Mid-Split, which delivers a long-term capacity runway, we can project that these DOCSIS 3.0 QAMs will be gradually reclaimed for DOCSIS 3.1 as those devices begin to dominate the CPE footprint, and base the analysis on the all-DOCSIS 3.1 case.

Thus, considering the use of Mid-Split supporting OFDMA, the DAA benefit of 3 dB can be used to estimate the impact on timing of a node split due to Compounded Annual Growth Rate (CAGR). For exemplary purposes, assume that the DOCSIS 3.1 upstream QAM format can be 256-QAM, where once it was 64-QAM for DOCSIS 3.0. The more powerful FEC of DOCSIS 3.1 delivers most of the additional QAM efficiency for “free” due to the extra coding gain. Improvements over time in the network RF quality that occur via node segmentation and noise funneling effects since the launch of DOCSIS 3.0 create additional dB benefits. Introducing the DAA gain to this situation, and the MER assessment above, suggest that we can improve to 512-QAM, or a 12% capacity gain.

An US CAGR of 25% is the current working assumption for business planners and capacity managers. This is a monthly growth rate of about 1.02%, and 12% capacity translates to 6-7 months. When considering the expense of tens of thousands of node splits per year, along with factoring in the time value of money, while it may be undramatic, this represents a quantifiable savings of capital by deferral to the node split budget.

In addition to capacity implications, this unavailable bandwidth must be accounted for when considering the penetration of speed tiers in a service group. As higher speeds are offered in the Upstream via Mid-Split, for example, increasing the peak burst, the remaining capacity during peak bursts must be shared by the other users in the service group. When the total available capacity is lower, there is less headroom above the peak, and this can impact the guidance for the number of users that can be added at the peak available tier. In addition, empirical data shows that when a subscriber upgrades to a higher speed tier, there is a predictable increase in their average utilization. As an example, the net result of these phenomenon, based on modeling for 300 Mbps US speeds, suggests that up to 6 users can have this speed tier in a Mid-Split system if utilization is below 60%. However, if the available payload is 9% lower, this number may be reduced to 4 or 5 users in a SG.

Furthermore, increasing US speed tiers has historically also increased utilization by freeing up latent demand. This again puts a premium on maximizing available capacity.

6.2. Customer Experience Metrics

Now consider Figure 14, which shows a common maturation trend observed when introducing the DAA platform followed by increasing production volume over the course (in this example) of a year. Statistically, the majority of customer trouble calls (TCs) are associated with issues that occur within the home, not the network. However, changes of any type on the network often results in interim jumps in TC activity. It is typical that as the DAA ecosystem components – vCMTS and RPHY Nodes – are rolled into new areas, there is an activation and support learning curve before returning to a “business-as-usual” state of health with respect to the customer experience. In some cases, this manifests itself as a temporary uptick in TC rate, which over time, as shown, returns to the mean.

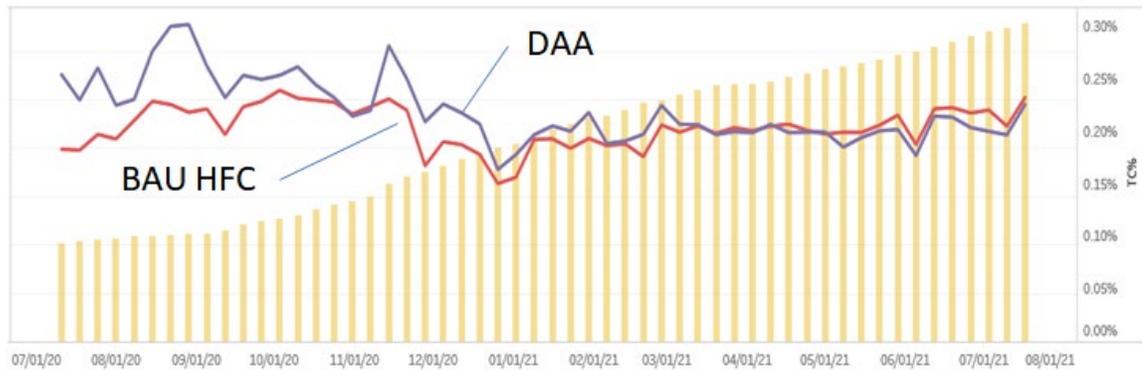


Figure 14 - Operational Metrics Proceed to Maturity as the DAA Platform Grows in Scale

As the DAA footprint continues to expand, availability and customer experience metrics are being carefully tracked and itemized so that any network-related causes are understood and addressed. Decreasing network maintenance activity and associated operational costs over time is cautiously expected, and saving targets established to quantify the costs removed from the operation. Observing network-related ticket root causes, improvements are anticipated for several reasons:

- Analog fiber links, prone to thermal variations and RF alignment sensitivity, become digital
- Large and more bug-prone monolithic SW releases of I-CMTS platforms become more incremental and agile
- The virtual implementation lends itself better to proactive measures and self-healing
- The virtual implementation lends itself better to smaller blast radii
- More robust DOCSIS 3.1 signals take over more of the spectrum allocation in both directions

We will continue to update the industry as DAA continues to scale across the enterprise, improving the network everywhere it is deployed.

6.3. Network Power Consumption Metrics

Returning to the topic of power consumption, one important original question about DAA networks, and in particular the very challenging N+0 DAA network – was what the impact would be on network power usage. RMD devices, with the added processing capability, distributed into the node, would increase the power consumption challenge further.

A power study was devised to measure the true impact of the DAA architecture on a section of an N+0 DAA network built in Denver in 2020. Five nodes were selected which contained 18 network power supplies. The 18 power supplies contained a mix of Alpha XM2 and XM3, 15 and 18 amp, 120 and 240 VAC variants. At the input and output of each power supply, measurements taken pre- and post-construction:

- Power
- Voltage
- Current
- Power Factor

All measurements were taken with a FLUKE 345 PQ CLAMP METER.

The N+0 architecture took advantage of the high RF output levels previously discussed. Four specific drop models were developed based on a majority of residential builds. All taps include the capability to shape the RF to provide an optimized DS and US RF profile.

Table 1 depicts characteristics of the five nodes selected. The power boundaries are restricted to the node boundaries in the design effort, so power distribution was not changed from a pre and post construction point of view allowing for an apples to apples comparison.

Table 1 - Node Statistics

| PS | Node ID | Model | VAC | HFC Legacy Design | | | | | xNET Design | | AER HHP | UG HHP | HHP | Aer Miles | UG Miles | MILES | % ports | HHP/MILE | |
|----|---------|-------|-----|-------------------|----|----|------|---------|-------------|-------|---------|--------|------|-----------|----------|-------|---------|----------|-----|
| | | | | node | LE | MB | WIFI | ACTIVES | NODES | Wi-Fi | | | | | | | | | |
| A | 11 | 915 | XM2 | 120 | 0 | 3 | 6 | 2 | 11 | 3 | 2 | 340 | 43 | 383 | 1.5 | - | 1.5 | 92% | 255 |
| B | 11 | 918 | XM3 | 120 | 1 | 9 | 6 | 4 | 20 | 3 | 7 | 343 | 241 | 584 | 1.0 | 0.1 | 1.1 | 100% | 510 |
| C | 11 | 915 | XM2 | 120 | 0 | 4 | 4 | 1 | 9 | 4 | 0 | 382 | 190 | 572 | 1.6 | 0.0 | 1.6 | 81% | 349 |
| D | 11 | 915 | XM2 | 120 | 0 | 7 | 3 | 0 | 10 | 2 | 0 | 187 | 0 | 187 | 0.8 | - | 0.8 | 50% | 238 |
| | | | | | 1 | 23 | 19 | 7 | 50 | 12 | 9 | 1252 | 474 | 1726 | 4.9 | 0.2 | 5.1 | 83% | 340 |
| A | 12 | 915 | XM2 | 120 | 0 | 8 | 4 | 0 | 12 | 2 | 2 | 342 | 7 | 349 | 0.9 | 0.1 | 1.0 | 100% | 335 |
| B | 12 | 915 | XM2 | 120 | 1 | 6 | 8 | 2 | 17 | 2 | 0 | 317 | 16 | 333 | 0.6 | 0.3 | 1.0 | 100% | 343 |
| C | 12 | 915 | XM2 | 120 | 0 | 6 | 5 | 2 | 13 | 4 | 1 | 529 | 137 | 666 | 1.4 | 0.1 | 1.5 | 100% | 441 |
| | | | | | 1 | 20 | 17 | 4 | 42 | 8 | 3 | 1188 | 160 | 1348 | 2.9 | 0.6 | 3.5 | 100% | 383 |
| A | 17 | 915 | XM2 | 240 | 0 | 14 | 5 | 0 | 19 | 4 | 0 | 490 | 30 | 520 | 2.4 | 0.8 | 3.2 | 100% | 164 |
| B | 17 | 915 | XM2 | 120 | 1 | 7 | 4 | 1 | 13 | 1 | 1 | 25 | 13 | 38 | 0.8 | 0.2 | 1.1 | 100% | 36 |
| C | 17 | 918 | XM3 | 120 | 0 | 12 | 5 | 1 | 18 | 5 | 0 | 464 | 87 | 551 | 1.9 | 0.9 | 2.8 | 100% | 198 |
| D | 17 | 915 | XM2 | 120 | 0 | 7 | 10 | 1 | 18 | 4 | 1 | 167 | 294 | 461 | 0.9 | 0.5 | 1.5 | 88% | 316 |
| | | | | | 1 | 40 | 24 | 3 | 68 | 14 | 2 | 1146 | 424 | 1570 | 6.0 | 2.4 | 8.5 | 96% | 186 |
| A | 8 | 915 | XM2 | 120 | 0 | 4 | 3 | 1 | 8 | 3 | 0 | 438 | 11 | 449 | 1.2 | 0.2 | 1.4 | 92% | 322 |
| B | 8 | 915 | XM2 | 240 | 1 | 11 | 7 | 3 | 22 | 2 | 0 | 274 | 0 | 274 | 0.7 | 0.2 | 0.9 | 100% | 292 |
| C | 8 | 918 | XM3 | 120 | 0 | 5 | 3 | 3 | 11 | 4 | 5 | 353 | 34 | 387 | 1.5 | 0.4 | 1.9 | 88% | 204 |
| | | | | | 1 | 20 | 13 | 7 | 41 | 9 | 5 | 1065 | 45 | 1110 | 3.4 | 0.8 | 4.2 | 92% | 262 |
| A | T0 | 915 | XM2 | 120 | 0 | 12 | 7 | 0 | 19 | 4 | 0 | 0 | 440 | 440 | - | 2.0 | 2.0 | 94% | 223 |
| B | T0 | 915 | XM3 | 120 | 0 | 10 | 3 | 0 | 13 | 3 | 0 | 0 | 268 | 268 | - | 1.6 | 1.6 | 75% | 171 |
| C | T0 | 915 | XM2 | 240 | 0 | 8 | 8 | 0 | 16 | 6 | 0 | 0 | 802 | 802 | - | 2.8 | 2.8 | 92% | 36 |
| D | T0 | 915 | XM2 | 120 | 1 | 20 | 6 | 0 | 27 | 5 | 0 | 0 | 459 | 459 | - | 1.9 | 1.9 | 70% | 242 |
| | | | | | 1 | 50 | 24 | 0 | 75 | 18 | 0 | 0 | 1969 | 1969 | - | 8.2 | 8.2 | 83% | 239 |

One of the most significant items from the data in **Table 1** is the reduction in the number of actives, which is facilitated by the high RF levels utilized in the fiber deep design. This reduction creates several operational benefits in the plant, including lower maintenance, less failure points and lower power consumption.

The results of this study are shown in **Figures 15-19**. Assessing the results of the trial, there was a significant reduction in overall utility power usage in the plant. We can see a net reduction of input and output power from most of the power supplies studied and a net reduction in input and output power in every original node boundary area. This is most welcoming news. While using less power is always an excellent result, another practical point of view exemplified by this result is that it has once again been shown that over time we can accomplish much more for a fixed amount of power, as network services continue to be added. Effectively, the bits-per-second-per-kilowatt-hour (bps/kwh) continues to improve. This has led to the deployment, for example, of outdoor WiFi Access Points (APs) or APs supporting CBRS (Citizens Broadband Radio Service) coverage.

Additionally, while cutting the total number of actives is the cause of the most significant decrease in power usage in this study, the effect of reduced transmission power loss in the cable plant is also to be noted. When current is drawn through the coax network, power losses defined by Joule's Law ($P=I^2R$) are incurred based on current (I) and loop resistance of the network (R). As the number of actives is reduced, and the current required to run actives in the plant is also lessened, so is the total power dissipated by transmission. Reducing this dissipated power is key to overall plant efficiency because, unlike power driving nodes or other actives, the power consumed due to transmission creates no value.

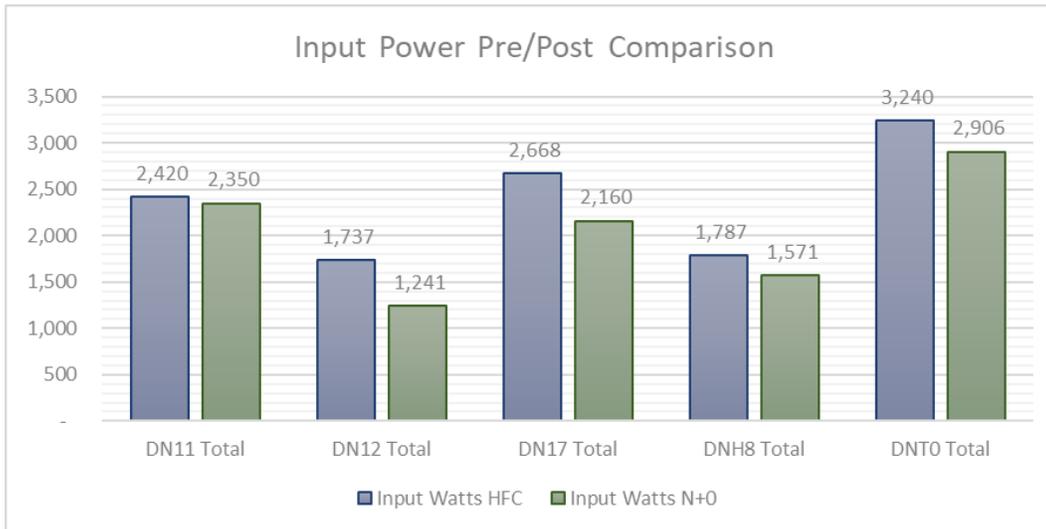


Figure 15 - Pre/Post Input Power Comparison by Original Node Boundary

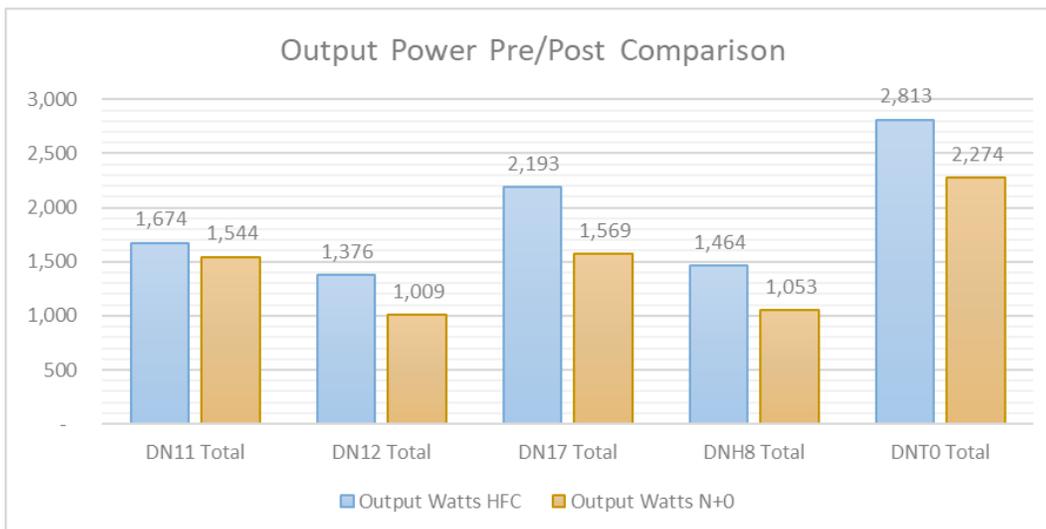


Figure 16 - Pre/Post Output Power Comparison by Original Node Boundary

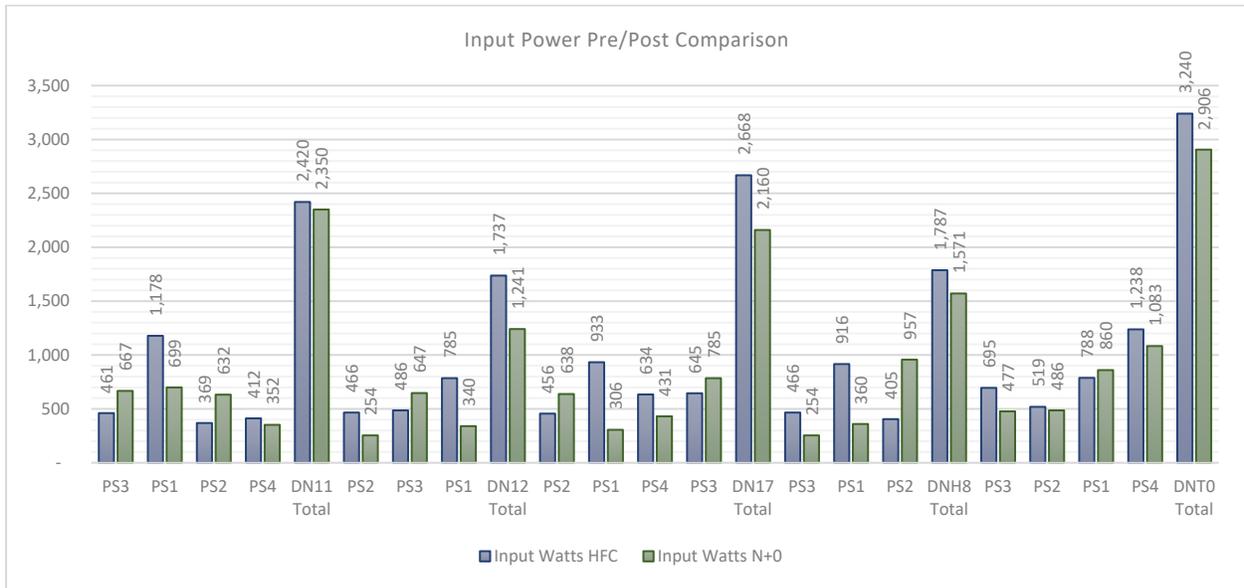


Figure 17 - Pre/Post Input Power Comparison by Power Supply

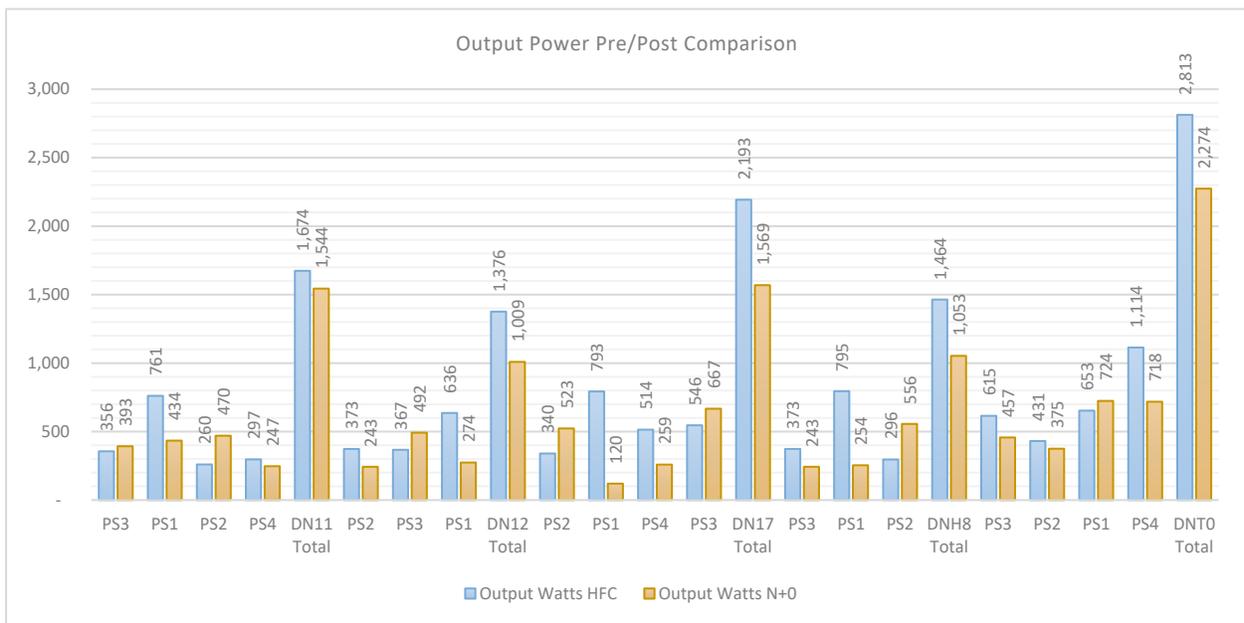


Figure 18 - Pre/Post Output Power Comparison by Power Supply

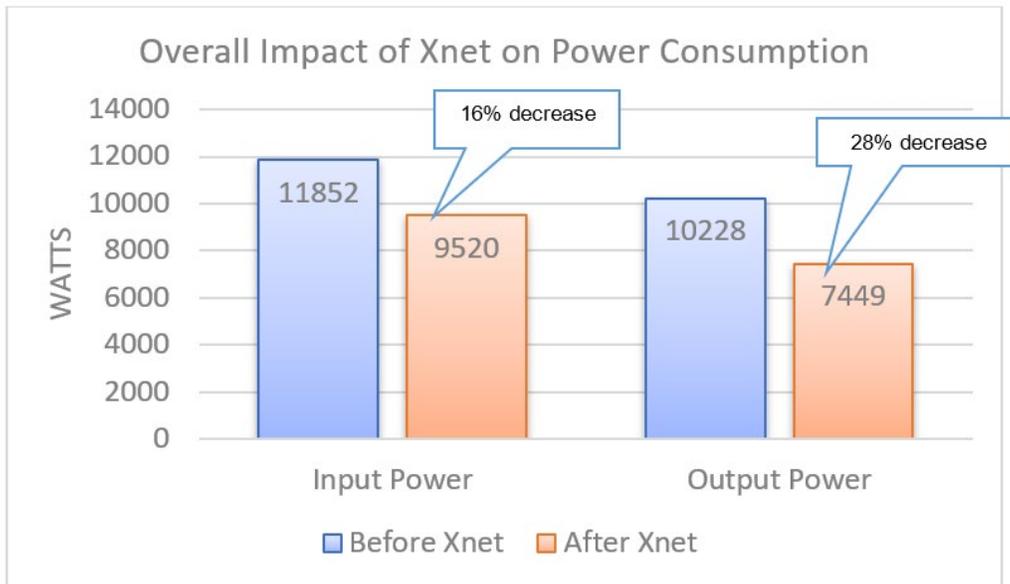


Figure 19 - Pre/Post Input Power Comparison Consolidated

An additional benefit derived from a reduction of plant actives and the resulting power savings is the positive impact on plant backup times. Using the powering data collected, and a 3-battery string of 114 Ah batteries, the sites in this study would experience an average of a 30% increase in battery backup time. This is shown in **Figure 20**. This additional runtime will reduce operational costs of additional backups to these sites during outages and decrease the risk of these outages becoming customer affecting.

Alternately, where current backup times are adequate, there may be opportunity to realize savings by either deploying a smaller, lighter, more cost-effective battery or by reducing battery string counts.

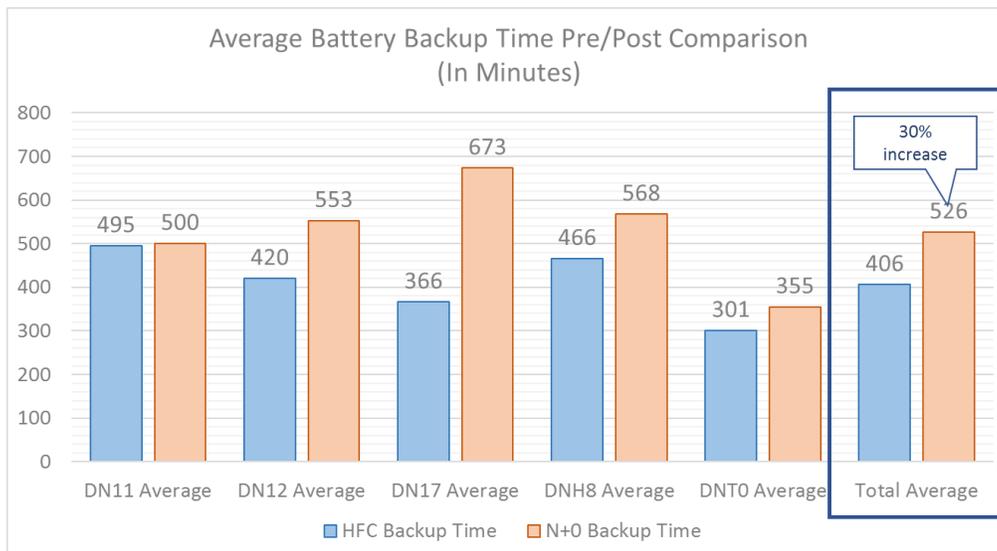


Figure 20 - Average Battery Backup Time Pre/Post Comparison

Finally, beyond the savings in network power, the reduction in network actives decreases the number of active failure points in the network. As with any system, reducing the number of potential failure points within the system, reduces the overall likelihood of system failure and, in this case, impact to customers.

7. Conclusion

Comcast began its DAA journey into the production network in mid-2018. By 2019, in every location that a Tier 1 or N+0 upgrade took place, it was done with vCMTS and DAA nodes. Remote PHY went from a CableLabs standard and PowerPoint to reality. There are now three production OEM partners producing DAA nodes in this interoperable ecosystem.

The scale of the deployment is now significant, yielding scale metrics of statistical significance, and validating projected performance improvements. We observed MER improvements that were anticipated with DAA, driving more capacity into the plant when it is combined with the more bandwidth efficient capabilities of DOCSIS 3.1. The capacity gains translate both to deferred costs of node splits, or other augmentations, that are driven by high utilization thresholds.

This added bandwidth can also translate to service speed scalability. While the added capacity technically could represent an opportunity for higher speed offerings, the difference is relatively modest for it to be a meaningful difference. For example, does it make very much difference to the customer or to the business if a service tier is 250 Mbps vs 280 Mbps (a 12% difference)? Probably not. However, the added capacity headroom will mean that more users can be added onto a higher speed tiers before a plant augmentation must occur to support the penetration – again, savings by deferring cost.

One of the initial uncertainties, and ultimately better than anticipated result, is how the introduction of DAA via RPHY nodes, as part of an N+0 upgrade, is, as it turns out, actually reducing power consumption of the serving area. This is a combination of the evolution of available technology replacing much older equipment in the field and incorporating efficient power network design as a priority from the outset. In the case of RF amplifiers that are eliminated in an N+0 build and effectively replaced by nodes, the amplifiers being replaced can be over 20 years old, and thereby relatively inefficient technology by today's standards.

Comcast is still relatively early in the DAA journey, but the engine is in 5th gear. The data from these upgrades is and will continue to be being collected and analyzed. This will drive optimizations and investment decisions going forward, such as the path to deployment of 10G and the evolution of the access network edge for alternative last mile technologies. Comcast will continue to keep the industry abreast of the observations and findings from the world's largest cable DAA network, fed by the world's most subscribed virtualized broadband platform.

Acknowledgments

A big thank you to the two Comcast technicians: Seth Williams, CommTech 5, Network Maintenance and Mike Torrez, CommTech 5, Network Maintenance. These talented gentleman collected all of the data for the power consumptions study.

Special thanks to Maher Harb and Ramya Narayanaswamy from Comcast's Next Gen Access Networks team for their support in providing dashboard readouts for the DAA production fidelity and customer experience metrics.

Abbreviations

| | |
|---------|---|
| AM | Amplitude Modulated |
| APs | Access Points |
| BAU | Business-As-Usual |
| CBRS | Citizen Band Radio Service |
| CMTS | Cable Modem Termination System |
| COTS | Commercial-off-the-Shelf |
| DAA | Distributed Access Architecture |
| EOL MER | End-of-Line Modulation Error Ratio |
| FMA | Flexible MAC Architecture |
| I-CMTS | Integrated CMTS |
| N+0 | Node plus Zero Amplifiers |
| N+X | Node plus X Amplifiers |
| OFDMA | Orthogonal Frequency Division Multiple Access |
| OSP | Outside Plant |
| PFC | Power Factor Correction |
| PMA | Profile Management Application |
| PON | Passive Optical Network |
| PP | Power Pack |
| QAM | Quadrature Amplitude Modulation |
| QSW | Quasi-Square Wave |
| R-CCAP | Remote Converged Cable Access Platform |
| RMD | Remote MAC Device |
| RPHY | Remote PHY Device |
| vCMTS | Virtual CMTS |

Bibliography & References

Howald, Dr. Robert, Aboard the Technology Wave: Surf Report, 2016 SCTE Expo, Philadelphia, PA, Sept 26-29.

Nandiraju, Dr. Nagesh, Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access, 2016 SCTE Expo, Philadelphia, PA, Sept 26-29.