Powering The Future:
Next Generation Access Networks

A Technical Paper prepared for SCTE*ISBE by

Todd Musat
Director, Critical Network Infrastructure
Shaw Communications Inc.
2728 Hopewell Place NE Calgary AB T1Y 7J7
1.403.930.9634
todd.musat@sjrb.ca

Chuck Carroll
General Partner
Saras Energy Consulting
1006 Monroe Ave River Forest, IL 60305
1.360.201.4205
ccarroll@saraspartners.com

Lew Rakowsky, Contributor, Saras Energy Consulting

Rene Spee, Contributor, Saras Energy Consulting

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

As bandwidth capacity continues to grow for existing hybrid fibre coax (HFC), so does the flexibility to utilize new technologies within HFC networks. Establishing practices which support new technology platforms, while fully employing the large investment in the HFC plant is important. As the coaxial cable is capable for shorter distances of carrying high bit rate traffic, MSOs generally implement a fiber deeper approach, where node splits allow fiber penetration further into the network, and closer to the customer. Often, an N+0 approach is considered, which has no active devices on the "last mile" coax between the optical node and the customer premises. However, cost considerations often require compromises, and in many cases, an N+1 (one amplifier between node and customer) or even N+2 approach is implemented — still a vast improvement over the historical N+4 or higher architectures.

HFC networks traditionally terminate node traffic on cable modem termination systems (CMTSs). Integration of CMTSs with other transport functions in the head-end and hub has driven implementation of the converged cable access platform (CCAP). CCAP combines the functions of the CMTS and the edge quadrature amplitude modulators (QAMs), allowing digital data processing prior to analog conversion for transport on the access network. Although CCAP itself drove more power and space efficient HFC edge network solutions, in a fiber deep architecture, the number of nodes will increase significantly, and with it, the requirement for additional CCAP hardware in the hub. This, in turn requires additional space, power and cooling in the facility.

In the distributed access architecture (DAA), the constraints on the headend are ameliorated since the physical and media access control (PHY and MAC) can be placed further in the access network (e.g. directly with the optical node). There are multiple implementation possibilities for DAA. Currently, the most popular and only fully CableLabs specified approach is Remote PHY, where the CCAP PHY layer is moved to the node, while the MAC layer continues to reside in the hub. The facility improvements in terms of space, power and cooling are moderate, as the CMTSs are not being eliminated.

Other possible implementations of DAA are Remote MAC/PHY and Remote CCAP. Here, both physical and MAC layers are moved to the node. Many functions can be virtualized and this in turn can lead to a drastic reduction of facility space, power and cooling requirements.

In 2019, as an integral part of evolving the network to 10G, Shaw began the process of introducing DAA into its network. As with other MSOs, the key driver of this effort was to move CCAP PHY processing from the hub to the node, replacing RF optics with metro ethernet (metroE) optics. The chosen strategy evolved DAA into the network in a controlled manner, implementing DAA in greenfield nodes and node splits, but not proactively in existing brownfield. As this was a different approach to DAA implementation, the company wished to better understand how a migration to DAA in this manner would affect energy usage across their footprint.

In conjunction with Shaw, Saras Energy Consulting worked to assess the energy use impact of the company's approach to DAA across its footprint. The workplan focused on the following investigative objectives:

- Establish a hypothesis about the impact of DAA on facility and outside plant (OSP) energy environments, based on the DAA design and implementation plan as per Shaw standard.

- Collect and analyze any relevant DAA architecture, equipment, and deployment data to model energy usage over time across the Shaw footprint.
- Use the model to prove and/or adjust the hypothesis and answer critical energy usage related questions.
- Utilizing the analysis, work with Shaw to draw conclusions with respect to the energy usage profile related to the Shaw DAA deployment plan.
- Analyze OSP powering capacity and architecture, specifically related to the ability of the power architecture to accommodate DAA evolution, as well as additional IoT, small-cell and other non-HFC connectivity alternatives.

To evaluate DAA impact on the network, the team first developed a model detailing energy usage for the business as usual (BAU) network architecture. It should be noted that the term BAU is used in this paper to refer specifically to Shaw's traditional plant build as currently defined. The team then developed a DAA model, based on Shaw's specific DAA implementation process, to compare against BAU. Conclusions were then drawn from the modeling work.

2. Business as Usual Network

The business as usual network has been built and developed over time in a manner consistent with MSO practice. The diagram below provides a view of the BAU network

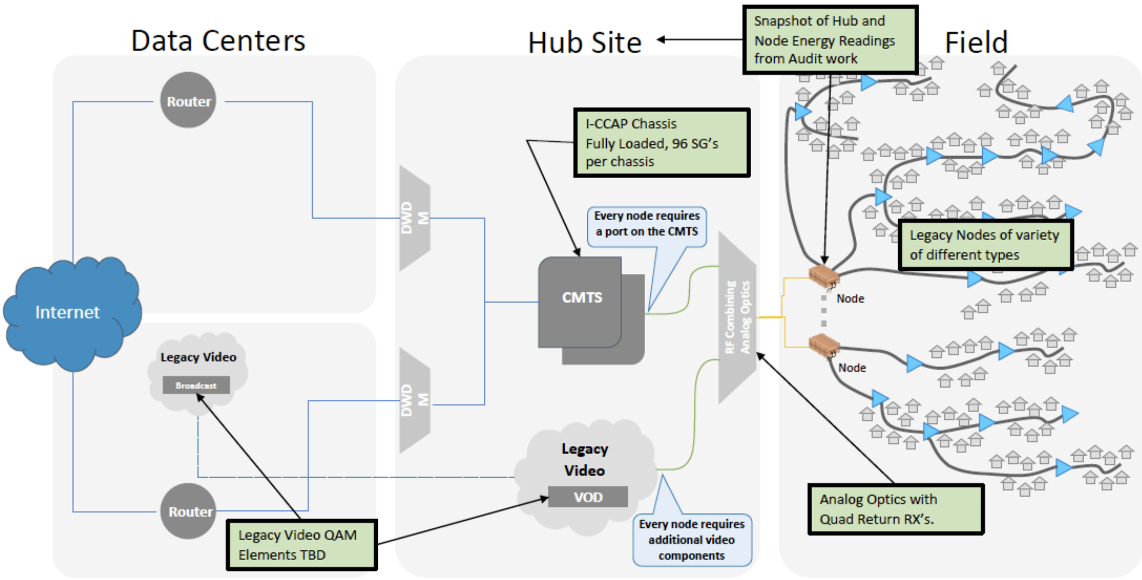

Figure - 1 Business as Usual Network Deployment

Shaw implements a tiered facility structure, with the top tier data center/core facilities located in key regions served and connected by a national mesh ring network. From the regional core sites, rings are used to connect to hub edge facilities in the region. Hub/edge facilities support a load of approximately 50,000 homes passed (HP) with capacity to go to 80,000 HP. The HFC distribution network stars out from the hub/edge facility. In the facility, the company connects nodes to an integrated CCAP (I-CCAP) device, with standard RF optical devices and connectivity infrastructure in place to support the forward and return bandwidth to/from the nodes.

With respect to Shaw's HFC infrastructure, it is N + X, with current attention to reduce or maintain node size and cascade lengths. Congestion relief in the downstream is achieved by node-splitting, where the node is physically divided to create two nodes, adding an additional optical path to the head-end, effectively doubling the dedicated downstream capacity to the same area. Upstream congestion relief is via a combination of node splits and conversion to a higher return path bandwidth of 85 MHz, techniques which increase dedicated upstream capacity for the area being served.

The company is committed to evolving fiber deep over time, with the ultimate goal of moving towards DOCSIS 4.0 in preparation for a 10G future. Evolution to N + 2 architecture is the first step to that future [1]. In support of this, greenfield plant and any brownfield upgrades are built to N + 2.

3. Approach to DAA

As noted in the introduction, an important part of Shaw's evolution to 10G is implementation of DAA in its network. DAA itself is a cable industry initiative, spearheaded by CableLabs, and forms a key part of the cable industry's drive to 10G in the future. Network benefits related to DAA implementation include:

Network Efficiency

- Increased network capacity
- Better end-of-line signal quality, higher modulation rates, higher bitrates
- Better spectral efficiency, more wavelengths per fiber

Operational Efficiency

- Reduced headend power, space and cooling requirements
- Hub consolidation
- Simplified Hub Architecture by leveraging digital optics (eliminates RF combining)

Enabling Convergence

- Extends IP network into the field
- Ability to leverage common IP Network for multi-tenant applications
- Alignment with FTTx build-out
- Required steppingstone to Virtual CCAP (vCCAP), 10G and the Edge Cloud

A high-level view of the Shaw approach to DAA implementation is shown in Figure 2 below.

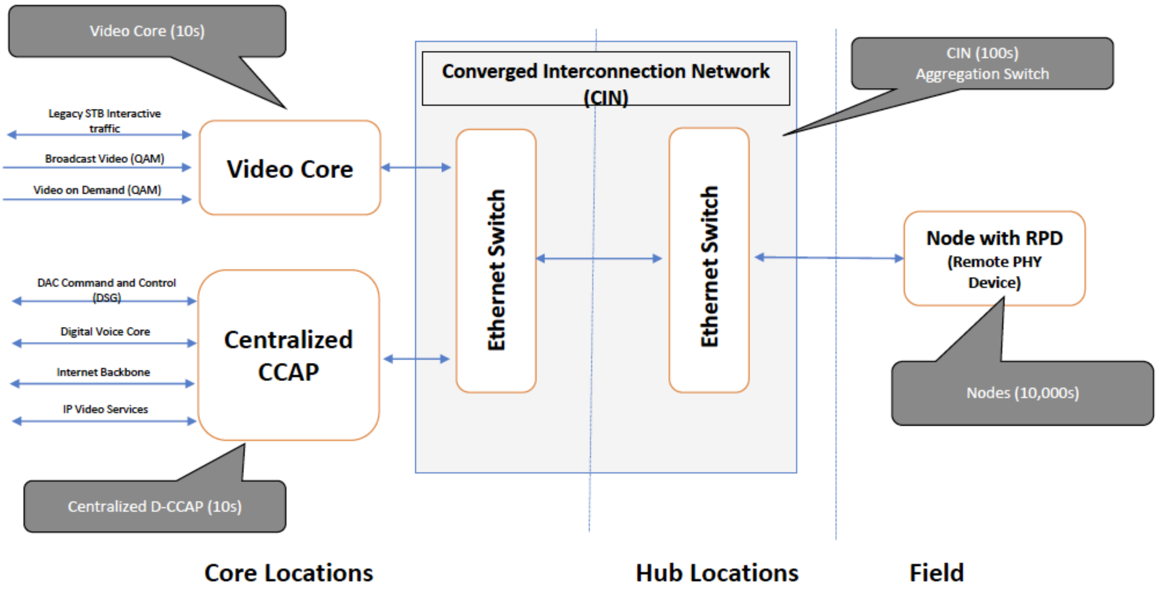

Figure 2 Shaw DAA Implementation

Key elements of the DAA infrastructure as defined by CableLabs in [4] and to be deployed by Shaw include:

- **Remote Phy Node (RPD):** The RPD is a device located in the node in the network which implements the Remote-PHY specification to provide conversion from digital Ethernet transport to analog RF transport.
- **Converged Interconnection Network (CIN):** The network (generally gigabit Ethernet) that connects a CCAP Core to an RPD.
- **CCAP Core:** A CCAP device that uses MHAv2 protocols to interconnect to an RPD. This device could be a DOCSIS Core, Video Core, OOB Core, RPD Controller (or a combination of these roles). Distributed CCAP (D-CCAP) is the implementation of CCAP Core the Shaw network uses and is used to represent the CCAP core function in this paper.

4. DAA Transition

Although Shaw's overall approach to DAA is consistent with that of CableLabs and the industry, a full DAA implementation will be controlled and measured. DAA will be implemented to meet node growth related to greenfield build and node-splits, as well as very selectively in existing brownfield upgrades. To support this evolutionary approach, initially a D-CCAP core and its associated CIN devices are to be placed in Shaw regional core sites. The infrastructure will be used to support the initial DAA node requirement for any hubs served from that core. D-CCAP and CIN network will grow with DAA node need, with D-CCAP chassis' and CIN leaf switches ultimately evolving outward to the hub sites as and when demand in DAA nodes dictates.

With respect to the nodes, Shaw's transition strategy will not entail pro-actively changing existing brownfield nodes to DAA with limited exceptions. Instead, DAA transition will cap legacy node placement as much as practicable, and use DAA nodes for new greenfield builds, as well as node growth related to node splits. Key attributes of this approach include:

- Focuses on building greenfield plant primarily with DAA nodes.
- In brownfield, implements node splits to relieve congestion primarily in the following manner:
    - Existing node split and new node created with node split will be DAA nodes in the field.
    - Both nodes from the node split will be placed in D-CCAP chassis on 10 Gbps connection via CIN.
    - Existing coax network beyond the nodes stays consistent with traditional HFC node split practices. This is not a change to existing coaxial distribution plant other than to use DAA nodes vs. traditional nodes in the field.

Once a node split has been completed and the new nodes connected to the DAA infrastructure, transition reaches completion with removal of any elements of the BAU I-CCAP infrastructure that has been taken out of service. But as this is an evolutionary process, this happens only after all in-service elements on the legacy devices are completely transitioned to D-CCAP and CIN network:

- I-CCAP chassis removal does not occur until all nodes are removed from the chassis. As the geography of where node splits occur is relatively random and based on many factors, this may be a few years away, given the spread of I-CCAP devices in the network and average number of ports in use on those devices.
- RF transmitters and receivers will be removed and powered down as and when all nodes connected to the chassis have been moved to the CIN. This may happen more quickly, potentially in the first few years after node split, as the number of nodes per chassis is smaller.

At this point, it is not Shaw's intent to actively work to re-arrange BAU node connections to vacate I-CCAP chassis' and/or RF optic modules as utilization on them lessens, but to let natural attrition determine when devices are vacated and can be removed.

Figure -3 below shows detail of Shaw DAA strategy.

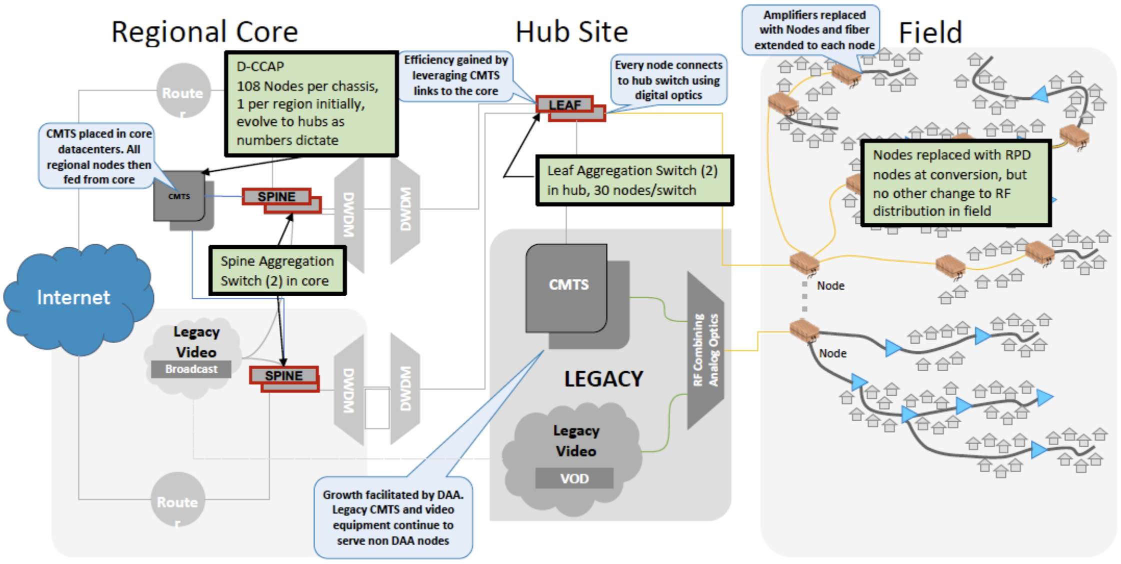

Figure -3 Shaw DAA Implementation

To summarize, with this approach, legacy I-CCAP infrastructure is largely capped in the footprint. The DAA infrastructure operates in parallel to the legacy I-CCAP and RF infrastructure, taking on all growth in nodes due to either greenfield build and/or node splitting. Legacy BAU I-CCAP equipment and supporting RF infrastructure is vacated over time by attrition, with equipment removed as and when service is no longer provided on it.

5. Modeling Energy Usage

As with other operators, Shaw's BAU networks grows energy usage year-over-year, due partly to continued growth in greenfield network homes passed, as well as increases in bits transported due to subscriber consumption related to internet use, streaming, etc. Greenfield growth drives additional nodes and actives in the field, along with the associated CMTS and RF optical components to support them. Increasing usage by subscribers creates node splits, along with the CMTS, RF optics, and new nodes required to support them. As subscribers' appetite for data consumption does not appear to be slowing, space as well as energy constraints in existing edge facilities are moving operators towards solutions which solve and/or mitigate these challenges. Evolution to DAA is driven by the needs of the cable operators like Shaw to solve the future space and energy usage problems growth of the BAU network creates if it continues at current pace.

Although in the past, work had been done on energy usage comparing BAU vs. the ***full*** implementation of DAA in a facility footprint [5], energy usage comparing BAU to a more evolutionary approach to DAA such as the one used by Shaw has not been modelled and/or analyzed. The following two sections of this paper detail this modeling, as well as results from the modeling work.

5.1. Initial Energy Usage Hypothesis

After initial review of Shaw's BAU scenario and DAA implementation plan detailed above, it was hypothesized that during the transition period from BAU to DAA, initially total energy usage would increase in comparison to BAU, before ultimately decreasing over time. This was due to three key reasons:

- Initially DAA D-CCAP, switch, and router equipment is added in core and hub sites to facilitate DAA, duplicating BAU network resources servicing nodes today.
- Removal of I-CCAP equipment will not occur until an individual chassis is completely vacated. This could take years.
- The DAA nodes in OSP are higher power devices than the BAU nodes they replace.

Figure -4 shows initial high-level modeling of percentage change to BAU associated with DAA implementation and transition done to form the initial hypothesis.

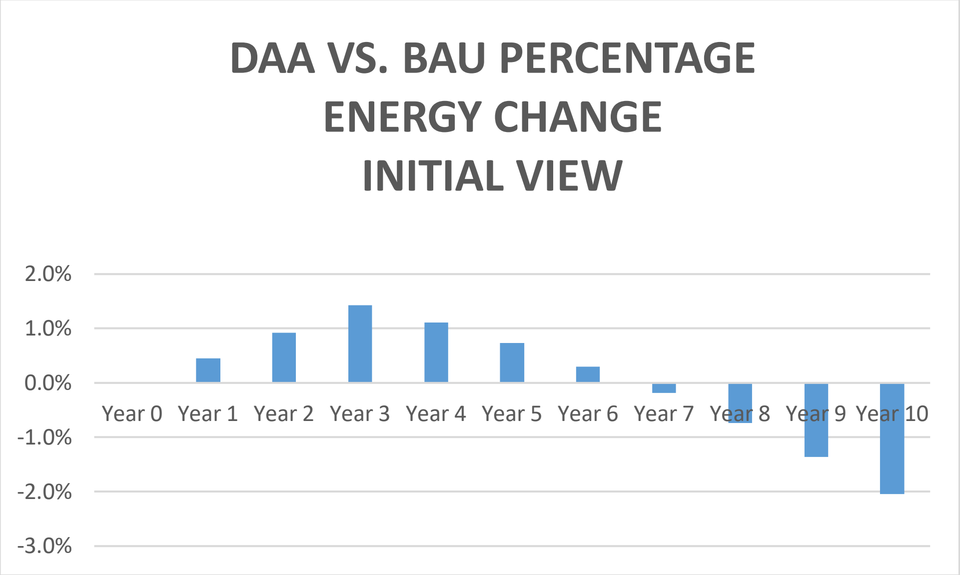

Figure -4 DAA vs. BAU Percentage Energy Change Initial Hypothesis Modeling

The initial high-level modeling represented in Figure 5-1 supported this hypothesis, showing an increase in energy usage in the early years related to deploying DAA in this manner. Ultimately, over time as DAA predominates and more BAU equipment is removed, DAA power does reduce, and go lower than BAU. That high-level modeling indicated timing of the cross-over point for DAA producing lower overall energy usage than BAU was a function of how fast DAA was implemented in the network, as well as how quickly I-CCAP infrastructure was removed from the network.

5.2. Energy Model Assumptions and Structure

To validate the initial hypothesis, a more detailed model was created. The detailed model starts with a base model for energy use in the BAU network. Layered on top of the BAU base is a model detailing changes related to implementing DAA in the plant, incorporating the Shaw transition assumptions. This allows comparison of BAU vs. DAA specific to the company approach. The base BAU model as well as the DAA overlay use company supplied and industry data, as well as current company practices for developing and growing the network.

Model Year 0 was specifically designed to align with Shaw current plant and energy information garnered from energy audit work performed in 2018. Key elements included:

- **Network footprint volume:** Used Shaw overall plant data in 2018 for HP, nodes, etc. as a starting point for the base BAU model.
- **OSP power per node:** Calculated OSP power per node using OSP PS data from audit performed in 2018 on calendar year 2017 data, combined with total node numbers from the plant data.
- **Hub/Facility power per node**: Calculated hub facility power per node using data from a sample of stand-alone hubs with known number of nodes connected.

From this data, total OSP and facility power across the Shaw footprint was calculated for Year 0, specifically matching audit data.

Moving in time from Year 0, the model changes homes passed and number of nodes year-on-year into the future using key drivers related to greenfield expansion and node splits. These include:

- **Greenfield expansion**: 1% increase in HP per year. This assumption is based on typical industry data for greenfield in developed markets. This was used to drive node growth due to greenfield build in the model.
- **Percentage node splits in year due to TB growth**: 5% of nodes from previous year. This was based originally on industry data, although a sample of Shaw data supported this assumption as reasonable. This was used to drive node growth in the model due to node splits.

The node growth resulting from the above drivers changes year-by-year energy usage in OSP and hub power based on the following key assumptions:

- **Facility equipment power:** RF optics power, CMTS chassis power, etc. as provided by Shaw for specific devices used in the company network. These numbers were used to calculate energy use changes in critical facilities for node growth related to plant extension and node splits in the year.
- **CMTS chassis utilization**: Model assumes one node per port per Shaw current practice, and 66% port utilization based on experience/industry norms. CMTS chassis' in the model grew as a function of node growth in combination with this utilization assumption.
- **Facility PUE**: Facility PUE was assumed to be 2.0, based generally on conservative industry norm. This is needed in the model as it drives changes for non-equipment energy (i.e. HVAC, lights, etc.) in facilities.
- **Brownfield node power:** Model uses per node power calculated from 2017 PS audit through-out the model.

- **BAU greenfield node power:** As greenfield nodes target N + 2, in theory they will be smaller than the brownfield assumption around node size and power. As such, greenfield node power is assumed to have W/HP consistent with a brownfield node.

- **BAU node split power**: After a node split, total power for the sum of the two split nodes is incremented by the power difference between the new node and the amp it is assumed to be swapped for.

For the DAA model, many of the base assumptions were similar to BAU. As with BAU, DAA model uses Shaw overall plant HP, nodes, etc. to model a Year 0 starting point similar to BAU today. Key growth drivers around greenfield growth percentage, node split percentage, PUE, and chassis utilization year-on-year are the same, as well.

The DAA model differs largely due to implementation and transition assumptions. Implementation drives changes to per node energy usage in the following ways:

- **OSP power per node:** DAA power per node is adjusted by the increment in power associated with replacement of the standard node with a DAA node.

- **Hub/Facility power per node**: Facility power is adjusted using DAA equipment elements instead of legacy equipment for the added nodes.

The DAA model itself is driven by the Shaw transition strategy. Key elements include:

- **DAA nodes only deployed in greenfield plant**: All plant extension greenfield nodes will be DAA nodes and therefore connected to D-CCAP chassis.

- **Node Splits:**

    - Existing node and the new node created from a split will be DAA nodes in the field.

    - Both nodes from the node split will be placed in D-CCAP chassis on 10 Gbps connection via CIN.

    - Existing node being split will be removed from I-CCAP chassis.

    - RF Optics connected to existing node will be disconnected.

- **Legacy I-CCAP**: I-CCAP infrastructure is capped – all node “adds” due to greenfield build and/or node splits become DAA and are placed on D-CCAP chassis.

- **BAU I-CCAP and RF Optics removal**: These items are removed but only after all BAU services are completely transitioned. Following assumptions are made in the model made for removal timing:

    - RF Optics removed in the same year I-CCAP node is moved to D-CCAP as part of node split.

    - Full I-CCAP chassis’ are removed three years after initial node moved to D-CCAP (i.e. starting year 4 in the model).

- **Evolutionary Strategy**: Strategy does not include targeting hubs and/or CMTS chassis/RF optics for complete conversion, and/or potential transfer of nodes to BAU ports to clear equipment for removal.

6. Findings from Modeling

Modelling work focused on a ten-year view of energy usage for each of the two options. Ten years was chosen because it is far enough in the future to see trends, but not beyond the point where the assumptions made could still seem applicable. Output was generated with respect to OSP plant energy usage, hub/facility energy usage, and combined usage, across the whole Shaw footprint. Data and analysis from that work is contained in this section.

6.1. BAU

BAU modeling shows future growth in energy usage in the footprint as assumed new greenfield build and node splits drive additional nodes. Figure -5 below shows the changes year-on-year.

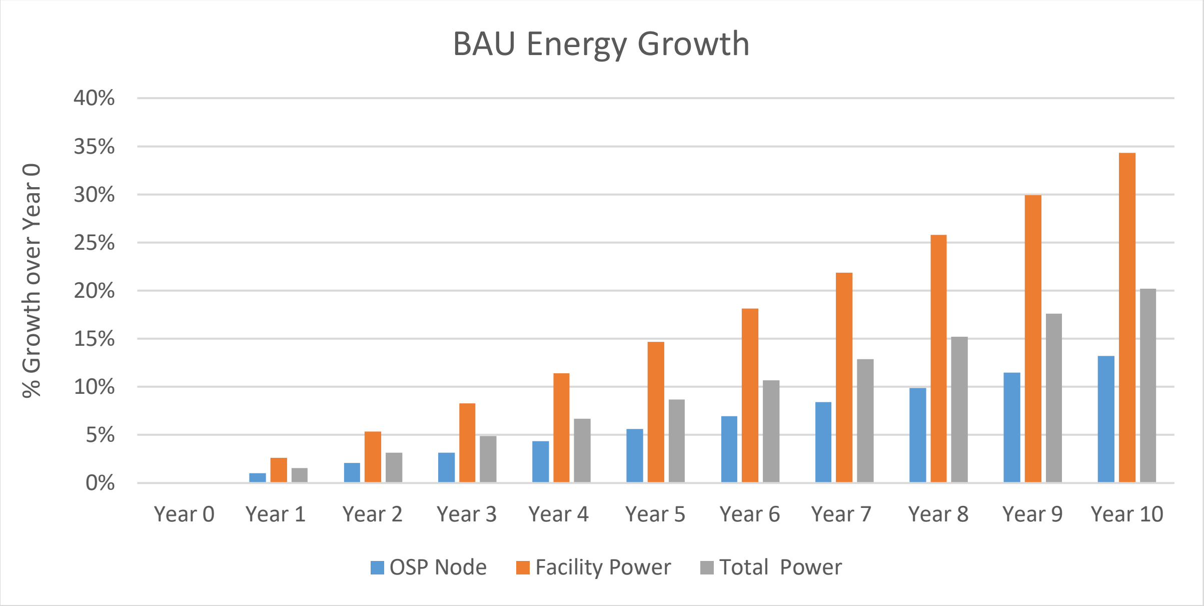

Figure -5 BAU Energy Growth

BAU OSP growth in the overall footprint is driven by a combination of greenfield expansion, as well as growth due to node split. Key impacting items associated with this were:

- Additional 1% of greenfield plant built in the footprint. This was the primary driver of BAU OSP energy use growth.

- The slightly smaller per node power specifically on the new greenfield nodes as at N + 2 they are smaller and require fewer active elements to be powered.

- The slightly higher power in split nodes due to swap of an amp with a node.

With respect to the individual components, OSP energy grows at a compound annual growth rate (CAGR) of 1.2%, making it 13% higher in Year 10 in comparison to Year 0.

Facility/Hub power grows at a much higher CAGR of 3.0% over 10 years. This is driven almost entirely by the addition of I-CCAP chassis', RF optics and infrastructure needed to support added greenfield nodes, as well as the nodes added via node splits. In comparison to Year 0, Year 10 usage is 34% higher in BAU.

Combining the effects of OSP and facility/hub growth, overall energy growth for BAU implementation had an energy usage CAGR of 1.9%. In comparison to Year 0, Year 10 usage is 20% higher. This combined view forms the basis of the BAU situation for comparison to DAA.

6.2. DAA

DAA modeling shows future growth in energy usage as well. As with BAU, new greenfield build and node splits drive this, tempered by the DAA transition drivers. Figure -6 below shows the changes year-on-year.

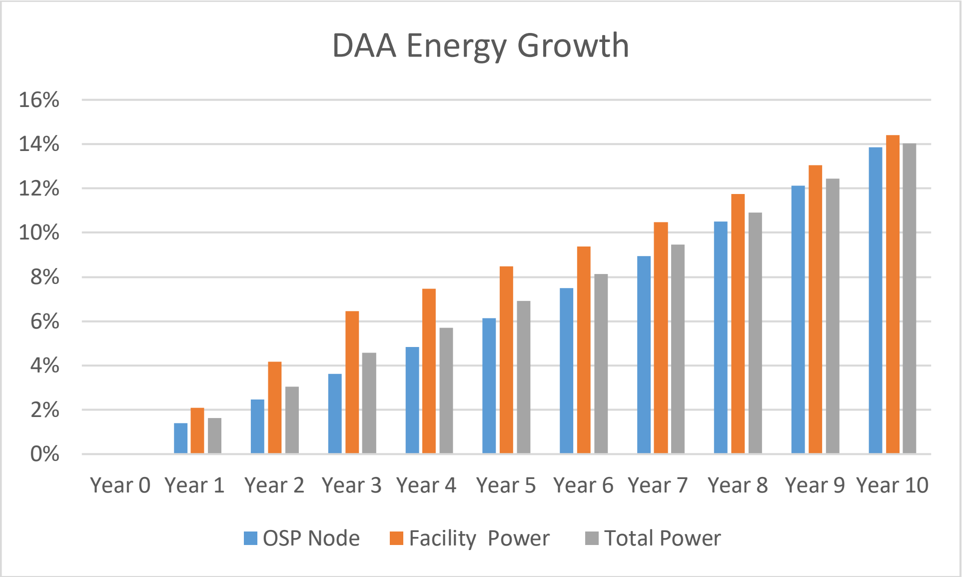

Figure -6 DAA Energy Growth

In relation to the individual components, DAA OSP energy is impacted largely by the same key drivers as BAU. Energy usage CAGR is 1.3% over the ten-year period, making energy usage 14% higher after 10 years. Both figures are slightly higher than BAU due to the swap out of legacy nodes with higher powered DAA nodes.

Facility/hub power growth undergoes much more meaningful changes from DAA implementation. Usage per year is higher in early years (2.1% - 2.2% per year), due to the addition of D-CCAP infrastructure coupled with the delay in removing legacy equipment related to DAA transition strategy, particularly the I-CCAP chassis' equipment. Starting in Year 4, the model begins the process of eliminating legacy I-CCAP chassis', and DAA growth approximately halves. The combination of removing I-CCAP chassis' from the network, and taking up port growth from Year 1 with lower power D-CCAP chassis', works to minimize yearly growth from that time to the 1.0% per year range.

As such, the facility energy use CAGR is 1.4%, with overall growth over the ten-year period slightly above 14%, both much lower than BAU.

The combined effect of the DAA transition strategy yields an energy use CAGR over 10 years of 1.3%, driving growth for the combination of OSP and facility to be 14% higher in Year 10, both lower than BAU. Even with the migratory approach, ultimately the swapping out of I-CCAP for D-CCAP, and swapping RF optics for CIN components, overcomes the added energy for DAA nodes in the OSP, leading to lower energy usage in comparison to BAU. Using the Shaw DAA transition strategy, energy usage will increase in future due to growth in plant and continued node splits – but DAA works to lessen that growth in future in comparison to BAU.

After modelling the Shaw DAA transition strategy, questions arose around how this might compare to energy usage for a complete conversion of a facility/hub to DAA in Year 1, as the company was contemplating this for a hub due to reasons related to an existing facility,. Although previously published work [5] on this topic was written from a more generic perspective, modelling of the Shaw specific use case for implementing DAA immediately in a hub provided some guidance.

Using energy audit data, along with homes passed, node and I-CCAP chassis data for facilities in the Calgary area, year one full DAA was modelled. The model used the following assumptions, yielding the following results

- **OSP DAA immediate implementation:** To calculate this, it was assumed all nodes within each of the hubs were swapped with DAA nodes. This drove a 4% increase in OSP power, solely related to the immediate swap out of lower powered existing nodes with higher powered DAA nodes.

- **Hub/facility DAA immediate implementation:** To calculate this, all I-CCAP chassis' in the hubs were swapped for D-CCAP. Additionally, all optical transmitters and receivers were eliminated, replaced with a CIN network to which the DAA nodes were connected. This resulted in a 42% decrease in hub/facility power, primarily due to the I-CCAP for D-CCAP platform swap.

- **Overall DAA immediate implementation:** The combination of these two effects drives a 5% decrease in overall energy usage. This relatively modest overall percentage reduction is a function of the heavy weighting towards OSP energy in the sample areas, which at ~80% is the majority of the total energy usage. Even though energy savings in the facilities for this approach would be meaningful, because of the large OSP component, they have a more limited impact on overall energy usage.

It should be noted that although not specifically a part of the work, high level estimation was that in addition to the significant savings in facility energy, fully implemented DAA would reduce space to somewhere between a third to a half of what it was pre-DAA.

Although meaningful energy and space savings would be achieved with immediate conversion to DAA in a facility, it is unlikely the savings associated with that change would offset the capital required to accelerate DAA implementation on its own. In special cases, lease challenges and/or space/energy limitations might necessitate such a change. The ability to use an all DAA implementation to meet a particular building and/or facility challenge is an important tool an operator can use. Although there are real benefits in energy and space savings in comparison to BAU to a full DAA implementation on Day 1, other rationale beyond those savings would be required to justify such a conversion.

6.3. DAA vs. BAU

Figure -7 compares year-on-year percentage difference in energy use change between BAU and DAA.

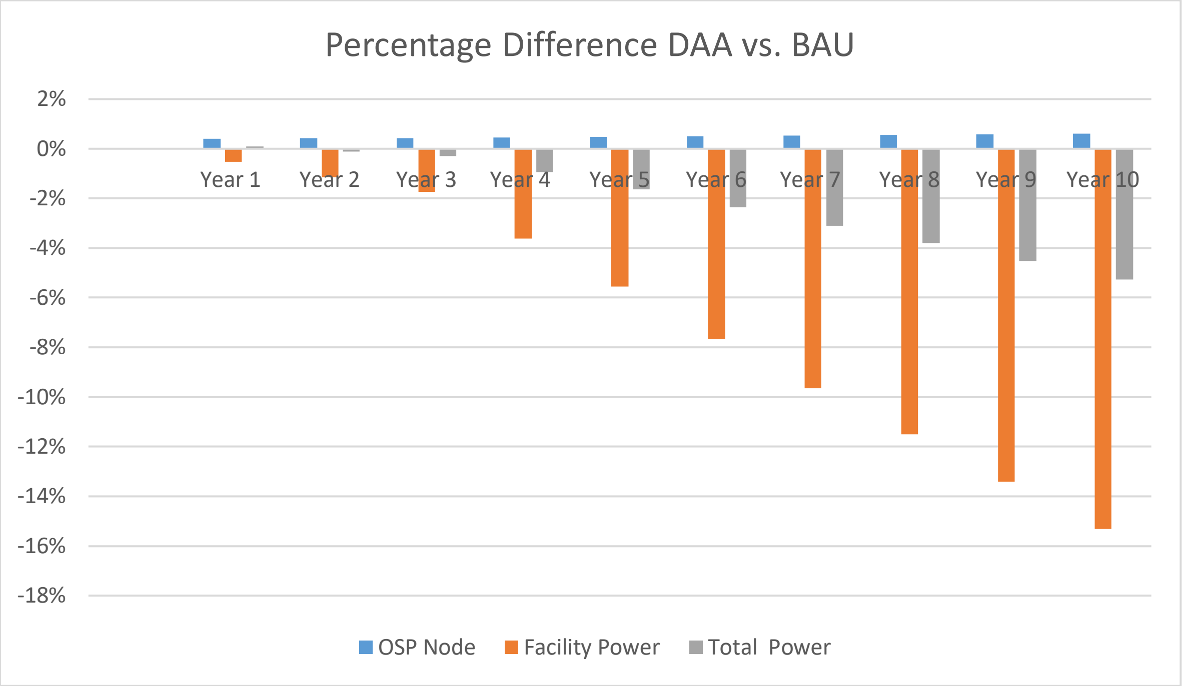

Figure -7 Percentage Difference DAA vs. BAU

This comparison of DAA energy growth to BAU energy growth shows the Shaw DAA transition strategy, even with its evolutionary approach, delivers energy usage reduction against BAU over time. Overall energy usage stays roughly the same in the early years, with the small decrease in facility energy offset by a slight increase in the much larger OSP component. This changes in the later years, as ultimately removal of I-CCAP components over time accelerate the energy usage reductions of DAA versus BAU.

Learnings from the modeling and subsequent analysis associated with it include:

Early years impact of dual DAA and BAU infrastructure: Section 5.1 hypothesized DAA energy use would increase in the early years with placement of dual DAA and legacy infrastructure for a period of time, reducing in comparison to BAU after I-CCAP removal started. In fact, the model showed the impact of this was negligible. Analysis indicated capping the placement of much higher power I-CCAP devices, and adding new capacity needs on newer, lower power D-CCAP devices in the DAA model, mitigated the added power of the dual infrastructure. A learning from this work was that the key to avoiding an early-year increase in energy usage related to this strategy is to cap I-CCAP and put all node growth in facilities on D-CCAP infrastructure.

Timing and size of DAA reduction – I-CCAP removal assumption: As expected, the size and timing of the DAA reduction varies with the assumption as to when I-CCAP chassis' would start to be removed. From a modeling perspective, earlier removal improves DAA in comparison to BAU, later removal lessens the energy improvement of DAA. The assumption of beginning the removal of chassis' in Year 4 equal to the number of ports vacated in Year 1 is an estimate trying to approximate when the nature of node split locations would fully vacate I-CCAP resource over time. If a more exact estimate of this is required, future work could look at adding depth to this part of the modelling based on real world experience, as well as producing quantitative analysis as to impact of varying I-CCAP removal starting point assumption in the model.

Timing and size of DAA reduction – improvement via consolidation: As noted in section 5.2, the DAA transition strategy does not include targeting lowly utilized equipment chassis' for complete conversion, to clear for removal. In fact, the timing and size of the DAA savings would be improved with a strategy for setting a minimum utilization point after which legacy architecture would be consolidated. Because a nominal spend would be required to accomplish consolidation, scheduling should be done strategically to minimize costs. One would not want to see this DAA implementation strategy lead to the continued use of a number of rarely-utilized I-CCAP and RF optic chassis' for a long period of time.

Facility/hub space savings: Although this work focused on energy usage, it should be noted DAA would lessen space requirements in the facility/hub locations. The DAA transition strategy slow-rolls the I-CCAP to D-CCAP evolution, so space savings will be evolutionary in much the same way energy usage evolves. Although conversion from I-CCAP to D-CCAP platforms yields minimal space savings for the chassis' themselves, elimination of RF optics, as well as the associated RF combining equipment over time, would free up significant space in the facility. The high-level estimate was that, with fully implemented DAA, facility space would be in the order of half to a third of the space used today. Future work specifically related to space would be recommended to better understand and quantify this.

Future of vCMTS: DAA and D-CCAP are both part of Shaw's network evolution strategy. Additionally, Shaw is evaluating the potential to evolve the D-CCAP architecture to vCMTS as and when that evolution makes sense. vCMTS holds the promise of housing the CMTS functionality in more space and power efficient white-box storage and server devices, potentially further improving energy and space efficiency in facility/hub devices. Future work to assess energy impact of vCMTS would be recommended as and when evolution to the platform is being planned.

7. Outside Plant Power Network Capacity Analysis

As noted in Section 2, Shaw BAU OSP network is standard HFC with ~N+X architecture with plans to maintain or reduce node sizes and cascades lengths. Consistent with HFC deployment norms, the company has implemented a distributed power infrastructure in the OSP to power nodes and active devices in the HFC network. In addition to the nodes and actives, the OSP power infrastructure also supports additional non-HFC devices e.g. small cell, Wi-Fi AP's, etc.

As a part of network evolution of taking fiber deeper, the company intends to use existing power supplies and power supply locations as much as possible. In addition to accommodating network evolution, the company expects the current OSP power infrastructure to be capable of incorporating more of the additional loads related to non-HFC equipment in the future. To date, Shaw has connected in the range of 10,000 Wi-Fi AP's to the coax network, and as an MNO, have connected 4G LTE small cells and are anticipating connecting 5G, both as noted in [3]. The table below outlines the approximate power draw for different types of devices which may be connected to the power network.

Table -1 Non-HFC Device Typical Power

Device	Power Consumption (per device)	Quantity per Site	Total Power/Site	Range
Small Cell	60-180W per device,	1-3	60 - 500W 300W-320W/site typical	~250M (3.5GHz CBRS)
Wi-Fi AP	15-35W	1	25W	~200M (2.4GHz)
IoT	60W	1	60W	2-4Km
Gateway	Modem (15W-20W) + Load	1	80W – 140W	For Cameras, etc. - depends on device under load

Source: Industry Data

To better assess the ability of the OSP power infrastructure to support these goals, as well as understand spare capacity in the power network, the energy audit work performed on the OSP power supplies was used to analyze utilization of the power infrastructure.

The audit found that Shaw’s OSP PS population totaled a little over 20,000, of which 95%+ were analyzed as a part of the audit work. With respect to voltage, data indicated ~75% of the PS’s were 90V with the vast majority of remaining PS’s at 60V, yielding an average plant voltage of 83V. Status monitoring data yielded an average current of 5.8A. Both the 90V/60V split, as well as the ~6A average current, are typical of MSO peer operators [2].

Power supply output voltage and current as acquired via the status monitoring system was used to analyze the PS network. Load investigations were based on Volt-Amperes (VA) calculations from this data, since the output power factor is not known. Key findings from the data included:

- Average output VA across the footprint is 482 VA. This is the average usage for each of the PS’s in the network.
- Average capacity for the mix of PS’s in the power supply universe is 1286 VA. This implies primarily 15A PS’s in the network, with some mix of smaller and larger variants.

This leaves on average ~800VA available across the footprint. As noted in [2], given the majority of the PS’s in the footprint are 15A capacity, good practice dictates power supplies target to run close to 80% of load, or with 3A available. This would equate to a minimum spare capacity requirement for standard network of ~250 VA. At an average of 800 VA available capacity, the power supply network appears to have more than enough headroom for future growth.

While averages provide important data, distribution of the count of power supplies around the average can provide operators a more informed view. Distribution of the power supply count by available VA is as shown in Figure -8 below.

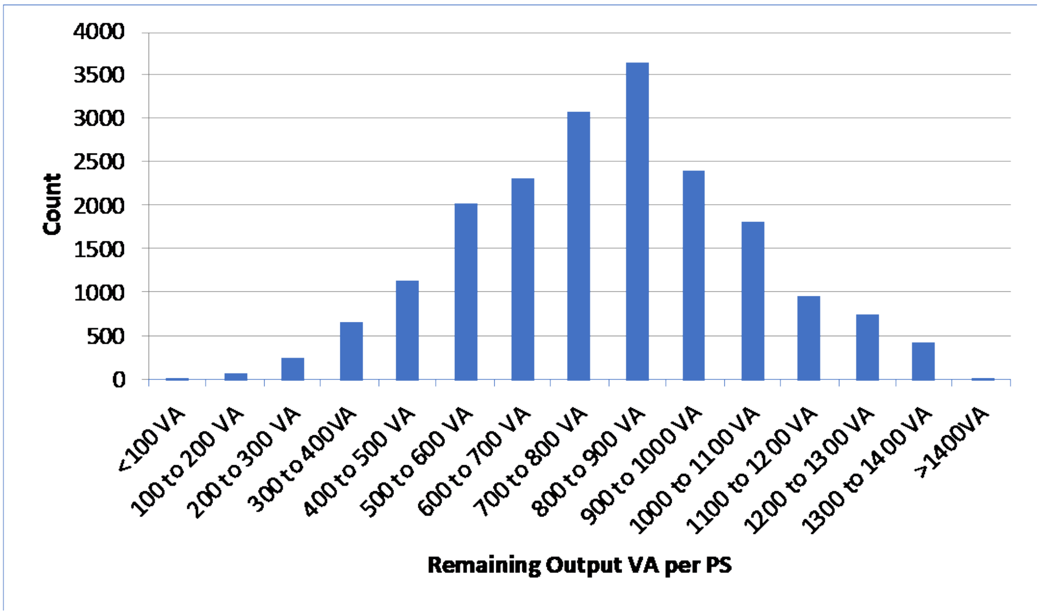

Figure -8 PS Distribution of Remaining Output VA

As demonstrated above, the vast majority of power supplies have in excess of 200 – 300 VA of capability minimum as headroom recommended. Figure -9 re-structures the data above with respect to the percentage of power supplies in each category.

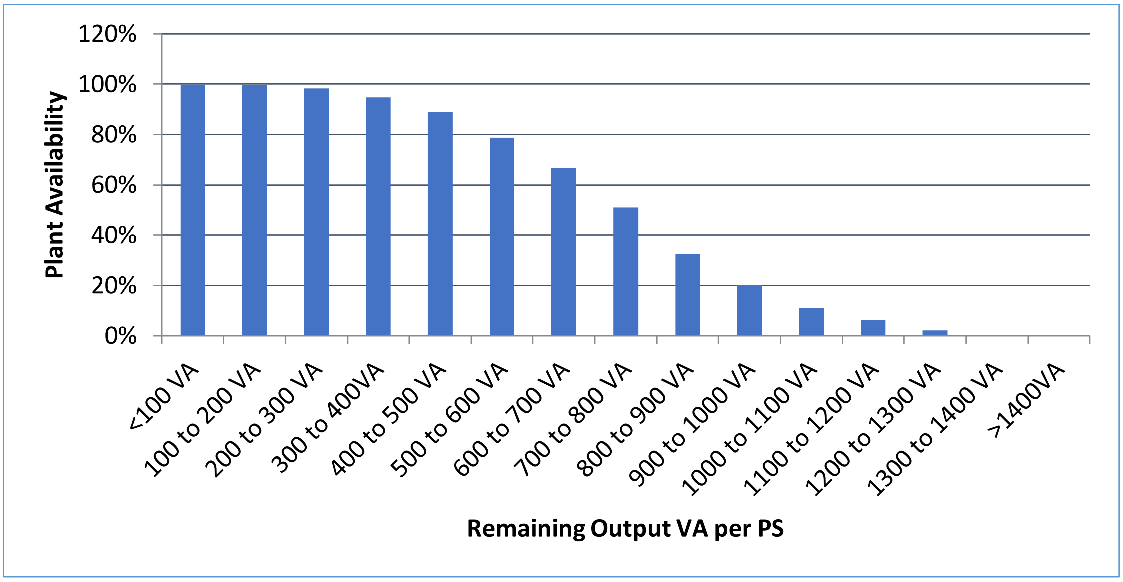

Figure -9 Distribution PS Remaining Output VA %'s

Key takeaways from this analysis include:

- 95%+ Power Supplies have minimum recommended headroom of 250 VA.
- 95% of Power Supplies have 300 – 400 VA or more, providing minimum headroom plus some added capacity for almost all PS's.

- 67% of Power Supplies have 600 – 700 VA or more, leaving minimum 350 VA over recommended minimum for additional loads.

- Average 800 VA of available capacity implies on average 550 VA over recommended minimum for additional loads.

Although the addition of DAA and continued node splits as Shaw evolves to N + 2 architecture will have an effect on reducing headroom, as noted above, it is the company's intent to manage any additional power needs related to network evolution and the addition of non-HFC elements from the existing power supply capacity. This analysis concludes that Shaw's OSP power architecture generally has the capacity to continue supporting the evolution of the HFC network, as well as additional small cell, IoT and gateway devices into the future.

8. Conclusion

The journey to 10G will push Shaw's network to evolve. Implementing DAA will be a key element of the network of the future. To properly balance timing and cost with needs and benefits, the company has chosen an evolutionary strategy for DAA implementation. This strategy will embed DAA into the network in a controlled manner, focusing implementation in greenfield nodes, as well as using DAA nodes in node splits at first, and proactively implementing DAA in existing brownfield only in very limited cases.

Even though it is an evolutionary approach, the Shaw DAA migration still provides energy savings over time in comparison to BAU implementation. Modeling work performed on the company's DAA implementation vs. BAU showed the following:

- DAA OSP energy slightly higher than BAU at 14% vs. 13% cumulative ten-year growth, due primarily to higher powered DAA nodes replacing existing nodes.

- DAA facility energy decreasing to 14% cumulative ten-year growth vs. 34% for BAU, due primarily to swap of I-CCAP with D-CCAP, as well as swapping RF optical devices with CIN elements. This is where the savings of the DAA transition are.

- DAA overall increases energy 14% vs. 20% cumulative ten-year growth for BAU. The Shaw evolutionary approach to DAA implementation slows but does not eliminate BAU infrastructure energy growth.

Capping I-CCAP growth and moving node growth to the lower energy DAA platform mitigates the initial power increase of dual D-CCAP and I-CCAP infrastructures. Real energy reduction in comparison to BAU begins when I-CCAP infrastructure is fully vacated and removed from facilities. Timing is impacted by speed of DAA evolution – if the process is accelerated, energy usage benefits will accelerate, and if it is slowed, energy usage savings will similarly slow. Energy usage reduction speed is also impacted by any actions that drive a less random, more geographically targeted DAA implementation. The faster I-CCAP and RF optical equipment can be fully vacated and eliminated, as opposed to lingering in service at low utilization, the faster energy usage benefits will accrue.

The benefits of immediate conversion to DAA were examined. Modeling from Calgary data indicated immediately converting a hub geography to DAA reduces facility power significantly (~42%), but overall impact is more modest (-5%). This is largely due to a large proportion of total energy (~80%) was OSP energy in the sample areas used – other areas where OSP energy was a more typical 70/30 split would see greater benefit. The small OSP increase also played a minor role. Although facility improvement in both energy usage and space reduction would be meaningful in an immediate transition of a hub to DAA, on their own it would be assumed they would not justify the cost. Full DAA implementation in a hub geography, however, could be a useful tool for solving specific space and power challenges, should they arise.

Finally, energy audit data was used to examine headroom in the OSP power infrastructure to handle DAA evolution, as well as additional non-HFC devices intended to be connected to the power infrastructure in the future. The data showed:

- Average headroom adequate to support evolution to DAA and fiber deeper, as well as multiple additional devices.
- ~2/3rds of nodes have headroom for even the largest anticipated single load device to be added and still maintain minimum headroom.
- 97%+ nodes have minimum headroom in place (20% capacity).

Although each case needs to be examined on its own merit, in general, the Shaw OSP power infrastructure has ample headroom to accommodate DAA evolution and the addition of non-HFC elements to the power infrastructure.

Abbreviations

4G / 5G	fourth generation / fifth generation
BAU	business as usual
CAGR	compound annual growth rate
CI	critical infrastructure
CIN	converged interconnection network
CMTS	cable modem termination system
DAA	distributed access architecture
D-CCAP	distributed converged cable access platform
HFC	hybrid fiber coax
HP	homes passed
I-CCAP	integrated converged cable access platform
IoT	internet of things
LTE	long term evolution
metroE	metro ethernet
MNO	mobile network operator
MAC/PHY	media access control layer / physical layer
MSO	multi-system operator
OSP	outside plant
PS	power supply
PUE	power utilization effectiveness
QAM	quadrature amplitude modulation
RF	radio frequency
RPD	remote phy device

TB	terabyte
VA	volt-amperes
vCMTS	virtual cable modem termination system
Wi-Fi AP	wi-fi access points

Bibliography and References

[1] Nadar Foroughi, “HFC Evolution – The Best Path Forward”, *Proc. Of SCTE Fall Technical Forum*, October 2018, Atlanta.

[2] Todd Loeffeholz, “Opportunities and Challenges of Implementing Wireless – Small Cell/Wi-Fi/IoT”, *Proc. Of SCTE Fall Technical Forum*, October 2018, Atlanta.

[3] Jennifer Andreoli-Fang, et al, “Blueprint for Mobile XHaul over DOCSIS – How Low Latency Xhaul (LLX) and Other Technologies Make DOCSIS and Ideal Solution for Mobile Xhaul”, *Proc. of SCTE Fall Technical Forum*, October 2019, New Orleans.

[4] “Data over Cable Service Interface Specifications DCA – MHAv2 Remote PHY Specification”, CM-SP-R-PHY-I14-200323, CableLabs

[5] John Ulm, Zoran Maricevic, “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures”, *Proc of SCTE Fall Technical Forum*, September 2016, Philadelphia

[6] Rajesh Abbi et al, “Powering the future 10G access networks – An End to End Perspective”, *SCTE Journal of Energy Management, Volume 5, Number 1*