Comparative evaluation of the impact of alternative CCAP implementations on facilities

Integrated CCAP versus Virtual CCAP, CCAP with Remote PHY, and CCAP with R-PHY and RF Analog Overlay

Strategic White Paper

First published at SCTE Cable-Tec Expo 2016 in Philadelphia, PA

The cable industry continues to assess how to evolve hybrid fiber-coaxial (HFC) networks to meet growing customer demands. Distributed access architectures (DAAs) are of particular interest. A comparative analysis of centralized networks utilizing integrated Converged Cable Access Platforms (CCAPs) and three commonly discussed DAAs shows that a virtual CCAP architecture delivers the best power and space reductions (> 85 percent).

Contents

Introduction

A primary objective of the cable industry's Converged Cable Access Platform (CCAP) initiative was to increase quadrature amplitude modulation (QAM) channel density. In 2010, promoters of the CCAP predecessor, the Converged Multiservice Access Platform (CMAP), announced that the proposed chassis could deliver twice the number of channels in half the space. Further efficiencies were envisioned from the reduction of combiners and separate QAM modulators. Power savings were estimated at more than 50 percent.

These efforts dovetailed with the industry's growing focus on energy efficiency. Unveiled by the Society of Cable Telecommunications Engineers (SCTE) in 2014, the "Energy 2020" campaign promoted innovation and set ambitious consumption and cost reduction goals to help cable system operators meet them.

Both initiatives expanded. As products emerged, multiple system operators (MSOs) began placing orders for integrated CCAP equipment. Meanwhile, network functions virtualization (NFV) and software-defined networking (SDN) technologies had made it possible to reimagine the monolithic CCAP head-end and reorganize the network at a functional level.[1] Leading equipment vendors collaborated with CableLabs to assess and specify new architectures that pushed parts or all of key head-end or hub CCAP equipment further out into the network.[2] Figure 1 illustrates CableLabs' comparison of centralized versus distributed CCAP architecture of integrated and modular CCAP/cable modem termination system (CMTS) versus three distributed architectures as shown in Figure 2.[3] In the end, the industry set their sights on two of the three distributed architectures, specifically Remote MAC+PHY and Remote PHY.

1 "Cable Networks at a Tipping Point – How Virtualization and Software-Defined Networking Can Supercharge the HFC Infrastructure," Samir Parikh, 2014 SCTE Cable-Tec Expo, SCTE.

2 "Distributed CCAP Architectures Overview Technical Report," CM-TR-DCA-V-1-150908, Cable Television Laboratories, Inc., 2015.

3 Figure 1 and Figure 2 are from ibid., pp. 7 and 20. The three DCAs that are discussed in the CableLabs report are R-MAC-PHY, Split MAC and R-PHY. Instead of the Split MAC, another form of R-PHY is analyzed here; for example, one that implements an RF Overlay to transport video.

Figure 1. Centralized versus distributed CCAP architecture

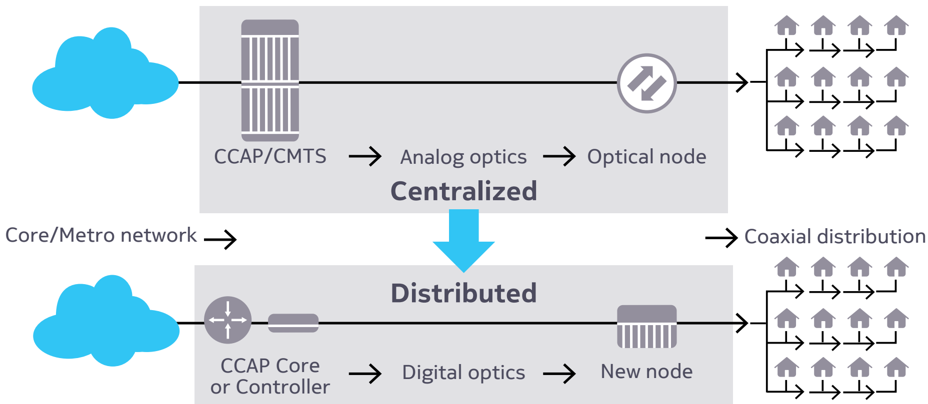

Figure 2. Three forms of D-CCAP architecture

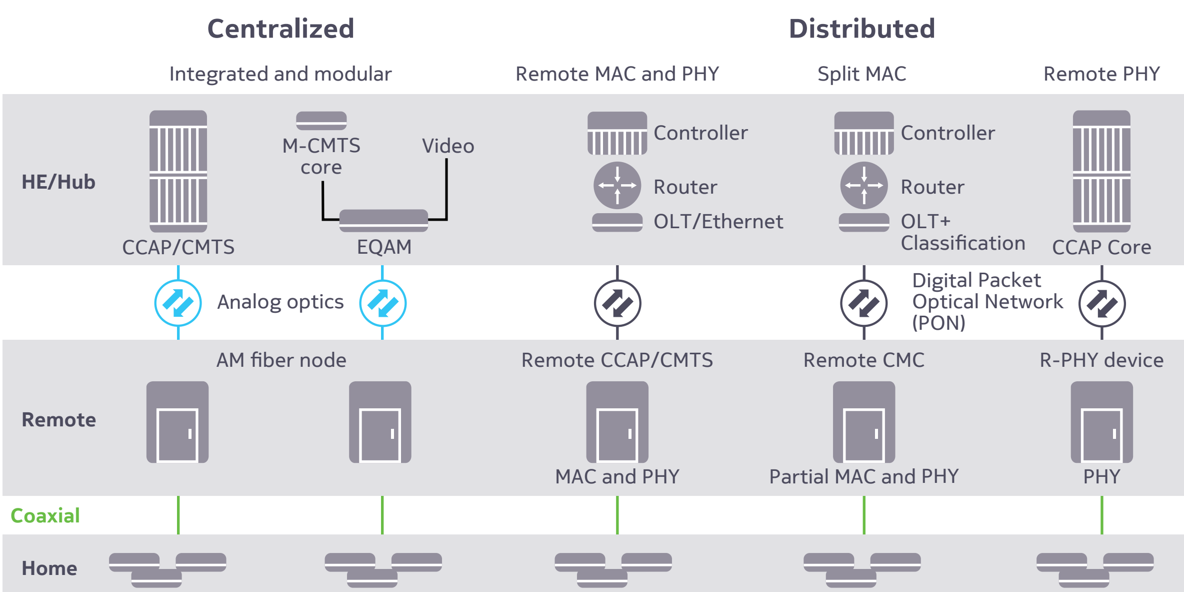

Comparative evaluation

1. Test and hypothesis

Proponents of distributed CCAP architectures (DCAs), also commonly referred to as distributed access architectures (DAAs), aim for two targets: higher levels of network performance and proportionally greater savings. The latter goal is of concern here. Our method of addressing power and space savings is to set up integrated CCAP as a point of comparison for three leading DCA options: Virtual (V)-CCAP, a type of Remote (R)-MAC-PHY; CCAP plus R-PHY; and CCAP plus R-PHY with RF Overlay. (The third option occurs when, for reasons discussed below, video is not integrated within the CCAP core, but instead is deployed using additional RF transport to the node.) Using typical configurations and publicly available specifications, we then calculate power consumption and RU (rack unit) space over a logical hybrid fiber-coaxial (HFC) deployment scenario.

Figure 3. CCAP architectures evaluated in this paper

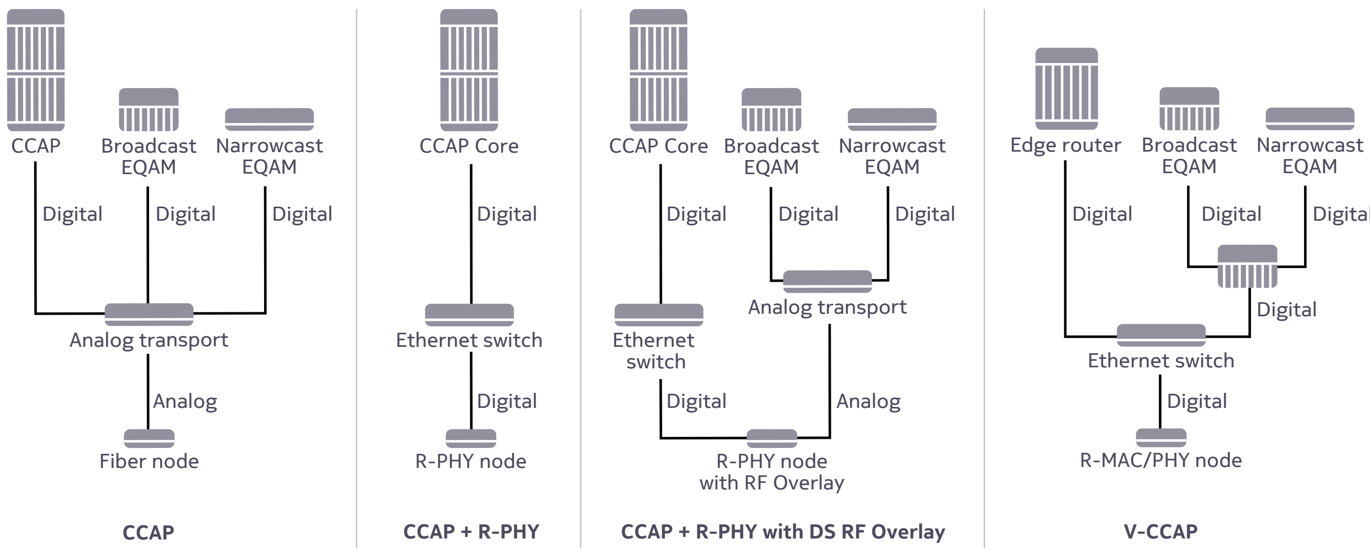

Our hypothesis is that these technologies affect facilities in different ways. That suggestion may not sound controversial, but not all stakeholders would agree. Some have contended that the total power consumption associated with DCAs should be "quite similar."[4] While the power consumed by DCA nodes is indeed roughly equivalent — 160 W to 180 W each, the other differences between these architectures are significant. A key factor is the location of equipment. Rack mounted electronics that reside in head-ends or hubs require power not only for operation, but also for cooling. Calculating that impact calls for a cooling power conversion coefficient, one of several constants assumed in this analysis. We will address those constants in more detail after first discussing our node selection scenario and other configuration inputs.

2. Node scenario selection

This analysis focuses on the impact of DCAs on network facilities, but the outside plant also matters to the extent that the placement of fiber optic nodes relative to households passed (HHP) impacts service group sizes, which in turn affects the deployment of head-end gear. In North America today, a common configuration is one node followed by as many as six active devices. Cable operators though are intent on reducing the number of actives, with node-plus-zero (N+0) serving as a cornerstone for industry-leading network evolution plans.

4 "Looming Challenges and Potential Solutions for Future Distributed CCAP Architecture Systems," Tom Cloonan et al., 2015 SCTE Cable-Tec Expo, SCTE, p. 18.

The reason for this reduction is the need for capacity. To expand bandwidth, operators have two primary tools:

1. Converting spectrum from video to DOCSIS

2. Reducing the number of users sharing the existing DOCSIS spectrum

Most operators will pursue both options, but they are limited to how fast they can convert spectrum. To keep up with demand, operators must also shorten the last mile. Just as with other service providers who have pushed fiber deeper, MSOs will continue to split nodes, pushing fiber closer to the home.

MSOs have their own node splitting strategies. Some will take small steps; others may make dramatic leaps. In either case, N+0 — no active devices between node and customer, remains the logical end game. It serves as a basic assumption of this analysis.

3. Node configuration inputs

In this exercise, the N+0 fiber deep architecture serves 50,000 HHP, requiring an estimated 800 nodes. (By comparison, N+6 would require 100 nodes and N+3, 200). That translates to 63 HHP per node. Moreover, after a node reaches approximately 60 homes, physical topology makes additional balancing or splitting unlikely, which in effect turns the node into a single DOCSIS serving group (SG).

Other assumptions that we hold equal over all four calculations involve spectrum, QAM and DOCSIS channels, narrowcast video and aggregation ratios (see Table 1). While not impacting the math, the top spectrum is assumed to be 1.2 GHz. For the long term, operators may go to 1.8 GHz or 2.0 GHz, but that spectrum remains beyond the reach of today's centralized head-ends because existing lasers cannot support launch levels needed to overcome the associated non-linearities.

Table 1. Shared configuration inputs

Top frequency	1.2 GHz
Node configuration	1 SG
HHP	50,000
Narrowcast video (VoD + SDV) QAMs	16
Broadcast video	30
D3.0 DS	32
D3.0 US	5
DOCSIS SG per node	1
DOCSIS SG to NC video (VoD + SDV) SG	8
Broadcast video ad zones	1
Aggregation ratio (EN to ER)	10%

As for high-speed data, we are assuming DOCSIS 3.0, but the calculations for power and space would not change with the addition of DOCSIS 3.1 spectrum. (In any case, CCAP hardware is assumed to be 3.1 capable.) Overall, to serve

50,000 HHP, we estimate needing 32 DOCSIS 3.0 downstream (DS) channels and 5 DOCSIS 3.0 upstream (US) channels, which is an industry standard for a 42 MHz split and a 1 Gb/s downstream service.

Video delivery requires 30 broadcast and 16 narrowcast video QAM channels. The broadcast channels comprise a single ad zone, limiting the requirements in video head-end equipment. With 63 HHP per node, multiple nodes can be aggregated into a single VoD SG. A target of 500 HHP per VoD SG translates to 8 DOCSIS SGs for each VoD SG.

A final configuration input involves aggregation ratios. The challenge for distributed architectures is that each node has a 10 Gb/s link, but instead of dedicating 10 Gb/s of throughput to each node, operators must manage an oversubscription of demand. In this case, we assume a 10 percent aggregation ratio. For every 10 nodes connected to the Ethernet aggregation switch, there is a single port going to the Ethernet router.

4. Test constants

A second group of assumptions brings us closer to assessing the impact of these architectures on network facilities. These statements concern power and cooling, head-end-based serving groups and RU numbers (see Table 2).

Table 2. Test constraints

Cooling power conversion	1.1594
Analog transport SG	32
CCAP SG	48
CCAP Core SG	72
CCAP RU	16
Analog transport RU	3
CCAP power	5,000 W
Analog transport power	1,078 W
Power cost	$0.15
CCAP core power	5,000 W

Among this set of constants, the cooling power conversion coefficient is critical for any accurate comparison of real costs. Derived from estimates of how many watts are required for cooling to offset the heat generated by electrical equipment, this figure more than doubles the amount of power that these devices require. By the same token, every watt associated with equipment that is moved out of the head-end or hub into an environment that is passively cooled in effect lowers power consumption by more than a watt (1.16 W, rounded up) in cooling expenses. That is on top of any other gains that a DCA might realize through new technologies or design.

The power requirements and RUs associated with analog transport and integrated CCAP devices are derived from publicly available data sheets for representative products. To begin with analog transport, one widely deployed platform drives 32 transmitters and receivers, while consuming 3 RU of space and 1,078 watts of electricity. For this study, we associate each analog system with 32 SGs (one transmitter and receiver for each).

For CCAP with RF connectors, a typical 16 RU deployment is limited to 48 SGs.[5] In a CCAP Core or R-PHY configuration in which RF ports are replaced with Ethernet ports, the assumption is that the chassis can scale up to 72 SGs. Limiting factors, at least on first-generation products, include RF connectors, media access control (MAC) processing and backplane bandwidth.

For CCAP power, data sheets indicate a wide divergence between nominal and maximum levels. Part of the challenge is that as rooms become hotter and fans spin faster, the amount of heat that can be removed narrows relative to the power consumed in doing so, which drives up power consumption and inefficiency. For present purposes, we have used a conservative, nominal number of 5,000 watts, applicable to both integrated CCAP and CCAP Core. At the same time, we realize that 5,000 watts is considerably — even dramatically — lower than what some MSOs believe CCAP devices will grow to consume as more services are enabled in the future.[6]

Finally, there is the question of electrical power costs. Location plays a large role in pricing. In regions with abundant hydropower, $0.12 per kWh or lower may be normal; elsewhere, rates could average as high as $0.18 or $0.20. For this discussion, we peg the rate in the middle of that spectrum, at $0.15 per kWh.

5. Calculated impact on power and space

In the four architectures assessed here — V-CCAP, CCAP, R-PHY, and R-PHY with RF Overlay — the initial numbers used for calculations are largely identical. We begin with the same numbers for fiber optic nodes, HHP per node, DOCSIS SGs, DS/US DOCSIS channels, narrowcast video SGs, and a narrowcast/broadcast channel (see Table 3).

Table 3. Shared calculation figures

Fiber nodes (FNs)	800
HHP/FN	63
HHP/SG	63
DOCSIS SGs	800
D3.0 DS channels	25,600
D3.0 US channels	4,000
Narrowcast video SGs	100
Narrowcast video channels	1,600
Broadcast channels	30

5 The average CCAP DS of 7 slots x 8 ports per card = 56, less N+1 protect card, = 48; average US of 5 slots x 12 ports per card = 60, less protect card = 48.

6 "Energy Efficient Cable Plant Facilities: Strategies to Increase Density through Capacity Reclamation, Site Configuration and Subscriber Based Financial Modeling," Daniel Marut, 2015 SCTE Cable-Tec Expo, SCTE, Tables A1 and A3, p. 35. The author, a principal engineer for Comcast, includes total wattage figures for the ARRIS E6000 and Cisco cBR8 of 8,000 and 11,970, respectively.

When our analysis reaches video, the R-PHY approach diverges. The CCAP design allows for video QAM integration, but a number of challenges, especially surrounding encryption and back-office integration, have impeded the execution of that option. In a pure R-PHY implementation, however, the lack of an analog laser and combining network makes video QAM integration a necessity. That means that the R-PHY approach has no need of the six Edge QAM (EQAM) modulators and one linear broadcast system required by the other solutions[7] (see Table 4). Space and energy for each EQAM modulator are 2 RU and 785 watts, and 4 RU and 1,450 watts for the broadcast system.

Table 4. Broadcast and narrowcast system requirements

	V-CCAP	CCAP	CCAP + R-PHY	CCAP + R-PHY with DS RF Overlay
EQAM systems for narrowcast	6	6	0	6
EQAM systems for broadcast	1	1	0	1

What distinguishes the V-CCAP system further are its centralized video modulator, V-CCAP system components, including the controller and other software systems, and Ethernet switch/router network. Taking up 7 RU of space and requiring 2,775 watts to power, the V-CCAP video modulator, or “engine,” modulates 384 6 MHz video QAM carriers to generate roughly 15 Gb/s of traffic, with each 6 MHz carrier becoming a 40 Mb/s Real-time Transport Protocol (RTP) stream. In this test network of 50,000 HPP, 100 individual streams align with each of the 100 VoD SGs, and one additional stream delivers the broadcast lineup to the entire market.

Other parts of the V-CCAP system are not associated with power or space. But handling 1.2 terabits of traffic requires other components that do incur physical costs. In particular: one 8 RU edge router that consumes 3,405 watts; and ten 2 RU 10 Gigabit Ethernet (GigE) switches that require 263 watts each.

On the other hand, instead of a 10 GigE based video modulator and Ethernet transport, the other three approaches require numerous CCAP systems and (in two cases) analog transport as well (see Table 5). Integrated CCAP, R-PHY and R-PHY with Overlay call for 17, 12, and 12 CCAP systems, respectively, each system 16 RU in size and requiring the conservatively estimated 5,000 watts of power. Deploying CCAP and R-PHY with Overlay also requires 25 and 13 analog transport systems, respectively. (R-PHY only requires forward analog transmitters.) Each 3 RU system needs 1,078 watts of power.

7 Six EQAM modulators are derived from 100 VoD SGs, divided by an assumed 18 VoD SGs per modulator.

Table 5. CCAP and analog transport system requirements

	V-CCAP	CCAP	CCAP + R-PHY	CCAP + R-PHY with DS RF Overlay
CCAP systems	0	17	12	12
Analog transport systems	0	25	0	13

6. Output and results

What remains is to run the numbers, multiplying the unit quantities by their rack sizes, power requirements and kWh cost, where applicable, and adding the cooling power coefficient for a grand total. These results are summarized in Table 6. To view the space and power savings generated by V-CCAP, R-PHY and R-PHY with RF Overlay relative to integrated CCAP, see Table 7 and Figure 4.

Table 6. Total outputs – power and RU

	CCAP	V-CCAP	CCAP + R-PHY	CCAP + R-PHY with DS RF Overlay
V-CCAP + Ethernet routers and switches				
Head-end space (RU)	–	35	–	–
Head-end equipment power (W)	–	8,810	–	–
Broadcast + narrowcast – VoD/SDV EQAM				
Head-end space (RU)	16	16	–	16
Head-end equipment power (W)	6,160	6,160	–	6,160
CCAP and CCAP + R-PHY components				
CCAP systems	17	–	12	12
Analog transport systems	25	–	–	13
CCAP head-end space (RU)	272	–	192	192
Analog transport head-end space (RU)	75	–	–	39
CCAP head-end power (W)	85,000	–	60,000	60,000
Analog transport head-end power (W)	26,950	–	–	14,014
TOTAL				
Head-end space (RU)	363	51	192	247
Head-end equipment power (W)	118,110	14,970	60,000	80,174
Head-end cooling power (W)	136,937	17,356	69,564	92,954
Power costs for 1 year	$335,131	$42,477	$170,247	$227,490

Table 7. DCA savings relative to CCAP

	CCAP	V-CCAP	CCAP + R-PHY	CCAP + R-PHY with DS RF Overlay
Head-end space (RU)	–	312	171	116
Head-end equipment power (W)	–	103,140	58,110	37,936
Head-end cooling power (W)	–	119,581	67,373	43,983
Power costs for 1 year	–	$292,655	$164,884	$107,642
Power costs for 5 years	–	$1,463,274	$824,422	$538,208
Head-end space (RU)	0.0%	86.0%	47.1%	32.0%
Head-end equipment power (W)	0.0%	87.3%	49.2%	32.1%
Head-end cooling power (W)	0.0%	87.3%	49.2%	32.1%

Figure 4. DCA savings relative to CCAP

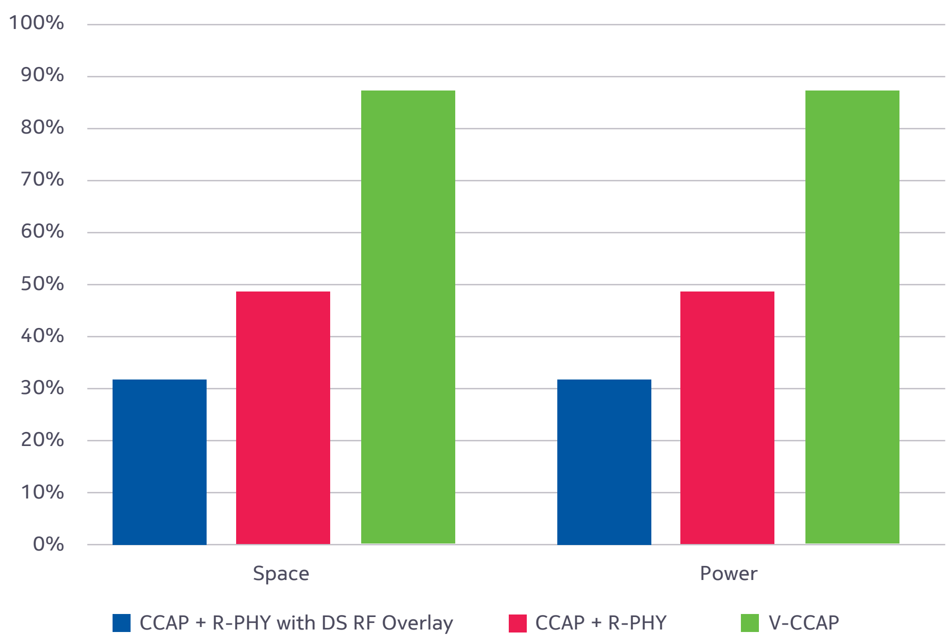

Summary

These results confirm our hypothesis that distributed CCAP architectures differ in their impact upon network facilities. Within the framework of this deep fiber 50,000 HPP test case, they differ widely. As compared with the impact of integrated CCAP on head-end or hub facilities, these DCAs perform as follows:

- CCAP plus R-PHY with an RF Overlay requires 68 percent as much space and power
- CCAP plus R-PHY requires 53 percent as much space and 51 percent as much power
- Virtual CCAP requires 14 percent as much space and 13 percent as much power

The wide variance between V-CCAP and the two R-PHY approaches is noteworthy. Given our conservative estimate of CCAP device power consumption, the gap could be much wider. These results should be encouraging to those looking to avoid space constraints, reduce operating expenses (OPEX) and mitigate the environmental impact of network electronics.

As for space, each cable system has its own limits. Those with large enough facilities may incur minor incremental costs, but the reality is that facilities are often small and crowded, with little room to expand.[8] Given the costs associated with expansion, it is useful to see the different impact on the RU measured space of these DCA architectures.

The reductions in OPEX achieved within these architectures and superiority of the V-CCAP approach are obvious. A total economic value assessment of DCAs would also assess the benefits enabling better signal quality and a more capable and flexible delivery framework.[9] Our results here would be highly complementary to such a study.

Among the industry's Energy 2020 objectives are reducing unit-based consumption and costs, reducing grid dependency and optimizing technical facilities and data center footprints. The efficiencies revealed here refer to the latter goals, which reflect a nationwide concern.[10] The CCAP initiative has already reduced inefficiency within the cable industry. Distributed CCAP takes that effort further, and V-CCAP, much further.

8 "Energy Efficient Cable Plant Facilities," ibid., p. 19, gives a description of how a site expansion from 20 to 35 racks at the Comcast Noblesville, Indiana, hub required the building of a new facility. That project, which collapsed another hub into the new site, reduced kWh/sub by 31 percent, which suggests the high level of aggregate savings attainable by collapsing hundreds of hubs with V-CCAP.

9 For more on overall advantages, see "Distributed Architectures and Converged Access Network," Jorge Salinger, 2016 Spring Technical Forum; CableLabs, NCTA, SCTE.

10 A widely circulated report estimated that the aggregate waste and inefficiency in US data centers at large would increase by 2020 to the equivalent of 50 large (500 MW) power plants. See "Data Center Efficiency Assessment: Scaling Up Energy Efficiency Across the Data Center Industry," NRDC, August 2014.

Acronyms

CCAP	Converged Cable Access Platform
CMAP	Converged Multiservice Access Platform
CMTS	cable modem termination system
DAA	distributed access architecture
DCA	distributed CCAP architecture
DOCSIS	Data over Cable Service Interface Specification
DS	downstream
EQAM	Edge QAM
FN	fiber node
Gb/s	gigabit per second
GHz	gigahertz
GigE	Gigabit Ethernet
HFC	hybrid fiber-coaxial
HHP	households passed
kWh	kilowatt hour
MAC	media access control
Mb/s	megabit per second
MSO	multiple system operator
NFV	network functions virtualization
OPEX	operating expense
PHY	physical layer
QAM	quadrature amplitude modulation
RF	radio frequency
R-MAC-PHY	Remote MAC-PHY
R-PHY	Remote PHY
RTP	Real-time Transport Protocol
RU	rack unit
SCTE	Society of Cable Telecommunications Engineers
SDN	software-defined networking
SDV	switched digital video
SG	service group
US	upstream
V-CCAP	Virtual CCAP
VoD	video on demand

References

"Cable Networks at a Tipping Point – How Virtualization and Software-Defined Networking Can Supercharge the HFC Infrastructure," Samir Parikh, 2014 SCTE Cable-Tec Expo, SCTE.

"Distributed Architectures and Converged Access Network," Jorge Salinger, 2016 Spring Technical Forum, CableLabs, NCTA, SCTE.

"Distributed CCAP Architectures Overview Technical Report," CM-TR-DCA-V-1-150908, Cable Television Laboratories, Inc., 2015.

"Energy Efficient Cable Plant Facilities: Strategies to Increase Density through Capacity Reclamation, Site Configuration and Subscriber Based Financial Modeling," Daniel Marut, 2015 SCTE Cable-Tec Expo, SCTE.

"Looming Challenges and Potential Solutions for Future Distributed CCAP Architecture Systems," Tom Cloonan et al., 2015 SCTE Cable-Tec Expo, SCTE.

"Data Center Efficiency Assessment: Scaling Up Energy Efficiency Across the Data Center Industry," National Resources Defense Council (NRDC), August 2014.

Nokia Oyj
Karaportti 3
FI-02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Product code: PR1608022175EN (September)