Distribute or Die: Hybrid Fiber Coax Hanging by a Cable

Samir Parikh
Gainspeed, Inc.

Abstract

The hybrid fiber coax (HFC) networks currently in place were designed to carry broadcast analog video, not the wide ranging high-speed data (HSD) services, IP video products, teleconferencing, cloud computing and other business services they need to deliver today.

There is enough capacity potential in today's hybrid fiber coax (HFC) networks to satisfy bandwidth demand, even with a 40% to 60% annual growth rate. However, a new architectural approach is required in order to unlock this potential.

This paper will explore how cable operators can re-imagine their networks and extend CCAP to its next logical step by virtualizing and distributing functions and services. Doing so will dramatically reduce power and space consumption in the hub and headend, lower costs, speed up service velocity and enable new technologies such as DOCSIS 3.1.

INTRODUCTION

Cable networks are being overwhelmed by a deluge of IP Video traffic. As a result, cable operators face escalating costs to deliver bandwidth and are forced to invest heavily in equipment and infrastructure just to maintain network performance levels. Hybrid fiber coax (HFC) networks were designed to carry broadcast analog video, not the wide ranging high-speed data services, IP Video products, teleconferencing, cloud computing and other business services they need to deliver today. New services and growing bandwidth demands are driving a 40% to 60% average annual compound growth in capacity requirements.

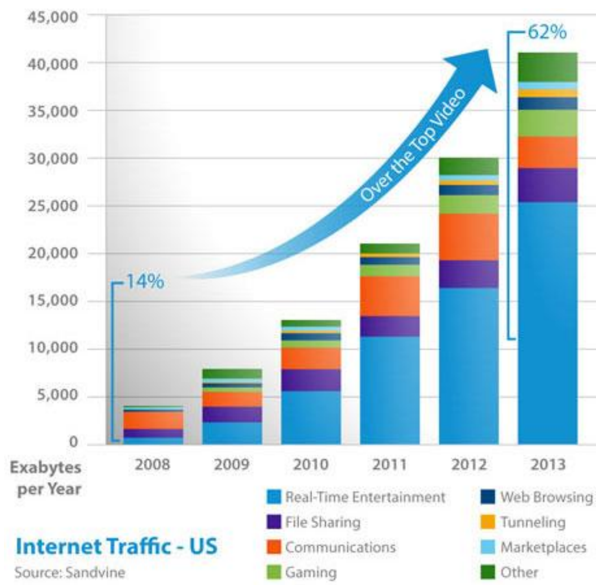

Figure 1: US Internet Traffic Profile

IP Video now accounts for over 62% of all Internet traffic – 67% of downstream traffic – and growing. These trends show no sign of fading and have radically changed everything.

At the time the Converged Cable Access Platform (CCAP) specification was conceived, operators were seeing tremendous demand for both IP services as well as legacy video services (video on demand and switched digital video). CCAP, which combines the physical edge QAM and CMTS devices to support both of IP and legacy video services on a single platform, made a lot of sense. However, just as platforms based on the CCAP specification are beginning to come to market, everything has changed. Growth in legacy video has dropped dramatically and IP Video (both operator provided as well as

Over-the-Top (OTT)) has become the dominant consumer of Internet bandwidth. Pressure from OTT is driving an accelerated shift away from QAM video to IP video, pushing the limits of DOCSIS capacity available in the network today.

The good news is that with DOCSIS 3.0, cable operators have close to 6 Gbps of raw IP capacity on a 1 GHz HFC plant. (DOCSIS 3.1 will further increase this.) The bad news is that there are numerous hurdles in unleashing this capacity:

- Existing CMTS have very limited scale – 8-16 downstream (DS) channels per port. Even the newer CCAPs only scale up to 32 DS channels per port. Scaling a service group (SG) to 128 or 158 DS channels will therefore require a lot of equipment in the headend. Perhaps in a couple of generations, CCAP systems will come to fulfill the promise of a single port per SG supporting full-spectrum services.
- Most networks only support up to 860 MHz of spectrum, and the bulk of this spectrum is allocated for legacy video services (analog video and digital QAM video – broadcast and narrowcast). Upgrading the network to 1 GHz or higher is a costly undertaking.
- Spectrum allocation changes require manual processes. RF combining and splitting networks need to be adjusted. Power levels need to be re-balanced in the headend and in the fiber node. This acts as a natural rate-limiter for capacity increases in the network.

SCALING CHALLENGES

Spectrum conversion to DOCSIS in conjunction with node-splits has been the modus operandi for operators for years now. But continuing down this path is unsustainable without a major change in the architecture. With each node split, the strain on headend resources (power, cooling, space) is being exacerbated.

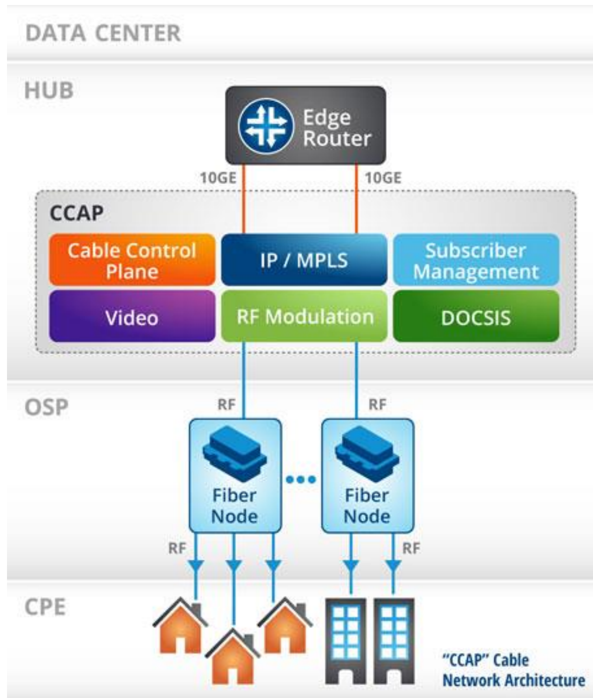

Figure 2: Network Architecture with CCAP

Let's take a very simple example to illustrate the problem operators face in managing facilities. In this example, an operator is going to add 1,288 new SGs in a given headend over a 7-year period. Figure 3 below shows the cumulative space required over this 7-year period.

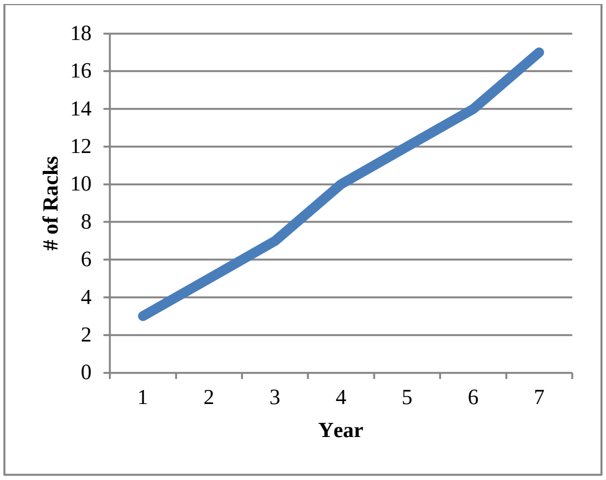

Figure 3: Cumulative Rack Space for Centralized CCAP Deployment

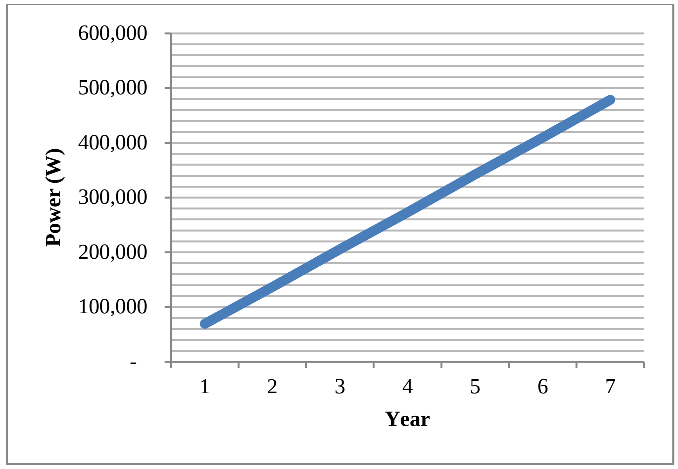

Figure 4: Cumulative Power Consumption for Centralized CCAP Deployment

A total of 17 racks are fully utilized after 7 years to accommodate this growth, and this is a very optimistic scenario as it makes the following assumptions:

1. The CCAP system supports 64 SG per device.
2. The CCAP supports full-spectrum services per DS RF port so that there is no need for RF combining and splitting. This requires the CCAP to support:
    a. 158 DS channels per port
    b. Flexible allocation of services to DS channels
3. The CCAP supports narrowcast and broadcast video encryption – the CCAP ECMG system is in place to enable support for Motorola Digicipher II and Cisco PowerKEY encryption schemes.
4. Operators collapse *all* services onto the CCAP device and retire existing QAM and CMTS platforms.

Achieving all of these requirements really isn't possible today, so the actual space needed will be a lot greater. Figure 4 shows the cumulative power consumption for this example, which again is optimistic given the constraints listed above.

Table 1 below has the raw values for the charts in Figure 3 and 4.

	YEAR						
	1	2	3	4	5	6	7
# of Racks	3	5	7	10	12	14	17
Power (kW)	69	136	205	272	342	409	478

Table 1: Raw Numbers for CCAP Power and Space Requirements

In order to address this challenge, the industry needs to develop and embrace a new architecture which minimizes power, cooling and space requirements in the headend. The architecture must leverage the power of software to provide operators with a flexible mechanism to reallocate spectrum. Furthermore the architecture must not force operators to rip and replace the legacy video EQAM infrastructure and integrate a new solution into the video back-office. It must also eliminate analog transport and use standards-based Ethernet to lower transport costs and improve fiber utilization. Finally, the architecture must lower costs (CAPEX and OPEX) and increase service agility and performance. The industry needs **Virtual CCAP**.

VIRTUAL CCAP

So what is a Virtual CCAP? At its core, it is a CCAP. It delivers all of the features and functions of a CCAP – all existing services are supported without modification, all existing customer premise equipment (CPE) works without modification, and existing back-office systems remain in place and unmodified. However, Virtual CCAP takes CCAP to the next logical step by distributing functionality and virtualizing control and management.

Rather than viewing CCAP as a physical platform, it can be viewed as a collection of base functions that support the services cable operators offer. These base functions include:

1. Cable Control Plane
2. IP/MPLS Control and Forwarding Plane
3. Subscriber Management
4. Video (QAM) Processing
5. DOCSIS Processing
6. RF Modulation

Software Defined Networking (SDN) technologies enable us to move away from the notion that all of these functions must be co-located in a monolithic system that is centralized in the headend.

Taking these base functions and bringing them together in a virtualized environment enables operators to leverage best-of-breed products for each function. The resulting eco-system will provide operators with the most feature-rich and reliable, yet lowest cost, solution.

Placement of these base functions can now be rearranged, moving each to the device and location in the network that optimizes the overall deployment.

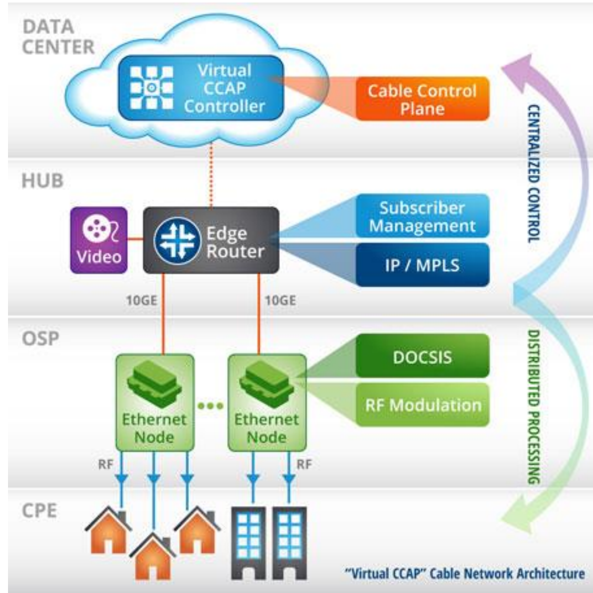

Figure 5: Network Architecture with Virtual CCAP

Cable Control Plane functions are collected into a single package and incorporated into the Virtual CCAP Controller. The Virtual CCAP Controller can be placed in any data center in the network, as it is a control and management entity only.

IP/MPLS and **Subscriber Management and processing** are collapsed onto the Edge Router in the hub. The Edge Routers commonly deployed in operator networks have all the native capabilities to support this; they just lack any cable awareness. This is addressed by the Virtual CCAP Controller.

Video (QAM) Processing stays in the hub, but migrates to Ethernet to take advantage of digital transport to the Ethernet Node. This enables operators to maintain the infrastructure that is already deployed, maximize the usable life of the equipment and avoid the costly integration of a new platform into their video back-office systems.

DOCSIS Processing (MAC and PHY) as well as **RF Modulation** is moved into the node, creating a new class of nodes – Ethernet Nodes. This eliminates analog optical transport and moves everything to digital, Ethernet transport. This improves RF performance, reducing customer calls and truck rolls and setting up the network for deployment of DOCSIS 3.1. The data-carrying capacity of the fiber is increased and operators can now converge their HFC transport systems with their Carrier Ethernet transport systems.

Looking at the same deployment scenario described earlier, it is clear that the Virtual CCAP solution dramatically changes the space and power equation.

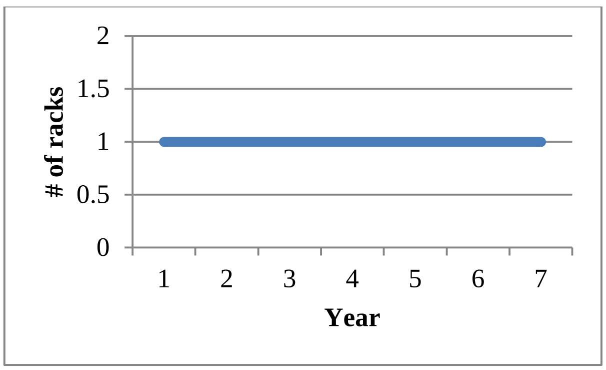

Figure 6: Cumulative Rack Space for Virtual CCAP Deployment

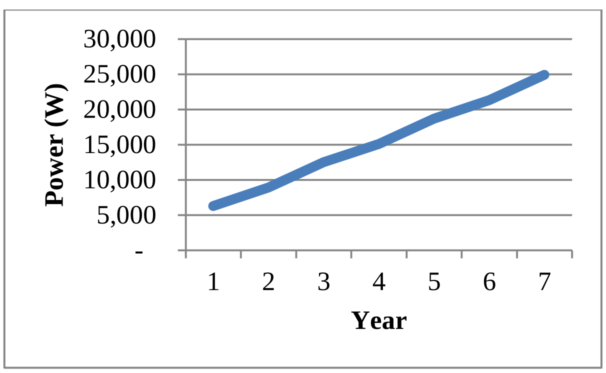

Figure 7: Cumulative Power Consumption for Virtual CCAP Deployment

	YEAR						
	1	2	3	4	5	6	7
# of Racks	1	1	1	1	1	1	1
Power (kW)	6.3	8.9	12.5	15.1	18.7	21.3	24.9

Table 2: Raw Numbers for Virtual CCAP Power and Space Requirements

Comparing the Virtual CCAP deployment against the Centralized CCAP deployment, operators can realize the savings described in Table 3 below.

	YEAR						
	1	2	3	4	5	6	7
Cumulative Power Saving (thousand's $)	82.9	167.6	254.1	338.7	425.2	509.9	596.4
Space Savings (Racks)	2	4	6	9	11	13	16

Table 3: Power and Space Savings of Virtual CCAP

This dramatic reduction in power and space enables operators to consolidate headends / hubs. Operators can eliminate smaller facilities and centralize the equipment in larger headends. This not only saves hundreds of thousands of dollars of costs related to ongoing site maintenance, it eliminates the need to spend millions of dollars to augment smaller facilities to accommodate the growing power and space demands of CMTS/CCAP devices.

DOCSIS 3.1

A Virtual CCAP clearly provides enormous financial and technological benefits compared to a traditional CCAP for DOCSIS 3.0. The difference becomes ever more pronounced with DOCSIS 3.1, especially with respect to the ease of transitioning.

DOCSIS 3.1 is designed to enable 10 Gbps downstream (DS) and 1 Gbps upstream (US) transmissions on the existing coax plant. There are a host of challenges in achieving this:

1) Improving end-to-end signal-to-noise (SNR) in order to achieve 4096 QAM modulation in the downstream and 1024 QAM modulations in the upstream
2) Increasing the top-end spectrum in the DS to 1.2 GHz
3) Changing the upstream split from 42 / 65 MHz to 200 MHz
4) Reallocating spectrum from existing services (analog and digital video, DOCSIS 3.0) to use for D3.1

The single largest impediment to achieving the necessary SNR for ubiquitous support of 4096 QAM in the DS and 1024 QAM in the US is the analog optical transmission system. The specific values vary based on a host of factors (age of system, top-end frequency supported, number of wavelengths carried on a single fiber, length of fiber, etc.), but replacing the analog optical transmission system with Ethernet provides a 3-7 dB improvement in end-to-end SNR. This is crucial to achieving the higher order modulation rates to make the cost of migrating to D3.1 worthwhile.

A traditional monolithic CCAP doesn't address the issues with analog optical transmission. It actually requires the use of the existing analog optical transport systems, whereas a Virtual CCAP transitions the optical transport to Ethernet.

With a traditional monolithic CCAP deployment, there are multiple changes that operators must perform in order to deploy D3.1 with the new frequency splits in the US and DS.

1) Replace existing splitter-combiner networks
2) Replace the analog optical transport systems
3) Replace the fiber nodes
4) Replace amplifiers
5) Replace linecards in CCAP to D3.1 capable versions

By contrast, in a Virtual CCAP deployment, all of the headend work is eliminated. A simple diplexer adjustment in the Ethernet Node changes the US split. Amplifiers will still need to be replaced to support the US and DS frequency range changes. The rest is just a system configuration change to reallocate spectrum to D3.1. The Ethernet transport infrastructure does not need to be touched at all; it already has the capacity required to support D3.1.

Alternatively, operators can also choose to deploy D3.1 within the existing US and DS frequency plan in order to take advantage of improved spectral efficiency while minimizing the outside plant changes.

In this scenario, for a traditional monolithic CCAP deployment, the following changes would need to be done:

1) Replace linecards in CCAP with D3.1 capable versions
2) Manually adjust RF power levels at the analog optical transport system in the headend and the fiber node in the outside plant to accommodate spectrum allocation changes.

For a Virtual CCAP deployment, a simple configuration change re-allocates spectrum to D3.1 and adjusts the RF power balance. No hardware needs to be replaced, no need to visit the node to manually adjust power levels.

CONCLUSION

As operators evolve their network infrastructure to address ever-growing and changing consumer demand, they need a new network paradigm to satisfy bandwidth requirements. Existing products limit the network architecture, increasing costs, limiting flexibility and impeding service deployment.

Technology advances in SDN and NFV have opened up new network architecture options. Virtual CCAP is a prime example of that. A Virtual CCAP deployment enables operators to support all of their existing services while lowering overall CAPEX and OPEX. It addresses the facilities challenges faced by operators, streamlines capacity additions to the network and sets the network up for a simple migration to DOCSIS 3.1. This is done in a way that seamlessly integrates with existing back-office systems and processes, and works with all current and proposed CPE.

It provides the foundation for operators to evolve their business and address the growing threat from FTTH and OTT competitors.