SERVICE AND MANAGEMENT ORCHESTRATION IN DISTRIBUTED HIGH SPEED DATA NETWORKS

Dan Torbet, Tom Cloonan, and Andy Aftelak
ARRIS
Steve Krapp
MaxLinear

Abstract

High speed data networks to the customer's premise are evolving rapidly, driven by the need to innovate quickly and to reduce the cost of both operation and capital spending. Distributed DOCSIS architectures, whether remote PHY or remote MAC/PHY, will create more intelligent nodes in the network that require management and create valuable data on network performance and effectiveness. Mixed distributed and centralized network Architectures will create management domains with different requirements, straining traditional management infrastructures and organizations. Management paradigms from the management of data centers, such as software defined networking (SDN) and network virtualization (NV or NFV) are being applied to telecom networks promising agility and innovation in the creation and monitoring of network services. Managing and innovating on these evolving network platforms will be key to both maximizing the return on past and future investments, enabling service velocity, and containing the operational costs of the network.

We describe an approach of using a service orchestrator that evolves with the network architecture. Using SDN principles, we show how such a cloud based orchestrator in conjunction with purpose built applications can begin to unburden the existing network elements in the collection and management of network data, and to create value added services for network configuration, network optimization, and ultimately network planning. The purpose built applications in conjunction with the Service Orchestrator can be well placed to orchestrate network management between traditional and distributed equipment domains, and play an essential part rapid configuration and control of the entire network.

INTRODUCTION

The cable industry is currently experiencing a wide range of rapid changes in its services as it moves into the future. These changes encompass not only the access networks that have been leveraging the Hybrid-Fiber Coaxial (HFC) plant, but how those networks are architected, provisioned, and managed. The HFC plant continues to carry a wide variety of services that depend on multiple transmission methods including traditional broadcast, IP HSD, video, and telephony over DOCSIS and PON, and cell backhaul over fiber based Ethernet. Increasingly, operators are migrating into a highly service oriented methodology of enabling, managing, monitoring, and provisioning services over this HFC architecture As a result, the HFC plant as a last-mile delivery system will undoubtedly change immensely in the next several years to support these changes in services.

The changes in the access network and HFC network will be driven by many advances, including:

Video Upgrades

- A transition from Standard Definition/High Definition video to Ultra-High Definition (4K, 8K) video
- A transition from MPEG-2 and H.264 encoding to HEVC encoding

- A transition from QAM-based video distribution to IP-based video distribution

High-Speed Data Upgrades
- A transition to 1+ Gbps downstream service level agreements (SLAs)
- A transition to symmetrical data services
- A transition from SC-QAM-based DOCSIS 3.0 channels to OFDM-based DOCSIS 3.1 channels
- A transition from predominantly 750 MHz downstream systems to 1.0/1.2 GHz downstream systems
- A transition from 42/55 MHz upstream systems to 85/204 MHz upstream systems
- Transitions from integrated and purpose built CMTS equipment to distributed devices where specific functions can be unbundled

Alternative Wireless Services
- The addition of cell tower back-haul support
- The augmentation of existing commercial Wi-Fi support
- Wi-Fi services such as voice
- IoT

While some of the above changes are definitely aimed at reducing the bandwidth requirements on the HFC network (ex: HEVC encoding, OFDM), most of these changes will push the network in the other direction and will likely force much higher bandwidth capacity requirements on the network equipment of the future. Each of these changes also requires that for the operator to be successful, the overall introduction and velocity of services must evolve.

Multiple System Operators (MSOs) are planning several actions and investigating many new technologies and architectures as they prepare to provide for this increased bandwidth capacity. As an example, High-Speed Data traffic engineering studies are currently underway to predict the future bandwidth capacity requirements as MSOs look towards the 2020 decade and beyond. These studies indicate that by the year 2020, the average downstream bandwidth per subscriber (in the busy-hour) may grow from its current level of 1 Mbps to ~5 Mbps, and the average upstream bandwidth per subscriber (in the busy-hour) may grow from its current level of 100 kbps to ~330 kbps. Maximum downstream SLA levels that are (in many cases) at 200 Mbps today will grow to be greater than 1 Gbps by 2020, and maximum upstream SLA levels that are at ~20 Mbps today may grow to be 75 Mbps by 2020. Some MSOs are even beginning to explore the likelihood of providing symmetrical 1 Gbps services by the 2020 time-frame.

All of these bandwidth trends will place a heavy strain on the existing HFC infrastructure and will lead to capitalizing on the higher spectral efficiencies and larger spectral widths supported by newly-arriving DOCSIS 3.1 equipment. While DOCSIS 3.1 improvements will undoubtedly help MSOs with increased bandwidth capacities, other techniques for supporting the bandwidth requirements will also be required. In particular, MSOs are looking at many new architectural approaches in an effort to provide even more bandwidth improvements.

One approach is to continue to perform node-splits to reduce the number of subscribers that share the bandwidth within a Service Group. As Service Group sizes decrease, the number of Service Groups will obviously need to increase. Future I-CCAP systems are being developed with much higher Service Group densities than exist today. However, another approach that permits MSOs to increase the density of Service Group support in the headend is to move toward Distributed CCAP Architectures.

DISTRIBUTED CCAP ARCHITECTURES

Distributed CCAP Architectures (DCAs) are

a new class of architecture for providing broadband digital services (such as high-speed data, video, and voice) over HFC plants. These DCA architectures move some (or most) of the CCAP functionality (including PHY and optionally MAC processing sub-systems) into the fiber nodes that exist within the outside plant. They therefore require the use of Ethernet-based or Passive Optical Network (PON)-based digital optics (instead of analog optics) for transmissions across the fiber within the HFC plant. DCAs offer many potential benefits to MSOs, including:

a) Reducing the required space and power requirements in the headend or hub

b) Increasing the number of simultaneous lambdas that can be wavelength-division-multiplexed on a single fiber (for support of future node-splits) based on the switch from amplitude modulation (AM) optics to digital optics

c) Circumventing nonlinear optical noise effects and increasing the End-Of-Line Signal-to-Noise Ratios (SNRs) and increasing the corresponding spectral efficiencies of DOCSIS 3.1 transport by moving the transmission of the signals closer to the device that a given node serves

Two types of Distributed CCAP Architectures are currently being explored for DOCSIS based access networks. They are:

1) Remote PHY- This architecture moves the DOCSIS PHY (Downstream and Upstream) into the fiber node or a remote PHY shelf and keeps the DOCSIS MAC in a CCAP core in the headend

2) Remote MACPHY- This architecture moves the DOCSIS MAC and PHY into the fiber node

Each of these architectures has its own pros and cons, but both suffer from some common challenges. One of these common challenges is the increase in the number of intelligent devices (managed elements) within the network. Whereas an integrated CMTS or CCAP would be managed as single managed element, the same system built using Distributed CCAP Architectural approaches might have hundreds (or even thousands) of managed elements (including CCAP cores and Remote PHY devices or Remote MACPHY devices). This is because fiber nodes will take on these traditionally integrated functions in a DCA. This increase in the quantity of managed devices will create new challenges that existing management paradigms will struggle to address. A new approach to management will be required. The remaining sections will discuss some techniques that could simplify the management of the large number of DCA elements.

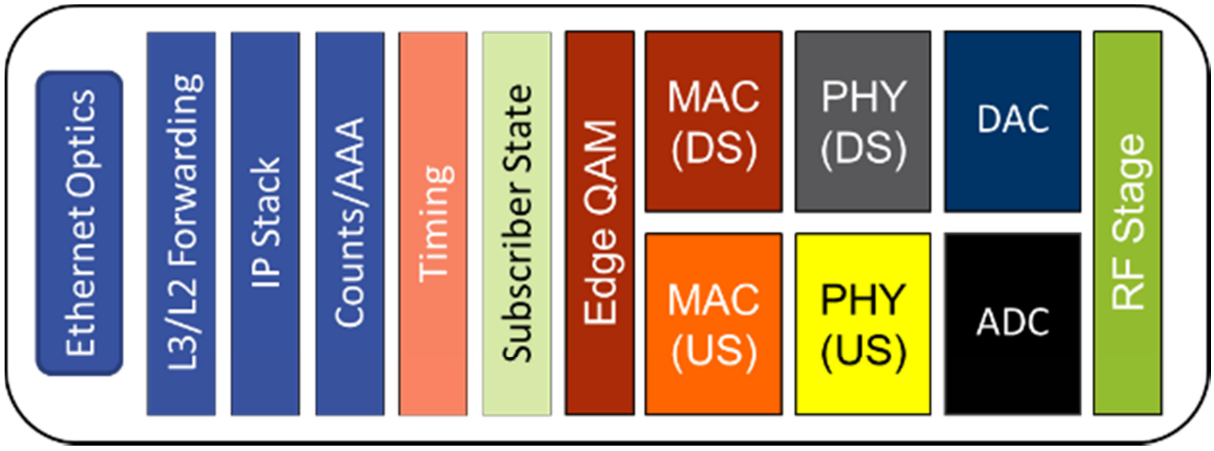

Figure 1: Basic CCAP Data Plane Architecture

Figure 1 shows a very basic data plane structure of a typical Converged Cable Access Platform (CCAP). A cable modem termination system (CMTS) would be similar minus edge QAM functionality. All of the components that are needed for a CCAP or CMTS to forward a given packet to and from the customer premises equipment (CPE) devices are integrated into a single enclosure and located in a headend. Figure 2 depicts the same basic breakdown when a Remote PHY architecture is implemented.

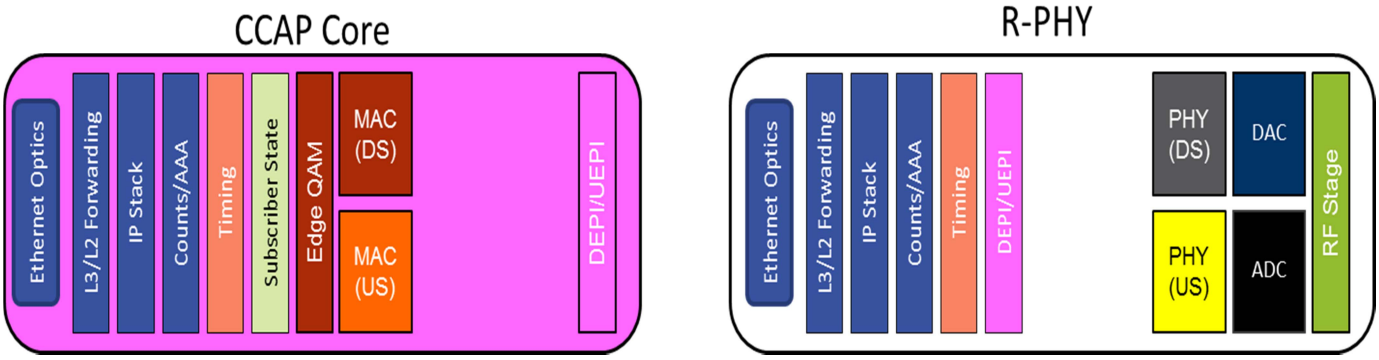

Figure 2: Remote PHY Architecture

In this scenario, the packet handling and DOCSIS MAC related functions are in a CCAP core element (traditionally located in a headend or hub site) and the downstream DOCSIS and video channels as well DOCSIS upstream channels are moved into the optical nodes in the outside plant that service customer homes. Where a single I-CCAP chassis may have serviced 40 fiber nodes, a CCAP core that supports the remote PHY architecture might support many times that number of fiber nodes.

In order to manage this explosion of devices, each of which needing to be provisioned and managed over their lifespans, operators are increasingly looking to evolving technologies like Software Defined Networking (SDN) and Network Function Virtualization (NFV) to help them cost effectively roll out these networks and

the services that run over them. In the next section we'll discuss SDN and NFV in more detail.

SDN, NFV AND ITS APPLICATION TO ACCESS INFRASTRUCTURE

Software Defined Networking (SDN) as a networking paradigm has been evolving since its inception the early 2000's along with the OpenFlow Specifications which were developed in 2008 and 2009. This paradigm is based on the notion that traditional networking equipment has been based on an integrated control and data plane built on proprietary hardware and software modules that are expensive and require long cycles for new features. SDN seeks to address two main goals. The first is the separation of the control plane from the data plane. As a result of this modification, several benefits can be realized, including:

- Networking equipment can become more operationally cost effective
- Commercial off-the-shelf (COTS) servers can be utilized, resulting in the ability to capitalize on the elasticity of COTS servers

The second goal of SDN is the enablement of a higher degree of network programmability, which results in a reduced time to market for features and capabilities. Where a new routing protocol or filtering method may have taken many months for a given vendor to implement on their hardware and software, an SDN approach may deliver that same functionality in a fraction of the time and, potentially, at a fraction of the cost. These benefits are driving the desire to migrate as quickly as possible to architectures that can adopt the new SDN paradigm.

A second technology trend has emerged alongside the SDN efforts. This effort is called Network Function Virtualization (NFV). It has taken advantage the notion that with the separation of the control and data planes, portions of the functions that were previously performed by proprietary equipment and servers could be moved to servers within the network and scaled separately from the hardware and control planes as virtual network functions (VNFs).

Where SDN provides the basic framework for increased programmability of the network and its equipment, Network Function Virtualization (NFV) takes this concept another next step forward by providing capabilities to house these data plane functions on hardware/servers that can be independently scaled to match the needs of the services they support. While SDN and NFV can be deployed independently of each other, the combination of the two technologies makes for a powerful combination.

Increasingly, MSOs and other broadband network providers have been working diligently to determine how and where these technologies can map into their business plans and how to apply these concepts to the access network and data center architectures. Multiple international standards organizations like the BroadBand Forum (BBF), European Telecommunications Standards Institute (ETSI), Internet Engineering Task Force (IETF), and the Open Network Foundation (ONF) have taken up the call to develop standards for their respective members.

Indeed, many telco providers like AT&T and Verizon have already made significant headway in the deployment of SDN and NFV in their networks and already realizing significant gains from it. In the cable space, several MSOs have been very active in defining how they would like to see these concepts deployed in their network. Some MSOs have chosen to focus SDN and NFV operations on the data center before moving out to the access network equipment while others have been equally

focused on the access network and the home CPE.

As mentioned previously, the DCA brings with it several challenges. One challenge is the management and configuration of all of the Remote PHY Devices (RPDs) that are serviced by a given CCAP core. Each of the RPDs is connected to a given CCAP core using an interconnected network of some size. RPDs can either be directly connected to the CCAP core, or there may be a network consisting of one or more network switches / routers between the CCAP core and the RPDs it services.

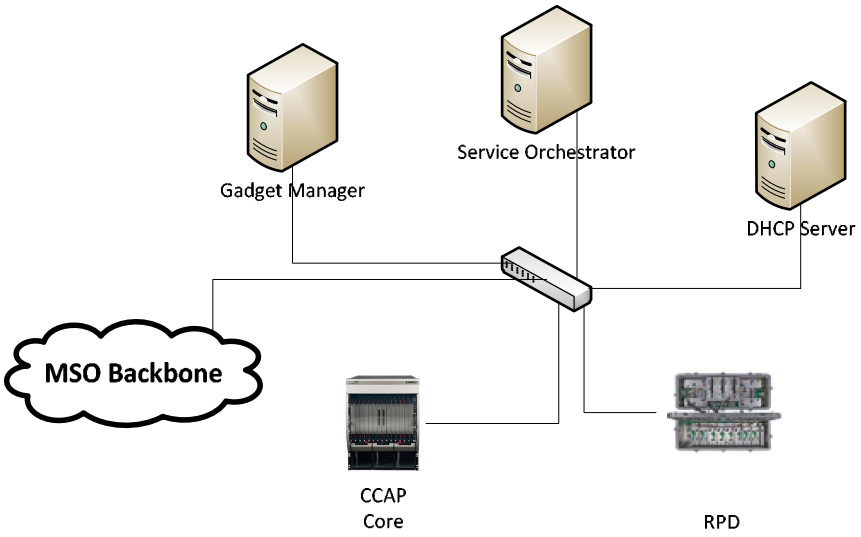

Figure 2: Generic Physical Network

As shown in Figure 4 below, when an MSO migrates from an integrated CCAP chassis to a DCA based architecture, the number of provisioned and managed devices increases very quickly. Where a DCA node in the diagram may have represented a MAC domain and a port on an upstream or downstream line card, the DCA node is an entity that has a MAC domain in the headend / Hub (perhaps virtualized across several servers in the headend /Hub) and then an RPD and the intervening network switches between the node and the hub site.

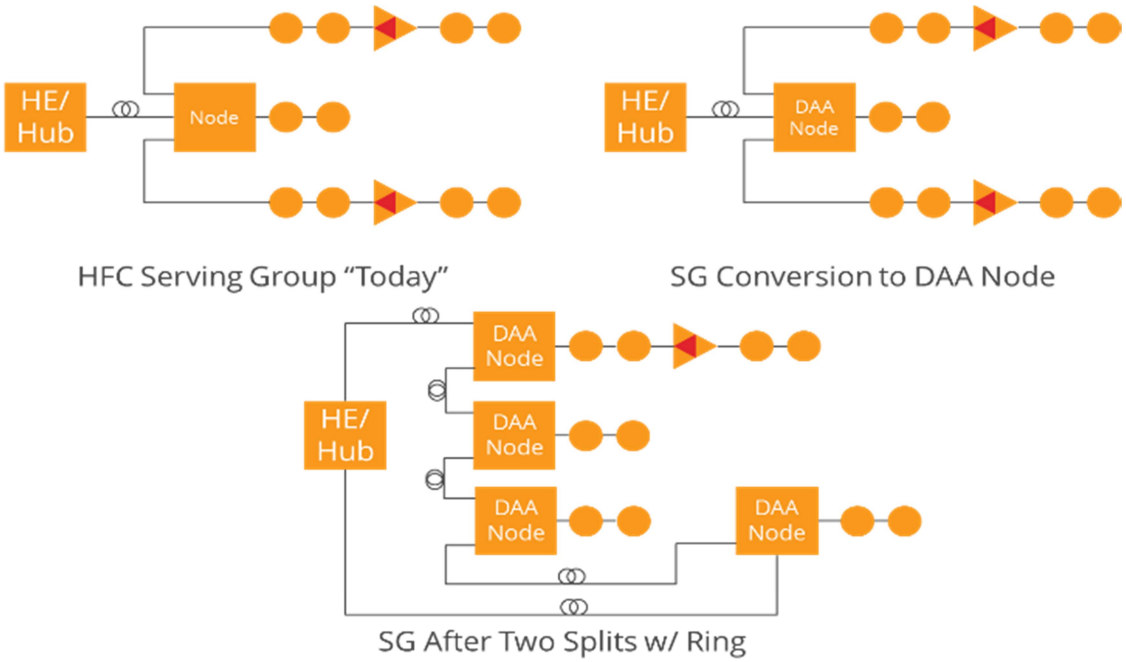

Figure 3: The Managed Device Explosion Brought on by DCA

The increased number of devices will make a CCAP Manager a necessary reality for MSOs. One of the benefits of migration to an SDN configuration model is that centralized control can provide programmatic deployment of systems and services. The ability to manage, configure, and update/expand network capacity utilizing an API based configuration model, rather than SNMP and command line interface (CLI) based model, provides significant benefits to the operator. Eventually the multiple functions integrated in the CCAP today can be pushed to VNFs and these can be located in the network where they make the most sense.

SERVICE ORCHESTRATION - AN EVOLUTION PATH TO A VIRTUALIZED NETWORK

What is Service Orchestration?

Service Orchestration is a concept that has evolved as part of the SDN and NFV efforts. It is the name for the function that facilitates the creation of a given end-to-end service among different Virtual Network Functions. It includes not only the virtual services, but also touches the physical network elements – like CCAP or CCAP Core, Remote PHY devices, and cable modems. The Service Orchestrator along with a NFV Orchestrator (NFVO) can be utilized to implement both the end-to-end connectivity, as well as the virtualized network functions and physical network equipment needed to fully realize a given service that a service provider wishes to make available. The NFVO in this context manages the virtual machine infrastructure that houses the virtualized functions – like RADIUS and related authentication servers, the DHCP services for the RPDs, and related functions.

To better understand the role of the Service Orchestration function in the context of the evolving Distributed CCAP Architecture, we need to look at several issues that are present today. One of these issues is the configuration, or on-boarding, of the RPD with its set of CCAP and Auxiliary Cores. In the current DOCSIS Remote PHY specifications, this is done using

the DHCP server and DHCP Options. In a network where a Service Orchestrator is present, the Service Orchestrator would be able to assign existing or new CCAP MAC and Auxiliary Cores dynamically and provide that information to the DHCP server function which can then populate the values in the DHCP reply messages. Should the Service Orchestrator need to create additional CCAP MAC or Auxiliary Cores, it can contact a NFVO to create additional core functionality and have it ready for the RPD to connect to it. This process shows the power of this dynamic provisioning or on-boarding model – the ability to add in capacity when new devices are discovered. Likewise, functions like the authentication of the RPD and the CCAP cores can be virtualized and those resources managed by the NFVO.

The Service Orchestrator can also be teamed with an SDN Controller in a given location that will help to manage the custom forwarding and message handling that is employed during the RPD initialization and alert the Service Orchestrator that a new or recognized RPD has been discovered. This helps then further distribute the load generated by the on-boarding of the RPD. In the past this was managed by the wiring of nodes to RF cards in the Integrated CCAP and EQAM network in a hub.

The Service Orchestrator can be further augmented by a Gadget Manager application as described in the next section.

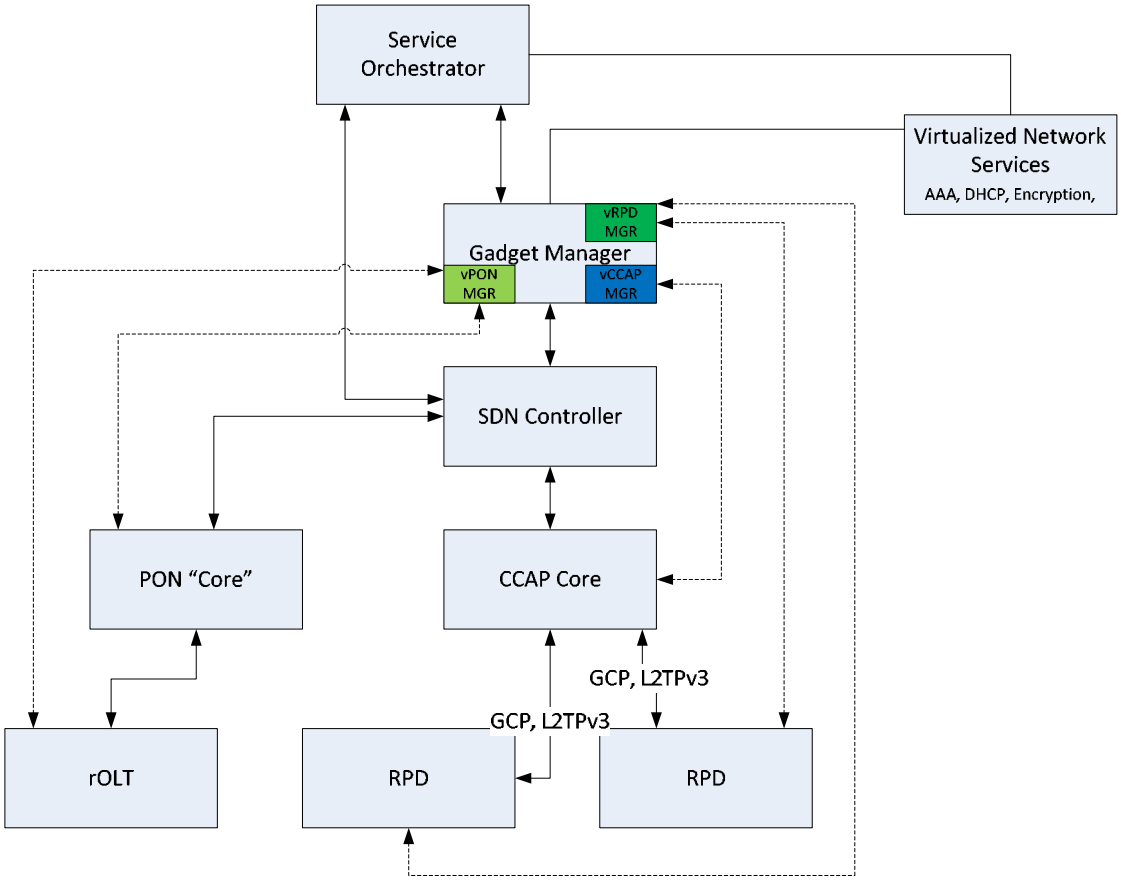

Figure 4: A Logical View of the Network

The Gadget Manager

As discussed above the advent of DCAs it has become apparent that a more centralized and programmatic method is needed to manage the large number of distributed access devices that would be deployed throughout MSO networks in the coming years. This presents both a challenge and an opportunity. The challenge is to build an Element Management System (EMS) that can manage thousands of devices and give the operators the control and monitoring required to support and maintain such a network. The opportunity is the realization that this system can be more than a very elegant EMS with a nice graphical interface. Rather, can be both the User Interface (UI) and the machine-to-machine interface for the system. Next, since this system doesn't live on the access device itself, rather it lives in a virtual compute architecture, it can be used to abstract the access layer from both management and control, thus realizing one of the main goals of SDN – the separation of the control plane logic from the data plane. Ideally the "Gadget Management" function can be used for management and presentation for any access technology including: I-CCAP, DCA, PON, Managed Ethernet, or even G.fast.

This new system needs a name that will communicate that the system is more than simply an EMS and that it can be used to manage any flavor of access device. At the time of the gadget manager's conception the industry had not yet settled on whether the DCA world would be an R-PHY world or an R-MACPHY world. "Distributed CCAP Access Device" could have been used but this seemed too formal, and would potentially create a name that was cumbersome. Also it would not communicate the idea's access agnostic nature. At the same time a small group was referring to DCA devices as "gadgets." A gadget can be just about any device or function, and the term was found to be endearing as it was used by the Manhattan Project as a handle for the design of the first atomic bomb. The atomic bomb's existence radically changed the world, and there is a belief that DCA, access agnostic architectures, and this new EMS that is more than an EMS, will radically change the cable and access worlds. Thus this new EMS has been dubbed the "Gadget Manager" (GM). The GM would be access agnostic striving to eventually manage any access gadget.

The GM is the Presentation Layer

"Presentation Layer" is a fancy name for an API, but in the case of the GM this idea can be extended to be any interface to CCAP system including both well-defined APIs of the DOCSIS OSS system and loosely or un-defined or even proprietary interfaces such as the Command Line Interface (CLI) with which users interact when managing CCAP devices. Thus the GM will "present" as a CCAP to the outside world weather that outside world was a back office management system, or was a human being banging on a CLI.

This idea tracks well with the virtualization efforts that are being discussed in the industry today. For example, consider a collection of R-MACPHY nodes. In the non-virtual world, each would have its own IP address, each would be managed as a single device, each would have its own set of DOCSIS MIBs to be managed and monitored, and each would have its own CLI. In the GM world each still has its own IP address, however users in this world don't use this address to interact with these devices, rather it's the IP address (or one of the IP addresses) associated with the GM with which users and devices will interact.

This is a virtual system in few ways. First, the world is interacting with an intermediate application, and not the devices themselves. The devices are hidden behind a veil of interfaces that make it look and feel as if directly

connected to the device. Second, the logical partitioning of gadgets to "virtual" CCAP is flexible. Meaning the operator can group, un-group, and re-group the DCA devices based on resource constraints that are not related to where the devices physically live. Lastly, the GM itself is virtualized and runs on virtual machine in a compute farm.

The example above can be further filled in. Consider an R-PHY deployment that consists of 100 such gadgets and the operator decides to partition these devices into four virtual CCAPs. The question becomes how these gadgets should be grouped. One method would be to simply split them into four groups of 25 nodes per CCAP core. In this case, one group of 25 nodes could come from then north end of town, 25 from the south, 25 from the east and 25 from the west. This is a simple partitioning that maps the physical world almost directly to the virtual world.

However there may be other methods for partitioning that are more advantageous. Perhaps the number of subscribers attached to each of the 100 gadgets is not uniform. Maybe 50% of the subscriber population lives on 20 of the gadgets. In this case, it may be desirable to balance the subscriber population between the various virtual CCAP instances. Here, we may end up with 30 gadgets associated with one virtual CCAP, another with 10, and two with 20. They may or may not be geographically partitioned. It could be that the 30 gadget V-CCAP consists of the least used gadgets for all parts of town, East, South, North, and West and the two 20 gadget V-CCAP's live only in the South and West parts of town where the bulk of the customer traffic is located.

It should be noted that this system is not virtual in some respects. It may be argued that a virtual CCAP as a system that has had its data plane processing virtualized. Here, with the GM, only the control plane has been virtualized. It may be that the deployment being managed by the GM is an existing physical I-CCAP or OLT, a physical CCAP-Core, or even a virtualized CCAP-Core. The GM is agnostic to the implementation of the data plane. It will be shown later that the benefits of the GM can apply to both physical and virtual data planes.

Purpose Built Gadgets are Compute Bound

Gadgets, whether they are in I-CCAP or a Remote PHY or Remote MACPHY node are all built on special purpose hardware and this purpose built hardware has an upper bound to the amount of processing power, or compute, of which it is capable. This in itself is not a bad thing as these systems are designed with excess compute power. The system designers target is to have at least enough compute – both in terms of control plane and data plane for the last day the system will be deployed. Operators would be reticent to purchase without the belief that these systems will last for a reasonable length of time.

It is difficult to know ahead of time exactly what an operator is going to demand of a system on the last day of deployment during the design phase of a product. It should be noted that the design phase typically occurs 10 years prior to the last day of deployment. Even if a designer did know exactly how an operator would use a system 10 years after it is first deployed, it may not be cost effective with the current generation of hardware to implement.

Thus designers and operators typically live in a world where they must live with the amount of compute that is available, and many times the amount of control plane compute is strained over the life cycle of a product.

The GM is Elastic (Not Compute Bound)

This is crux of the opportunity of the GM. Moving the presentation layer off of the gadget

and onto a virtual machine makes it possible to refactor functions that were once integrated into a more distributed environment. Separating data plane and control functions allows for them to be independently scaled. In a compute bound world an MSO may have wanted to poll all MIBs every five seconds on an integrated CCAP, but was unable to do so due to the compute limits of the system. With the GM an operator will have the ability to do this type of rapid polling; however, it now becomes a cost to the operator to add enough compute to the GM to do so. The operator can thus weigh this cost against any perceived or real benefit.

A more real world scenario is that an operator's needs grow over time. With elastic compute the operator can add compute "just-in-time." This means that at the start of a project, less compute can be deployed initially than would have to be deployed with compute bound solution that cannot be changed over the life cycle of the deployment. In the elastic compute world as more compute is needed it can be added. Also of benefit is that late added compute can take advantage of the performance gains that are occurring throughout the time of the deployment. This has two affects. The first is that the costs are spread out through the duration of a deployment, and second the total cost of the compute is reduced assuming the price of compute goes down during the deployment.

Challenges to Elasticity

While the GM is elastic it is still dependent upon the compute capabilities of the gadget. This is because the GM presents a view of the gadget to the world that is built upon data collected from the gadget. Thus the poll all MIBs at a one second interval example above is only meaningful if the GM can get all MIB data from the gadget every one second. There is also the cost in terms of bandwidth to deliver this data at this rate from the gadget to the GM. Therefore the problems of limited gadget compute and limited network bandwidth available for gadget data collection need to be addressed in order for the GM idea to be successful. A potential solution is presented next section.

Connectionless Push and Gadget Specific Protocol Design

As the gadget Presentation Manager, the GM has to be able to access each gadgets configuration and operational data store. Gadget implementations that support traditional IPDR protocols are possible. However, given a strong MSO desire for each gadget to be low in cost or specific in functionality, there may be limited compute capacity within each of the gadgets. As a result, it may become critical that the communication path between the GM and the individual gadgets be as lightweight as possible. This will help us reduce the work load required of the gadget to processes data and minimizing the amount of data that must traverse the network. Thus a connection-less protocol is preferred over a connection oriented protocol. The reason for this is that with a connectionless protocol, no state is needed for the protocol itself.

By going connectionless the sending of IP data grams can be simplified and even implemented with simple logic in an FPGA or an ASIC. This has a secondary benefit in that the sending of data can be distributed within a gadget to different functional blocks where the data being collected actually lives and since the protocol is connectionless these functional blocks do not have to coordinate with each other, nor does their respective data need to be curated by a central control function within the gadget. All of this simplifies the design and reduces overhead. It would also allow for individual blades such as a multi-bladed I-CCAP device to send data grams. Lastly, previously sent values would not need to be

maintained either in shadow memory or on either volatile or non-volatile storage.

The only downside to a connectionless protocol is the potential for out-of-order or lost packets. One way to mitigate this is to have a sequence number associated with each datum or set of data that is produced by the same packet source within a gadget. A time stamp could be used but could incur undue overhead to distribute and synchronize it. All data sent must be sent as absolute values. Thus packets coming out-of-order to the GM can be re-ordered by the GM which has elastic compute, and any lost packets would be detected, and would only represent a gap for the specific datum or data lost, and only for one collection period.

If that temporal data collection gap due to lost packets is deemed to be undesirable, then alternative protocols (such as those based re-transmission requests or those based on multicasting of the collected data to redundant, check-pointed GM receivers) can also be considered.

Another design consideration that one should take into account when looking at the design of this protocol – should the data transfer be a push between the gadget and the GM or a pull from the gadget? So called pull protocols, like SNMP, poll the gadget at regular and random intervals and pull data from the data store. Push based protocols, like IPDR, are based on a set of predefined service definitions that include elements from the operational data store that are pushed to the GM at regular intervals. IPDR also allows for ad-hoc requests to be made of the operational data store. REST based interfaces like RESTCONF are more API driven and thus suitable for both push (post) or pull (get) and can be implemented using the north and southbound APIs that can be accessed using and SDN controller. In the end, the GM may utilize both types of protocols when gathering the needed data from a given gadget. Likewise, the GM must be able to provide data to operator management stations using both push and pull protocols that SNMP and IPDR so that these systems can continue to function and gather data for applications that rely on that data.

Following the above principles will allow for a fast and robust data collection between gadgets and the GM. While gadgets are typically thought of as new devices that are yet to enter the network, the GM can be deployed with existing I-CCAPs and with the same principles of connection-less, non-prescriptive, and single collector, to improve data collection rates in today's systems.

THE GADGET MANAGER

We have discussed the challenge facing network evolution and the applicability of SDN and NFV technology and the principle of service orchestration- all of these technologies enable network evolution and provide productivity and OPEX gains in managing the network. In applying these methods, we propose the idea of a Gadget Manager. The genesis of the Gadget Manager is that it can immediately address the issues of managing the explosion of devices in a distributed access architecture by managing the deployment and configuration of these new devices. The Gadget Manager can centrally manage devices in operation. It will also become an aggregator of data from remote devices, legacy CCAP and CCAP core, reducing the burden on the infrastructure to collate key performance parameters and establishing a central repository for operations critical data and key performance indicators. The Gadget Manager also plays a vital part in configuring the network as DOCSIS 3.1 is rolled out, and will be the basis of a service orchestrator and SDN controller.

Figure 6 shows the conceptual locations of the Gadget Manager.

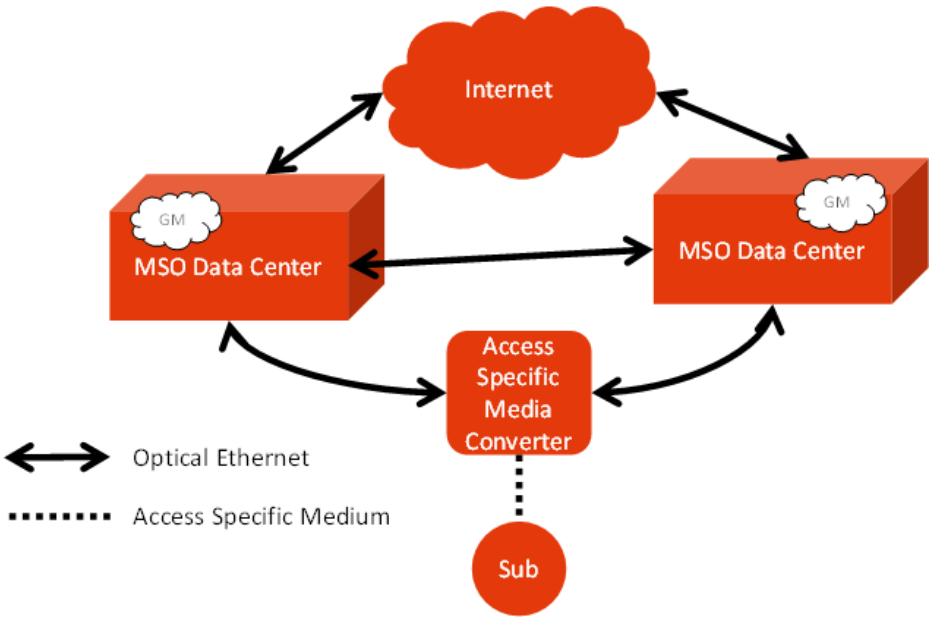

Figure 5: Gadget Manager Architecture

Alternatively the Gadget Manager can be configured as a hosted application for Operators that may not have the resources or choose to operate in a more centralized fashion. This architecture is depicted figure 7.

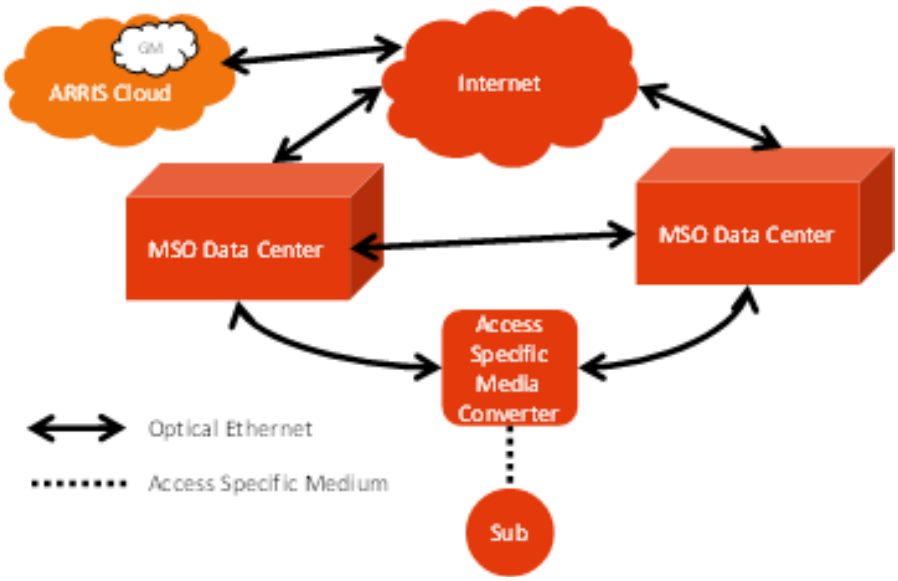

Figure 6: Gadget Manager Hosted Architecture

The Gadget Manager is developed as a cloud-based software that provides:

- An aggregate Interface between the MSO and multiple remote gadget devices
- Manages remote gadget versions and configuration
- Optimizes gadget monitoring
- An aggregate Interface between the MSO and legacy CCAP and new CCAP cores
- Fits into the existing MSO network and allows the MSO to provision and implement the same services that they provision and implement on CCAPs today
- Value added functions such as data aggregation, CCAP core and Remote device license management and reporting, back office integration, device provisioning and traffic management tools
- Eventually, a controller that uses SDN and SON to manage forwarding rules and plumb in NFV functions

Evolution of a Gadget Manager

The evolution of a Gadget Manager could occur in a few phases.

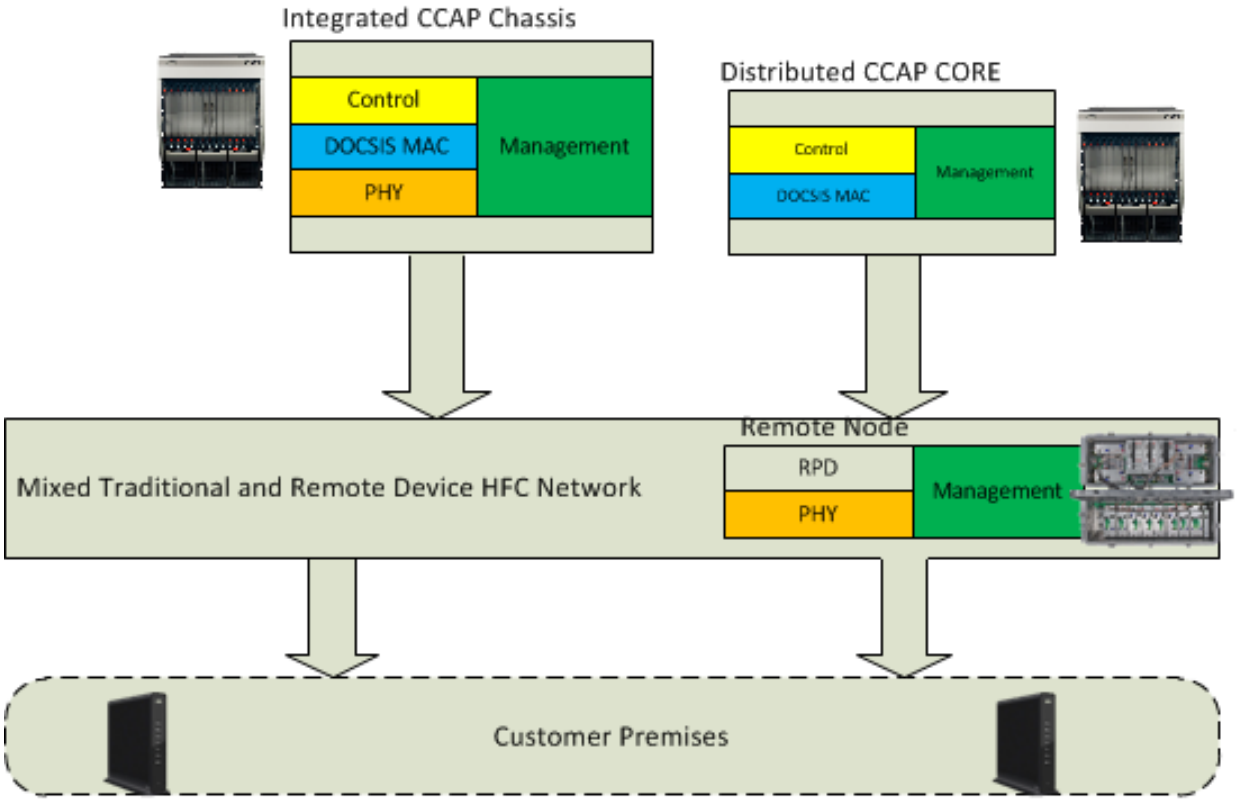

Figure 7: Current Day Network View

Phase 1: the Gadget Manager hosted in the data center provides value added network services in the deployment of DOCSIS 3.1 and remote devices. In this phase, the Gadget Manager can start to provide value added functions like DOCSIS 3.1 profile management and the ability to manage the on-boarding of RDPs to CCAP Cores and discussed previously. This is depicted in figure 9.

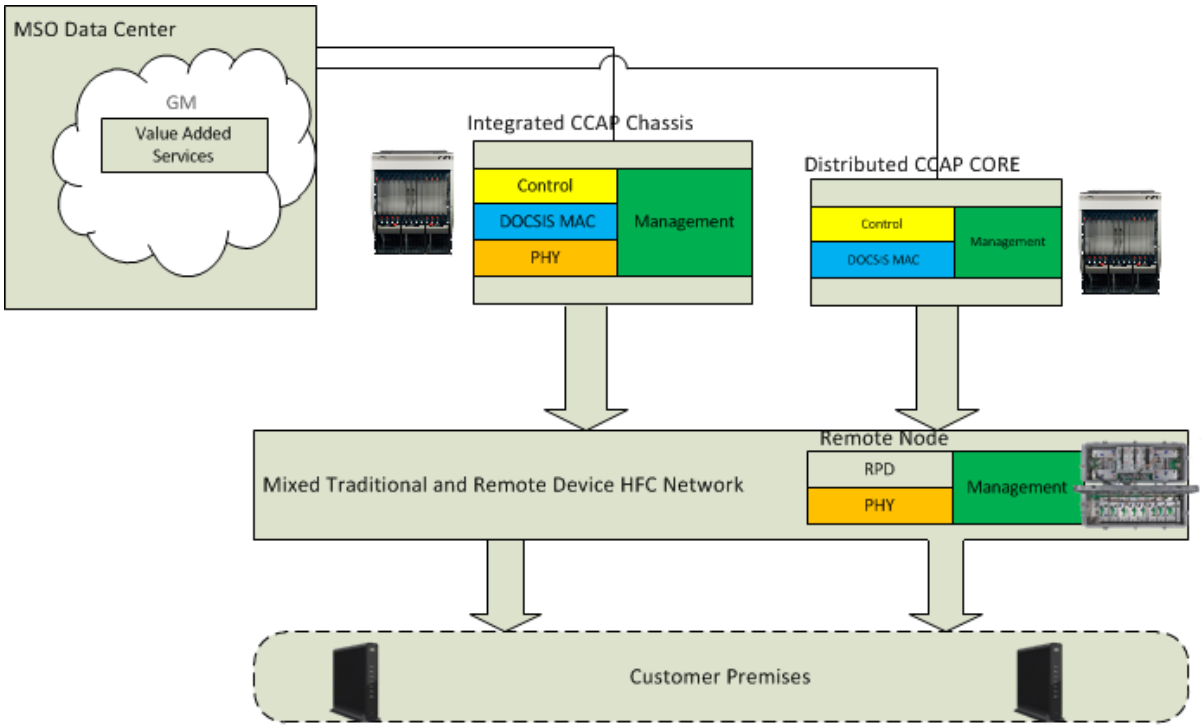

Figure 8: Introduction of the Gadget Manager

Phase 2: In this phase, the Gadget Manager begins to take on more and more of the management functions for the managed devices. This includes the beginnings of the Gadget Manager serving as the management entity for the CCAP and RPD resources by gathering and managing all the performance and configuration data for the core and RPD devices. The GM also provides the configuration interface for the mapping of individual RPDs and CCAP and Auxiliary cores. License management and additional performance applications are also integrated under the Gadget Manager umbrella or interface with the GM to perform their work. This is phase is depicted in Figure 10 below.

The Gadget Manager takes the role of SDN controller for CCAP cores in a distributed access architecture. Its ability to provide control and management of networks with full CCAP, CCAP core/DCA deployments and virtualized deployments protects and extends the life of the investment made in CCAP, providing the benefits of separated control and data planes, simplified management and rapid service creation without having to undertake the risk laden move to NFV before data centers can provide the performance required to effectively run a full CCAP core.

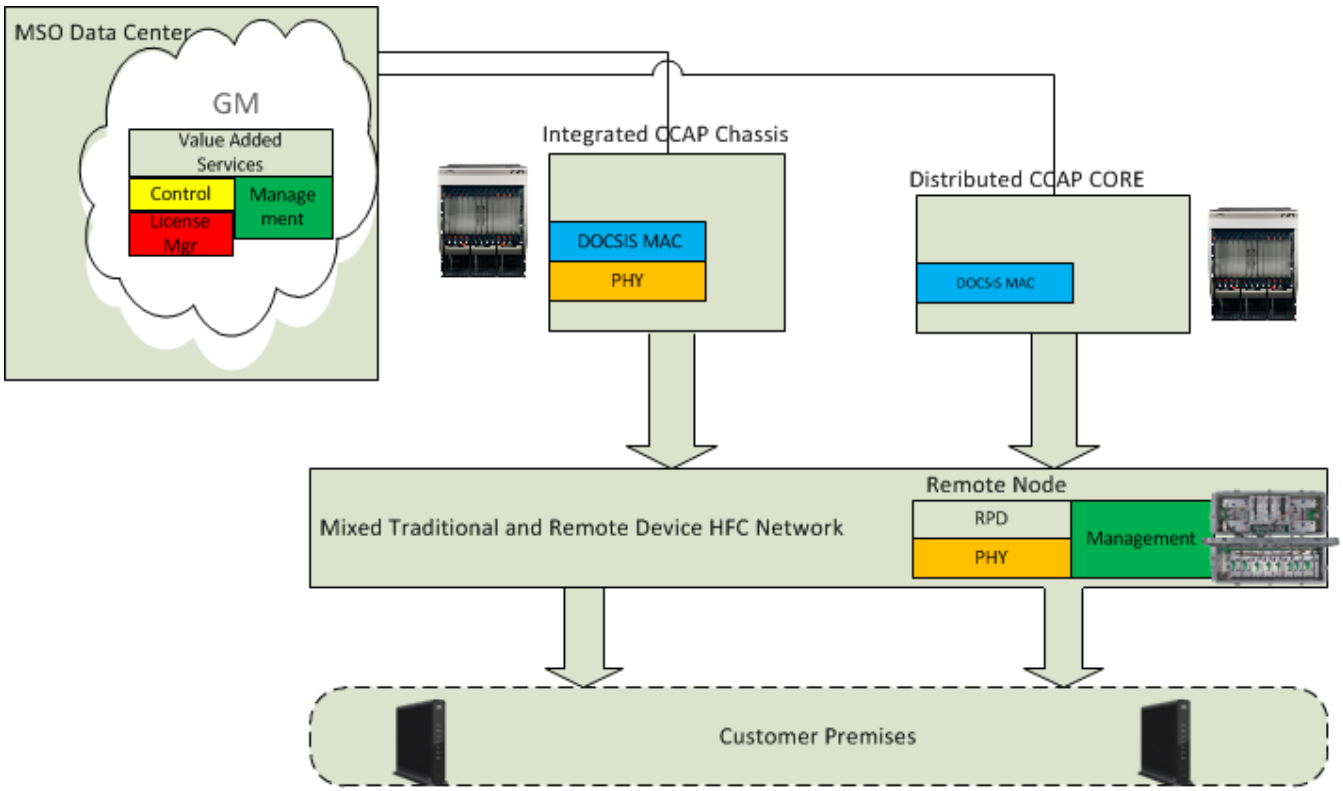

Figure 9: Migration of Functionality to GM

As more and more MAC Core functionality is virtualized, the Gadget Manager becomes more and more integral in managing and configuration of the mapping of the mac core virtualized functions and related traffic to the RPDs and customer homes. Eventually, the Gadget Manager and the Service and NFV orchestrators are enabled to dynamically manage and on-board resources as those resources are needed. This end state is depicted in Figure 11.

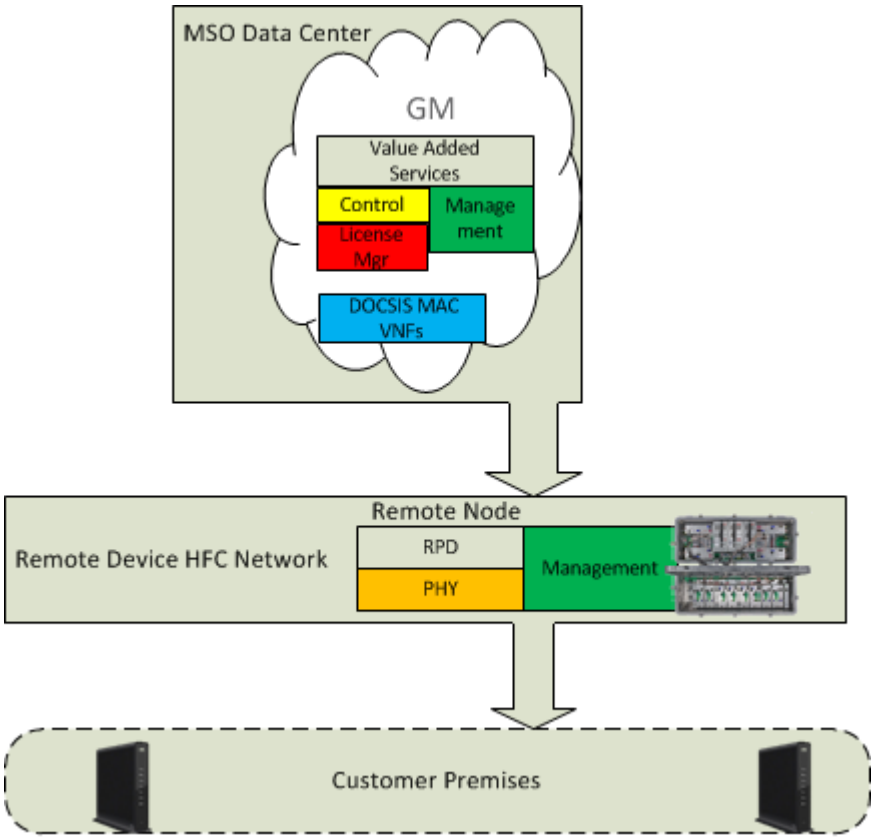

Figure 10: Gadget Manager Fully Realized

SUMMARY

High speed data networks to the customer's premise are evolving rapidly, driven by the need to innovate rapidly and to reduce the cost of both operation and capital spending. These changes encompass not only the Access networks that have been leveraging the Hybrid-Fiber Coaxial plant, but impact how those networks are architected, provisioned and managed. Distributed DOCSIS architectures, whether remote PHY or remote MAC/PHY, will create more intelligent nodes in the network that require new levels of management and will also create valuable data on network performance and effectiveness that must be collected and processed. Mixed distributed and centralized network Architectures will create management domains with different sets of requirements, straining traditional management infrastructures and organizations. Operators will increasingly look to the emerging paradigms like SDN and NFV to help provide a more services-based approach to the provisioning and management of the overall service delivery including the HFC based access network. Using technologies like the Service Orchestrator and a Gadget Manager to configure and manage the increased numbers of distributed HFC gadgets that MSOs will deploy over the next several years, the MSO can realize the benefits of migrating their current service provisioning environment to a more elastic and dynamic model. This paper has shown one possible migration strategy that can be used to migrate from the current generation of provisioning and service enablement to a much more dynamic, cloud based model. This model will permit MSOs to easily change the form and scale of their management system to match the ever-changing demands of the evolving network infrastructure being managed.

ACKNOWLEDGEMENTS

The authors would like to acknowledge and thank some subject matter experts for their invaluable inputs and insights. In particular, Jeff DeMent, Bill Hanks, and Erich Arnold (of ARRIS) were instrumental in giving guidance on several sections within this paper.

ABBREVIATIONS & ACRONYMS

API	Application Programming Interface
BBF	BroadBand Forum
CCAP	Converged Cable Access Platform
CLI	command line interface
CM	Cable Modem
CMTS	Cable Modem Termination System
COTS	Commercial off-the-shelf
CPE	customer premises equipment
DAA	Distributed Access Architecture
DCA	Distributed CCAP Architectures
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data Over Cable Service Interface Specifications
EMS	Element Management System
EQAM	Edge Quadrature Amplitude Modulation
ETSI	European Telecommunications Standards Institute
HFC	Hybrid Fiber Coax
HSD	High Speed Data
I-CCAP	Integrated-Converged Cable Access Platform
IETF	Internet Engineering Task Force
MPEG	Moving Picture Experts Group
MSO	Multiple System Operators
NFV	Network Function Virtualization
NFVO	Network Function Virtualization Orchestrator
NV	Network Virtualization
OFDM	Orthogonal Frequency Division Multiplexing
ONF	Open Network Foundation
PON	Passive Optical Network
RPD	Remote PHY Devices
SDN	Software Defined Networks
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratios
SO	Service Orchestrator
UI	User Interface
V-CCAP	Virtual CCAP
VNF	Virtual network functions