Managing Distributed Access Devices with SDN

Centralized Operations for FIB Management

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Steven Krapp
Product Management Director
ARRIS Enterprises, Inc.
2400 Ogden Ave Suite 180 Lisle, IL 60532
Office: +1(630) 281-3142
Steven.Krapp@arrisi.com

Chris Busch
Product Management Director
ARRIS Enterprises, Inc.
3871 Lakefield Drive, Suwanee, GA 30024
1-678-473-8401
Chris.Busch@arrisi.com

Jeff DeMent
Principle Systems Engineer
ARRIS Enterprises, Inc.
2400 Ogden Ave Suite 180 Lisle, IL 60532
Office: +1(630) 281-3149
Jeff.DeMent@arrisi.com

Overview

Distributed Access Architectures, including Remote CCAP (R-CCAP) and Remote PHY(R-PHY), will create new networking and operational challenges as the number of managed devices increases by two orders of magnitude. Operators will need to address basic questions in regard to service activation, device deployment, remote device provisioning, and network resource management including IP address pool allocation. Fortunately new tools are becoming available both as part of the CCAP specification in terms of XML based provisioning and Software Defined Networking that can help address these challenges. However, a fundamental shift in back-office systems will need to occur in terms of logic process, data sharing, and visualization of how remote architectures will influence our view of topology and assigned resources.

This paper will focus on how SDN can be used as part of a remote CCAP architecture to both provision and manage remote devices. It will also demonstrate how centralized SDN controller co-ordination with today's back office platforms can be used to orchestrate network resources enabling rapid service deployment throughout the network.

Contents

List of Figures

Distributed Access Architectures (DAA)

What is DAA?

For Hybrid-Fiber Coaxial networks, two main camps are forming in the industry for DAA[1] . These are Remote PHY (R-PHY) and Remote CCAP (R-CCAP). Other options have been considered, but these two options are emerging as realistic candidates for actual widespread deployments [2][3] [7] .

Basics Idea of DAA

This paper assumes that the role of the service provider is to provide its customers access to information. In the case of cable, the service provider is the MSO, the customer is the cable subscriber, and the information is voice, video or data, any of which the subscriber desires to consume. Traditionally the MSO has provided this information via a Hybrid Fiber-Coax (HFC) network. HFC networks cost effectively deliver information through the use an Ethernet to coax to fiber to coax to Ethernet set of media conversions. This can be seen below:

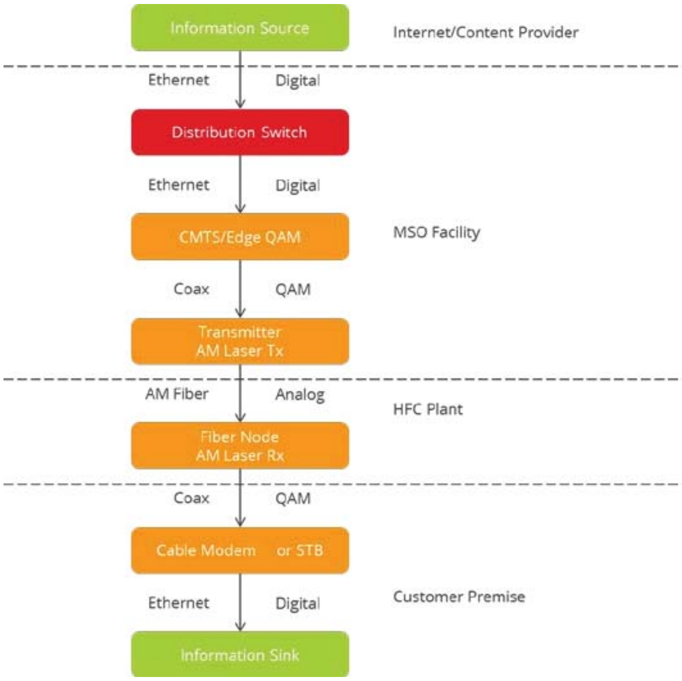

Figure 1. Media Conversion in a MSO Network

The basic premise of DAA is to move the QAM generation as close to the point of consumption as possible. The motivation for why an operator would want to do this is discussed in the *Why Deploy DAA* section that follows below.

R-PHY and R-CCAP

For the intents of this paper R-PHY is a modular architecture that splits the CCAP into a MAC core that lives in an MSO facility, and a PHY portion that lives in fiber node enclosure that resides outside of the MSO owned facility closer to the customer premise. Whereas R-CCAP is an integrated architecture that moves both the MAC and PHY into a fiber node enclosure also outside of the MSO facility closer to the customer premise.

Figure 2 shows the same network as in Figure 1, but using DAA. R-CCAP is on the left and R-PHY is on the right:

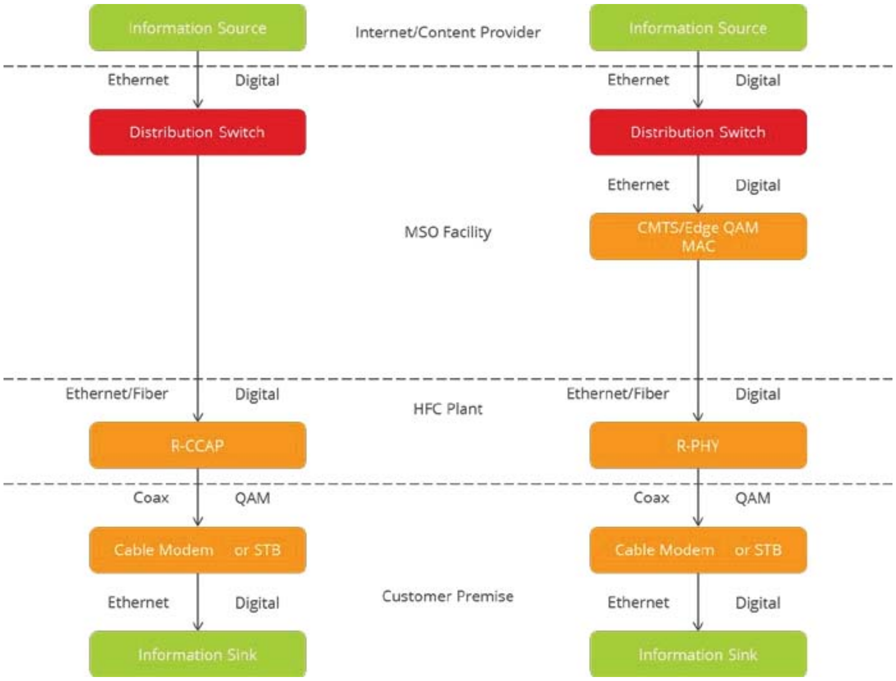

Figure 2. Media Conversion in DAA networks

From examining Figure 2 it can be seen that QAM generation is moved out of the headend and into the node for both R-PHY and R-CCAP. It should also be noted that R-PHY and R-CCAP are evolutionary rather revolutionary designs. Both still keep the Coax and its QAM signaling to preserve the existing “at premise” interface to the cable plant. The goal is to provide gigabit services as cost effective as possible.

For further reading on R-PHY and R-CCAP there are several industry papers[4] [5]

[6] that explain these concepts in much greater detail than what will be explored here. Also, for an excellent breakdown of what OSI stack functions live where for the various architectures, Emmendorfer and Cloonan explain this in excellent detail in their 2013 paper[1] .

Plant Topology in DAA

As can be seen from the citations above much work has gone into figuring out how DAA will work in the HFC network for both R-PHY and R-CCAP. One question that remains open is what will the topology of a DAA network look like, what kind of connectivity will be required, and what kind of cost effective redundancy can be achieved to maintain and even improve the reliability and availability of end user services.

Topology distribution options include:

1. Point-to-Point
2. Point-to-Multipoint via Daisy Chain (Also called “Tree”)
3. Point-to-Point via Star
4. Point-to-Multipoint via PON
5. Multipoint-to-Multipoint via Ring

Point-to-Point

Point-to-Point is the simplest topology. It consists of one fiber feeding a single DAA node device (R-PHY or R-CCAP).

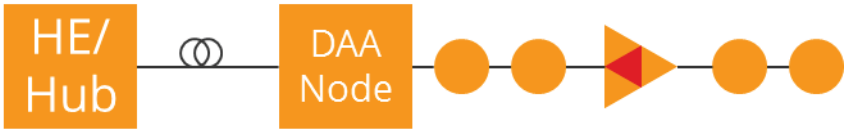

Figure 3. Point-to-Point Distribution

Figure 3 shows an MSO facility, in this case either a headend or hub. It is connected to the DAA node via a 10 Gbps Ethernet fiber. The node then serves two drops (the circles), an amplifier (the triangle), and two addiotnal drops.

The advantages of the point-to-point design are straightforward. Simply replace the existing node with a DAA node, reuse the fiber that previously was used for the AM laser for the 10 Gbps connection.

Daisy Chain

The basic idea of the daisy chain is that each DAA node has two 10 Gbps inputs.

In order to create the daisy chain, one 10 Gbps is connected to the MSO facility, and the other interface is connected to the next DAA node in the chain. This is repeated as many times as capacity demand will allow.

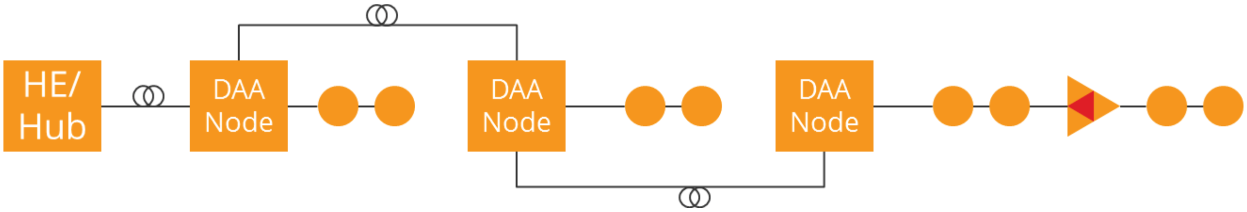

Figure 4. Daisy Chain Topology

Figure 4 shows exactly this. In this case there are three DAA nodes in the chain. The first two simply have two drops each, and the third has 4 drops and an amplifier.

The daisy chain topology is interesting as it is a simple way to support multiple DAA nodes from a single 10 Gbps feed from the MSO facility. The negative to the daisy chain is that each DAA node and each link between DAA nodes represents a single point of failure. If a single daisy chain serves a large number of homes this may not be acceptable; however, it may be acceptable in some deployments where a single daisy chain serves the same area as that served by a single analog node prior to the DAA deployment. An example of this is where a single analog node is replaced not by one DAA node, but by a daisy chain of DAA nodes. This would be the case if an operator substituted DAA nodes in place of existing amplifiers in the network. Thus the serving are of the daisy chain would be the same as that of the analog node prior to the DAA deployment, and like the DAA node, the analog node has single points of failure, namely the AM laser, the optical receiver, the amplifier in the node, etc. Thus in case such as that described the availability of the daisy chain may be no worse than the current HFC availability.

Star Topology

The star topology simply consist of multiple point-to-point legs all emanating from the MSO facility.

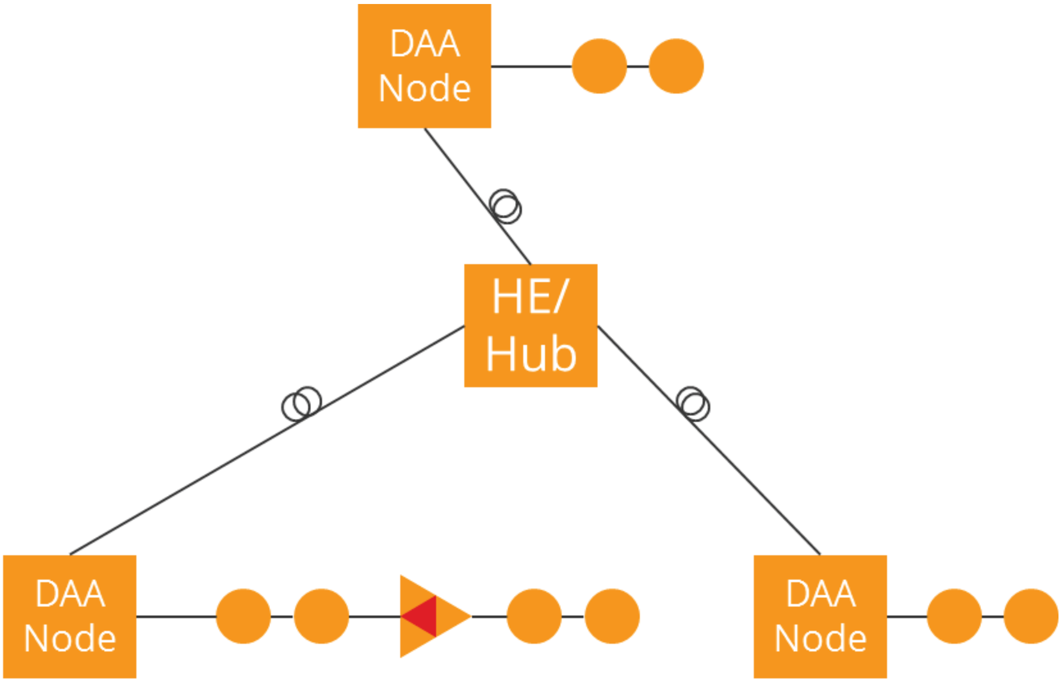

Figure 5. Star Topology

Figure 5 shows a basic star topology with three legs. Each leg is essentially a point to point. Unlike the daisy chain, in the event of a link or node outage, only a single node or link would be affected. This is the same for the simple point-to-point topology.

Where the star topology becomes of interest is in regard to how a DAA network may evolve. Salinger explores this in detail in his 2014 paper *Remote PHY: Why and How.* While he focuses on R-PHY, the basic network topology will work for either R-PHY or R-CCAP. In his paper, R-PHY is deployed incrementally in steps starting with a single R-PHY (or RPN as described by Salinger) in place of the existing fiber node in a point-to-point network. As demand for capacity grows, capacity is added by adding additional R-PHY devices. Initially these new RPNs are fed from the first RPN in a daisy chain fashion. Additional capacity is added by providing additional fibers or possibly λs, to provide each of the RPNs with their own 10 Gbps feed. Thus the culmination is a network of RPNs fed from the MSO facility via a star topology.

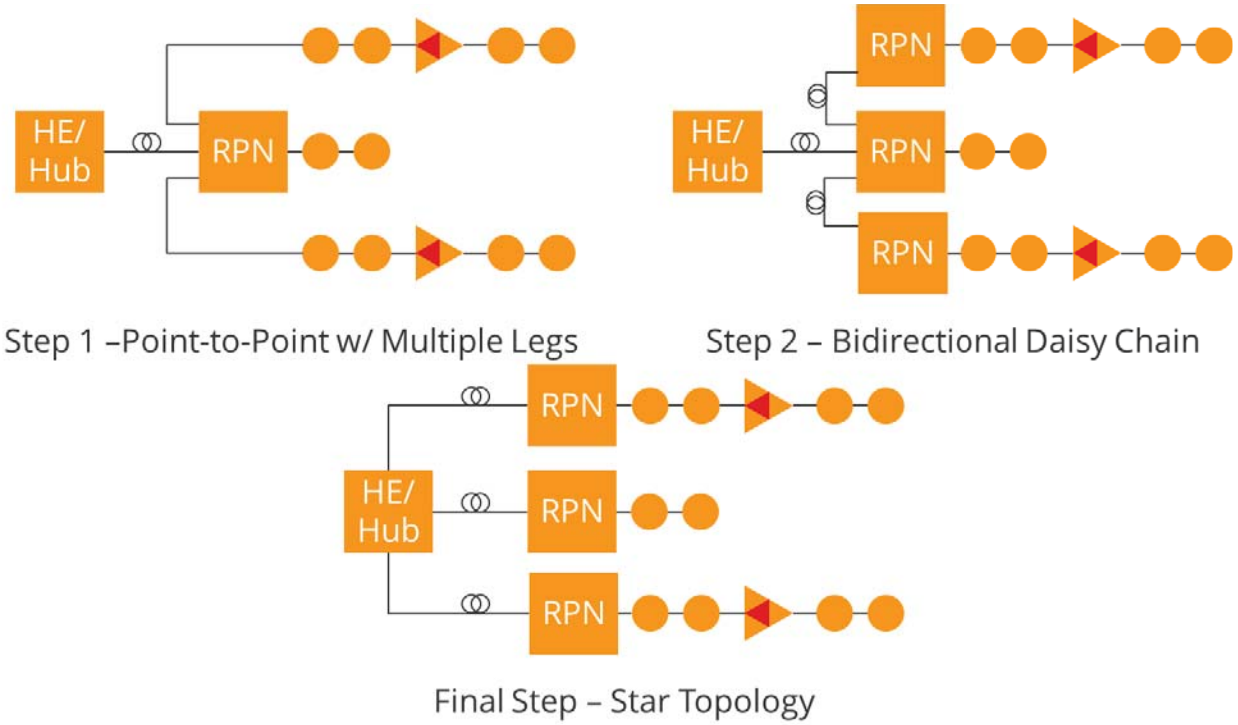

Step 1 –Point-to-Point w/ Multiple Legs

Step 2 – Bidirectional Daisy Chain

Final Step – Star Topology

Figure 6. Salinger's Network Evolution[2]

Salinger shows five steps in the evolution. Three were shown here for expediency. The last step shows a star topology where each of the RPNs has a dedicated fiber link (or possibly λ) emanating from the MSO facility (Headend or Hub in Figure 3).

PON Topology

Up to this point all the topologies have used 10 Gbps Ethernet to interconnect the headend and the DAA nodes and the DAA nodes themselves. 10 Gbps is inexpensive, can support a large number of λs (up to 150), and can be deployed at distances up to 80 km. All of these attributes make these topologies a good choice for DAA interconnect. However, if we assume that a single 10 Gbps link can support multiple DAA nodes, as with the Daisy Chain topology, then it seems logical that instead of dedicated 10Gbps Ethernet interconnect, it would be possible to support some number of DAA nodes via a PON distribution network. Figure 7 shows a possible PON distribution topology:

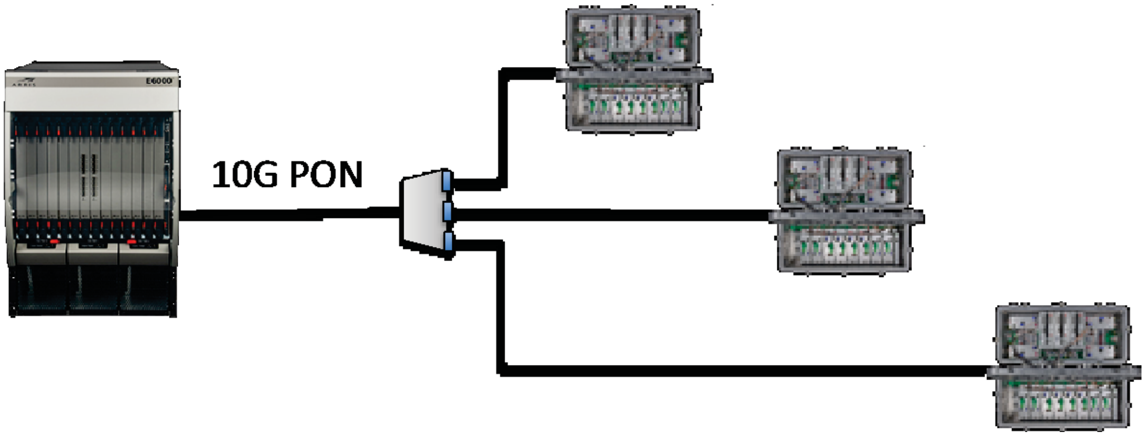

Figure 7. 10G PON Fed DAA Nodes

As shown it is a 10 gigabit PON, either EPON based or GPON based. The PON OLT module is located in the MSO facility and the ONU module resides within the DAA node.

Others have suggested the use of GPON[10] for distributed CMTS distribution. This may work for just the CMTS technology, but with only 2.5 Gbps of data available GPON is not sufficient to support full spectrum remote CCAP or R-PHY deployments, which could consume more Gbps than the 2.5 Gbps available.

A PON topology may or may not be more cost effective than 10 Gbps Ethernet for DAA interconnect. The factors that go into determining this will be similar to those used by an operator to determine should it deploy a PON topology or point-to-point Ethernet topology when deploying network services. For example the relative cost of the OLT and ONU modules are typically greater than the cost of a 10 Gbps aggregation switch plus the cost 10 Gbps SFP+ modules. However, this may be offset by the fact that only a single OLT port is required to support multiple ONU's in a PON deployment. Compare this to the star topology where each customer will require a separate termination point in the operator's facility. Thus the PON topology may be a less costly solution than the star topology as it would not require unique termination for each DAA node at the operator facility. However, like PON, the daisy chain topology also only requires a single termination in an operator's facility. Thus the daisy chain topology may be lower cost compared to the PON topology as it too has a single termination but can make use of ordinary 10 Gbps Ethernet links.

Lastly distance will be a factor in cost as long distance PON networks require long distance capable optics in both the OLT and the ONU. If the number of ONUs is large the higher costs of the long distance optics can become onerous.

Compare this to a long distance daisy chain where the link between the headend and the first DAA node in the chain would require long distance optics, but the distance between each the DAA nodes may be relatively short. In this case less expensive shorter reach optics for the links can be utilized between each of the DAA nodes.

Ring Topology

Ring Topologies have a bad name in the HFC world. This topology conjures up high cost, difficult to manage right-of-way, and multiple high cost transmitters/ and receivers in the MSO facility where only half are active at any given time. However, 10 Gbps Ethernet is much different technology than traditional AM optics. First Ethernet is packetized and each packet has addressing. Second each DAA node has two bidirectional 10 Gbps Ethernet links. The plan is to use SFP+ based interfaces. Each SFP+ comes with both transmitters and receivers. Lastly each DAA node will have a built in L2 or L3 switch (most likely L2, but for this discussion switching is the key, not the layer upon which the switching is decided.) What all this means is that a ring can be built inexpensively from the daisy chain by simply pulling an additional fiber between the last DAA node and the MSO facility.

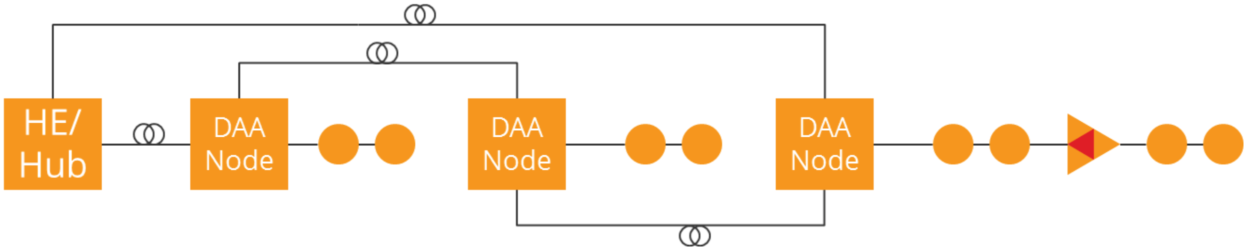

Figure 8. Daisy Chain with Ring

Figure 8 shows a ring built from the daisy chain in Figure 4. When fiber is constructed to all of the DAA nodes in the daisy chain, additional fibers can be “pulled” for essentially zero additional labor cost. All that is incurred is the price of the fibers themselves.

Unlike the other topologies discussed, the ring topology offers protection in the event of failure.

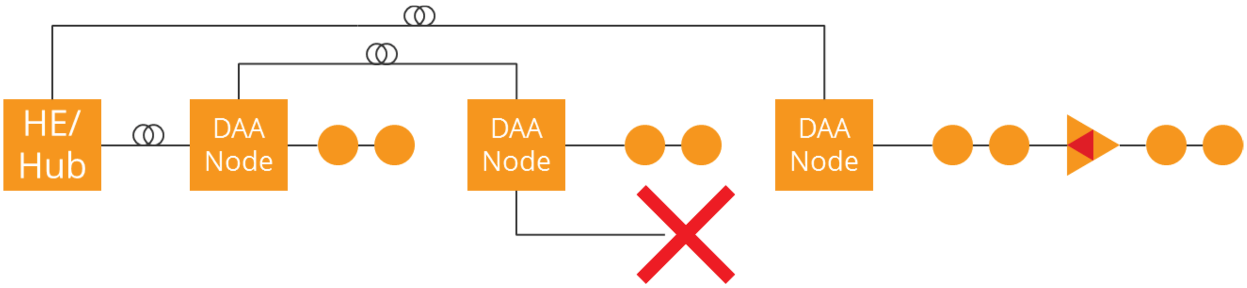

Figure 9. Link Failure in Ring Topology

It can be seen in Figure 9 that despite the cut in the fiber between the last two nodes in the ring, there is still reachability from the MSO facility to all of the DAA node devices.

The ring topology is very attractive. It provides network path redundancy for the cost of an additional fiber, it can provide geographic path redundancy if the needed right-of-way is available or can be made available, it only requires two terminations per ring in the operators facility, and it can make use of ordinary Ethernet optics including short distance optics when DAA nodes are located in relatively close proximity to each other.

Why Deploy DAA

Operators are looking toward DAA as a likely solution to help solve issues that are expected to arise as both service group count and channel count continue to increase.

The problems that operators are anticipating include:

1. Spectral Efficiency
2. Access Technology Costs
3. Facility Space and Cooling costs

DAAs may offer help to address one or more of the above issues.

Spectral Efficiency

DAAs help to improve spectral efficiency by moving the media conversion, Ethernet to QAM modulation, closer to the home. This move means that the implementation loss of traditional AM optics is negated, and thus the MER at the node can be on par with that of facility based equipment. If the improvement is at least 3 dB then in conjunction with DOCSIS 3.1 LDPC FEC signaling the maximum QAM modulation order that is sustainable on a given plant can be increased by at least one step.[12] [2] . For example consider an HFC plant that has a MER of 27 dB. This plant can support 512 QAM utilizing DOCSIS 3.1 LDPC FEC. In this scenario a 3 dB improvement to MER will mean that MER would increase from 27 dB to 30 dB and that the maximum modulation order would then increase to from 512 QAM to 1024 QAM. This means that spectral efficiency in terms of bps/Hz would increase by roughly 11%.

The counter viewpoint of this is that with current headend equipment MER and AM optic implementation loss it is expected to be able to support up to 8K QAM in most systems[6][8]. The idea here is that despite the fact that node based synthesis improves MER, it doesn't matter if the MER of facility based synthesis is good enough.

Access Technology Cost

The next issue is Access Technology costs. Even if we assume the cost of Analog Modulated (AM) optics is equivalent to the cost of digital optics, the reality is that for the same distance, more digital DWDM wavelengths can be supported than the number of AM DWDM wavelengths on a single fiber.

Related to the issue of cost is facility consolidation. This has been the Holy Grail of operators as a way to reduce operational costs. Simply put fewer facilities should equate to lower operational costs. Digital Optics in conjunction with DAA may offer a path to cost effectively bypass today's small facilities, connecting large facilities directly to the equipment in the neighborhood.

The counter view point here is that at 10 km or less the difference is negligible, and that for even up to 40 km distances as many as 44 λs can be supported with reasonable MER[8] .

Facility Space and Cooling

DAAs help with facility costs in multiple ways as Moore's law is enabling the ability to add some or all of the CCAP functionality into the node, or onto the pedestal. With or without DAAs, more service groups will eventually lead to more nodes in the network. However since some or all of the CCAP functionality can be fit in the node, zero or very close to zero headend real-estate will be required to support these new service groups and the new equipment is deployed on the strand. This means that facilities can either stay the same size, or possible even shrink. In addition to space savings, since the DAA element lives outside of the MSO facility and it will be passively cooled, the MSO does not bear the burden of cooling the DAA and its environment.

The counter viewpoint of this is that Moore's law will not only permit increased functionality in the node, but it will also increase service group density in headend based equipment, thus negating the need to change the current HFC deployment architecture [5].

Summing Up the Why's of DAA

DAAs aren't a magic bullet. Improvements will continue to be made in facility based Centralized Access Architectures. But there is enough evidence to suggest that due to the improvements of MER, increased number of λs, for at least some portion of the network DAAs will have an advantage. Therefore it is worth considering what kinds of tools will be required to manage this latest architecture. However, as distances increase, the difference becomes significant. Assuming the number of service groups increases by an order of magnitude, there is a large potential cost difference between digital and AM optics.

The Problems Created by DAAs

How Many DAA Devices Will be Needed

Assuming Distributed Access Architectures will be deployed, they will create new networking and operational challenges as the number of managed devices increases by one or two orders of magnitude.

This may seem like a stretch, but consider an example of a medium sized operator that today has roughly 10,000 nodes in its network. If we assume that there are 60 nodes per CMTS on average, then that operator manages roughly 165 CMTS today. Next if we assume 20% of the existing service areas are converted to a DAA system and for each of the serving areas where DAAs are deployed each service group is split two times – meaning where there is one node today, there will be 4 DAA devices tomorrow.

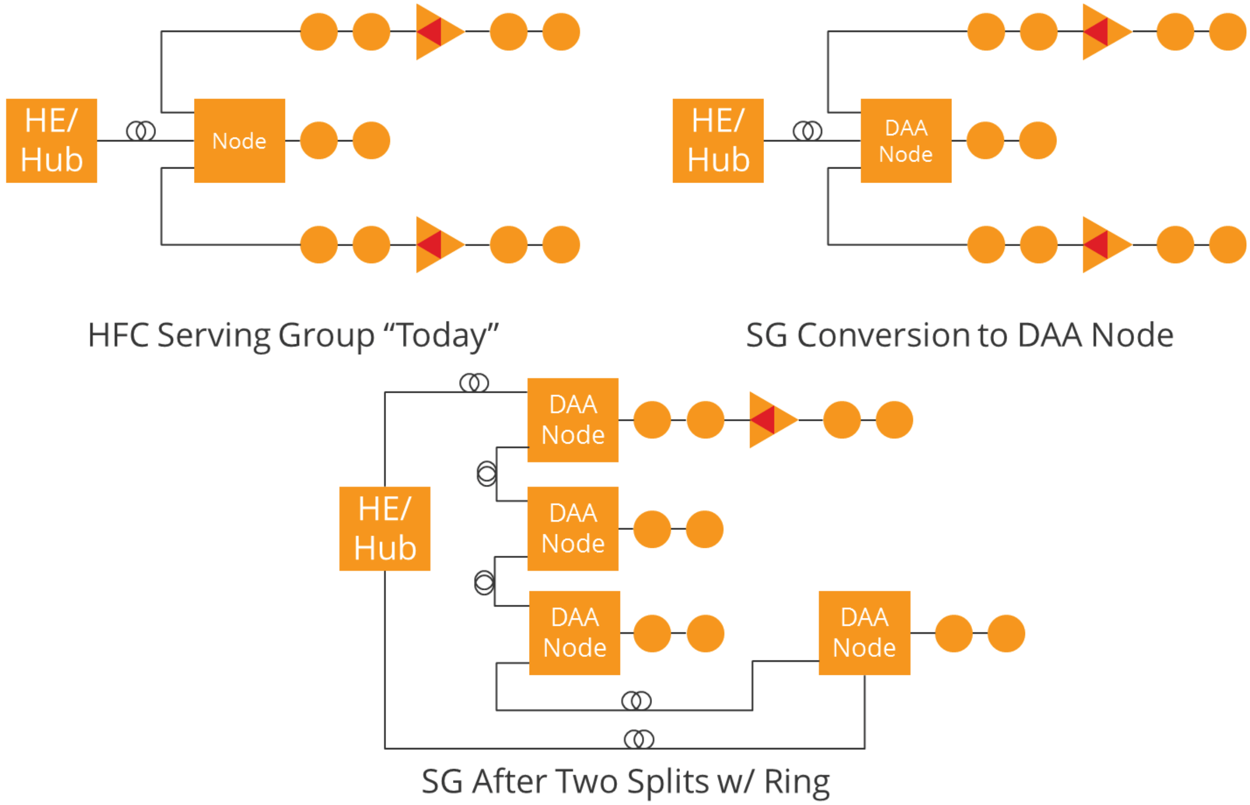

HFC Serving Group "Today"

SG Conversion to DAA Node

SG After Two Splits w/ Ring

Figure 10. Two Split DAA Network Evolution Example

This additional split seems reasonable, as the move to a DAA deployment would most likely be triggered by the need to increase system capacity, or it may occur over time. Thus eventually the operator will have 8,000 DAA node devices, and each will be a network host and require configuration and monitoring. This is in addition to the 132 remaining CMTS for a total count of 8,132 managed devices. This represents almost 50 times the number of devices currently managed. If DAAs are wildly successful and this operator decides to convert 100% of its network, it would mean 40,000 devices to manage, or over 200 times the number of managed devices than exist in the network today.

The point is that while remote architectures represent opportunity for better noise performance and possible facilities consolidation, they also bring many operational challenges. These include physical, environmental, and device management challenges.

Operators will need to address basic questions in regard to service activation, device deployment, remote device provisioning in this expanded network.

The operator will need to evolve its operational procedures to deal with the increased number of devices in its network. There are multiple paths that an operator may choose for this evolution. They include:

1. Software Defined Networks
2. Independent Network Element Monitoring and Management
3. Standalone Provisioning for Network Elements

This remainder of this paper will focus on SDN solutions to this problem.

The Solution: Software Defined Networks (SDN)

SDN is a technology that seems like a good fit to manage the explosion of managed devices due to DAA in MSO networks. An SDN control plane must support the same functions and services that are supported by today's CCAP deployments.

What is a Software Defined Network

Software Defined Networking (SDN) concepts were conceived only a few years ago among several academic initiatives examining the control and configuration of today's network infrastructure. One of the principle concepts of SDN is separation of control plane decisions from the forwarding element. This separation creates a network 'Controller'. The Controller has knowledge of all network elements and is able to provide a number of services to the underlying network infrastructure. The Controller can provide topology discovery, network element provisioning, client forwarding policy, and effectively transform the network into an Application Programming Interface (API) centric model for these actions.

The SDN Controller represents an application centric view of the network. With direct knowledge of forwarding control, the Controller may now support de-coupling several traditionally local forwarding decisions of routers, switches, and CCAPs, into common compute, open interfaced, cloud scale northbound platform. The Controller, for example, also enables movement of legacy DHCP services into the Controller directly. This 'closer to network' approach will de-layer the current provisioning back office for cable operators.

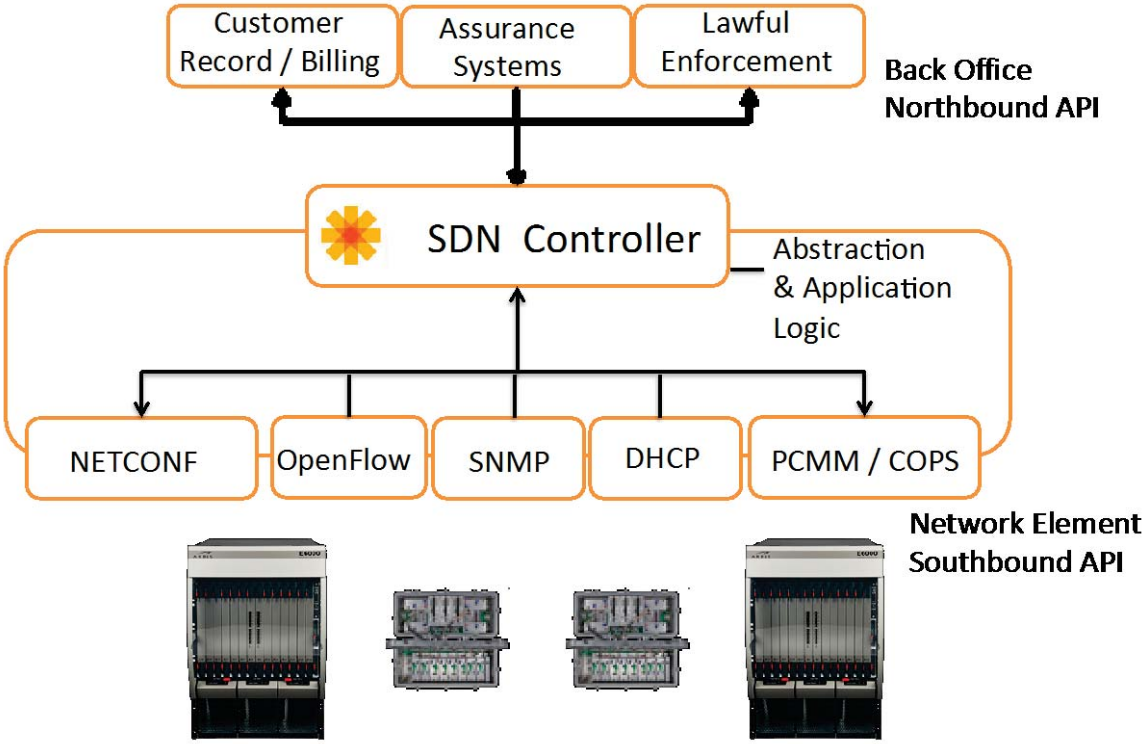
Customer Record / Billing
Assurance Systems
Lawful Enforcement
Back Office Northbound API
SDN Controller
Abstraction & Application Logic
NETCONF
OpenFlow
SNMP
DHCP
PCMM / COPS
Network Element Southbound API

Figure 11. SDN Network

A significant concept in SDN is flow based networking. The most predominant significant protocol in flow-based networking is OpenFlow. The OpenFlow protocol was the first standard interface of SDN supporting separation of traffic control from forwarding within network elements. SDN however is not OpenFlow. SDN is a network architecture enabling the decoupling of network devices from the control layer and exposing the network and its control to northbound applications using standard APIs. OpenFlow enabled SDN is a dynamic infrastructure of forwarding, control, and per customer network programmability.

With the entire network topology now discovered and known operationally by the SDN Controller, a number of disruptions to Assurance systems are ahead. The SDN Controller enables a simplified collection of network statistics in terms of traditional reporting. As will be discussed later, OpenFlow based SDN networks implement a rich set of interface, flow counters and QoS instruments which Assurance systems now no longer need to collect from each network element individually. This may now be performed with a simple RESTFul call over the northbound API to the SDN Controller. In fact the entire known topology of the network is now a simple REST call.

One of the most significant business reasons for SDN adoption in the industry is this separation of forwarding from control. The reason for this is the increasing demand for network density toward the edge. DOCSIS and other service provider access networks worldwide are shrinking the ratio of access mux to subscriber served. R-CCAP and WDM-PON are two examples of such innovations. As this trend continues, the increase of access network interfaces directly increases aggregation network density demands. Without a new approach to managing the dramatic increases in network elements year over year, operators may soon be faced with spending more money operating their network than the network is generating in revenues.

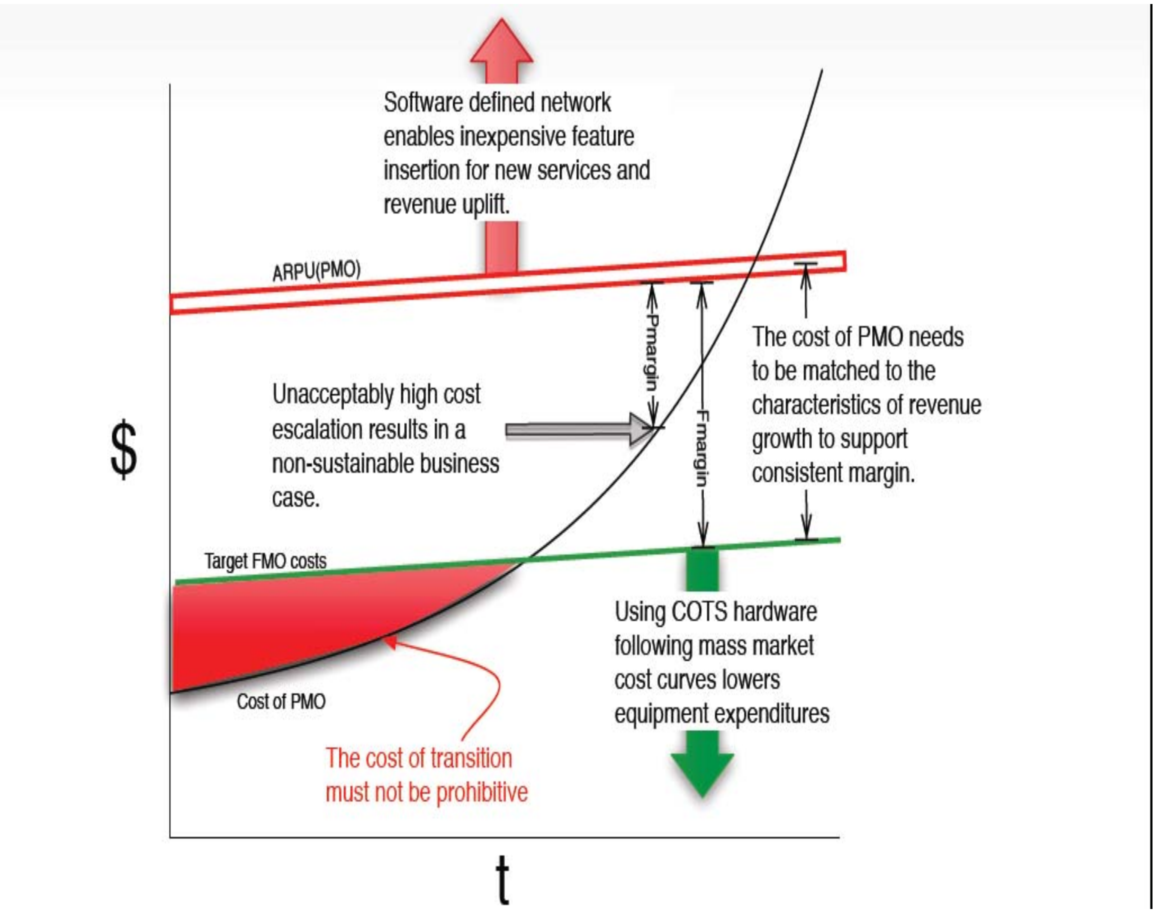

Figure 12. Verizon N2 GN Business Case[9]

In the above figure, Verizon was suggesting two key points. Verizon had predicted a cross over point where the ability to keep pace with year over year increases in bandwidth demand meant annual investments in more access, aggregation and core network elements become unsustainable. The traditional model of answering this growth, using vendor closed architecture platforms, means the fixed capital costs of growth remain unchanged. Further, the traditional network model demands an increasing operational cost per deployed network element that eventually out pace revenue.

For further reading, Finkelstein et al, provide a nice summary of SDN and NFV concepts in their 2014 paper[3] .

Today's CCAP Services

The CCAP environment today consists of either I-CCAP or M-CMTS deployments. MPEG/IP convergence occurs at the downstream RF port either on the I-CCAP box or on a universal edge QAM in the case of M-CMTS system. However from an IP forwarding point of view and as far as this paper is concerned the video functionality in either type of system is not very interesting. While an edge QAM is essentially an IP to MPEG bridge, it can be better thought of as an IP host.

The IP services offered by a CCAP system include the following:

1. Cable Bundling
2. Layer 3 Abstraction/Layer 2 "Wall"
3. DHCP Relay with subscriber and interface identification
4. Cable Source Verify
5. Cable Proxy ARP/ND
6. Layer 2 and Layer 3 Virtual Private Network Services
7. Virtual Router Forwarding
8. Policy Based Routing
9. Layer 2 and Layer 3 Access Control Lists
10. Protocol Throttling

Other services that will be interesting that are not specifically IP related include:

11. Classification
12. Policing
13. Rate Shaping
14. Subscriber Management
15. Subscriber Usage Reporting via SNMP or via IPDR

So the challenge for the operator when deploying a DAA is to ensure that each of these services that are provided by today's advanced CCAP solutions remains to be supported in equal or better fashion.

Provisioning CPE in CCAP Today

Current CableLabs provisioning is dependent on the DHCP protocol. CCAP elements provide DHCP Relay and DHCPv6 Prefix Delegation services to enable cable modem device initialization and IP addressing process to occur.

DOCSIS service groups provide cable modems with a MAC domain for forwarding of traffic. This forwarding supports the ability for broadcast traffic, such as DHCP DISCOVER packets to reach the virtual router in the CCAP. The virtual router terminates the Layer-2 MAC domains in the CCAP. It is at this point where Layer 3 services are provided such as DHCP Relay, IP Routing, and DSCP QoS marking.

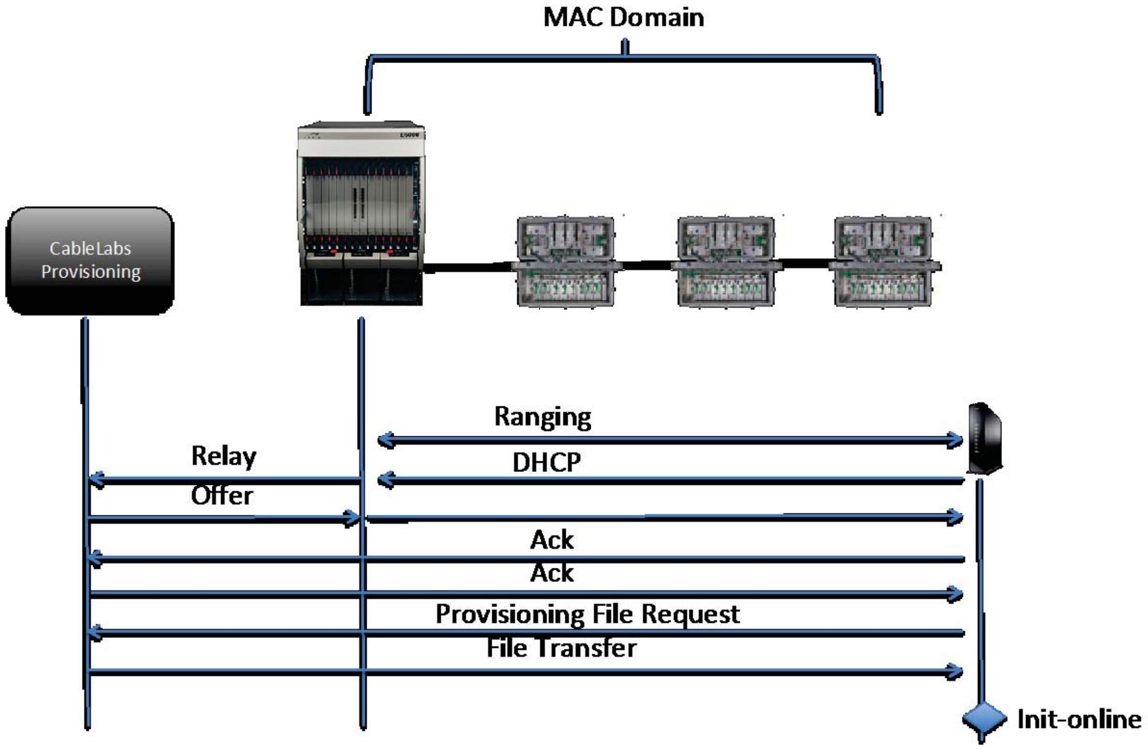

Figure 13. Current DHCP MAC Bounce Chart

Layer 3 Today

Today's Hub based CCAP platforms are advanced processors of DOCSIS PHY/MAC, Video QAM services, encryption, lawful interception, usage based billing records, access control, Layer-2 and Layer-3 forwarding including an array of protocols from which to learn neighbor and forwarding details with. Essentially today's Hub based CCAP network element is a multi-service platform operating over multiple different physical interfaces.

One key aspect of the CCAP today is Layer-3 forwarding. The Virtual Router inside the CCAP builds a neighbor database from a variety of gateway routing protocols. The Virtual Router also builds a local forwarding table tracking device attachment to interfaces. This local Forwarding Information Base (FIB) is the table of which port of a MAC domain owns which cable modem mac addresses.

When we discuss various R-CCAP topologies we often work to replicate our understanding of the current approach to Layer-2 and Layer-3 forwarding as it currently exists in the hub based CCAP.

For reasons we will demonstrate, it is impractical to replicate all the Layer-3 processing that exists within the hub based CCAP in remote CCAP nodes. The hub based CCAP can expand to support the emerging R-CCAP logical topologies, in the process the hub CCAP serves the R-CCAP nodes as the northbound or upstream service router.

Provisioning CPE in R-CCAP DAA

We previously discussed the relationship of Layer-2 broadcast forwarding as being a function of the MAC domain. MAC domain termination by a Virtual Router in the CCAP provides Layer-3 services and it is this interface providing DHCP Relay enabling cable modem provisioning.

However the MAC domain in an R-CCAP has segmented the former MAC domain into several smaller serving groups. While the benefit to smaller service groups, terminated closer to the attached cable modem has benefits in spectral performance, the logical network has now changed where initial cable modem provisioning is concerned.

It is impractical to suggest operators manage and allocate individual subnets to R-CCAP nodes, within what was a cascade of a single DOCSIS service groups served from the Hub. This would be both wasteful in IP allocation and operationally complex. Modems will continue to rely on DHCP, therefore the R-CCAP needs to become DHCP aware, one approach to do this is snooping packets.

One method of determining where in the segmented cascade of R-CCAP nodes a modem has attached is to shim additional information into the DHCP packet.

Effectively the R-CCAP operates a Layer-2 DHCP Relay Agents as specified in IETF RFC3046. A Layer-2 network element may operate between client and a DHCP relay agent. The purpose of this Layer-2 agent is to encode meaningful option 82 relay information objects. This approach is well established from Broadband Forum TR-101 and used in both xDSL and PON networks today to convey subscriber identity to attachment point during DHCP operations.

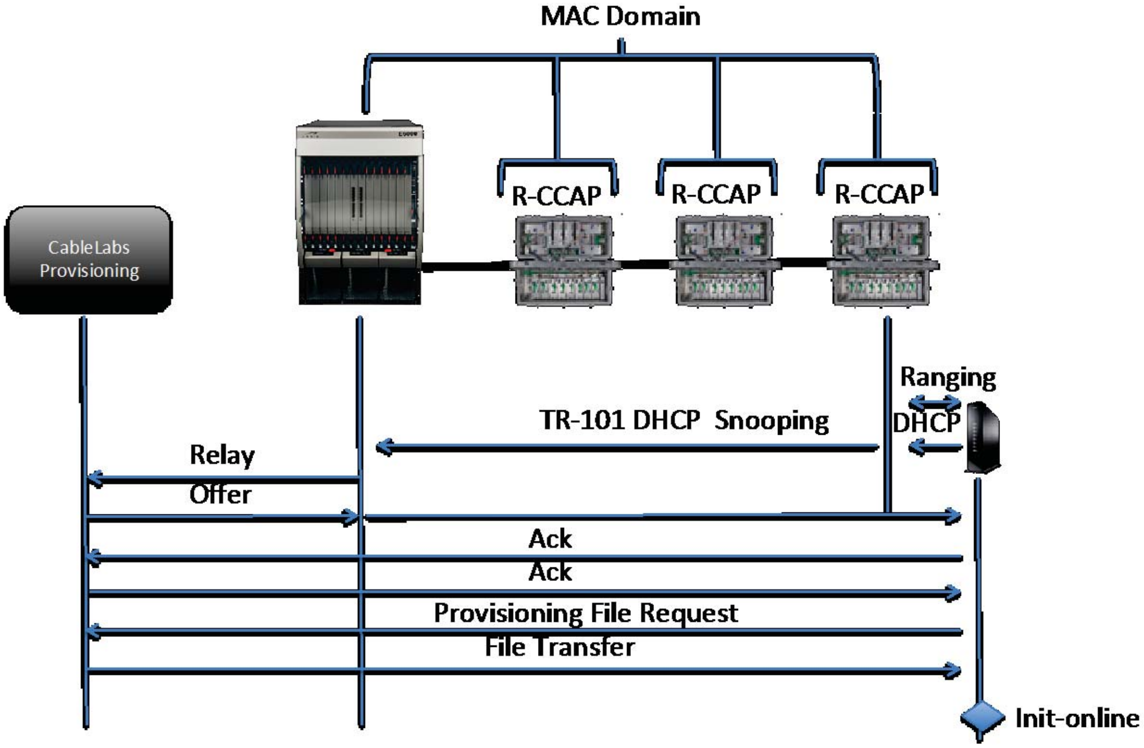

Figure 14. TR-101 DHCP Relay Snooping R-CCAP

In this scenario the R-CCAP remains a Layer-2 device with the exception of a northbound management IP interface. Client forwarding is operated at Layer-2.

The selection of Layer-3 addressing is based on a few factors. In addition to the incoming Relay GI-Address the upstream router has inserted as the DHCP Relay, the option 82 sub option 1 Circuit-ID has been populated with R-CCAP node name and port name. The TR-101 defined general syntax is “relay-node-identifier eth slot/port:[vlanid]” assuming Ethernet based access node identification. The exact method to determine slot and port identification is left as vendor determined which supports the conversation for R-CCAP of inserting network information to help steer logical topology. Sub option 2 remains populated with the cable modem MAC address as the remote-ID (RID) value.

When evaluating the incoming DHCP packet from relay, the DHCP service parses the 82.1 ASCII string to apply rule matching, which may dictate more specific addressing.

Each R-CCAP is receiving a tagged 802.1q trunk from the upstream router. In essence the current model of Layer-2 forwarding is no different than when the MAC domain was served from a Hub based CCAP. The main difference is the R-CCAP is switching PHY-MAC layers, adding DHCP snooping in the last mile for location identity.

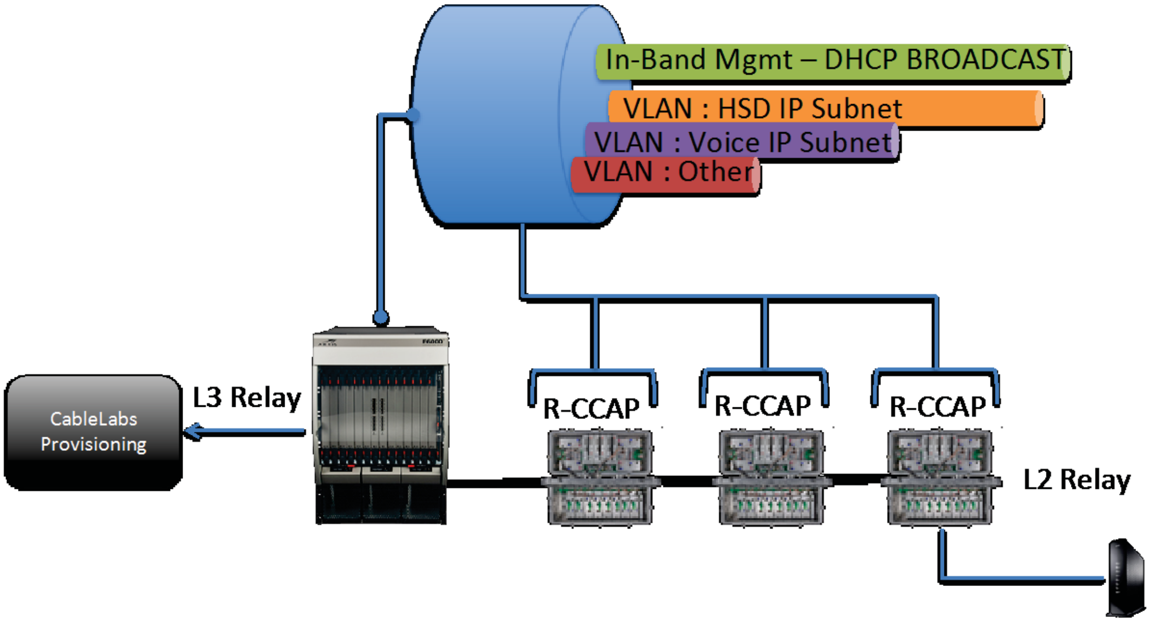

Figure 15. Logical Topology R-CCAP Upstream Router

SDN Support of CCAP Services

The following sections will outline how SDN can support CCAP services. Specifically these sections will address Network Provisioning, R-CCAP discovery, and CPE provisioning.

SDN Provisioning – The Network

The SDN Controller presents a centralized network provisioning function enabling each R-CCAP element to be provided local forwarding rules, or even per port instructions either physical port or logical interface such as MAC Domain based port(s).

As we have discussed, the DOCSIS network edge in DAA architecture will create a one or two magnitude increase of managed network elements to serve the end customer. It is entirely impractical to conceive of the legacy per box mechanisms of management, configuration, and local forwarding control when faced with such a massive increase in managed network elements. Traditional CLI approaches may not entirely disappear, however when introducing SDN to the DAA, we will inherit a new network that is not only more spectrally efficient, we will also have a fully programmable network offering agility for the next wave of IP service delivery.

In work already underway at CableLabs, a new network element configuration schema has been described. Within the CM-SP-CCAP-OSSI specification the CCAP XML file based and NETCONF based approaches to network configuration are described.

The principle difference between XML file based and NETCONF based configuration is the processing of configuration. In an XML file based CCAP, the full or partial configuration of the CCAP is transferred to the CCAP as an XML file confirming to the schema requirements. This transfer is enabled via SCP, or HTTP/s or similar secure delivery. On receipt of the XML configuration, the CCAP reads in file and processes all configuration changes. Effectively, the enumeration of the XML described objects is the application of the changes to the CCAP configuration and completes with merging the new operating configuration with the CCAP stored startup configuration.

In the NETCONF approach, a configuration session is opened with the CCAP where configuration changes are communicated as Remote Procedure Calls (RPCs). With each RPC there is error and status handling. Similar to the file based CCAP provisioning method, NETCONF protocol uses an XML based encoding for configuration data. The NETCONF protocol defined by IETF RFC6241 specifies its client operations over an RPC layer. The client session is negotiated from the provisioning system to the CCAP using YANG modeling permitting server to respond in a mutually agreed manner to the RPC events. YANG is defined by IETF RFC6020 and provides a data modeling language for NETCONF RPC operations.

The idea of NETCONF as a mechanism to enable network wide provisioning is not new, nor was it conceived by the efforts from SDN initiatives. Similar to CableLabs CCAP OSSI specification, SDN programs have chosen to inherit the benefits of NETCONF and YANG to offer an abstraction to large-scale network wide provisioning.

SDN and R-CCAP Discovery

Earlier the micro-segmentation of the cascade into unique R-CCAP elements suggested basic operations of DHCP would be augmented if the traditional approach to network service and topology were to be maintained.

CableLabs CCAP OSSI specification suggests the CCAP will initially be configured using CLI via a direct access means such as serial console cable. This certainly applies to the hub based CCAP today.

Field personnel in the outside plant will deploy R-CCAP. These teams typically do not locally configure strand or node based devices. The ideal path would be R-CCAP auto-initialization using common approaches available from the current back office. For example, if we think about the traditional IP network, a R-CCAP might initially start on a Gigabit or 10-Gigabit interface with a default 802.1q VLAN and initiate DHCP. This management VLAN DHCP traffic would be uniquely treated from any other logical network from the perspective of the DHCP service. The R-CCAP would likely be supplied an IP address and could be supplied location of a provisioning system similar to how the DHCP method of directing TR-069 location for ACS provisioning systems is specified by Broadband Forum using DHCP option 43 sub-options.

There are a few ways to consider the next steps of R-CCAP management and operations. If the implementation of SDN is purely for configuration management then it is possible to assume the R-CCAP, having been directed to its SDN Controller, will initiate a connection identifying itself in a unique manner. The SDN in this scenario would have pre-determined knowledge of the unique identification of the R-CCAP and process either a file based XML or RPC based NETCONF provisioning process. The result for the field technician is a status light on the R-CCAP changes indicating provisioning success.

In this approach, provisioning of R-CCAP was enabled, and continues to depend on high layer forwarding decisions within each R-CCAP. This essentially preserves the traditional switch-routing infrastructure known today.

However, if the DHCP services are part of the SDN Controller, and the R-CCAP itself implemented OpenFlow, the ability to uniquely treat network traffic in a programmatic and subscriber aware manner can be realized.

If the R-CCAP implements OpenFlow, the supplied location of the SDN Controller enables OpenFlow protocol communication. The R-CCAP in this scenario will initiate a TCP connection on port 6633. The conversation with the

Controller from the R-CCAP will be to enumerate each port in the R-CCAP associated with the internal OpenFlow switch. The Controller may now direct the R-CCAP per port with flow based forwarding instructions.

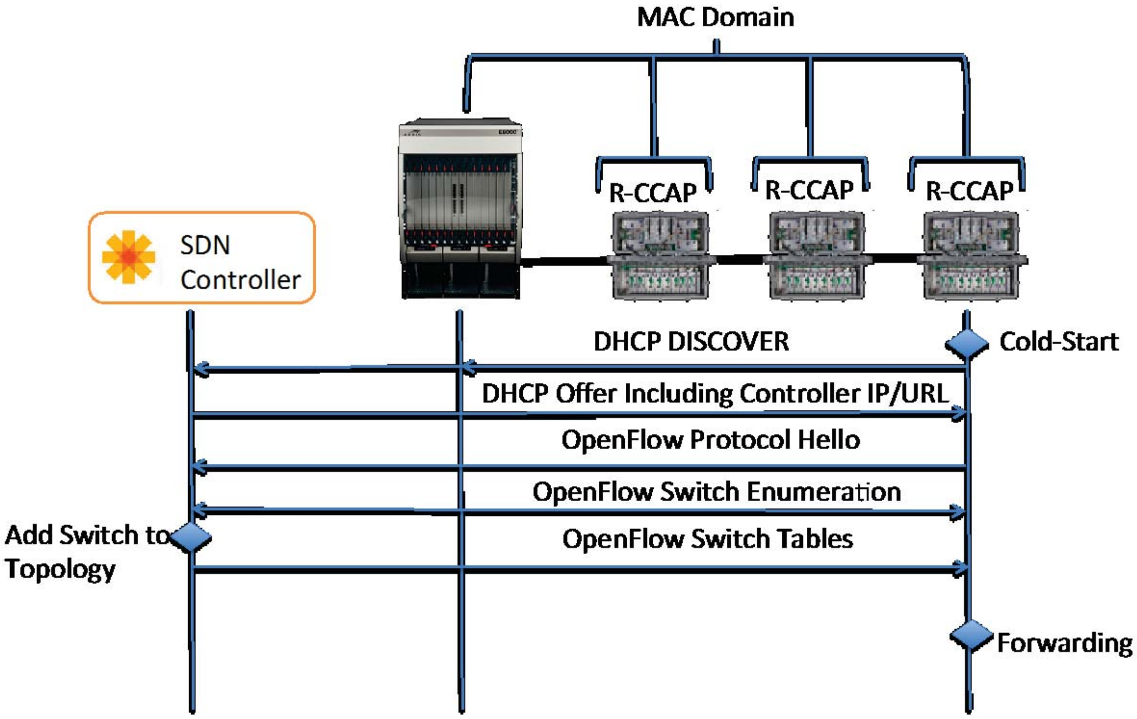

Figure 16. R-CCAP OpenFlow Initialization

Pipeline processing is applied to packets as they enter an OpenFlow ingress port. The ingress port may be used to match and classify packets. The OpenFlow pipelining process preserves knowledge of the ingress port per packet. Decisions of forwarding packets may be made by local flow table policy loaded into the OpenFlow switch from the SDN Controller, or may be sent to the Controller for decision and optionally stored for a period of time locally. OpenFlow supports physical and logical ports that support the needs of an R-CCAP or hub based CCAP MAC domain concepts.

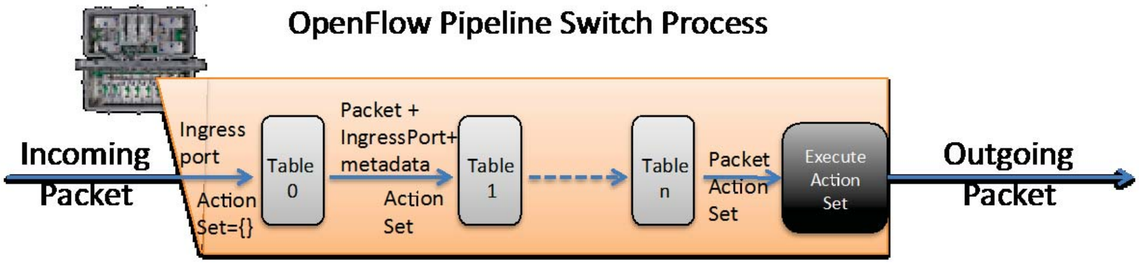

Figure 17. OpenFlow Pipeline Switching

OpenFlow switch tables are processed in sequential order, starting with highest priority matching table based on the packet. When the OpenFlow switch processes a packet the ingress port is used as a condition to initial match actions. Metadata may also be considered updated or cleared and passed to the next table if applicable based on the prior match table. When the packet has traversed all matching tables, the complete instruction set is executed. Unless the action set contained instruction to 'GoTo' another flow table, the packet is forwarded from the switch.

Match Fields	Priority	Counters	Instructions	Timeouts	Cookie

Figure 18. Flow Table Structure

If an incoming packet fails to be processed by the OpenFlow switch table, meaning the switch has no known actions based on the incoming packet, a 'table-miss' occurs. In this case, the R-CCAP or hub CCAP OpenFlow switch could simply drop the packet, send to a table specifically for table-miss or using the CONTROLLER reserved port, asks the SDN Controller for processing instruction.

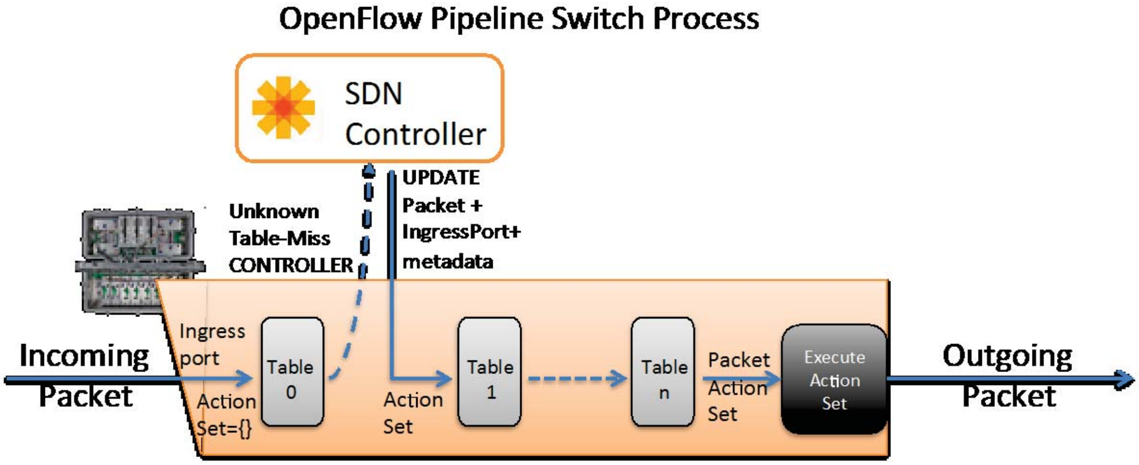

Figure 19. OpenFlow Pipeline Switch Process

If the Controller chooses to add a flow table for this type of packet, the packet is acted upon locally by the R-CCAP or hub CCAP switch. The Controller may have chosen to indicate when this flow table action should expire, meaning the decision is not necessarily permanent.

As part of the flow table processing, one table that may be updated is a Meter Table. Meter tables are where OpenFlow pipeline processing performs QoS. The Meter Table can record the rate of packets matching its table that is certainly of interest to recording performance of packets through the network. The Meter Table may act on packets applying QoS functions such as prioritization and rate-limiting in conjunction with per port queue techniques the R-CCAP or hub CCAP may be capable of implementing.

Further, in an OpenFlow topology, the minimum bandwidth that any queue should afford a packet as part of its flow path may be communicated during flow setup. This enables an end-to-end bandwidth setup known in OpenFlow as 'network slicing',

OpenFlow counters support the recording of multiple traffic levels and are implemented on a table base, with port, queue, group, meter and bucket basis. Counters enable visibility to active flow entries in the switch, lookups and matches. Flow based packet counters and the bytes or bandwidth per flow are optionally exposed. Per port received and transmitted bytes are required, as are queue packets. Many other counter objects are supported in the standard with exposure likely to vary by vendor for the optional elements.

The output of the OpenFlow pipeline process is to apply the action set. This is the combination of all table processing. This might be as simple as forward the packet out a specific port. There may be action to mark with QoS value, apply VLAN or MPLS or PBB header information or any possible combination as specified in an action list.

At this stage, FIB management is centralized while forwarding remains de-centralized for scale. The results are topologies that support autonomous inter-communication among common SDN Control. Ingress port table classifications create actions on forwarding decisions including egress ports.

A hybrid of OpenFlow and traditional forwarding model is also possible. The R-CCAP might be provisioned to enable a sub-set of its ports to OpenFlow and define other ports as being members of the traditional switch-routing facilities. Hybrid switches typically also support packets moving from OpenFlow pipeline to the traditional switch pipeline based on the action set to the packet.

This would be a common scenario in border node elements where the OpenFlow topology is interfaced with the current network infrastructure. Again as part of OpenFlow ingress packet table processing, if a packet for example were destined to participate in an MPLS topology, it may be so marked by the OpenFlow switch and preserved at the border node for seamless transport.

In this hybrid scenario the R-CCAP remains dynamically discovered by the back office platforms north of the SDN Controller. The forwarding decisions, on a per hop, per packet basis in the traditional network, are preserved where the non-OpenFlow ports are concerned. For the OpenFlow enabled ports, the SDN Controller will have complete control of the decisions for forwarding, policy, failover, bandwidth, QoS, and subscriber control.

SDN Provisioning – The CPE

Returning to the example of a DOCSIS cable modem entering the R-CCAP after ranging completes, the modem attempts DHCP. The ability for the cable modem to send traffic at all is a result of the cable modem MAC being admitted to the R-CCAP or hub based CCAP forwarding information base (FIB). The FIB admitted a new MAC address into the associated table based on the MAC domain. In an OpenFlow 'FIB', the cable modem MAC address may be learned as part of MAC learning associated to the MAC Domain or ingress port. The OpenFlow switch may have a default table allowing learned MAC addresses to pass DHCP and TFTP traffic. The modem without having realized it is communicating with an OpenFlow switch process initializes DHCP Discover message.

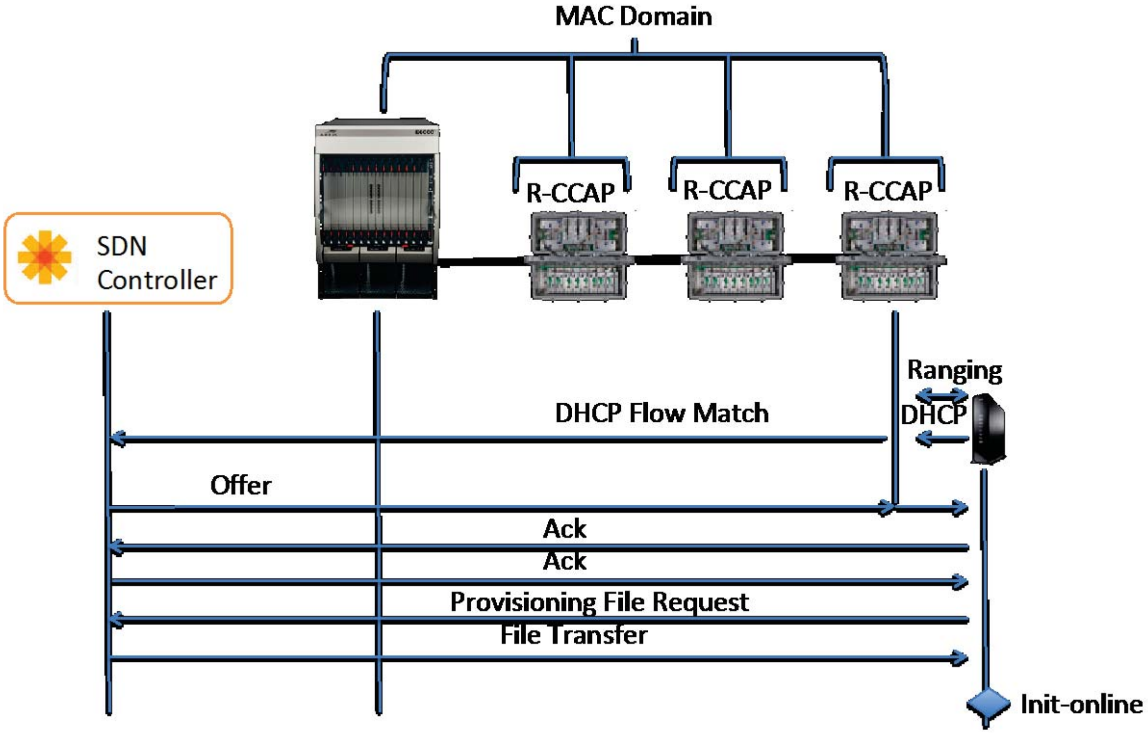

Figure 20. SDN CPE DHCP

Provided a flow table based on DHCP protocol exists the R-CCAP may have an egress action that forwards to one or more DHCP servers in the network. The R-CCAP action-set could be programmatically defined to preserve the cable modem MAC address known as remote-ID (RID), and write the equivalent of TR-101 legacy style DHCP location information as a Layer-2 DHCP action. However assuming the DHCP services exist within the SDN Controller, the visibility to the OpenFlow switch could be passed to the address allocation functions of the SDN Controller and thus eliminate the need to provide TR-101 legacy style DHCP location information as the SDN Controller is now the source of truth for all topology information.

This further reduces the operational burden on both the edge of the network in terms of maintaining configuration alignment, as well as co-ordination with DHCP systems in the back office as now the topology is dynamically inherited for provisioning and addressing concerns. The hub CCAP or R-CCAP could be further segmented and the back office would not need to be altered.

Another use case that highlights the significant change when provisioning systems are collapsed within the SDN Controller is in commercial services and static addressing. With commercial services today, cable operators must orchestrate the configuration of L2VPN or MPLS paths through the network itself while also coordinating the use of the L2VPN 802.1q VLAN identifiers or MPLS Pseudo-Wire identifiers within the cable modem provisioning file. While it may still be beneficial to implement the Type Length Value (TLV) modem configuration approach, many operators simply dedicate a cable modem to the commercial service. When the modem equals the service, an SDN Controller need only apply an 802.1q VLAN to all packets from the cable modem MAC. Alternatively the SDN Controller could update the push of an MPLS tag onto the packet as it egresses flow table processing northbound of the R-CCAP. Conversely, the packet would have the VLAN or MPLS additional information popped off before egressing southbound over the RFI interface to the cable modem.

This centralized topology awareness present in SDN is a significant step forward from current practice of coordinating local route descriptors (RD) for MPLS ingress interface processing at each CMTS, in addition to the label switch path (LSP) configuration necessary to encapsulate a path toward the core network which also must have broader topology awareness to align use of LSP's across the network. All the while these operations are controlled and secured using locally defined Access Control Lists (ACLs) per hop often increasing the manual network management burden.

For static IP customers, their commercial cable modem routers are assigned subnets that are painful to reconfigure when node splits occur and the MAC domain owning the IP sub-interface next-hop can no longer be found. As mentioned earlier, the OpenFlow concept of network slicing can be applied to solve for this operational issue. The commercial cable modem will be provisioned with its eRouter static IPv4 address. This will generate an ARP packet. OpenFlow network slicing permits match condition based on IP field of the ARP. Therefore the source IP of the commercial customer eRouter can programmatically be dynamically moved throughout the R-CCAP network as its egress table action will be to present the packet to the appropriate upstream router.

By comparison, managing the movement of a static IP subscriber to another cable interface can involve multiple manual touch points in the operator back office. There are two common approaches to static IP customers today. First is to provision the customer eRouter cable modem to advertise its subnet using RIP. The CMTS is then configured to restrict advertised networks via RIP learned over the cable interface using an ACL. For this approach to work means managing ACL's across all CMTS in the network which represents its own effort and often involves co-ordination with the provisioning and or billing back office systems. For operators who do not permit any route advertisement to originate from the eRouter device, the CMTS must have local knowledge of the destination subnet for commercial static IP customers for each cable mac domain. Therefore managing the static forwarding tables implies its own operational cost of management. Coordinating the movement of static IP subscriber when either of these two methods are in use represents multiple manual network touch points.

These are only two of many subscriber cable modem forwarding use cases that serve to highlight the significant operational savings possible with the introduction of SDN technologies and centralized FIB management.

Cable Industry View of SDN

While much of the early hype around SDN was focused on disruption of the traditional switching and routing network model with complete separation of the control and data planes, most of the service provider community and the cable industry in particular, have adopted the more evolutionary hybrid SDN approach described above. In this approach SDN capabilities are added alongside the established service provisioning and forwarding functions implemented in current generation CCAP devices, with a focus in simplifying the provisioning process and streamlining operations.

The choice of SDN controller technology is a key factor in determining how effectively MSOs can meet these operational goals. Inside CableLabs, as well as many vendor research labs, the OpenDaylight SDN controller has emerged as the primary choice for prototyping and evaluating SDN operations. Other commercial and open source controllers remain viable alternatives, but the OpenDaylight controller provides several key attributes needed for a successful cable network SDN implementation:

1. Programmability, in the form of an extensible data model, associated "northbound" APIs, and a service abstraction layer to isolate applications from low-level protocol details;
2. Flexibility to extend the data model with cable-specific attributes and new APIs;
3. Support for OpenFlow, NetConf, and SNMP, as well as a "plug-in" architecture to add additional "southbound" protocols to communicate with network elements.

For example, CableLabs has previously demonstrated a prototype OpenDaylight plug-in to support PCMM, a protocol used in MSO networks to provision new service flows with desired QoS attributes.

Much work remains to be done before SDN technologies are ready for widespread MSO access network deployment. Migrating provisioning and management functions to SDN-based applications while maintaining compatibility with existing back-office systems will be challenging. Troubleshooting in a hybrid SDN environment will require new tools and a significant investment in training. Other questions to be addressed include redundancy models, reliability, and scale. But the challenges of managing DAAs may be the driver to move SDN out of the research labs and into MSO networks.

Conclusion

DAA whether in the form of R-PHY or R-CCAP is attractive to MSOs as a cost effective way to extend the existing HFC network architecture such that it remains competitive into the 2020 decade and hopefully beyond. However, DAA does have challenges in that the number of network devices that need to be managed increases by up to two orders of magnitude. SDN and OpenFlow offer a starting point for the HFC community as tools to help build a network that can scale to these demands. SDN is a standards based eco-system of software enabling the network to be programmed thus promoting rapid service delivery. SDN has a growing user and vendor community, its support of dynamic service creation based on topology and subscriber awareness delivers the most significant disruption in IP network control and forwarding of the last 20 years.

Bibliography

[1] EXAMINING THE FUTURE EVOLUTION OF THE ACCESS NETWORK, Michael J. Emmendorfer and Tom Cloonan, NCTA 2013

[2] Remote PHY: Why and How
Jorge D. Salinger

[3] USING SDN AND NFV FOR INCREASING FEATURE VELOCITY IN A MULTI-VENDOR NETWORK
Jeff Finkelstein, Alon Bernstein, and Samir Parikh, NCTA 2014

[4] NEXT GENERATION - GIGABIT COAXIAL ACCESS NETWORK,
Michael Emmendorfer, Tom Cloonan, Scott Shupe, and Zoran Maricevic, NCTA 2010

[5] PREDICTIONS ON THE EVOLUTION OF ACCESS NETWORKS TO THE YEAR 2030 & BEYOND,
Tom Cloonan, Mike Emmendorfer, John Ulm, Ayham Al-Banna, and Santhana Chari

[6] A SIDE-BY-SIDE COMPARISON OF CENTRALIZED vs. DISTRIBUTED ACCESS ARCHITECTURES,
Michael J. Emmendorfer, Thomas J. Cloonan, John Ulm, and Zoran Maricevic, NCTA 2014

[7] Remote PHY for Converged DOCSIS, Video and OOB,
John T. Chapman, CTO Cable Access BU & Cisco Fellow

[8] Scaling Traditional CCAP to Meet the Capacity Needs of the Next Decade
John Ulm, Tom Cloonan

[9] Software Defined Networks: A Carrier Perspective,
Stuart Elby, Verizon, Open Networking Summit
2011 https://m.youtube.com/watch?v=xsoUexvljGk

[10] Unlocking the Potential of DOCSIS Through D-CMTS
Jeff Heynen, Infonetics Research, Huawei Communicate, Issue 71 November 2013

[11] RESTful Web Service

Richardson, Leonard; Sam Ruby (2007), O'Reilly Media, ISBN 978-0-596-52926-0, retrieved 18 January 2011.

[12] Breathing New Lifespan into HFC: Tools, Techniques, and Optimizations
Dr. Robert Howald; NCTA 2013

Abbreviations & Acronyms

ACS – Auto Configuration Server

AM – Amplitude Modulation

API – Application Programming Interface

ARP – Address Resolution Protocol

ND – Neighbor Discovery – IPv6 protocol for ARP functionality.

ASCII - American Standard Code for Information Interchange

CCAP – Converged Cable Access Platform

CLI – Command Line Interface

CMTS – Cable Modem Termination System

CPE – Customer Premise Equipment. Usually means equipment located on the premise to provide network access such as a cable modem

D-CMTS – Distributed CMTS

DAA – Distributed Access Architecture

DHCP – Dynamic Host Configuration Protocol – see RFC 2131

DHCPv6 – DHCP for IPv6

DOCSIS – Data Over Cable Interface Specification

DSCP – Differentiated Services Code Point

DWDM – Dense Wave Division Multiplexing

EPON – Ethernet Passive Optical Network

FIB – Forwarding Information Base. This is the table used by a layer 3 switch to determine where to send packets.

GPON – Gigabit Passive Optical Network

HFC – Hybrid Fiber Coaxial

HTTP/s – Hypertext Transfer Protocol

I-CCAP – Integrated Converged Cable Access Platform

IETF – Internet Engineering Task Force

IP – Internet Protocol

IPDR – Internet Protocol Detail Record

IPv4 – Internet Protocol version 4

L2VPN – Layer 2 Virtual Private Network

M-CMTS – Modular Cable Modem Termination System

MAC – Media Access Control

MER – Modulation Error Ratio

MPEG - Moving Picture Experts Group – usually refers to the framing format of the protocol used to communicate movies.

MPEG/IP – MPEG delivered over IP

MPLS – Multiprotocol Label Switching

MSO – Multiple System Operator

NCTA – National Cable & Telecommunications Association

NETCONF - Network Configuration Protocol

NFV - Network Functions Virtualization

OLT – Optical Line Terminal – This is the headend equipment to deliver PON

ONU – Optical Network Unit – This is the CPE equipment to deliver PON

OOB – Out of Band. Can be any out of band messaging. However in this context refers to SCTE 55-1 and SCTE 55-2

OSI - Open Systems Interconnection

PBB

PHY – Physical Layer

PON – Passive Optical Network

QAM – Quadrature Amplitude Modulation

R-CCAP – Remote Converged Cable Access Platform

R-PHY – Remote PHY – Refers to the PHY layer functionality of a modular CCAP network

REST - Representational state transfer

RESTFul – A web service can be characterized as "RESTful" if it conforms to the constraints described in the Architectural constraints section of [11] .

RF – Radio Frequency

RFI – Radio Frequency Interface

RID – Remote ID

RPC – Remote Procedure Call

RPN – Remote PHY Node[2]

SCP - Secure Copy

SDN – Software Defined Networking

SFP+ - Enhanced Small Form Factor Pluggable Interface

SNMP – Simple Network Configuration Protocol

TCP – Transmission Control Protocol

TFTP – Trivial File Transfer Protocol

TLV – Type Length Value

VLAN – Virtual Local Area Network

WDM-PON – Wave Division Multiplexed Passive Optical Network

XML – eXtensible Markup Language

YANG - Data modeling language for the NETCONF network configuration protocol

xDSL – Any flavor of Digital Subscriber Line