SCTE·ISBE CABLE-TEC EXPO 2018 | ATLANTA, GA OCTOBER 22-25

2018 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS

Node Provisioning and Management in DAA

What Changes Are Coming for Techs

A Technical Paper prepared for SCTE•ISBE by

Robert Gaydos
Fellow
Comcast
1701 JFK Boulevard, Philadelphia PA
215-286-8737
Robert_gaydos@cable.comcast.com

Mehul Patel
Senior Principal Architect
Comcast
183 Inverness Drive West, Englewood, CO, 80112
303-658-7826
Mehul_Patel@cable.comcast.com

Joe Solomon
Principal Engineer
Comcast
401 Wynkoop St Ste 300, Denver, CO 80202
303-242-7037
Joe_Solomon@cable.comcast.com

Table of Contents

List of Figures

List of Tables

Introduction

In traditional HFC deployments, the headend technician controls the services that are delivered to a fiber node because this is done via RF combining in the headend. CMTSs and Edge QAMs output RF channels at configured channel frequencies. This output is split and combined such that the appropriate services are delivered to the lasers going to the node; once the wiring in the combining network is complete, it is rarely changed.

By contrast, and as HFC infrastructure evolves toward a Distributed Access Architecture (DAA), consequent RPHY (Remote PHY) node combining is done virtually, with software. Each node must be virtually directed to appropriate service "Cores," expressed by frequency plan per service, and video multiplexes to be joined. Service cores span DOCSIS flows, linear and on-demand video, and legacy, out-of-band information.

More specifically, each node must be software configured to listen to its appropriate QAM broadcast and VOD feeds, as well as connect to the correct CMTS, and legacy Out-of- Band (OOB) components.

These advances mean that the industry's technical workforce needs to be able to program the node, along with Cores which also have to be configured to provide the right data. Pre-planning which physical node (and MAC address) will be installed at a given fiber location is improbable, because line technicians typically carry many nodes. As a result, mechanisms to map logical nodes with intended service configurations and physical instantiations are required.

At the same time, an increasingly intrinsic design goal for network design is to prevent "vendor lock in." In order to encourage a competitive cost and inovation environment, operators prefer and require multiple sources of components and to be able to pivot to new resources easily. In addition operators must deal with the realities of supporting differing QAM video conditional access (CAS) systems throughout their footprint.

If an operator the size of Comcast used a CMTS to provide all services and manage all aspects of a DAA node, for instance, it could yield as many as 18 permutations of nodes, CMTS, and CAS to test and integrate. Clearly, this is not sustainable. This led to the notion of applying separate "Cores" for broadcast video, VOD, out of band (OOB), and high speed data flows. Cores allow for best-of-breed product selection. Also, keeping video out of the HSD cores simplifies CMTS operations; more importantly, any call to pivot to a new CMTS, or multiple CMTS providers, would sidestep the need to re-integrate video services across six permutations of nodes and CAS systems.

The work related to disaggregating service flows into multiple Cores presented the next dilemma: Deciding which "Core" to make the "primary," or lead coordination Core. This led to the creation of a "primary core" that is, in essence, a vendor-independent orchestrator. This allows us to mix Cores and nodes at will, and to use our own internal software management processes and tools when needed. We call this software "GCPP," which stands for "Generic Configuration Protocol Principle." This internally-developed software performs the following functions, which will be discussed in the paper:

- Network access and steering – aligning the DAA node on the network with the appropriate Cores for configuration
- Configuration – providing QAM Video, VOD, legacy out-of-band configurations and other non-DOCSIS functions
- Service routing – orchestrating the multicast IP routing of video and out-of-band content to the DAA node across the IP network

- DAA node management – managing the versions of software and the code's signed certificates used by DAA nodes in the field

Yet another design goal was to do as little upfront design as possible, relying on auto-discovery instead of complex, pre-drawn wiring diagrams to connect nodes with switches and photonic muxes. This relates to the concept of the logical vs. physical node. The logical node has a name or ID known to billing systems and GIS systems; it has a known channel map and frequency plan. The physical node is the hardware that hosts the logical node. The physical node can be replaced because of hardware failure or natural disaster. Our system uses "late binding" of physical to logical node mapping, meaning that it happens at the time of install (via an app), thus allowing any node in inventory to host the logical node. (This alone vastly simplified construction processes.) In addition, we tried to ensure that the node's connectivity to a CMTS would be detected, rather than designed. This makes capacity management easier, ensures the databases are up to date, and gives greater visibility into the network for technicians..

This paper describes the software infrastructure used to manage the thousands of nodes that will be transitioning to DAA. The solution makes use of software defined networking (SDN), distributed cloud servers, and multiple Cores.

Overview of a DAA Control and Data plane Architecture

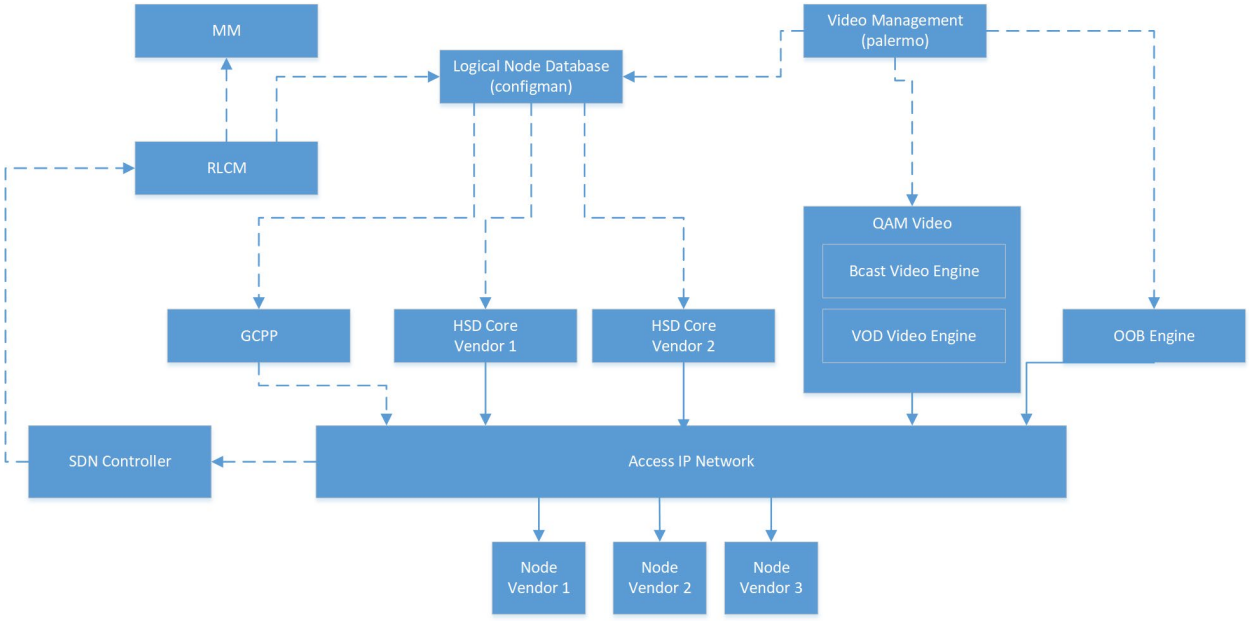

Figure 1 Comcast's DAA Architecture

Figure 1 illustrates the DAA architecture we developed to overcome the challenges covered in this paper. While we don't envision running Cores from different vendors in the same sub-section of the plant, it is possible. In fact, it is highly probable that nodes from multiple vendors will exist in the same sub-section of the plant. The following is a brief description of the components:

- SDN Controller – A software-defined networking controller that programs the access IP fabric for both multicast and unicast, detects new and dropped hosts in the network, and heals faults in the fabric.

- Access IP Fabric – A group of whitebox switches comprised of redundant leaf and spine switches that connect to a DAAS (Distributed Access Architecture Switch), which aggregates 10GE connections to remote nodes. There is at least one access IP network per headend and on average 10 individual networks controlled by their own SDN controller. Each network serves approximately 12,000 HHP (households passed.)
- GCPP – General Config Protocol Principle, an internally-developed software component that configures the RPD, receives alerts, and refers the RPD to auxiliary Cores, such as HSD Cores and possibly legacy OOB Cores[1] (55-2). GCPP receives dynamically-generated configurations for each node.
- RLCM – RPD LifeCycle Manager. This internally-created component is the workflow engine that performs the late binding configuration of nodes. It receives notifications from the SDN controller when a node joins the access IP network, which includes the switch and port number. The RLCM will also trigger a "match maker" function to pick an available V-CMTS Core and configure it. RLCM will also communicate with a server for 802.1x authentication, as well as an inventory management system to ensure the device isn't cloned or otherwise outside of our equipment inventory.
- Logical Node Database – hosts the frequency map (both HSD and video) for a logical node. Maps the video multicasts to static pseudo-wires -- for example tunnels and the frequencies that they are placed on.
- Video Management – "Palermo" is another internally-created tool that tracks desired video configurations on a per node basis. It also must be aware of or program the configurations of video and OOB engines.

DAA-related Challenges Facing Field Employees

1. Network Access and Steering

Each DAA node connects to the CMTS in the headend via Ethernet optics, rather than the traditional HFC network. While 802.1x-based network authentication verifies that the remote-PHY device (RPD) has valid, CableLabs-provided certificates, it does not ensure that the device is one from the operator's inventory, versus a cloned device. Additional measures to authenticate devices connected to the operator's IP network will be needed. As mentioned earlier, we implemented an app that binds a physical node to its logical configuration. No device is allowed to enter the network, meaning that the 802.1x request will not be granted and the port will remain "off" if this binding is not done. Furthermore, any attempts to bind multiple physical nodes of the same ID (MAC address) to different logical nodes, i.e. a cloned device attempts to enter the network, will fail because the backoffice prevents such security breaches.

Once connected and authenticated, the DAA node has to be directed to the appropriate service-providing devices (HSD Core, Video Core, OOB Core, etc.). When a DAA node obtains its IP address from the DHCP server, the DHCP server also provides the IP address of its Principal Core. The challenge for the operator is to determine how to direct that DAA node to the correct Principal Core, among the several IP reachable Cores that service that node's population. In other words, in a DAA world, several principle IP cores are IP-reachable, whereas before DAA only one CMTS was physically reachable, because of the

[1] Cores use l2tpv3 protocol to establish dynamic tunnels. It is a two-way control protocol; the RPD picks the tunnel ID. Engines multicast data using static tunnel IDs. They do not communicate via a control protocol with the RPD. Static tunnels are provisioned on the RPD via the GCPP.

wiring; what used to be 1:M, CMTS to nodes, is now N:M, depending on how many switches are connected together..

2. Configuration

Each Core contributes to the node's frequency plan, and for each RF channel in the frequency plan, the node has to be configured to join video multicasts and transmit the multiplex, or to stand up the channel as a DOCSIS service. Legacy video out-of-band (SCTE 55-1 and SCTE-55-2) upstream and downstream channels also have to be configured on the right frequencies. The channels that are present are determined by the geographical area the node serves (its logical configuration) -- but as the DAA node joins the network, it is impossible to ascertain its desired serving area.

When a new node joins the network, the only identifier that it provides is its MAC address. Again, pre-planning which physical node (identified by its MAC address) will be installed at a given location/connected to a fiber is improbable, because line technicians and construction crews typically carry several nodes on the truck at a time and should not be tasked with ensuring that node *A* is placed in location *Z*. For those reasons, a mechanism was required to map a logical node, with its intended service configuration, to its physical instantiation installed in the plant.

3. Service Routing

Because the service and control traffic is delivered via IP on pseudowires, the unicast and multicast routes must also be set up and managed on the network, between the Cores and the DAA nodes. The broadcast video content is multicast to multiple nodes, as they share video channel service groups -- but network capacity is not infinite. Therefore, operators need to ensure that video multicasts are present only on links that will use them. Pseudowires containing video-on-demand content may also be multicast to multiple nodes for efficient use of Edge QAM resources. Orchestrating this routing, especially as new nodes join the network, is complex.

4. Node Management

Once verified, the node must get the latest firmware version. Unlike RDK-based cable modems and gateways, where all possible drivers are distributed, or legacy set-top boxes with on-board agents that check for updates, the Primary Core must ensure that each node is running the latest software, and if not, instruct the node which software file to download and from where to retrieve it. The Principal Core needs to provide software details that are specific to the make and model of the DAA node, as firmware is vendor- and version-specific. Also, in a CI/CD (continuous integration/continuous development) software deployment, not all nodes will get the same software at the same time.

DAA nodes also need additional controls placed upon secure shell(SSH) access – because as shipped from manufacturers, these nodes have static admin usernames and passwords. Because the DAA node is an IP device in an unsecured location, stronger user authentication was needed and must be applied before the node goes into service. Updating thousands of DAA nodes on a service bench, before installation, is impractical, as nodes are shipped to warehouses across an operator's footprint. Ultimately, a solution was needed to allow the node to be updated in the field as part of the installation process.

Core Configuration

In our architecture, it is the responsibility of the Principle Core to configure the DAA node with the settings necessary to bring it up and to begin delivering video services.We developed a Principle Core that delivers no services, but instead supports such management via the generic control protocol, or GCP. The GCP Principle Core (GCPP) is responsible for providing the DAA node's initial configuration; the only interaction it has with the DAA node is via GCP. Once initial configuration is completed by the GCPP, the DAA node is "handed off" to an Auxiliary Core, which configures HSD services. That configuration process, and the tasks performed in each stage are described in the following subsections.

1. GCPP Configuration

The GCPP is responsible for configuring the following on the node:

- The DAA Node's Auxiliary Cores
- Non-service-specific operational configuration (precision timing server settings, RF ports, event management)
- Video services (described in section 1.1)

As described previously, the DAA node gets the IP address of the GCPP from the DHCP server when it receives its IP address. When the DAA node sends a configuration request to the GCPP, the GCPP does not know what configuration to apply to the DAA node, as it only has the node's MAC address and no other indication of where the node has been installed, or what service groups it serves. Essentially, the GCPP is not pre-provisioned with the configuration for a DAA node. Instead, the GCPP sends a request to a new back office element that serves as a repository for DAA node configuration files – the Configuration Manager. This request includes the MAC address of the RPD; the Configuration Manager, working with other back office systems, determines the appropriate configuration to apply to the DAA node. The process of associating a DAA node's MAC address with the appropriate configuration detail is discussed in more detail in the Video Configuration section.

While the video configuration details are determined by node's geographic serving area, the other items that are configured by the GCPP are the same for all DAA nodes it manages. Figure 2 depicts the GCPP's event flow, and Table 1 summarizes the configuration settings that are the same for all DAA nodes that the GCPP manages.

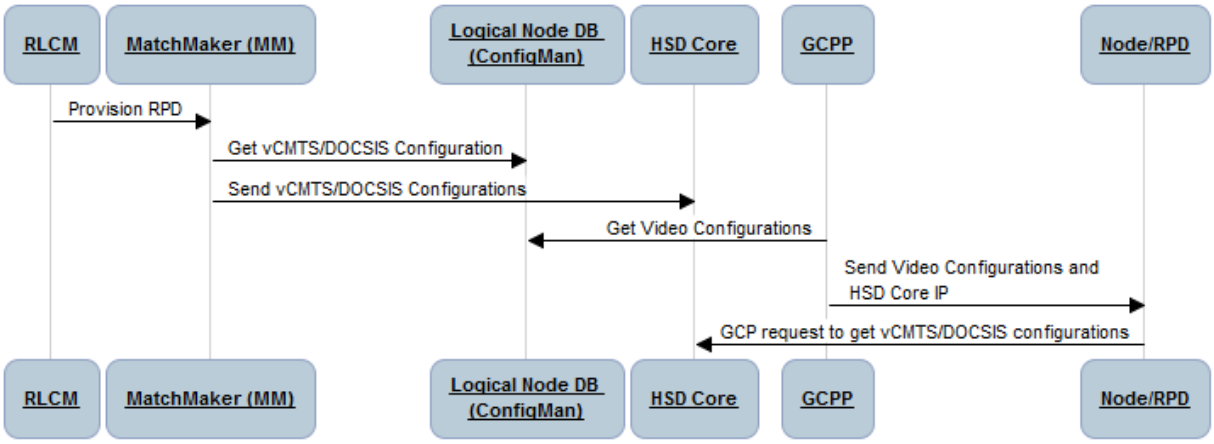

Figure 2 - GCPP Event Flow

Table 1 – GCPP Common Configuration

Settings	Purpose
Core details, including Aux Cores	Configures the GCPP Core details and the IP addresses of any Auxiliary Core to which the RPD will connect once the GCPP finishes configuration (e.g., CMTS Core, Video Out-of-Band Core, etc.).
Core Reconnect	Controls how long the GCP control plane can be idle before the DAA node considers the connection to the GCPP down; also determines the action the DAA node is to take upon connection failure.
PTP settings	Configures the DAA node's connection to the precision timing server.
Downstream and upstream ports	Configures the active RF ports on the RPD and associated power levels.
Carrier Wave (CW) tones	Configures the placement of CW tones in the downstream RF spectrum that can be used for RF leakage detection and automatic gain control.
Event reporting	Configures how DOCSIS-defined events are handled by the DAA node – logged locally or sent to the GCPP Core

The GCPP receives almost all of these configuration details from the Configuration Manager via a configuration file created for the specific DAA node. However, some of the configuration settings are local to the GCPP and Aux Core installation and are not known by the Configuration Manager or the back office – for example, the local IP addresses of the Precision Time Protocol (PTP) server and the Auxiliary Core. For these configuration parameters, the configuration file passed to the GCPP contains variables. The GCPP reads a local file, written in YAML (Yet Another Markup Language and/or YAML Ain't Markup Language), and replaces the variables with the values that are specific to the local installation. As part of the DAA installation process, the YAML file is updated with the appropriate values by the local technicians with access to that information.

1.1. Video Configuration

1.1.1. Logical Nodes

In the past, cable operators were only really concerned with the logical node since physical nodes were not addressable. The logical node ID is associated with the service addresses and is governed by the franchise agreement. It is this logical node ID that is referenced in billing systems, GIS mapping systems, and OSS tools.

The logical node ID associates all the configuration data needed to provision and manage services for a DAA architecture. The logical node ID is best conceived coincident with the initiation of the fiber design. The designer will use these logical node IDs when designing the node and fiber build out from the plant. From a design perspective, the logical node ID is used to represent the location of the physical RPD.

The logical node ID follows some rules and criteria from a node provisioning and management perspective, including:

1. The logical node ID can be an arbitrary value but must be unique across the entire footprint
2. The logical node ID is what will be used to associate with location/addresses
3. The logical node ID is what will be used to associate with subscriber accounts
4. The creation and assignment of the logical node ID must be completed prior to build out and activation of any RPDs
5. The association of the logical node ID must be one-to-one to a physical RPD MAC
6. The logical node ID will be used to create the Video and DOCSIS configurations needed to provision and manage services

1.2. Linear Service Group

The Linear Service Group (LSG) is a logical construct for grouping nodes that share the exact same binary multiplexes to their customers, i.e. they are in the same PEG zone, blackout zone, EAS zone, and ad zone. They all share a common channel lineup which is carried on the same channel map/frequency plan. A node can be in only one LSG. To create an LSG, a location and franchise-based community unit ID (CUID) is used as foundation, to help with identifying the specific video configuration to associate to a node ID and also helps with associating the right zones to the specific nodes.

The LSG will need to have configuration for the following:

1. Broadcast services - a broadcast services configuration, including channel lineups pulled from the video controller; the video controller is a possible resource for helping to identify the source IDs that can be used to create a LSG.
2. Public, Educational, and Government Access Channels (PEG) services.
3. OM (Out-of-Band Modulator) and VARPD – Downstream and upstream traffic configuration.
4. DSG Configuration

To help with collecting the configuration data for the broadcast and PEG services, we initially looked at the video controllers (DAC and DNCS) to help map the corresponding source ID and multicast feeds that feed the channel lineup for a downstream plant. The lesson learned was that for local access, the source IDs were duplicated and would cause issues when trying to build unique configurations for a specific lineup. For testing and trial purposes, the initial video configuration was incorporated into a spreadsheet, and a python script was created to build configuration files necessary for the GCPP and Node provisioning. The mapping of a logical node to an LSG was done manually and the LSG's configuration

information copied into the spreadsheet for every node in the LSG. By doing it manually first, we were able to quickly adapt as we learned which parameters were essential and which were superfluous.

The management of this configuration has now been placed into a video application internally called "Palermo," which is a tool to manage LSG configurations, map logical nodes to LSGs, and finally push the video configurations to the configMan. We do not keep LSG information in the Access Network platform. Rather, we let Palermo manage the definitions of an LSG in case the defintions change. This isolates video business policy from the enforcement provided by the access network. To the access network, it looks like each node has its own video configuration.

A Video On Demand (VOD) Service Group is a logical construct of grouping nodes together that share all the same On Demand video feeds. A node can only be in one VOD Service Group, but a VOD service group can contain multiple nodes. The number of nodes and the number of QAMs in a service group depends on historical usage, number of homes passed per node, and the amount of frequency the operator can afford to provide for the service. We chose to use four nodes per service group, with a shared pool of four QAMs.

The question then became how to map a node to a service group. In the past, this was done in the RF combining network. The VOD system had no knowledge of nodes and this still holds true. Palermo only needs to provide the frequencies for the VOD QAMs. Four multicasts from the VOD engine are mapped to four ports on the Ethernet switches connecting the RPDs to the access IP network, i.e. the DAAS. When a node is connected to the DAAS, its VOD service group is implicitly selected. The RLCM is notified of this port connectivity by the SDN controller. The configuration for the VOD portion of the node is then completed and stored in the ConfigMan, including the multicast addresses that the node should map to the frequencies given by Palermo.

2. Aux Core Configuration

2.1. vCMTS Configuration Collection

The original CableLabs architecture assumed DOCSIS as the primary core, with the option to hand control for certain services to other Auxilary cores. Our primary core is the GCPP, and all DOCSIS cores are auxiliary. The Aux Core is responsible for configuring the DOCSIS data paths for the RPDs. In our vCMTS architecture, each node is mapped to a single vCMTS instance, which is a set of containers managed by Kubernetes for that node only.

For configuration of the Aux Core, a "Remote PHY CRD" (Custom Resource Document) is created in a JSON format. One CRD is created for each vCMTS instance The CRD is comprised of many parameters, some of which are dependent on logical node while others are standardized or static for all nodes.

A CRD template was created that maintains all the static values. For the components that required specific variables, the template was broken out and specific "configlets" were created. This allowed specific modifications of the CRD using site-specific variables. The following configlets were created to allow site- and node-specific configurations:

1. Downstream (DS) Channel Configlets – Allow modification of the number of channels and site-specific center frequencies.
2. Service Class Name (SCN) Configlets – Allow site-specific tiers of service flows, where the name would be standard but the speeds can vary by site and node.

The DAA node will inherit the configuration of its parent i.e. the analog node that served the same set of customers prior to be converting to digital. The backoffice will retrieve the the specific DS channel configuration, the center frequencies used for those channels, and the SCN profiles and specific values for downstream and upstream speeds from the CMTS that hosted the parent node.

2.2. Node ID to RPD Configuration Association – An App for That!

The challenge of mapping a logical node to its physical host was solved by building a smartphone app specifically for the line techs doing the initial node install. Each RPD comes with a QR code that contains the serial number and MAC address of the physical node. Once the technician arrives at the location where the RPD will be installed, geolocation is derived and used to provide a list of nearby intended logical Node IDs. The technician, via the handheld application, associates the specific RPD MAC to a unique node ID by scanning the QR code on the RPD. After the association is made, a trigger is then sent to the RLCM, which in turn triggers the "match maker" to pick an available vCMTS core, generating most of the configurations for video and vCMTS. The RPD is then activated using the configurations from Configuration Manager. The final configuration of the RPD is then stored in a database.

Node Management

1. Software Version Management

Another significant shift facing the industry's technicians, is that DAA nodes require software, where analog nodes did not. Cable operators carry vast experience downloading code to embedded devices, such as set-tops, cable modems and gateways. However, there was always a "get out of jail free card" available if the download was not successful. That came in the form of the ability of the customer to restart the device. That's not an option here: Restarting a node hanging on a strand would necessarily require a bucket truck.

As with CPE devices, the software that ships on node devices is usually out of date by the time it goes into service. Therefore, it's almost a guarantee that the initial provisioning of a node will require its software to be upgraded. However, the latest version of software might be immature. The last thing anyone wants is to deploy new software to a new node. It would be impossible to ascertain if resultant customer issues were attributable to the new software, or the new plant changes. Therefore, the configuration management software must be capable of determining which code version a node should have, based not only on manufacturer and model, but also the logical node itself. At Comcast, this function is provided by the RLCM.

Ideally the initial push of a new software version would not go to nodes that are have had recent or ongoing issues. Perhaps other system software upgrades, i.e. new guide firmware, was recently pushed and for related reasons, the system is at risk of having a negative customer experience. Or, perhaps the node recently suffered a fiber cut or other outage. Therefore the configuration management system should be able to use heuristics to determine the appropriate software version for a node. The DAA node is told the version of software it should have at its initial connection to the GCPP. Given that the value changes on a per-node basis, static configuration files are not a viable solution. The configuration manager must dynamically determine the version of software that a node should use. When a software upgrade is rolled out, the nodes need to be notified of the change in a throttled manner to reduce the risk of having too many changes happening to the network at the same time.

2. Secure Shell (SSH) Lockdown

One of the first challenges we encountered was that nodes ship with a default set of accounts and passwords, for debugging purposes. This struck us as an unreasonable security risk. As a result, we elected to prevent any device from accessing the nodes directly; only certain servers, with a given set of keys, are able to access the nodes. Our internal term for this system is "autobahn." The nodes need the public keys of the "jump hosts" that are allowed to connect to the nodes. Upon initial provisioning, the default user name and passwords must be revoked or changed, and the current public keys associated w/ the jump host private keys must be installed.

The RLCM detects the presence of default accounts and absence of "autobahn" keys, and takes corrective action to change the access controls.

3. Controlling Multicast Video by SDN Controller

In our DAA architecture, we do not pre-plan to which DAAS or port a node is connected. We discover this information when the node joins the network. It's at this stage that the VOD service group is determined. Based on the physical-to-logical node mapping provided by the aforementioned app, and the port and switch information garnered by the RLCM, a multicast path can be created. We do not use SSM (Source Specific Multicast) in the node, because we have multiple sources of broadcast video. We also chose not to use multicast control protocols e.g. PIM or MLD, as this would defeat the purpose of using inexpensive white box switches. Instead, we build the multicast paths, and then monitor them. If the SDN controller detects a failure in the path, it reprograms a new path. Also, the multicast redundancy app that we created for the SDN controller will monitor the bit rates of the multiple video sources, and can cause drops of all sources but one. In this way, the system controls the sources,destinations and internal port flooding of multicast switches, based on real time information and port discovery.

4. How Do You Know It Worked

In the current architecture, headend techs can tell if their combining is right by tapping into the RF combining network and hooking up set-tops and cable modems. This verifies not only that the video is working, but that the correct channel map, PEGs, and ad zones are being included.

With DAA, signal generation is done for the first time in the node. Therefore, there is no way for techs to check their work, in terms of the soft configuration of the system. We approached this problem in the following manner: The first DAA node that we turned up wasn't actually done on a pole. We installed a node in the headend, which required installing its own power supply. Then we connected the RF-to-optical analog transmitters and recievers, which were connected to existing nodes in the field. This let us create a pseudo-mCMTS. It also let us roll back to an existing upstream and downstream port, on an existing non-DAA CMTS, as well as the legacy video RF combining network.

In the next phase of deployment (i.e. the first node on the pole), we built a portable test rack with a set-top box, cable modem, eMTA, etc., that can be carried on a truck. When the fiber to the node is spliced in, temporary power is applied and an RF cable is temporarily run to the truck. We dynamically provision the node, as described above, and ensure the node is working properly. Technicians can check that the video is correct. Then the node is spliced into the RF and power network. At this point, we are still somewhat blind as to what's in the network. That's why it is essential to apply real time telemetry, to see if the node is receiving and transmitting data on all of its channels. (That discussion is beyond the scope of this paper.)

Conclusion

The evolution of the industry's workhorse HFC architecture toward a Distributed Access Architecture brings with it several changes to node configuration and maintenance that will impact how headend, line and field technicians work.

In traditional HFC deployments, the headend technician controls the services that are delivered to a fiber node because this is done via RF combining in the headend. CMTSs and Edge QAMs output RF channels at configured channel frequencies. This output is split and combined such that the appropriate services are delivered to the lasers going to the node; once the wiring in the combining network is complete, it is rarely changed. Node combining in DAA is done virtually, with software; nodes are virtually directed to different service "Cores," expressed by frequency plan per service, video muxes, DOCSIS-related flows, linear and VOD, and legacy/out-of-band information.

These advances mean that the industry's technical workforce needs to pivot toward node programming, Core configuration and dynamic mapping, to link logical nodes with intended service configurations and physical instantiations. The learning curve will necessarily involve network access and steering, configuration, service routing, and node management.

Comcast worked extensively to develop automated, manageable, and monitor-able solutions for provisioning and managing DAA nodes deployed in the field. While the foundations for these systems are available via the standardized DAA protocols, additional work was required to simplify the provisioning and management of services on DAA nodes.

Abbreviations

CAS	Controller Access System
CRD	Custom Resource Document
CUID	Community unit ID
CW	Carrier Wave
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
DWDM	Dense wave division multiplexing
GCP	Generic Control Protocol
GCPP	Generic Control Plane Principal Core
ISP	Inside plant
L2TPv3	Layer 2 tunneling protocol, version 3.
LSG	Linear Service Group
Mcast	Multicast
Mplx	Multiplex – a collection of streams
OOB	Out of band
OSP	Outside plant
PTP	Precision Time Protocol
PW	Psuedowire – a data tunnel
RLCM	RPD Life Cycle Manager
RPD	Remote Phy Device
SCN	Service Class Name
SSH	Secure Shell

VOD-SG	Video On Demand Service Group

