# Leveraging Legacy Video In Digital Access Architecture Networks

A Technical Paper prepared for SCTE•ISBE by

**Wesley Weiss**
Network Architect III, Video Infrastructure
Shaw Communications Inc.
2728 Hopewell Place NE, Calgary, Alberta, Canada T1Y 7J7
587-393-0654
Wesley.Weiss@sjrb.ca

**Anjan Bajwa,** Co-writer/Shaw Communications Inc.

**Corwin Martens,** Contributor/Shaw Communications Inc.

 

# Table of Contents

**Title** **Page Number**

Table of Contents ..... 2
1. Introduction ..... 4
2. Drivers for Legacy Video Integration into DAA ..... 4
3. Principal and Auxiliary Core Overview ..... 5
4. Current State of Legacy Video (Pre-DAA) ..... 5
5. DAA Technology and Components ..... 7
   5.1. VUE Solution ..... 8
   5.2. APEX 3000-NE ..... 10
   5.3. Video Topology Manager ..... 10
   5.4. Legacy SCTE55-1 Interactive Signaling ..... 10
   5.5. DSG SCTE55-1 Interactive Signaling ..... 13
   5.6. Converged Interconnect Network ..... 13
   5.7. Timing ..... 14
   5.8. Remote PHY Device ..... 14
6. DAA Monitoring Solution ..... 15
7. Video and VOD Encryption ..... 15
8. Video Broadcast Acquisition Architecture ..... 16
9. Narrowcast Architecture ..... 17
10. IPv6 Addessing Standards ..... 19
11. Multicast, DEPI, and SessionID Relationships with Physical Frequencies on the RPD ..... 19
12. Regional Video Realities and Solutions ..... 20
13. Shaw Automation Strategy ..... 20
   13.1. Service Director ..... 21
   13.2. Resource Inventory ..... 22
   13.3. VOD Central Database ..... 22
   13.4. Automating the Video Topology Manager ..... 23
14. RPD Provisioning and Activation Process Through Automation ..... 23
15. Lessons Learned ..... 25
16. Conclusion ..... 26

Terms and Abbreviations ..... 27

References ..... 30

 

## List of Figures

**Title** **Page Number**

Figure 1 Legacy Hub Site Video Architecture .... 6
Figure 2 DAA Data Flow .... 7
Figure 3 Shaw DAA Rack Layout .... 9
Figure 4 RPD and vARPD Relationships .... 11
Figure 5 DAA Interactive Network (Physical Locations) .... 12
Figure 6 DAA Work Flow .... 12
Figure 7 DSG Interactive and Data Flow Traffic .... 13
Figure 8 CIN Network .... 14
Figure 9 DAA Monitoring Architecture .... 15
Figure 10 Broadcast APEX 3000-NE Data Centre Network Configuration .... 16
Figure 11 NE GUI Showing VOD Encryption .... 17
Figure 12 Narrowcast Architecture of the APEX 3000-NE Blade and VUE Configuration .... 18
Figure 13 SD DAA Interactive Map .... 22
Figure 14 Activation Sequence Diagram During the RPD Activation Process .... 24

 

## 1. Introduction

Digital Access Architecture (DAA) is driving the transformation in next-generation cable-access networks. In contrast, the effective integration of legacy video services and automation deployments can be a daunting, and an often-overlooked task. DAA technology is continually evolving, and these issues apply to many operators facing similar challenges. Creating a solid foundation and adopting DAA technology in stages allows for a scalable video architecture and a *one-touch* approach to deployment, thereby avoiding expensive upgrade costs to customers and infrastructure.

This whitepaper explores Shaw's operational and technical complexities of integrating legacy video into a DAA network and presents the best practices and lessons learned from this undertaking. A robust discussion around functional and automation transformations are covered to prepare operators for large-scale Remote PHY Device (RPD) deployment, modification, and monitoring in an evolving network ecosystem.

## 2. Drivers for Legacy Video Integration into DAA

Shaw believes that DAA is the path forward in upgrading to a robust, resilient, high performance, future-proof, hybrid fibre-coaxial (HFC) network. However, a large install base of legacy customer premise equipment (CPE) exists and needs to be considered. Shaw set out to create an architecture that fully supports in-place capabilities of legacy CPE while still being able to embrace a move to DAA, allowing for upgrading to DAA infrastructure without the added cost of CPE replacement. Furthermore, due to the realities of hub site to node combining, a non-trivial amount of power and space is being expended to the end of manual radio frequency (RF) combining. The ability to reduce entire racks of connected gear to a handful of 1RU switches and dense wavelength division multiplexing (DWDM) hardware could offer significant power, space, and cooling savings. As a result, reducing the requirements for real estate, possibly being able to house an entire hub site in a small fibre cabinet.

While initially researching architectures for legacy Video DAA, we endeavoured to utilize a solution that could re-use as much of the original legacy acquisition, encryption, and out-of-band (OOB) communication hardware. Leveraging as much virtual and Commercial Off-the-Shelf (COTS) hardware as possible to minimize vendor lock-in and maximize flexibility in the event of changing use requirements. Another significant consideration was compatibility with chosen RPHY nodes that were being tested and decided on by other teams. Taking all of these considerations into account, we built the architecture described in this paper.

The evolution of DAA infrastructure has allowed Shaw to utilize the spectrum more efficiently. It has already reduced combining losses at the plant level for installed RPDs by 6dB downstream and 3dB upstream. We have launched our first RPD-only hub site with no video RF combining. Once networking was present on-site, Shaw was able to turn this hub site up and achieve full RPD functionality in a matter of hours, due to the reduced RF combining.

 

## 3. Principal and Auxiliary Core Overview

When creating the initial architecture, the legacy video architecture was approached with the following three principles:

- Scalability
- Compatibility with current legacy systems
- Difficulty of automation

Having a separate auxiliary core in each region for legacy video services behind a principal core is the solution that most closely meets Shaw's current and future video needs. The auxiliary core is a secondary GCP controller core tasked explicitly with handling video RF and channel map configurations after handoff from principal core. The principal core is a dedicated GCP controller that handles initial configuration and handoff to auxiliary cores as well as all DOCSIS frequency configuration and control. This design's substantial benefit is the Cable Modem Termination System (CMTS), and video can be scaled separately depending on customer and business requirements. In the future, because of the modular architecture, the video auxiliary core can be eliminated once the transition to Internet Protocol television (IPTV) is complete with no impact to the principal core, aside from the removal of a few global settings on the CMTS. The video auxiliary core is lightweight, scalable, and relatively cost-effective compared to utilizing a CMTS video core in each region and affords a more finely tuned automation schema.

## 4. Current State of Legacy Video (Pre-DAA)

Using a headend/hub site architecture has various components distributed across the network. Most of the video components and RF combining happens in hub site facilities, which may have space and power limitations. Linear broadcast services are processed in core sites, i.e., a headend, and then combined in hub sites with narrowcast Video on Demand (VOD) and DOCSIS services to produce the full RF line up. Scaling activities and node splits can increase complexity when combining occurs in the hub site, potentially forcing the need for additional hardware.

 

An example of Shaw's original legacy hub site architecture is displayed in the following diagram:

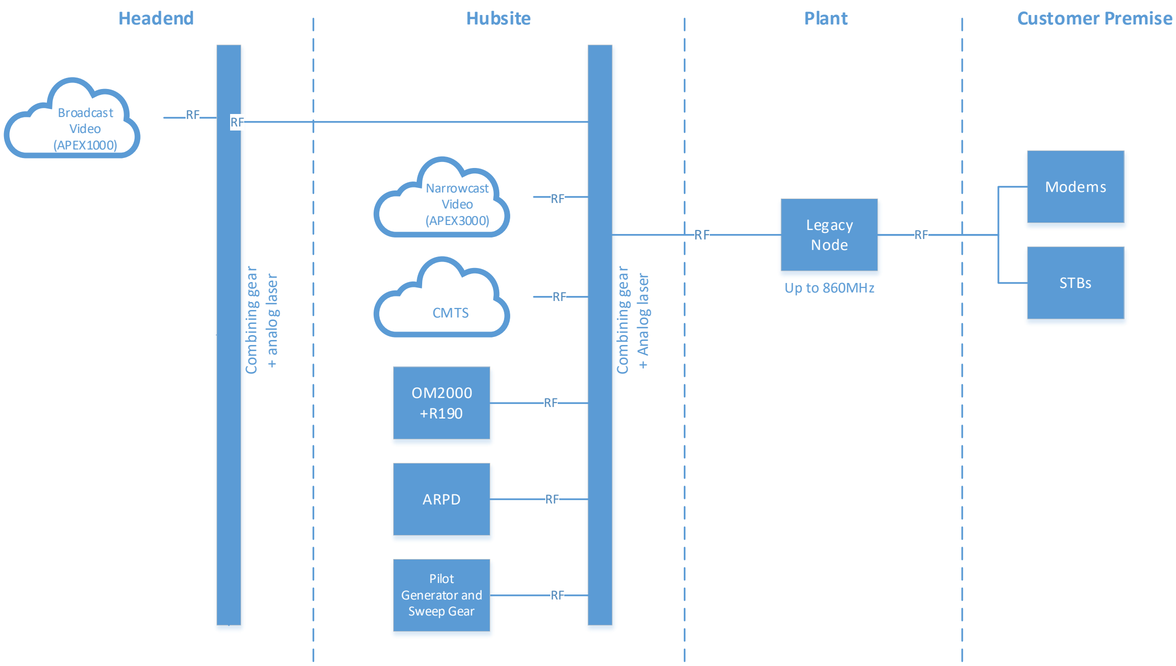


**Figure 1 Legacy Hub Site Video Architecture**

With this engineering, linear video is acquired at various signal acquisition facilities (SAF) and made available to video multiplexers—Digital Content Managers (DCMs) and edge Quadrature Amplitude Modulation (QAM) CommScope APEX1000s. Here it is distributed on Shaw's video wide area network (WAN) over an internet protocol (IP) multicast format.

Once multiplexed into Multi-Program Transport Streams (MPTS), edge QAMs in the headend subscribe to these video multicasts and modulate the signals to video QAM carriers over RF. This broadcast RF spectrum leaving the headend is then transported to attached hub sites, where the broadcast video gets combined with local narrowcast services (VOD), and put on the plant.

 

## 5. DAA Technology and Components

The DAA video core infrastructure is responsible for the aggregation and encryption of video services, and the configuration of all video, OOB, and SCTE-55 communication for DAA RPDs. In collaboration with CommScope, it was possible to lay a solid foundation for a genuinely scalable DAA video architecture that creates a reliable infrastructure for rapid expansion without incurring technical debt.

Considering the long-term goal of DAA and that DAA requires the same interactive components of a legacy node (RADD [Remote Addressable Danis/DLS Downloader], OM [Out-of-Band Modulator], ARPD [Advanced Return Path Demodulator], etc.), it was decided not to create net-new. Instead, Shaw opted to utilize existing equipment already in the hub site and leverage it in the DAA solution. Each net-new and replacement DAA video node is attached to an already existing and functioning analog node or hub site in a logical manner, therefore, enabling seamless and transparent migration from analog nodes to DAA nodes in the event of a node swap or split, providing the same valued customer experience across the entire footprint.

Housing the non-legacy component of DAA in the data centre, it is possible to reduce a hub site footprint to as small as a single Converged Interconnect Network (CIN) leaf router.

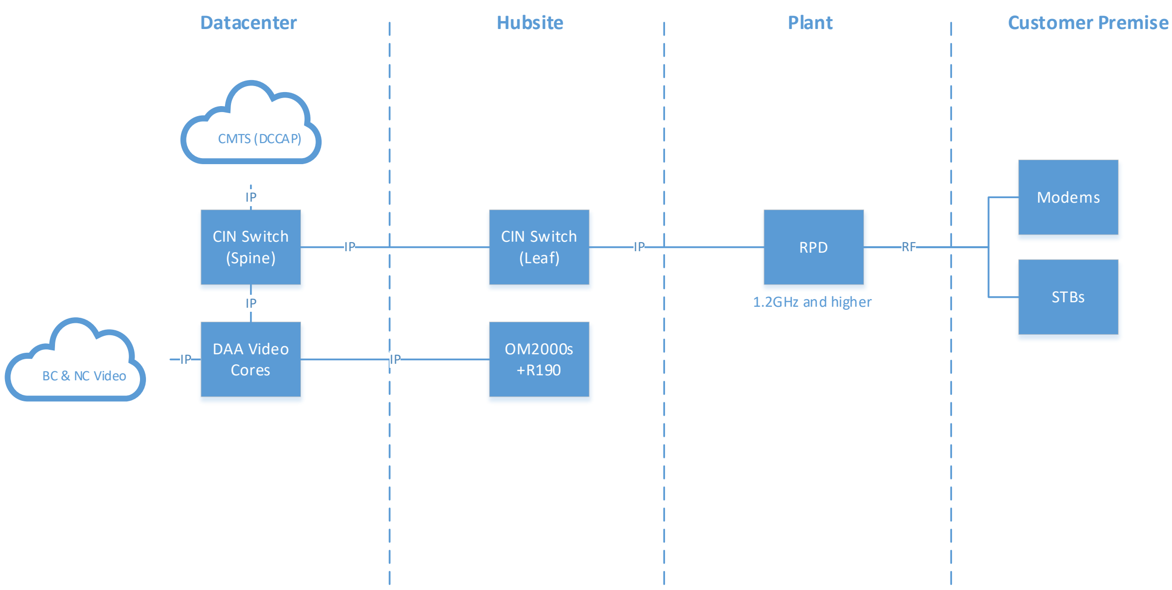


**Figure 2 DAA Data Flow**

 

### 5.1. VUE Solution

The Video Unified Edge (VUE) provides the network and RF information between the video sources and the RPDs. It is a highly independent scalable solution designed to replace the legacy headend video gear functioning as an auxiliary core for Shaw's DAA environment. VUE is a containerized application which can be used in both bare-metal and virtualized environments.

The following servers establish the VUE:

- Two video platform server (VPS) systems operating on separate virtual machines:
  - VUE Docker repo (VPS#1), and
  - VUE platform services (VPS#2)
- VUE application servers (pipe servers), which operate on bare metal chassis

The two VUE platform servers share the functions of managing and maintaining the VUE environment. VPS#1 deploys and controls the entire VUE OS/Docker solution and performs upgrades and maintenance. VPS#2 is responsible for the Generic Configuration Protocol (GCP) communication to the RPD through the Auxiliary Core Control Plane (ACCP) sub-system, which is in a dedicated Docker container. Because this system communicates directly with RPDs, this sub-system requires connectivity to the CIN in addition to a connection to our management network.

VUE application servers (VUE pipe servers) are responsible for the creation of Downstream External PHY Interface (DEPI) and Upstream External PHY Interface (UEPI) streams. These streams carry interactive and video MPTS multicast traffic to the RPD. In Shaw's infrastructure, each pipe on the VUE pipe server is in its own Docker container and carries a single 38.8Mbit MPTS.

Internal communication between all CommScope VUE servers (applications server and VUE platform servers) use either Secure Shell (SSH) or Application Program Interface (API) over our Operations Administration Maintenance and Provisioning (OAMP) network.

The following image displays the rack elevations for the DAA infrastructure. Each network encryptor (NE) blade on an APEX 3000-NE can output approximately 6.5G of traffic, allowing Shaw to build 42 service groups per VOD COR – a VUE pipe server designated for narrowcast, at four Annex B MPTS per service group. Annex B is a cable standard that modulates a MPEG-TS input into a QAM-256 output. Each service group is assigned four VOD carriers and four RPD nodes initially. This number can be scaled up and down to dynamically meet fluctuating capacity demands utilizing our back office and automation tools.

Each DAA rack with seven VOD pipe servers can feed 1176 DAA nodes with no physical combining involved. With the second rack added, a total of 3,528 nodes can be fed.

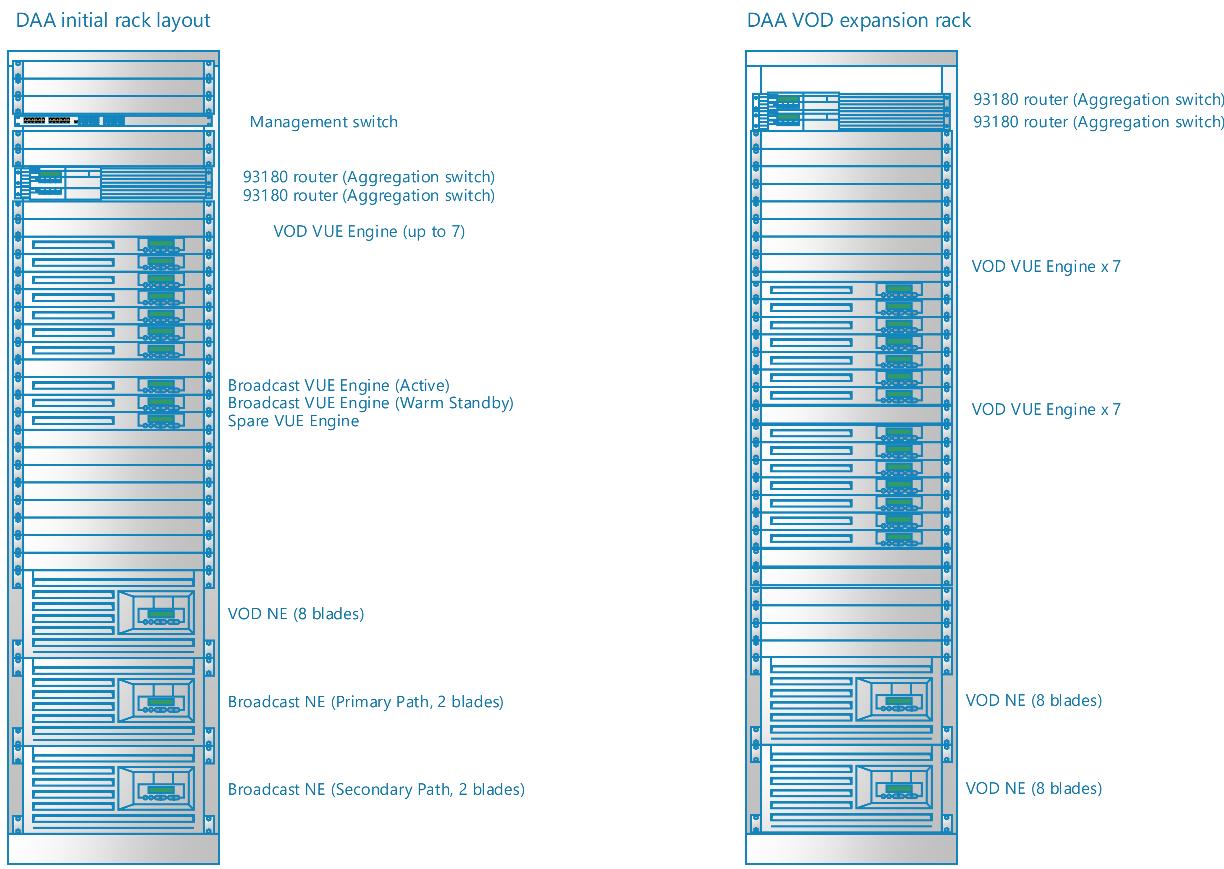


**Figure 3 Shaw DAA Rack Layout**

## 5.2. APEX 3000-NE

The APEX 3000-NE device is a high-density network encryptor capable of supporting and ingesting 12,288 unique program identifications (PIDs); up to 4096 multicast streams. It can support Single Program Transport Streams (SPTS) for VOD and MPTS for broadcast services. The APEX 3000-NE egresses encrypted traffic on the same physical port that it ingests unencrypted traffic, allowing for a denser installation requiring fewer physical switchports.

Features:

- Eight blades configured in an N+1 model
- 12 total 10GigE interfaces where four interfaces are the back up for the first eight interfaces
- Redundant host modules
- Redundant power supplies
- Support for User Datagram Protocol (UDP) port mapped VOD and includes support for Real Time Streaming Protocol (RTSP) provisioned VOD services
- Performs broadcast and SDV encryption
- Supports PSI generation and message insertion
- Support for SCTE-52 encryption for broadcast and VOD through Common Tier Encryption (CTE)

## 5.3. Video Topology Manager

The Video Topology Manager (VTM) is a software-defined management solution that configures and manages the video configurations in DAA. Service groups with specific settings are created in VTM for each of the video services offered and then assigned to applicable RPDs. Additionally, VTM interfaces with the Digital Addressable Controller (DAC) and receives all channel information for the region. This related information defines service group information for the RPDs configured by the regional ACCP sub-system.

A single instance of VTM can administer the entire DAA topology. Aside from using VTM to automate, this solution can also build serving groups and modify channel maps. In the future, there will be an option to use VTM as a dashboard that enables granular troubleshooting on a per RPD basis country wide.

## 5.4. Legacy SCTE55-1 Interactive Signaling

### 5.4.1. Virtual ARPD

For interactive traffic, the Advanced Return Path Demodulator (ARPD) functionality is on the RPD. RF streams are demodulated and encapsulated using the Remote Upstream External PHY Interface (R-UEPI) standard and routed to the broadcast application server (VUE pipe server). The application server identifies the upstream traffic by the Layer Two Tunnelling Protocol (L2TP) *sessionID* and *varpdSourceid*, then routes this data to the correct port on a legacy Network Controller (NC-1500). Each Virtual Advanced Return Path Demodulator (vARPD) – one per hub site, is its own logical entity in the DAC. It is configured identically to a physical ARPD on both the legacy Network Controller and the DAC.

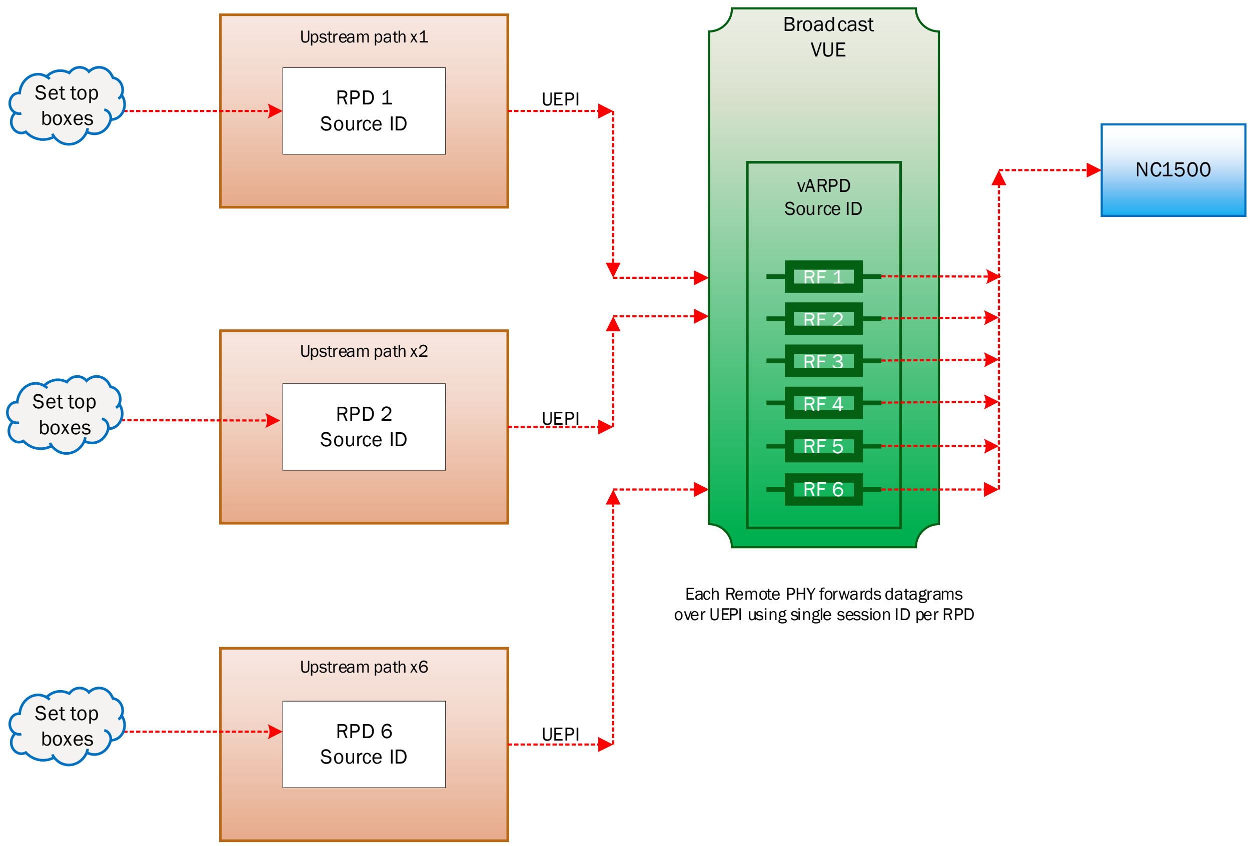


**Figure 4 RPD and vARPD Relationships**

### 5.4.2. Out-of-Band Modulator

Downstream OOB data is still handled by an existing OM2000 located in each hub site. The out-of-band modulator (OM) sends forward data streams containing guide data, code modules, channel map information, Emergency Alert System (EAS) messaging, DAC messaging, and app network data.

On the legacy OM2000, an IP interface is enabled pointing at the VUE application server (VUE pipe server) over the OAMP network. The OM2000 communicates with the VUE application server utilizing unicast UDP IPv4 traffic. The OM2000 packs seven, 188-byte MPEG packets into a single 1,500-byte packet and forwards to the VUE pipe server's IPv4 OAMP interface on a unique port. Each hub site OM's communication port was configured in an exclusive range per DAA region. For example, the UDP port range 4000-4100 is the Calgary region and port 4001 is specified for a particular hub site within that region and is not re-used; keeping the ports separate provides teams to discover and identify problems quickly.

### 5.4.3. NC1500

The NC1500 is a network controller that functions the same as it does in the legacy network; it works as both a DHCP server and a router which assigns IP addresses and routes upstream messages between nodes and RADDs. To complete configuration, the NC requires an IP address for each vARPD; therefore a virtual IP address is created and reserved in the Internet Protocol Access Management (IPAM) system and inserted into the NC configuration. The vARPD IP is a placeholder and not reachable on Shaw's OAMP network, so ping checks to these IP's will fail from the NC1500. Other upstream connectivity tests such as DAC refreshes are utilized instead of ping checks to verify return path functionality during troubleshooting.

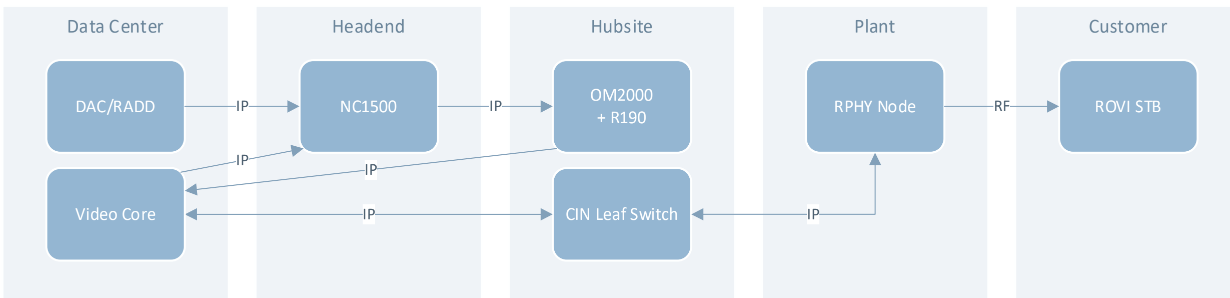


**Figure 5 DAA Interactive Network (Physical Locations)**

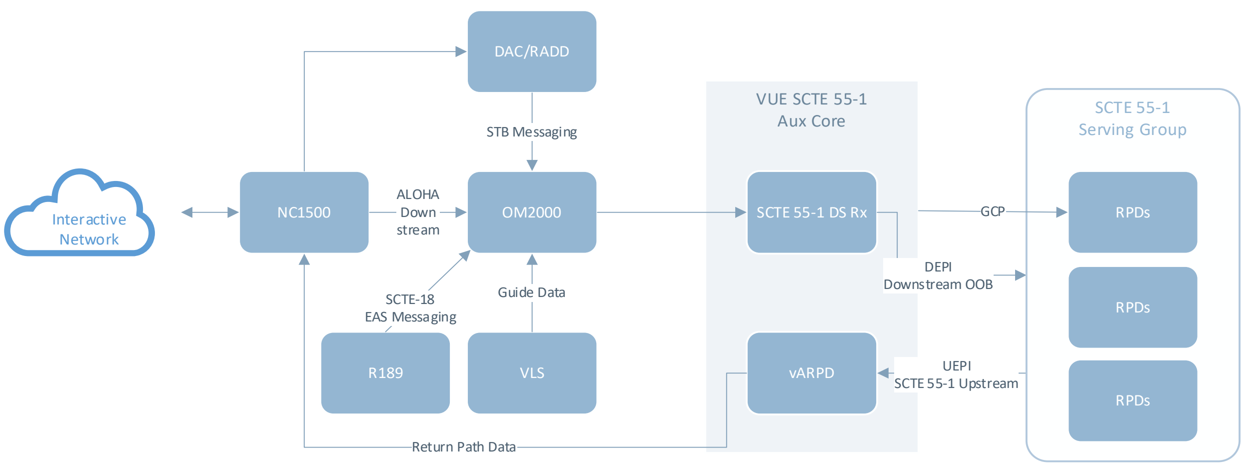


**Figure 6 DAA Work Flow**

## 5.5. DSG SCTE55-1 Interactive Signaling

For the DOCSIS set-top gateway (DSG) boxes, new DAA Converged Cable Access Platform (CCAP) cores have been deployed for their provisioning and control traffic. Targeting these tunnels is done at the Media Access Control (MAC) level and are configured using the existing Conditional Access (CA1 and CA2) and EAS SCTE-18 multicast streams the legacy DSG boxes already use.

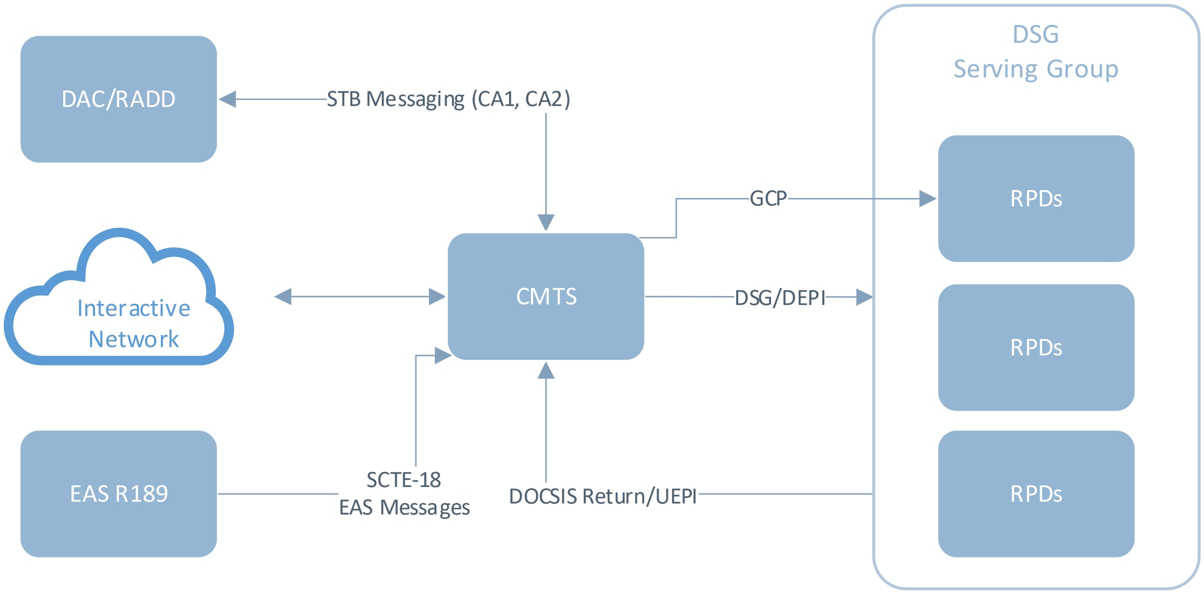


**Figure 7 DSG Interactive and Data Flow Traffic**

## 5.6. Converged Interconnect Network

CIN is a critical network that provides a redundant, highly available, scalable, and agile network solution for DAA and other Shaw services.

The CIN network leverages a spine-leaf architecture enabling a scalable DAA delivery. The spine (deployed in headend/data centres) typically has 100G interfaces and aims to maximize the fibre segments' throughput. The leaf is deployed to hub sites and aggregates traffic for multiple RPDs. All traffic carried on the CIN network is IPv6.

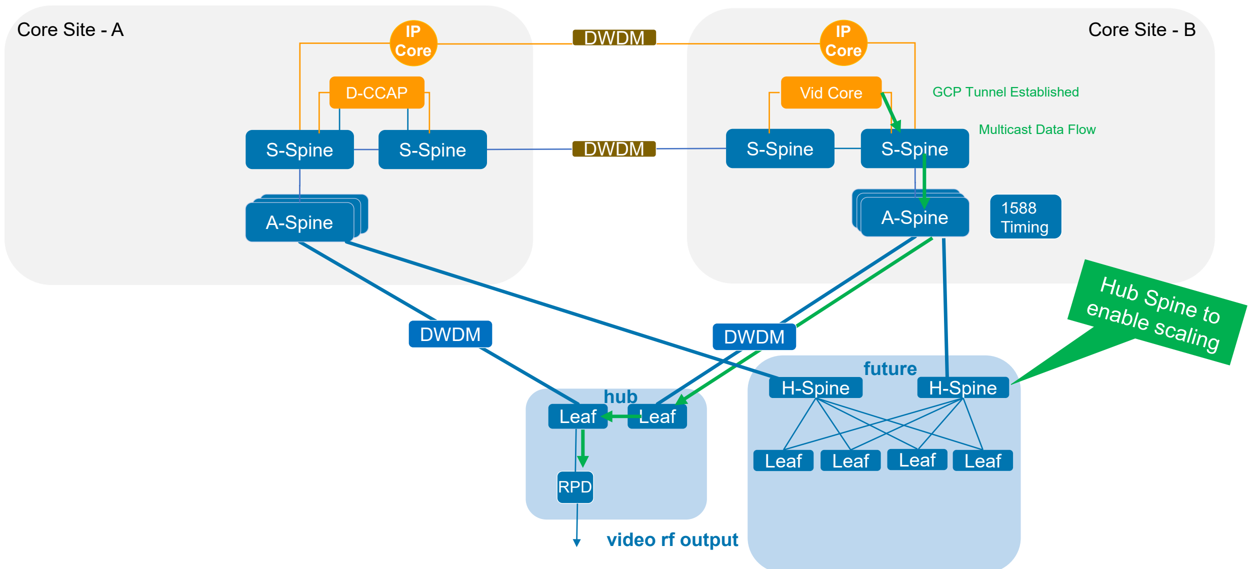


**Figure 8 CIN Network**

## 5.7. Timing

Shaw's DAA architecture utilizes the Precision Time Protocol (PTP), also referred to as IEEE-1588. The PTP protocol provides a method to synchronize all the devices connected to the CIN. PTP minimizes network jitter and synchronizes packet delivery across the CIN network. This is important because in current architectures, the DEPI pipes and RPDs have minimal jitter buffer. The principal core, the pipe servers, the RPDs, and any additional auxiliary cores are configured and synced to a PTP server. As with all other traffic on the CIN network, PTP is utilizing IPv6 addressing.

In Shaw's DAA design, the PTP master clock chassis syncs to the Global Positioning System (GPS) using a roof-mounted antenna, which synchronizes the two clocks deployed in each DAA region. These clocks are configured hot/hot to achieve necessary resiliency.

Currently, PTP clocks have a maximum allowance of 500 clients per clock or 1,000 clients per master clock chassis. In the future, Shaw will be transitioning to a boundary clock architecture where each RPD is assigned a more local boundary clock that syncs back to the PTP master chassis and presents itself as a PTP peer. This boundary clock architecture will allow Shaw to grow exponentially without utilizing dozens of PTP master clocks.

## 5.8. Remote PHY Device

The Remote PHY Devices (RPDs) replace traditional HFC nodes and assume the functionality of the edge QAMs in today's architecture. They have a network interface fed by fibre northbound, and RF interfaces southbound towards customer premises. The initiation sequence is described in the automation section 14. RPD Provisioning and Activation Process Through Automation.

 

# 6. DAA Monitoring Solution

To achieve an accurate monitoring solution, it was deemed necessary to install probes as close to the edge as possible. This accomplished two goals – end-to-end network verification and quality assurance for customer video feeds. There are probes in each region, with multiple probes in our largest region Vancouver. The probes installed in hub sites are on the same leaf routers that the RPDs are and acquire the same DEPI encapsulated multicast. The probes decapsulate multicast DEPI and then analyze the MPEG streams utilizing industry-standard methods. Shaw can currently monitor the entire broadcast channel lineup at multiple sites in real-time and trap any errors to a central dashboard. If bandwidth and licensing are concerns, narrowcast and OOB streams can be verified on an on-demand basis.

Because of DAAs regional and multicast architecture, the number of probes required can be minimized yet still maintain a high level of confidence in video quality to customers. An incumbent vendor has already integrated into the operational team's dashboard solution, so new probe outputs were added directly into an existing view. The Operations team required no further training on the monitoring solution because the probes present video quality metrics identical to the legacy MPEG probes.

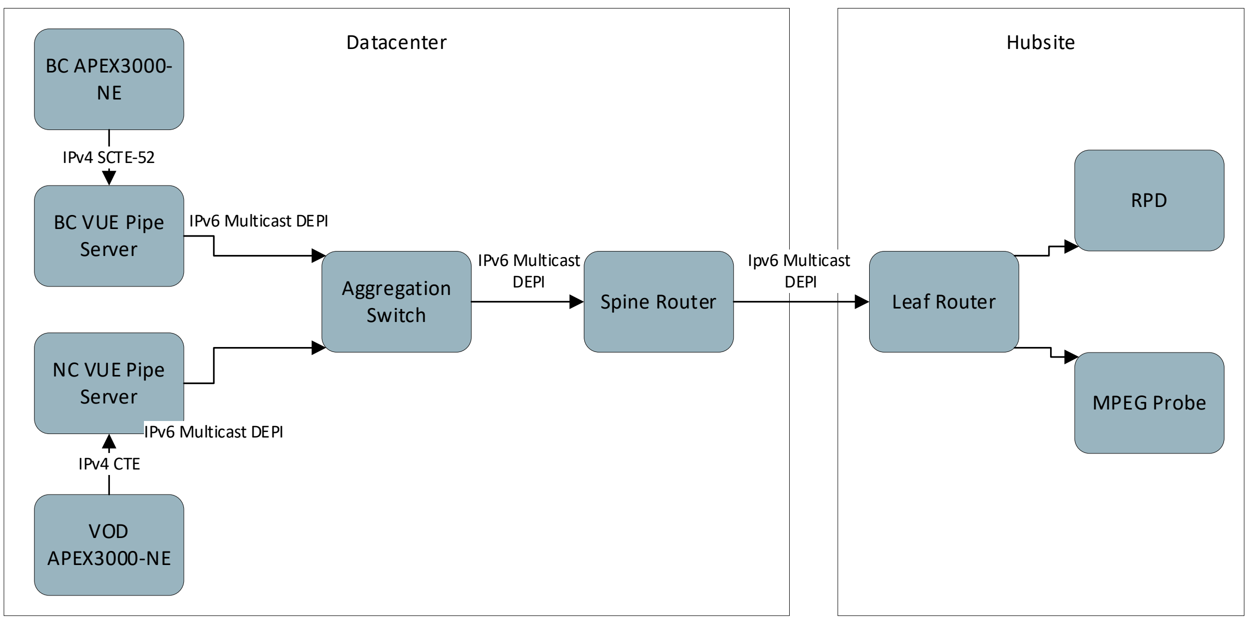


**Figure 9 DAA Monitoring Architecture**

# 7. Video and VOD Encryption

Shaw's conditional access system, the CommScope DAC, was initially configured to encrypt video services using DigiCipher II (DCII) – a CommScope proprietary encryption standard. As part of the DAA/auxiliary core project, Shaw migrated from this proprietary standard to the more open and industry-accepted alternative SCTE-52 (DVS-042). The conversion to SCTE-52 has allowed Shaw to deploy headend devices made by vendors who are not capable of licensing or using the CommScope DCII scheme.

As a sub-category of this process, all the VOD devices in production were converted to the new SCTE-52 encryption schema.

## 8. Video Broadcast Acquisition Architecture

Within Shaw, broadcast video delivery, encoding, multiplexing, and transport over the video backbone network have not changed for the DAA video solution. Instead of edge QAMs doing the encryption in the headend/data centres, video streams are encrypted by an APEX3000-NE. The NE is a product offered by CommScope based on the APEX3000 hardware with a bulk network encryptor license. Once the NE license is applied, it disables the RF ports on the chassis, receives the transport streams over IP, encrypts, and outputs over IP.

Two APEX 3000-NEs were installed at each DAA region to provide full chassis redundancy. Currently, the APEX 3000-NE cannot switch between streams. Therefore, it is impossible to utilize a single APEX 3000-NE for both primary (red) and backup (blue) broadcast video paths. Each APEX 3000-NE is connected to the primary and secondary video network. Both APEX 3000-NEs have two interfaces connected to an aggregation switch to provide link redundancy. Encrypted data streams from the APEX 3000-NEs are aggregated by the VUE pipe server, which evaluates the quality of incoming streams utilizing TR 101 290 and packet counters—in turn, selecting the best copy. The pipe server encapsulates the encrypted MPEG streams into DEPI (L2TP) tunnels.

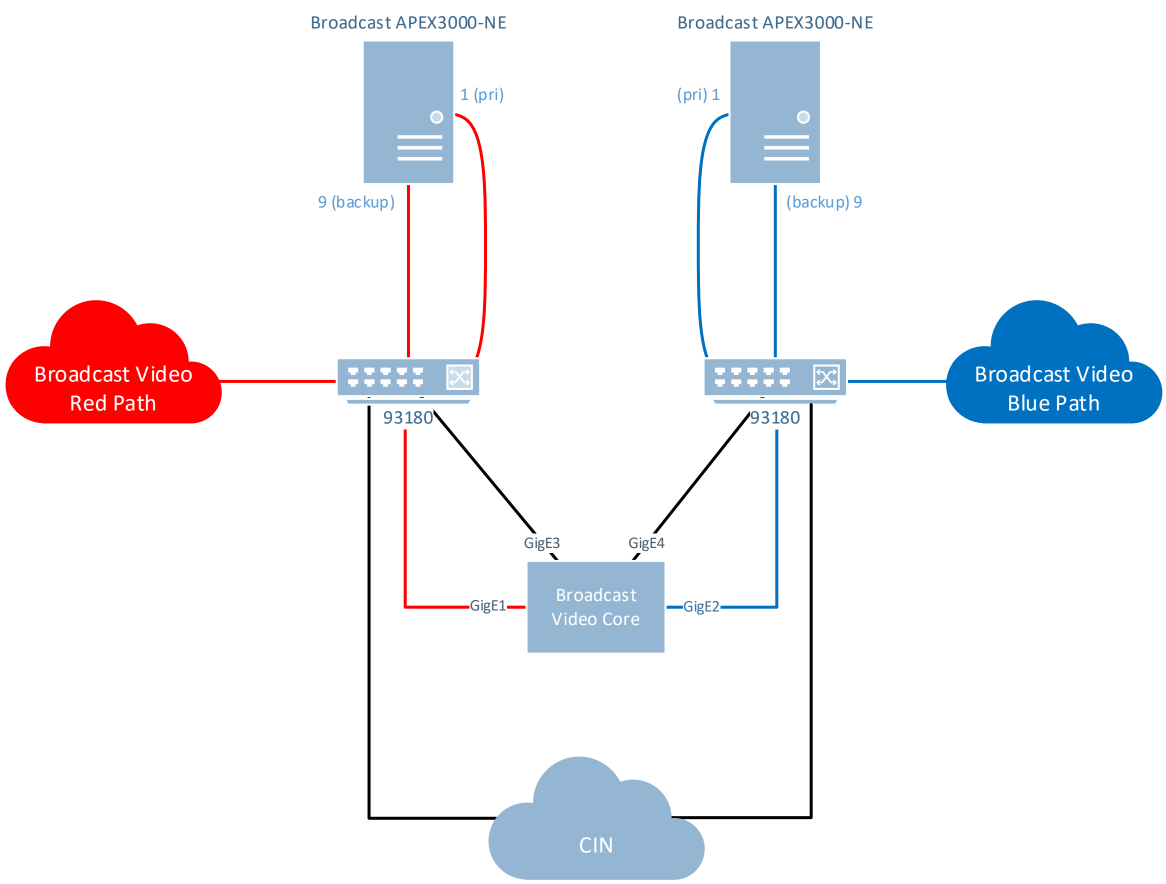


**Figure 10 Broadcast APEX 3000-NE Data Centre Network Configuration**

 

## 9. Narrowcast Architecture

Video-on-Demand (VOD) is delivered using previously existing VOD pumps that are now consolidated and centralized in the data centres in the three main regions – Calgary, Edmonton, and Vancouver. Shaw's narrowcast architecture is based on UDP mappings, whereby each RF QAM carrier is assigned a UDP port that has 21 programs (PIDs) attached to it. Program numbers 1-10 are for standard content, while programs 12 -21 are for adult content.

Currently, each QAM carrier has a unique Transport Stream Identifier (TSID) tied to it, and program 11 on each of the UDP ports is reserved for a Barker Channel embedded with TSID. This PID has been reserved in the VUE core and is automatically assigned to any port on the output UDP mappings.

A region unique TSID exists for every VOD carrier, on every VOD service group. When a VOD session begins, VOD back office (referred to as Edge Resource Manager (ERM)) maps the customer to a VOD encryptor, and calculates the next available UDP port. ERM then initiates a stream from the video server (VOD pump) to the VOD encryptor and VUE. The VOD back-office forces the STB to tune to the carrier and MPEG number mapped to the UDP port. If there is a TSID mismatch, ERM will attempt to tear down this first session and regenerate it to the correct location. This is referred to as TSID auto-correction. TSID auto-correction occurs if there is a billing system or VOD back-office configuration issue, or a physical wiring issue in the hub site. ERM is responsible for monitoring the bandwidth of the devices to ensure a VOD session can be established.

VOD encryption is also handled by APEX 3000-NE devices configured for full encryption. Each STB is authorized and entitled for VOD through our billing systems, the APEX3000-NE contains this package. Encryption is SCTE-52 and key distribution happens via the controller (DAC). The redundancy features for the APEX 3000 NE do not change from its legacy predecessor APEX 3000. No chassis redundancy has been implemented for VOD in DAA, as this mirrors currently existing VOD infrastructure on the legacy network.

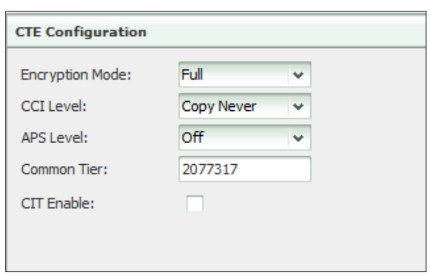


**Figure 11 NE GUI Showing VOD Encryption**

In the DAA architecture, each blade on the APEX 3000-NE is a one-to-one mapping with a VOD COR. Service groups in DAA are software-defined, and VOD sessions are streamed out over multicast. All RPD's under one single service group will subscribe to the same multicast.

A total of four interfaces are required, two each on the NE blade and VOD COR, respectively. To standardize and simplify our operations, all the four IPs selected are in the same IPv4 subnet. The output of the VOD COR propagates through the CIN network:

- Traffic interface on NE
- Virtual interface on NE
- Traffic interface on VOD COR
- Virtual interface on VOD COR

 

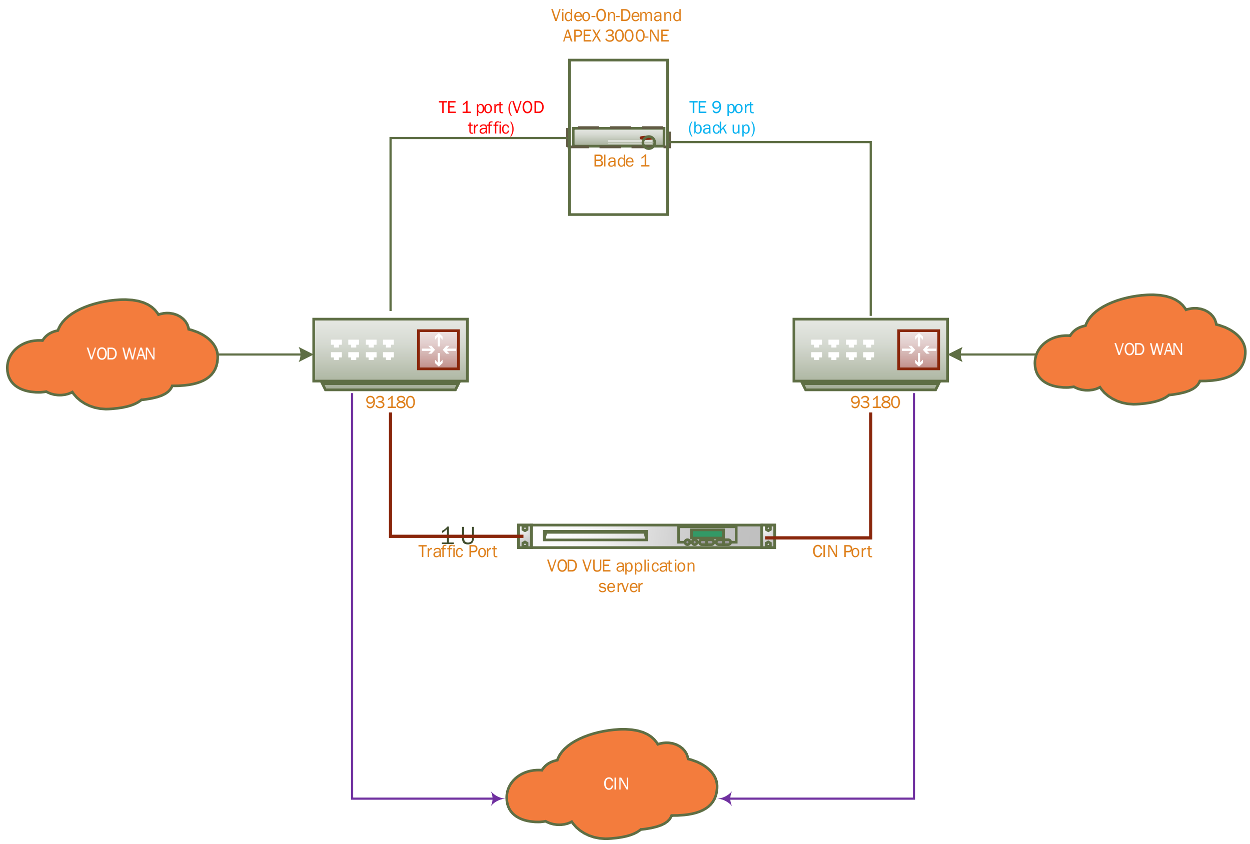


**Figure 12 Narrowcast Architecture of the APEX 3000-NE Blade and VUE Configuration**

 

## 10. IPv6 Addressing Standards

During the architecture process, it was decided to utilize strictly IPv6 for all addressing. This provided us with the room to logically and sustainably address RPD, cores, and pipe servers in a manner that would enable us to keep them logically separate from a regional and city perspective while also allowing for ample room for addressing. We feel that because of the number of addresses at our disposal, having separate hextets for cities and regions ease troubleshooting in the case of problems and allow the ability to be very specific when configuring network routing, assisting with keeping our network secure.

When choosing multicast addresses, a similar schema was used. Easily identifiable addresses outline OOB, local broadcast, system broadcast, and VOD just by looking at the address, while still giving us ample room to expand to as many different types of multicast as required—with no addressing space constraints.

## 11. Multicast, DEPI, and SessionID Relationships with Physical Frequencies on the RPD

RPDs assign physical frequencies directly based on the L2TP *sessionID* that DEPI tunnels deliver. In a broadcast video scenario, MPTS sessions are muxed into a DEPI (L2TP) tunnel and given the *sessionID* that points to the physical frequency that the MPTS is reconstituted onto. In our solution, *sessionID* is calculated as shown below:

```
session ID = baseMulticastSessionId + freqInMHz / 2
```

In our infrastructure, the *baseMulticastSessionId* is 0x80002000 in hex, or the integer 2147491840.

The following example calculation is utilized to determine the *sessionID* for 195 MHz in the main ACCP configuration file:

```
"channelType": "DsScQam",
"frequency": 195000000,
"multicastIp": "ff35:c531::daa:8123:1",
"tsId": 0,
"modulationType": "Qam256",
"interleaverDepth": "I128-J4",
"powerAdjust": 0
```

**freqInMHz** – 195

**sessionID** – 195 / 2 = 97

**chanObj.sessionID** – 2147491840 + 97 = 2147491937 or 0x80002061 in hex

The RPD receives the L2TP *sessionID* and reconstitutes it to 195 MHz on the plant.

 

Shaw has approximately 48 broadcast MPTS, and current RPD specifications only allow for 16 multicast streams. To streamline and facilitate a common architecture in all the DAA fed regions, Shaw has opted to create two unique broadcast service groups, each assigned a new multicast IPv6 address. These service groups are referred to as system and local. The system service group carries the channels common across the region, while the local service group carries channels specific to the city fed by the remote PHY nodes.

Future utilization of simultaneous substitution (simsub) content may be required (currently part of the local lineup), into a different multicast group. This schema remains within the RPDs 16 multicast session limit and allows for growth in the future.

## 12. Regional Video Realities and Solutions

Unfortunately, the method in which Shaw rolled out plant infrastructure via acquisitions and net-new builds did not adhere to any physical frequency standard regarding the transport stream to the Electronic Industries Alliance (EIA) channel. Physical frequencies are created on the RPD; therefore, the two solutions that came forth were to either significantly increase the broadcast auxiliary core hardware and network capacity or align all the hub sites to a single physical frequency/MPTS map.

Currently, Shaw is working on aligning system transport streams (TS)/physical channels across all regional markets. Once standardization is achieved, Shaw can carry non-simsub channels using a single multicast stream, utilizing one or more secondary MPTS DEPI feeds for any simsub and local content on a per-market basis. This process will save substantial traffic across the CIN network, and a non-trivial amount of server chassis as the bulk of our video traffic is not simsub.

Due to the distance and jitter limitations of DEPI, smaller micro-cores were designed to feed smaller regional sites that are otherwise outside of the regional pipe server footprint. As long as a hub site has routes to the RED/BLUE video WAN, a single APEX-3000 NE with four raptor blades can be utilized, and three initial VUE pipe servers to feed DEPI narrowcast and broadcast to customers supplied in that remote market. This maintains GCP signalling from the regional data centres as GCP is not latency and jitter sensitive, thereby alleviating any distance concerns for DEPI video. This also reduces traffic heading across the backbone between markets and services surrounding areas with DAA without building an entirely new full-scale DAA core.

## 13. Shaw Automation Strategy

Manually provisioning video services is time-consuming, error-prone, and a repetitive process. Shaw designed the DAA provisioning system with many small, discrete systems to handle specific tasks, rather than large systems that perform many tasks. This is a significant change from many of our past designs. Each domain (e.g. video, CCAP, CIN, DHCP) manages their provisioning components, and these systems primarily use HTTP/Representational State Transfer (REST) and SOAP as a means of communication.

Video automation has fulfilled the long-term objectives to:

1. Deploy data-and-software-driven automation to support legacy video services via RPDs.
2. Enable Shaw to deploy RPDs rapidly to meet our capacity needs to continue to offer excellent service to our customers.
3. Provide a framework that can be used to automate other CCAP components as part of any node change.

 

Shaw uses an in-house video automation system named VADR (Video Activation of DAA RPDs). VADR is a new system that handles the video provisioning and de-provisioning requests from Service Director (SD). Additionally, VADR supports a VOD dashboard, that assigns and re-allocates VOD service groups dynamically for RPD's depending upon their usage and the total number of nodes attached to them.

Due to the nature of DAA, multiple tools had to be either created or adapted to be functional in an RPD environment. Presently, VADR interacts with the following systems in the Shaw network to provision a node successfully. Each of these systems is either net-new or are modified to function with VADR:

- Service Director (SD)
- Resource Inventory (RI)
- VOD Central Database (CDB)
- Video Topology Manager (VTM)

### 13.1. Service Director

Service Director (SD) is a system that performs provisioning and de-provisioning of the Remote PHY Devices (RPDs). It interfaces with multiple network elements and Operations Support Systems (OSS) to complete the provisioning and de-provisioning processes.

The following two items are the OSS network elements that SD interacts with to extract information and passes to the VADR for RPD provisioning.

- **Asset Inventory System (ERP)**: This system provides the SD with the MAC address, serial numbers, and other RPD device-related information. SD uses the REST interface to query the ERP database for RPD parameters.
- **Resource Inventory (RI)**: Details the CIN, RPDs provisioned, and other information such as IPv6 prefixes on the network.

SD interacts with multiple network provisioning elements, using the OSS components mentioned above to activate the RPD:

- **BEANS** – Automated provisioning system used to configure distributed CCAP (D-CCAP).
- **CIN (NETCONF)** – Protocol used to interact with leaf routers (hub sites), aggregation spines (headends), and super spines (data centres).
- **DHCP/FLEXO** – Automation tool used to assign IPv6 addresses to the RPDs. The SD interacts with FLEXO using the REST interface.
- **VADR/VUE** – Composes the RPDs for video configuration by receiving the provisioning and de-provisioning requests from the SD. Interaction with VADR is enabled utilizing REST calls.

The following diagram displays the SD DAA interactive map showing its interaction with multiple OSS and network elements:

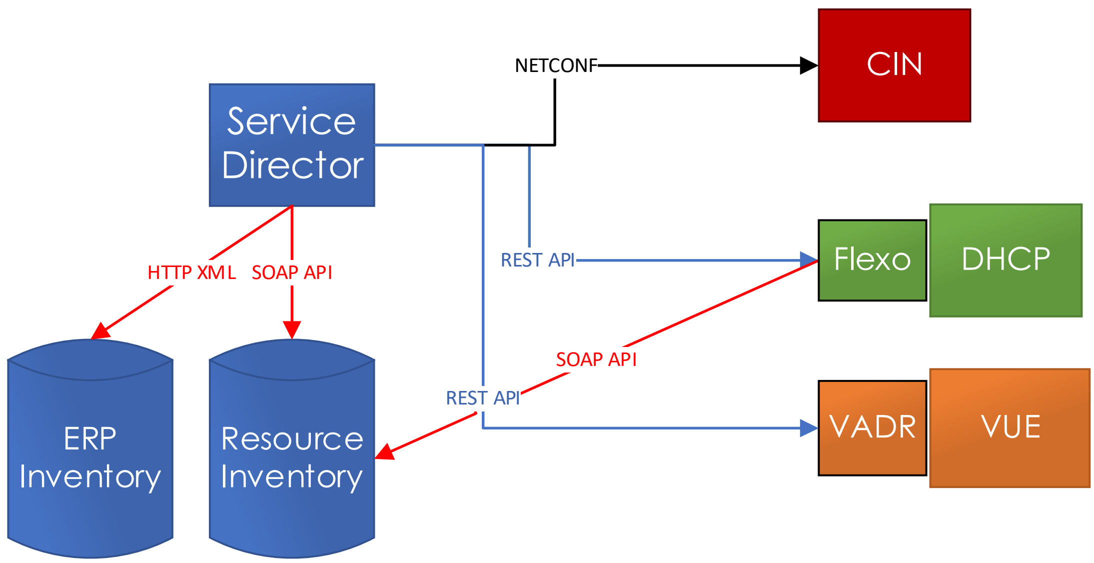


**Figure 13 SD DAA Interactive Map**

## 13.2. Resource Inventory

Resource Inventory (RI) provides information about the broadcast and OOB service groups. Currently, in RI, each region is classified as a Regional Video Centre (RVC), which encompasses the region's global configurations. Each hub site under these RVC's is assigned a local and OOB service group depending upon the simsub requirements for the hub site. In Shaw's case, due to the possibility that a hub site may have a different local serving group configuration, we have enabled RI to overwrite RVC configurations with hub site configurations if a hub site configuration exists. Otherwise, the RVC configuration is passed to the RPD.

As Shaw improves the automation processes further, RI will be configured to be the source of truth for VOD COR and APEX 3000-NE configurations for VADR.

## 13.3. VOD Central Database

The VOD Central Database (CDB) is one of the critical elements of the ERM infrastructure that assigns the RPDs to a specific VOD service group and its associated multicast based on the information received from VADR. It is the central location that keeps track of and assigns return paths to each new RPD that is activated in the field.

The VOD CDB is also responsible for creating and adding VOD APEX 3000-NE blades into the VOD database tables and its correlating VOD COR.

 

### 13.4. Automating the Video Topology Manager

VTM is a control plane and RPD configuration system that presents a standardized API into the VUE infrastructure. Shaw uses VTM to create and modify broadcast and narrowcast serving groups, administer RPDs, and monitor and modify regional VUE solutions.

VADR communicates with VTM to configure RPDs for video. VADR collects and interprets data from SD, RI, and VOD CDB then proceeds to set configurations utilizing the VTM API. During configuration, VADR adds the RPD with the proper MAC address and relevant serving group information into VTM. VTM then takes this information and configures the VUE ACCP server itself. This allows for a highly flexible and scalable single configuration endpoint allowing for ease of troubleshooting and administration over multiple regions and a large amount of RPDs.

## 14. RPD Provisioning and Activation Process Through Automation

Due to Shaw's efforts with automation and RI, once the hub site parameters are configured into the RI and CDB systems, the bulk of RPD configuration happens seamlessly at activation.

The following steps describe an RPD node turnup:

1. A technician arrives at a location and begins wiring the node physically.
2. The technician provides the MAC address of the RPD to the implementation team.
3. The Implementation team kicks off the SD provisioning process, which bridges each of the following services, provisioning and activating the RPD on the fly:
    - BEANS configures the principal core
    - SDs NETCONF service provisions the CIN devices
    - VADR sets the auxiliary video core via VTM
4. At this point, the RPD is considered provisioned. On completion of the physical wiring, the technician can power up the RPD.

When an RPD is first connected and powered up, it receives IP addressing information and the address of the principal core through Dynamic Host Control Protocol (DHCP). Once the communication with D-CCAP is initiated, it will receive the following configuration data from the device:

- A command to enable its PTP service.
- Network information on how to locate a PTP server.
- DEPI/UEPI and RF information (frequency, modulation, etc.), for the enablement of DOCSIS carriers.
- The address(es) of auxiliary core(s).

One of the auxiliary cores is the VUE auxiliary core. The RPD will send a startup-notification containing the RPDs MAC address to the auxiliary core, which will send the video DEPI and UEPI configuration for the RPD based upon the configuration the ACCP controller received from VTM. After the initial setup is complete, then the RPD will attempt sync with the PTP server.

The ACCP controller (VUE controller) requests a PTP sync status from the RPD after configuring all the channels and pseudowires. Once the RPD notifies VUE that PTP sync status is true, VUE then activates the previously configured channels by setting RF Mute to *Disable*, and the Admin state to *Up*. After activating all channels, VUE sets the RPD state to *Operational*. If the RPD has not yet achieved PTP sync, the channels being configured are set up in an RF Mute state, and the Admin state is set to *Down*.

 

If the VTM later updates channels on the RPD, the RPD has achieved PTP sync, and has been made operational, VUE will set up these channels with all of the configuration parameters; RF Mute state to *False*, and Admin state to *Up*, all at the same time.

Once the PTP sync message is successfully forwarded to the auxiliary core, all configured video RF channels are unmuted.

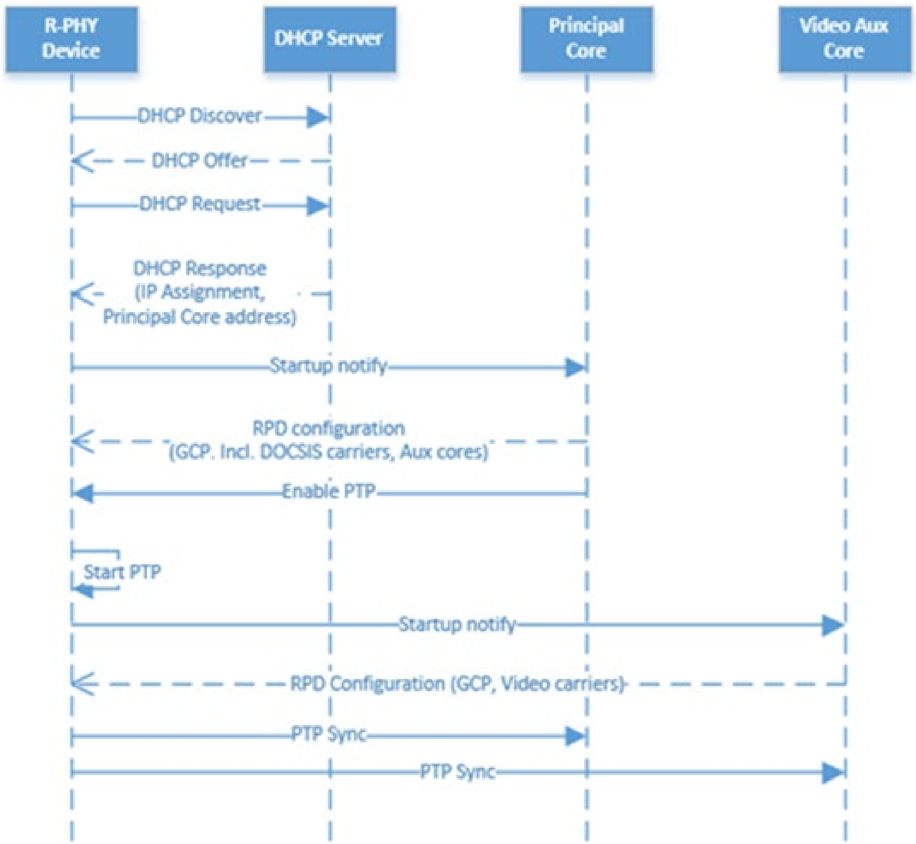


**Figure 14 Activation Sequence Diagram During the RPD Activation Process**

 

## 15. Lessons Learned

Architecting and integrating DAA with legacy video is a challenging task, and as an organization, many lessons were learned in architecting this solution.

During the architecture planning process, it was decided to utilize IPv6 for all RPD addressing. This enabled efficient use of the address space for RPD's, CIN network, and auxiliary core on a per region and hub site level while also allowing ample room for scaling out our DAA infrastructure. The sheer amount of IPv6 addresses has empowered us to have separate hextets for our regional video cores, which ease and assist with troubleshooting our network during outages. With the flexibility that IPv6 multicast addressing offers us, we decided to choose an addressing scheme that quickly identifies our out of band (SCTE-55), local, system and VOD service groups just by looking at the addresses.

We underestimated the bandwidth and multicast requirements required to service a footprint that had hundreds of unique TS/EIA frequency variants between each of our markets. Even though our channel maps reflected the same virtual channels between respective markets, this did not necessarily mean that the virtual channels resided on the same physical frequency. To be as economical as possible, we decided to align all of our TS/EIA schemas region-wide, in each region. This cut down on the duplicate multicast and the number of servers required to serve the duplicate multicast to separate markets, where the only difference is what EIA frequency the service resides on.

Early in our design phase, we came across specific markets built in a way that made sense at the time; however, as business needs or the market changed, the design and buildout of the market did not keep up. In these markets, specifically, there were conditional access requirements that differed in conditions to adjacent cities to roll out DAA video to the market. For these use cases, we had to either move the entire city between DACs or add a net-new set of DAA video pipe servers just for that city to adhere to our architectural requirements for DAA. As we advance, we take a much more holistic approach to conditional access architectures and how they interplay both now and in the future. Early in the project, the idea was that DAA should use a net-new DAC. In hindsight, it would have cost a little more, but as DAA technology improves, it may have proven to be more elegant.

Due to our sizeable footprint and the long distances between regional pipe servers and some of our more distant markets, we were outside the nominal range for DEPI. What we came up with, instead of having to create and license net new regions for smaller centers, was to create a "mini region," which would include pipe servers, but rely on the larger regional data centers for GCP command and control, as well as encryption. To facilitate this, we choose a central location for these pipe servers with a good network and a properly hardened datacentre facility and utilize a single "mini region" to feed as many small markets as possible. An excellent example of this would be our interior BC region, which has a large number of low population towns nearby, and good fibre availability, yet too much distance from a DEPI perspective to backhaul video from a fully-fledged datacentre.

From an automation perspective, early on in the project, we realized that our resource inventory system was inadequate to support our DAA initiative fully. Next was months of design and decision meetings, culminating in a multi-phased rollout to achieve full compatibility with automation tool requirements. Two of the major obstacles we had were the initial implementation of IPv6 addressing into our resource inventory system and specifying serving group, out-of-band, and VOD configurations to our automation system in a granular, sustainable, and information-complete way. As we advance with RI, we now put as much information as we have available in as specific and granular a format as possible, even if we are not currently using all the data. There may come a time when we may need to automate against it.

 

These are some of the more valuable lessons that we have learned while designing and implementing this solution, and we continue to learn and grow from our work in this technology daily. With features being added all the time, our architecture continues to shift and mature. The final lesson we have learned from this endeavour is to design the architecture to be modular. This avoids rip-and-replace, and duplicate efforts wherever possible.

## 16. Conclusion

As Shaw continues to embrace next-generation technologies, we feel our architecture is adaptable and scalable enough to maintain pace with any future considerations. Because of our automation efforts, we can run two legacy video infrastructures side-by-side with minimal impact, and no new resource requirements for operational teams. As we move towards an IP-only video solution, we can systematically and remotely free up plant spectrum for CMTS, until finally completing our transition to all-IP. Our auxiliary core legacy video architecture is an essential tool in enabling Shaw's network evolution going forward.

# Terms and Abbreviations

| Term | Definition |
|---|---|
| ACCP | Auxiliary Core Control Plane |
| Annex B | A cable standard that modulates an MPEG-TS input into a QAM-256 output |
| API | Application Program Interface |
| APS | Analog Protection System |
| ARPD | Advanced Return Path Demodulator |
| CA1 | Conditional Access 1 |
| CA2 | Conditional Access 2 |
| CCAP | Converged Cable Access Platform |
| CCI | Copy Control Information |
| CDB | Central Database |
| CIN | Converged Interconnect Network |
| CIT | Constrained Image Trigger |
| CMTS | Cable Modem Termination System |
| COTS | Commercial Off-The-Shelf |
| CPE | Customer Premise Equipment |
| CTE | Common Tier Encryption |
| D-CCAP | Distributed Converged Cable Access Platform |
| DAA | Digital Access Architecture. Evolution and decentralization of how data centres, headends, hub sites, and nodes feed video and data services to customers' homes. |
| DAC | Digital Addressable Controller |
| DCII | DigiCipher II |
| DCM | Digital Content Manager |
| DCT | Digital Cable Terminal |
| DEPI | Downstream External PHY Interface |
| DHCP | Dynamic Host Control Protocol |
| Docker | A tool designed to make it easier to create, deploy, and run applications by using containers. Containers allow a developer to package up an application with all the parts it needs, such as libraries and other dependencies, and deploy it as one package. |
| DOCSIS | Data-Over-Cable Service Interface Specifications. An international telecommunications standard that allows for the addition of high-bandwidth data transfer to an existing coaxial cable TV system. |
| DSG | DOCSIS Set-top Gateway |
| DWDM | Dense Wavelength Division Multiplexing |
| EAS | Emergency Alert System |
| ECM | Embedded Cable Modem |
| EIA | Electronic Industries Alliance |
| ERM | Edge Resource Manager |
| ERP | Enterprise Resource Planning |
| FLEXO | Control plane for OSS Dynamic Host Control Protocol systems |
| GCP | Generic Configuration Protocol |
| GigE | Gigabit Ethernet (GbE or 1 GigE) is the term applied to transmitting Ethernet frames at a rate of a gigabit per second (1 billion bits per second). The most popular variant 1000BASE-T is defined by the IEEE 802.3ab standard. |
| GPS | Global Positioning System |
| HFC | Hyper Fibre Coaxial |

| Term | Definition |
|---|---|
| HTTP | Hypertext Transfer Protocol |
| IP | Internet Protocol |
| IPAM | Internet Protocol Access Management |
| IPTV | Internet Protocol Television |
| IPv4 | Internet Protocol version 4 |
| IPv6 | Internet Protocol version 6 |
| L2TP | Layer Two Tunneling Protocol |
| MAC | Media Access Control |
| Mbps | Megabits Per Second |
| MPEG | Moving Picture Experts Group |
| MPTS | Multiple Program Transport Stream |
| NC | Network Controller |
| NE | Network Encryptor |
| NETCONF | A protocol defined by the IETF to install, manipulate, and delete the configuration of network devices. |
| OAMP | Operate Administer Maintain Provision |
| OM | Out-of-Band Modulator |
| OOB | Out-of-Band |
| OS | Operating System |
| OSS | Operations Support Systems |
| PHY | Physical Layer |
| PID | Program Identification |
| PTP | Precision Time Protocol. Also referred to as IEEE-1588. |
| QAM | Quadrature Amplitude Modulation |
| R-UEPI | Remote Upstream External PHY Interface |
| RADD | Remote Addressable Danis/DLS Downloader |
| RDS | Rights Data Server |
| REST | Representational State Transfer (RESTful API). API built using the rules of representational state transfer software architecture.<br>Web protocol that utilize much less overhead than SOAP. Allows for communication beyond XML. |
| RF | Radio Frequency |
| RI | Resource Inventory |
| Rovi | Rovi is a global leader in digital entertainment technology for some of the largest CE manufacturers, service providers and online, mobile, and application developers. |
| RPD | Remote PHY (Physical) Device. A device that receives Internet, telephone and video over IP and converts it to RF. |
| RTSP | Real Time Streaming Protocol |
| RVC | Regional Video Centre |
| SAF | Signal Acquisition Facilities |
| SCTE-18 | Emergency Alert Messaging for Cable (ANSI J-STD-42-C) standard by the Society of Cable Telecommunication Engineers |
| SCTE-52 | Data Encryption Standard - Cipher Block Chaining Packet Encryption specification standard by the Society of Cable Telecommunication Engineers |
| SCTE-55 | ALOHA protocol-based standard used by General Instrument and Motorola equipment |
| SD | Service Director |

 

| Term | Definition |
|---|---|
| SimSub | Simultaneous Substitution. Signal substitution occurs when a signal is temporarily replaced by another one airing the same program at the same time. Usually, an American signal (e.g. affiliates of ABC, NBC, CBS, Fox) is replaced with a Canadian signal. Sometimes, a Canadian signal from outside the area is replaced with a local signal. |
| SOAP | Simple Object Access Protocol |
| SOD | Shaw on Demand |
| SPTS | Single Program Transport Streams |
| SSH | Secure Shell |
| STB | Set-Top Box |
| TR 101-290 | Defines measurement guidelines for monitoring MPEG Transports Streams (MPEG TS) |
| TS | Transport Stream |
| TSID | Transport Stream Identifier |
| UDP | User Datagram Protocol |
| UEPI | Upstream External PHY Interface |
| VADR | Video Activation of DAA RPDs. Video Technology's service that receives provisioning/deprovisioning requests from Service Director and configures VTM using VTM APIs with data received from the various data sources. VADR also preconfigures VTM with the necessary Service Group information for VOD. |
| vARPD | Virtual Advanced Return Path Demodulator |
| VOD | Video on Demand |
| VOD COR | Video on Demand Core (pipe server) |
| VPS | Video Platform Server |
| VTM | Video Topology Manager |
| VUE | Video Unified Edge. Provides the network and RF information between the video sources and the RPDs. |
| WAN | Wide Area Network |
| XML | Extensible Markup Language |

 

# References


*Remote Downstream External PHY Interface Specification*; CableLabs

*Remote PHY Specification*; CableLabs

*CCAP Architectures Technical Report*; CableLabs

*TNO-F18-D-100 Distributed Access Architecture*; Corwin Martens & Wesley Weiss - Shaw Communications

*DAA-CIN Converge Interconnect Network Solutions Design*; Shaw Communications

*Service Director Documentation*; Shaw Communications

*Architecture Definition Document – LITE*; Darren Gamble