Creating The Intelligent Edge

Increasing DAA Velocity Using Service Orchestration

A Technical Paper prepared for SCTE•ISBE by

Chris Busch
Fellow – CTO Network Solutions
CommScope
2400 Ogden Avenue, Lisle Illinois
630-281-3150
chris.busch@commscope.com

Dave Baran
Fellow – CTO Network Solutions
CommScope
101 Tournament Dr. Horsham Pennsylvania
215-323-1973
dave.baran@commscope.com

Jeff Finkelstein
Executive Director Network Architecture
COX Communications
6205 Peachtree Dunwoody Rd Atlanta
jeff.finkelstein@cox.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

Distributed Access Architectures (DAA) open many new business opportunities at the edge of the network for the cable operator. Increasing service options in bandwidth terms, overall service quality, reduced cost of optics, and lower latency, are key values of the architecture. The infrastructure introduces all digital wavelength division multiplexed 10-100 Gigabit Ethernet from Hub to pole with highly precise timing support. This further creates a simpler path to 5G roll out for the MSO, in addition to residential and business services over any form of wireline access.

While Remote PHY is the most common distributed access technology today, there are multiple forms of DAA including Remote PHY, Remote MACPHY, and OLT-PON, in addition to emerging 5G backhaul networks.

Serving each of these platforms is an Ethernet system known as the Converged Interconnect Network (CIN). This digital infrastructure enables multiple services over each of these remote functions. These same CIN systems enable 5G backhaul opportunities that require access network sharing, sometimes termed network slicing functionality. Lastly, DAA brings with it support for the era of cable Network Function Virtualization (NFV). In fact, DAA is not any one element, it is all these platforms working together.

However, one fundamental role that the CCAP platform had must be introduced into the Distributed Access Architecture. The automatic discovery, provisioning, and telemetry for Physical layer and MAC layer functions in the topology must exist for DAA as they do today within any Integrated-CCAP chassis. Without this functionality in place, the increased management complexity due to the larger number of DAA devices can negatively impact operations and expected cost savings.

The authors seek to address the opportunity to add intelligence at the DAA edge using a select set of back office systems creating an SDN-based ‘Intelligent Edge’ for the DAA. Through this Intelligent Edge, operators may now see these distributed and often decoupled systems realized not at a network or element level, instead being brought together to deliver services in an end-to-end fashion, with dynamic provisioning effectively returning the same Zero-Touch onboarding and deployment a CCAP-based chassis provides for HFC-based PHY and MAC functions.

In this paper, we will review options for top down orchestration based on industry practice. Additionally, an approach to bottom up network discovery and considerations for automation and device provisioning from the edge of the network into the Hub and Headend are detailed to provide guidance on Zero-Touch provisioning for DAA.

Content

1. Goals of The Intelligent Edge

For the cable operator, introducing Distributed Access Architecture involves many stakeholders across the business. Engineering needs to understand the network bandwidth and interconnect needs for each DAA remote device type. Operations needs to understand how performance and troubleshooting data will be made available to preserve service assurance systems. Warehouse and Inside Plant Engineering teams have needs where logistics of inventory and assignment to outside and inside physical plants are concerned. Field technicians are concerned with time spent during maintenance windows and ability to close work orders for each installation. Service assurance as well has additional resources available for monitoring and managing the performance of the holistic HFC network.

Goals for most operators deploying DAA include:

- Simple inventory association
- Automated provisioning for any remote device
- Minimize remote device installation time
- Service assurance of remote device operations

From a high level, many if not all of these goals and stakeholders can be aided through an effective onboarding process for DAA remote systems.

Onboarding is the function of network element discovery and provisioning based on role and required service attributes.

To enable onboarding, multiple systems in the current back office, specifically the Operational Support System (OSS) evolve to support effective DAA deployment.

This evolution influences and creates the Intelligent Edge through:

- Inventory Process
- Device and Topology Discovery
- Service Activation
- Device Authentication
- Dynamic Device Addressing
- Device Provisioning
- Network Provisioning
- Field Installation Application

Perhaps more accurately, the Intelligent Edge is initially created through integrations among each of these systems to enable the goals operators seek in deploying DAA technologies as a solution.

2. Orchestrating the Intelligent Edge

The Intelligent Edge has knowledge of the resources in the DAA network and their relationships. Resources can be thought of generally as the various actors in the network involved with any DAA service and the values these actors must have configured for each DAA service and function to operate.

We can build up an effective relationship design for the Intelligent Edge by thinking of any DAA as a service that consumes available inventory in part or in whole. Services are defined at a top level for the operator and become associated to operating regions and entities across the network.

The approach presented here is a product catalog or so called 'model driven' approach to service creation whereby association to common elements minimizes manual entry and promotes high reusability of configuration data such that the least amount of configuration is required when a device is attached to the system. The idea of the catalog is to define standard products that are available for sale with standardized attributes that ensure that the provisioning/service delivery phase has all of the necessary information collected to ensure proper implementation.

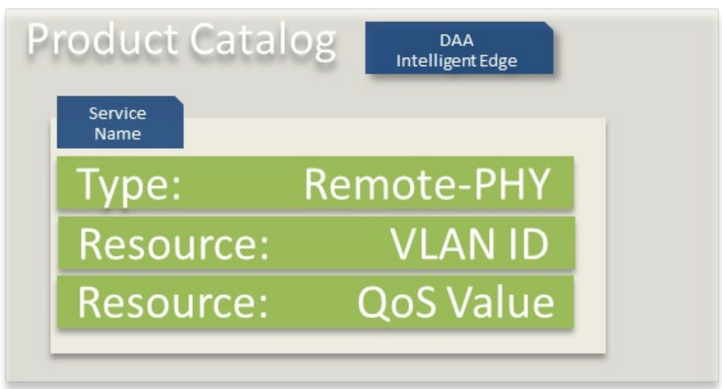

Figure 1 - Product Catalog Creation

Following that service model from a catalog approach, a Remote-PHY service type can be defined in the catalog to cascade down global configuration to members of the service. In this example, assume that the Remote PHY service includes a VLAN identifier and a QoS value for traffic separation on the network link. From this stage onward, other associations with increasing granularity may be defined and associated in a chain like manner.

For example, 'RPD_Service_A' is added to an entity such as 'North-East'. Within North-East, there may be multiple Headends. 'Headend_A' includes resources such as PNF (Physical Network Function) and VNF (Virtual Network Function) DOCSIS cores for Remote-PHY association. Headend_A also includes defined RF resources and Headend_A may have a Hub level resource below it that inherits from the entire chain described. The RF resources are a many to one asset, meaning they are defined as a named template and associated to an asset in the Headend or Hub, specifically the named service group RPD_Service_A which becomes a member of the Headend or Hub inheriting access to resources at that level.

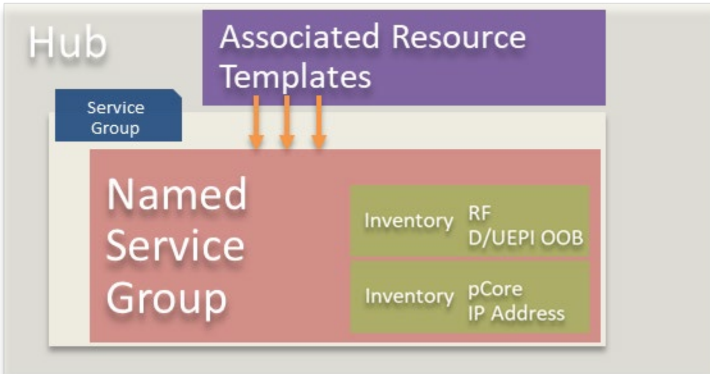

Figure 2 - A Single Service Group

When the product catalog for the RPD Service is associated to the Named Service Group, additional inheritance of functionality occurs. The product catalog is less specific in terms of Headend or Hub resources, such a RF forward and return or video out-of-band configuration values. Certainly, if the operator had the same DOCSIS RF resource plan and Video QAM RF resource plan across many Hubs then these could be configured up in the catalog. The relationship is purely one of re-use and ownership to facilitate orchestrating a service.

In the example demonstrated, all Remote-PHY service groups have a Remote-PHY VLAN ID value and QoS value the example here chooses to apply to all Hubs that have joined the service. By having this as part of the product catalog, the necessary information is obtained up front to ensure proper configuration of both the Remote-PHY device and the Secure Leaf (s-Leaf).

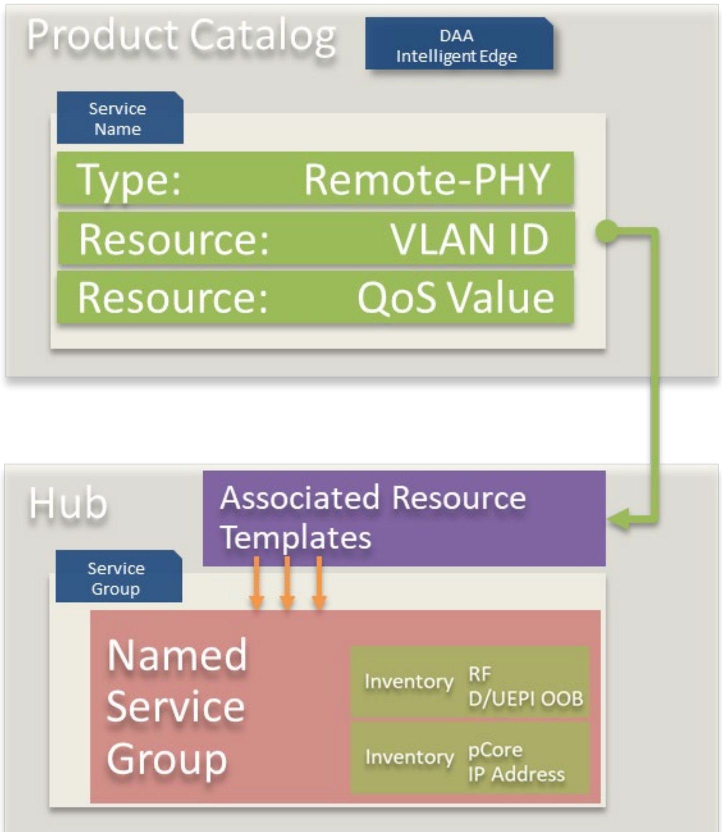

Figure 3 - Service Group to Named Service

2.1. Creating Service Groups

To create the service-based view of DAA systems, we have borrowed the concept of a 'Service Group' and use this as a simple entity to hold members of configuration objects both provisioned and learned. This 'Named Service Group', we could say 'Service Group 405', will have associated Resource Templates and be a member of various inventory objects in the network. These are the provisioned objects.

Resource Templates are values for configuration such as Video Out-of-Band, Downstream and Upstream channels, VLAN value, QoS value that service 'Service Group 405' may inherit. As inventory is associated, such as a Remote PHY DOCSIS device, the resource templates are applied by network function and executed in the service specific manner updating the device record held in our DAA Inventory function as indicated by the 'Service Group 405' record.

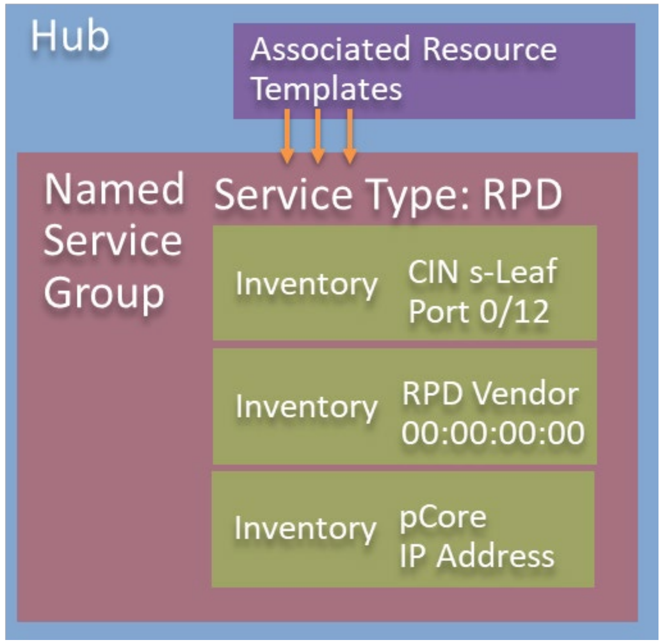

Figure 4 - Service Creation with Active Resources

Secure Leaf switch is also known as an s-Leaf exists at the edge of the network to accept a variety of distributed access device connections. A Secure Leaf is so termed given its role in supporting IEEE 802.1AE MACsec, and encryption standard for Ethernet physical layer.

When a physical RPD is attached to the network and joined to the service, its network address and s-Leaf attachment information becomes dynamically known to the Named Service Group, for example Service Group 405.

Service Group 405 may have simply held the RPD MAC address and the membership to the Hub. The Hub may have the pCore (Principal CCAP Core) for Remote PHY and therefore Service Group 405 inherits this pCore IP address.

As the system evolves, we will see how more of the service group may be constructed dynamically rather than through a buildup of static configurations pushed to any single device. Instead, we will have the ability to construct and deconstruct the needed provisioning updates at run time effectively when a device is discovered at the DAA edge using SDN control functions to execute these tasks on behalf of the service. Changing the provisioning model from "before use" to "at use" simplifies service creation and allows it to be executed without human involvement which is critical to success with a DAA architecture.

3. DAA Inventory

3.1. Static-Inventory

A Static-Inventory based system assigns the unique identification of a remote device to a service in the network. This may occur during a work order ticket for the field team as the truck is loaded for the day or may occur from other systems. When the activation process is driven from a Static-Inventory view, the field technician must use the specific remote device for that service order.

For Remote-PHY, this may mean that a pre-staging has occurred where the warehouse team have bench tested the RPD, or other DAA remote device type, prior to deployment and assigned it to its designed 'Service Group' where all that remains is for the physical installation in the field.

The steps involved in a Remote-PHY Service Group from Static-Inventory include:

- Association of RPD MAC address from Inventory to Service Group
    - Service Group inherits RF configuration and Core attributes
- Association of RPD MAC address to a Core
    - At minimum a Principal Core for DOCSIS services

Most DAA remote devices use some form of Ethernet backhaul such as 10 GB, 25 GB, or 40 GB Ethernet. This backhaul or access aggregation in the outside plant typically occurs over Wave Division Multiplexing (WDM) either Dense or Coarse (DWDM/CWDM). Others may use xPON or other technology (e.g. 5G, mmWave, etc.), but for the purpose of discussion we will focus on point-to-point Ethernet.

The two ends of this backhaul are pluggable optics, where the Hub or Headend side has an s-Leaf switch populated with Small Form Pluggable+ (SFP+) or a variant based on the PHY rate that is interconnected to a fiber mux tray at the top of the rack. At the other end in the outside plant (OSP), a Fiber Optic Splice Closure (FOSC) will present one or two fiber pair(s) to the remote device. The remote uses a pluggable optic to terminate the interconnect. The pluggable optics involved in this system may be fixed wavelength or dynamically tunable.

In the event the optics are on fixed wavelengths, the Static-Inventory system must include assignment of the SFP+ assets at both ends of the fiber path from s-Leaf to remote port(s).

The static model has a significant drawback as it produces distinct devices rather than fungible ones. At each stop along the installer's path, they have to get the right device and hope that it was provisioned correctly. In the event that a device doesn't work, the technician can't easily diagnose the problem by swapping another unit from the truck as that unit would be provisioned differently. The key to simplicity is to keep things interchangeable, which requires a dynamic model.

3.2. Dynamic-Inventory

A Dynamic-Inventory based system assigns the unique identification of a remote device to a service in the network at the moment the technician installs the device in the OSP.

This is made possible using several systems and application integrations discussed in further detail.

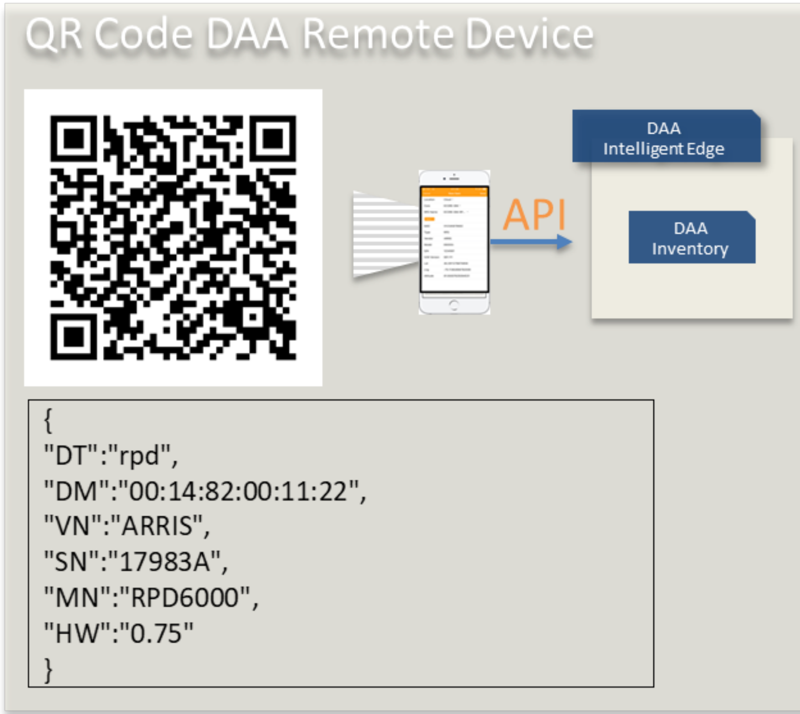

Figure 5 - Warehouse QR Scan

An effective system will integrate the warehouse view of physical inventory with the network view of actual in-service inventory. There can be subtle differences in how the operator may choose to assign inventory from warehouse to network. Specifically, whether a DAA remote device is assigned to a specific work order or destination service group, or whether the field technician will generate that assignment during installation. While the distinction is subtle, the implications are distinct to the back-office system.

3.3. Using a Mobile App

Onboarding for DAA Remote devices can be assisted from the field using an install app from a mobile device. Where Remote PHY devices are concerned, there is a QR code available that is associated to the DAA device module. The field technician at time of installation may scan this code with the mobile app to facilitate the onboarding process.

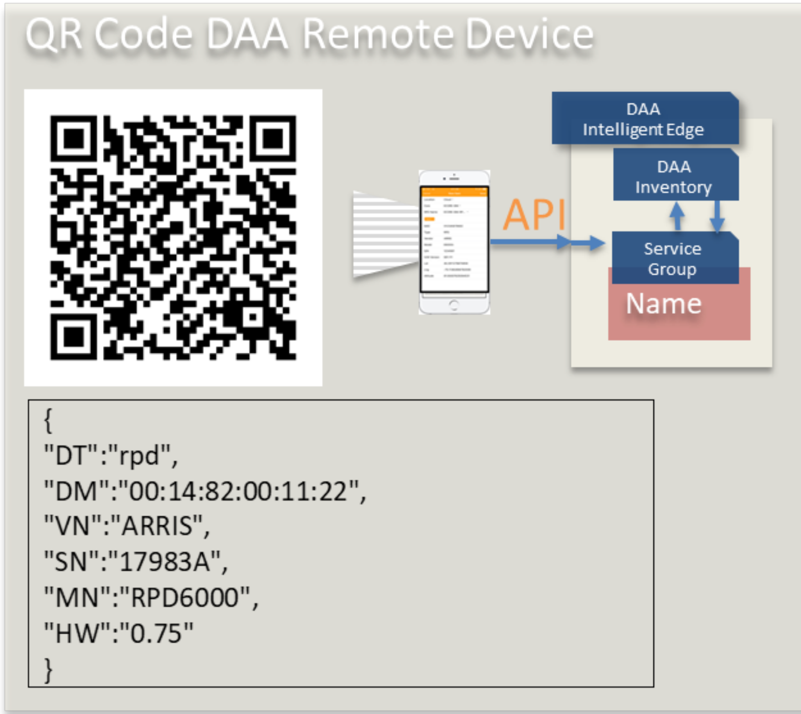

Figure 6 - Mobile App QR Scan Service Group to Inventory During Field Installation

The mobile app will also send to Inventory the geo-coordinates of the scanned device. These are sent as latitude and longitude, stored with the device record in the Inventory of the Intelligent Edge.

The field technician is prompted to select the Service Group to which this device is to be associated with. The Intelligent Edge will have returned a list to choose from or optionally the service group may be searched for within the mobile application.

Once submitted, an API is updated with the same inventory information from the warehouse, including geo-coordinates and now a named service group association to the DAA Inventory device record.

This is a key stage in the onboarding process. Association of a physical device to a 'Service' will now have impact on the edge topology discovery and the DHCP provisioning discovery stages.

4. DAA Inventory Provisioning

At the heart of the Intelligent Edge is the idea of an 'Active Inventory'. This 'Active' state is simply a closed loop of updates from SDN control functions to the Inventory database of the Intelligent Edge.

If we consider the Integrated-CCAP chassis model of line cards for PHY, MAC and routing, the fundamental of these interconnects was the chassis backplane. In a DAA system, the backplane is the CIN. As an important first step, enabling the CIN to itself be Zero-Touch onboarded is important.

The management network DHCP uses a lookup to the DAA Inventory to learn the configuration file value for the CIN element assuming the device is onboarded in a configuration file download manner. This approach is similar to 'golden config' as the returned configuration file will have globally significant configuration based on the device role either Spine or s-Leaf. In either case at a minimum two main configurations are enabled. First the direction to SDN Controller for ongoing management and Link Layer Discovery Protocol (LLDP) communications. Second access across all s-Leaf ports to the management network for DAA remote device onboarding. This includes the VLAN - VTEP association in all s-Leaf devices and the participation of EVPN forwarding from s-Leaf to Spine and Spine-to-Spine forwarding.

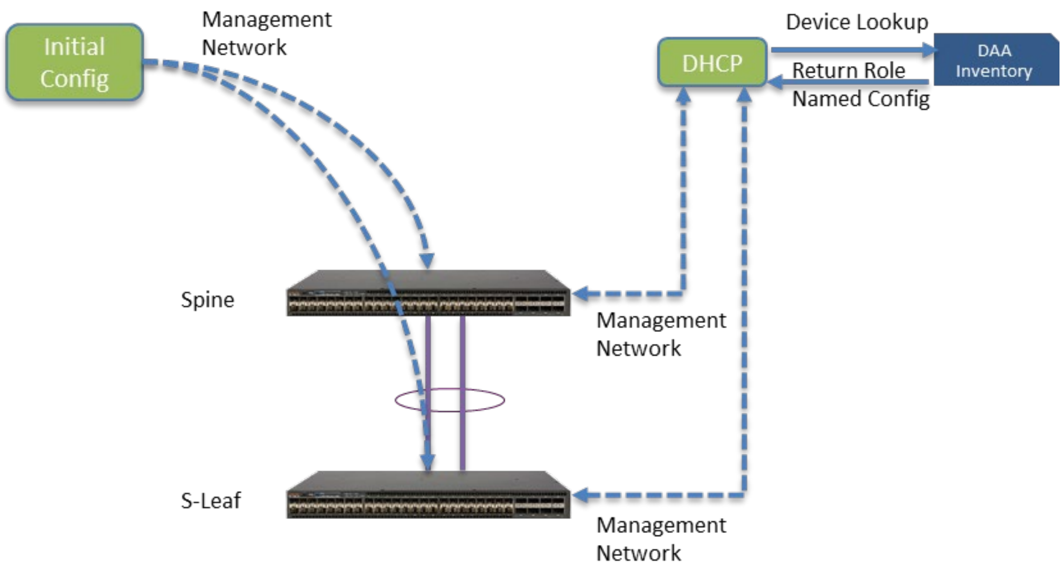

Figure 7 - Initial Switch Config

Next is the dynamic provisioning of service ports for each s-Leaf switch in the CIN. As each port should be capable of supporting any DAA service, we need to ensure discovery and provisioning at the edge are working perfectly. To achieve this, it is necessary to build a view of the complete topology.

Construction of the topology is possible through a graph of LLDP based information provided by all CIN devices. Several methods for asynchronously obtaining LLDP to a northbound system such as an SDN Controller are possible and include OpenFlow, sFlow, and NETCONF or RESTCONF streaming where each of these provide for the inherent MAC to Port information necessary to construct a dynamic view of the topology.

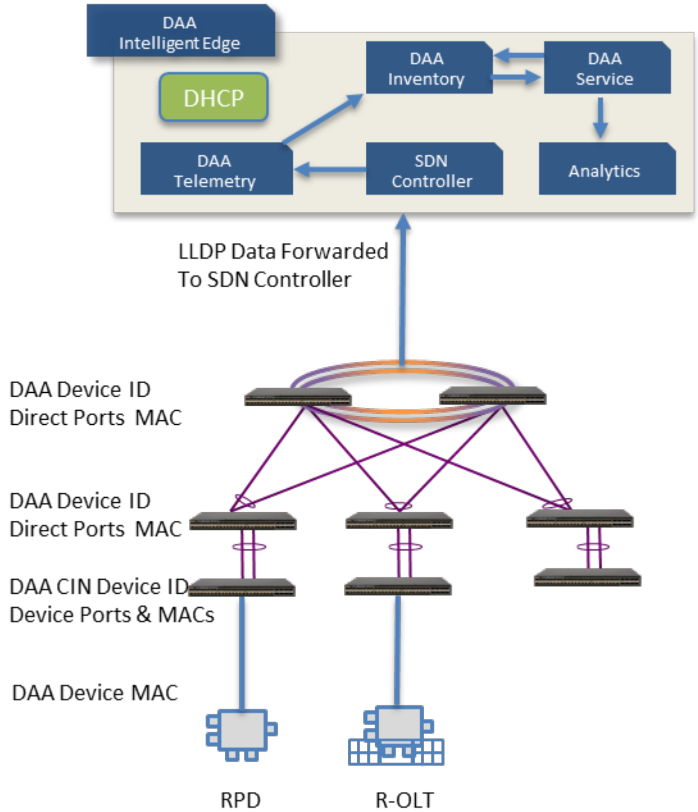

Figure 8 - OSS Consumption of LLDP

Additional resources in the example RPD Service Group 405 will need to associate with and involve the CIN. Specifically, a port on an s-Leaf switch.

The Intelligent Edge would also have knowledge then of the CIN elements, based on their assigned role as an 's-Leaf' in the network at any location such as Headend or Hub serving Service Group 405.

At this stage, the basic system logic is in place that will support either a 'Static-Inventory' or a 'Dynamic-Inventory' based approach to deploying Remote-PHY and essentially any DAA remote technology.

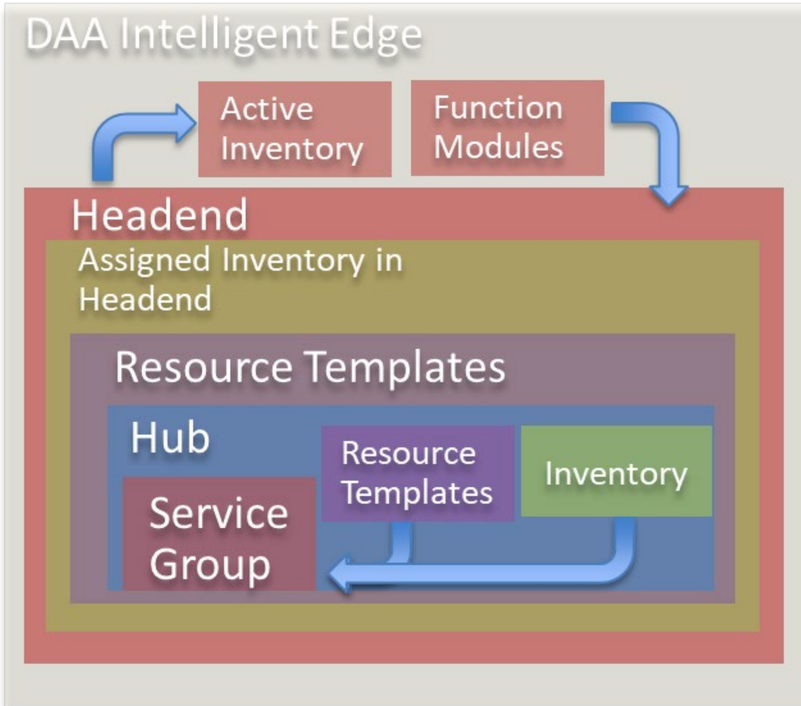

Figure 9 - Intelligent Edge with Dynamic Inventory

Using LLDP collection and SDN Controller to build a topology, the initial knowledge of edge port associations is made visible and accessible for automation.

The system design providing LLDP processing is quite robust. LLDP TLV data is exposed to the SDN Controller by any of the means described earlier. The SDN Controller transforms this LLDP packet data into a message bus topic and publishes this for consumption by the other functions of the Intelligent Edge.

In this system design, there is a DAA Inventory service always consuming topics from the message bus to remain in an 'active inventory' state for the entire network. As updates occur for device records, the Inventory system takes a decision if there is any change in the known data based on the information present in the message bus topic. If there is a change, the update is sent to a DAA Service Activation function which sends an API request to the SDN Controller. The SDN Controller then can modify the port profile configuration for the given service role communicated (such as RPD-Service or R-OLT Service or similar), which the will inherit specific configuration values during the process.

If there is nothing to update in the network, the device record is simply updated in the DAA Inventory with any additional values such as counters or other operational values. No actions to DAA Service Activation occur in this case.

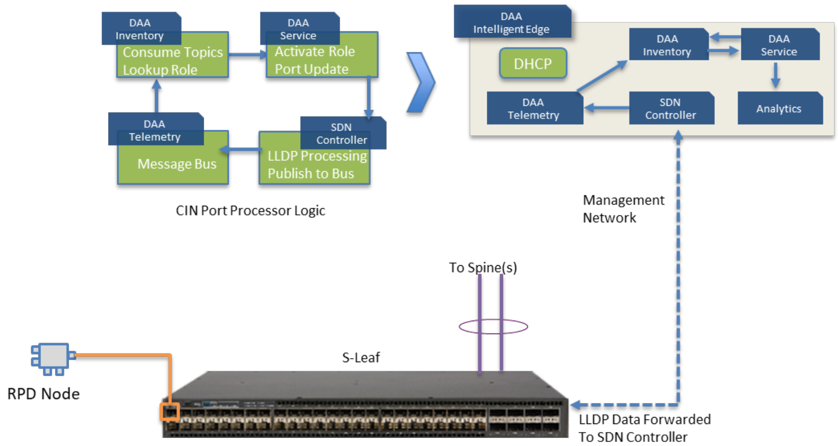

Figure 10 - CIN Port Dynamic Profile Provisioning

The result of this system design is all ports of the CIN are dynamically capable for any defined service known to the Intelligent Edge system as the CIN device record exists in Inventory and all ports have management network reachability. As any DAA device is attached, it is dynamically discovered before 802.1x or IP address stage such that Service Activation may occur just before the DAA device is provisioned.

For the field technician who may be responsible for deploying Remote PHY or Remote-OLT PON, or 5G Backhaul, the simplicity at the edge could not be better. Similar to how an Integrated CCAP chassis automatically detects and onboards line cards in the chassis based on their position and role, the CIN now provides the same logic in a highly distributed and in some cases disaggregated fashion.

For the field installer, attaching a DAA remote device in the Intelligent Edge system becomes as simple as adding a line card to a CCAP chassis where the CCAP chassis has a supervision card or function on a common card watching for interrupts on the backplane that occur when a card is inserted.

When such an event occurs in a CCAP chassis, the control card will determine the type of card inserted and prepare it for operation, possibly upgrading the firmware of the card, then configuring the line card for operation.

In a distributed system, the line card, in part or in whole, is the DAA remote device, the chassis fabric is the Converged Interconnect Network and the supervisor or control card function is an OSS back office application.

For each s-Leaf port the switch sends an LLDP message with switch unique identifier, port identifier, and the attached device MAC address to the controller. The controller emits this to a common message bus where the Inventory and Service applications work together to support dynamic edge discovery and service control by consuming the new message bus topic. It may then be determined if an inventory update is required. If an update is required, it will initiate the action sending the Service Activation the matching service profile from DAA Inventory, where the s-Leaf port config is then updated based on the service profile, initiated by Service Activation, and executed by SDN Controller.

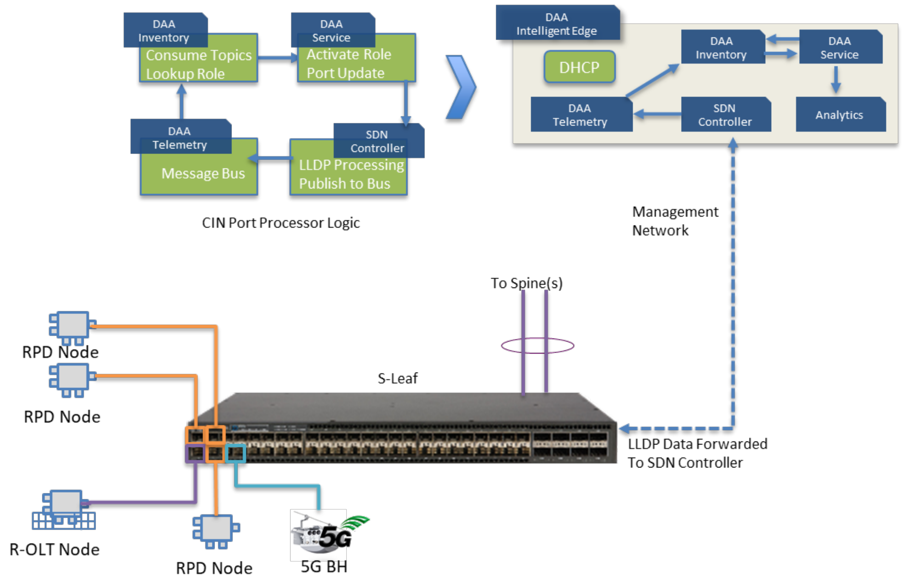

Figure 11 - Any DAA Service Any Port

5. CIN Toplogy

In our model, we have chosen to deploy a CIN based on Ethernet Virtual Private LAN (EVPN) and Virtual Extensible Local Area Network (VxLAN).

Ethernet Virtual Private LAN is a control plane for routing and forwarding Virtual Extensible LAN overlays. When these technologies are applied to the CIN, the back-office co-ordination tasks are greatly reduced as now the interconnect between s-Leaf switches over the Spine is handled within an autonomous system that acts like a forwarding fabric based on prefix identifier.

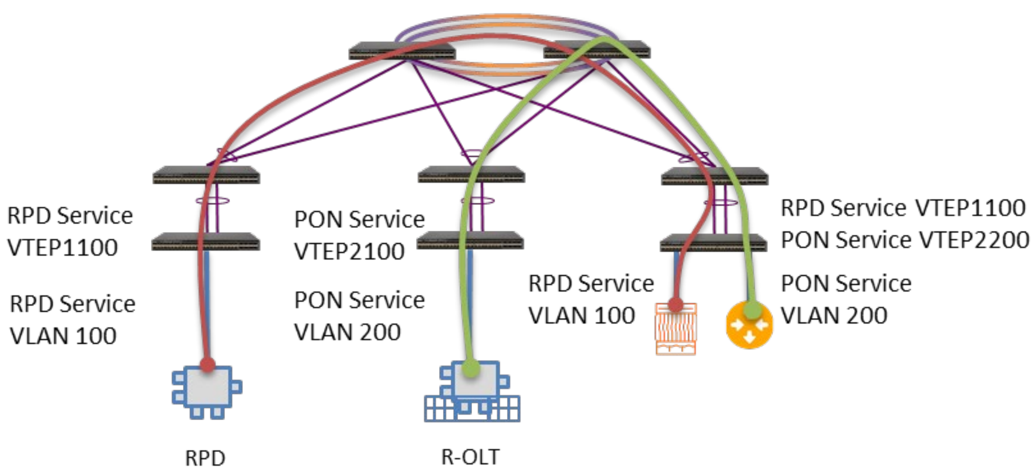

Figure 12 - EVPN VxLAN CIN

The EVPN with VxLAN design has the added benefit of isolating MAC address learning from the Spine. Virtual Tunnel Endpoints (VTEP) are exposed via the s-Leaf devices. Learning for MAC addresses, for example an RPD connecting to a Principal Core, is handled at each of those respective edges.

Essentially, the VTEP will encapsulate a tagged or untagged VLAN identifier from the incoming port of the s-Leaf and map it to a VxLAN identifier. The VxLAN then functions as an overlay network forwarded by prefix across the Spines for any other s-Leaf participant to join the created overlay service.

The Intelligent Edge may configure Spine devices with a uniform configuration to support all the s-Leaf devices, where the bulk of any CIN related CRUD (Create, Read, Update, Delete) actions will now occur to control VLAN to VTEP associations per device or service as required at the edge.

Multiple VLANs are possible per port, which enables the general theory of operation that management VLANs may be the default untagged VLAN for all ports of the s-Leaf in the CIN. Additional VLANs are likely to represent DAA subscriber service networks.

This management VLAN can reach the DAA Intelligent Edge OSS systems which include DHCP services for Dynamic Host Configuration Protocol IP addressing, as well as Authentication Authorization and Accounting (AAA) systems for 802.1x authentications in addition to the collection of topology related LLDP messages destined for the SDN Controller.

Additional VLANs may present per port based on the DAA remote service type, such as a remote OLT PON device with both a management interface for access to an OLT Manager in addition to a VLAN for subscriber IP internet services terminated by a Provider Edge (PE) router. In such a scenario, the PON Service would have two VLANs associated with two VTEP associations to associate forwarding across the Spine.

There are also further scenarios for remote OLT PON where multiple service providers may be served or a requirement to provision unique service separation, the same design of service profile and port configuration orchestration to VTEP associations are automated. The breakout of the PON service then occurs as a result of the PE element occupying a unique VTEP : VLAN Port appearance egressing an s-Leaf verses the management port egress to the back-office systems where the OLT Manager provides R-OLT provisioning and OAM functions with the remote device.

6. Dynamic Edge for DAA

6.1. LLDP Stage

A large portion of our intelligence in the Intelligent Edge is achieved through learned changes in device operating state from the moment a device is attached to an s-Leaf port. The system is dependent on Link Layer Discovery Protocol (LLDP) packets being exposed to the DAA Intelligent Edge SDN Controller.

Each Ethernet frame contains an LLDP Data Unit (LLDPDU) which holds type length values (TLV) for several objects. These may include ChassisID, PortID, Time-To-Live (TTL), Port Description, System Name, System Description, System Capabilities, and Management Address. Additionally, vendor specific LLDPDU TLVs are possible.

Using the device MAC address with the LLDPDU contents when processed from all elements in the CIN becomes incredibly valuable to the DAA Intelligent Edge.

The use of LLDP permits the calculation of the topology graph by the controller. The topology is now available for interrogation by other systems in the Intelligent Edge back office. This is particularly useful when seeking to understand attached location or physical or logical forwarding path based on DAA remote device by a variety of OSS systems.

For our purposes, workflow is triggered from our DAA Intelligent Edge orchestration based on the identity of the attached device to an associated service.

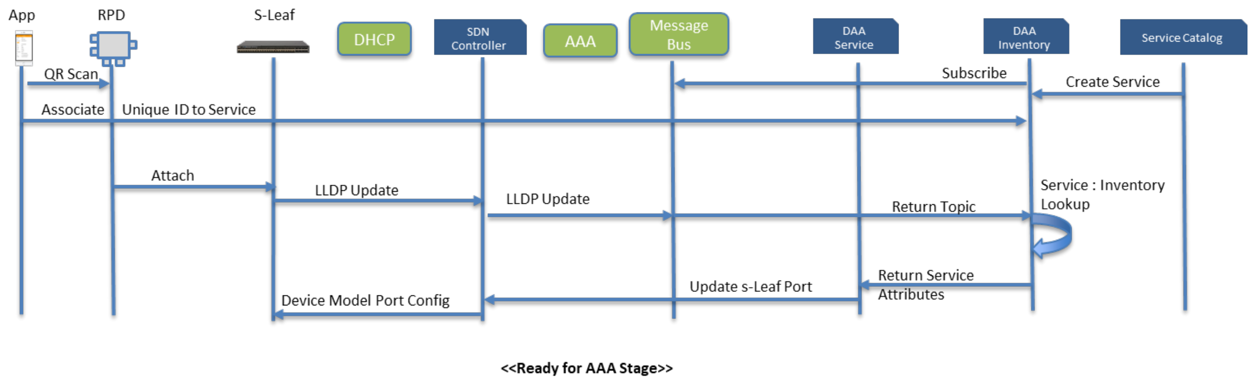

Figure 5 - Intelligent Edge Discovery and s-Leaf Provisioning Sequence

The Intelligent Edge Inventory holds the device records for all network elements in our DAA systems. Spine switches, s-Leaf switches, Remote PHY, Remote MACPHY, and Remote OLT PON devices all have device records in Inventory. By orchestrating the service to the device, we can now inherit dynamic provisioning of the DAA s-Leaf port.

Given the design assumption made in this paper has been a CIN based on VxLAN overlays with Spine forwarding using EVPN control, this means the association of a service based on the discovered device attachment to port permits 'service based' port configurations primarily consisting of VLAN and VTEP setup across the CIN.

6.2. 802.1x Stage

Given s-Leaf ports are configured by default to permit access to the management network, when CableLabs® standardized DAA devices attach, there is a requirement to process Extensible Authentication Process (EAP) to secure the edge network ports. EAP authentication in these terms is usually certificate driven meaning the CableLabs device will use its CableLabs X.509 based certificate or a valid vendor certificate chained from the Root-CA to complete the secure identity association.

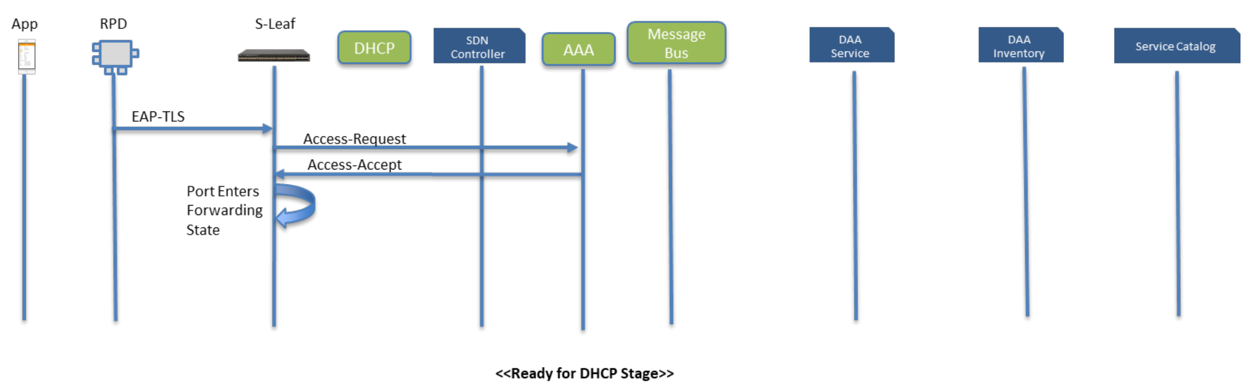

Figure 6 - AAA Sequence

To support this, each s-Leaf device provides an authenticator role to interwork with the AAA services in the back office. While the s-Leaf has informed the SDN Controller of the attached device by its MAC address and including the port information, Layer 2 traffic is not available until the authentication request is successful.

6.3. DHCP Stage

DHCP and DHCPv6 systems have formed a central role in the cable industry since the earliest dependencies on the protocol as effectively a control plane for dynamically provisioning cable modems. In DAA systems, we make use of DHCP and DHCPv6 to assist remote devices in the discovery of their management systems or some other functional actor that will complete the remote device provisioning stage.

For an RPD, this may be a Principal Core capable of configuring the RPD with DEPI, UEPI, and/or OOB RF resources over Generic Control Plane (GCP) protocol. For an R-OLT, DHCP may be used to discover the address of the OLT Manager system capable of provisioning the EPON remote OLT using some vendor defined means thereby supporting rapid deployment of CableLabs based DOCSIS Provisioning over EPON (DPoE) services.

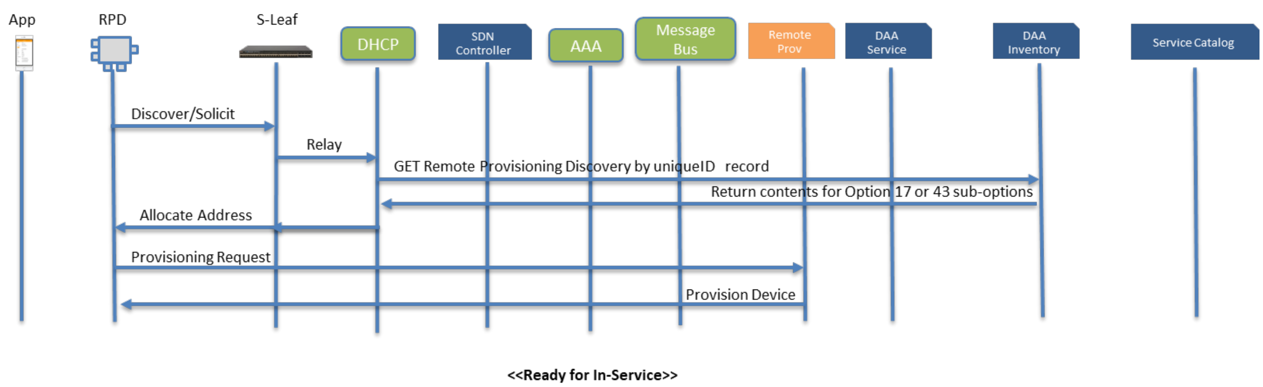

Figure 7 - DHCP Sequence

DHCP systems can often enrich and expose attributes of the network edge to back office systems given they understand gateway relays that have provided edge routing of DHCP discovery and solicit messages to reach the DHCP / DHCPv6 system itself. This provides a hint about the edge of the network a device has attached to, however, it often lacks further details.

When DHCP is part of the DAA Intelligent Edge, it can be said that DHCP systems may enrich the DAA back office and the DAA back office may also enrich the DHCP systems.

Conclusion

As exciting as the numerous new solutions are that came with these distributed system changes, DAA also brings with it a dramatic increase in managed elements in the network with relationships spanning edge into the eadend and beyond. Operators with early experience in DAA deployments have indicated challenges with Dynamic Host Configuration Protocol (DHCP) based back office integrations and a need to enhance the overall relationships of the back office to support the new distributed network in all its forms.

To achieve the business and technology goals of DAA, it will require automation and additional intelligence in the OSS back office systems that many would call Software Defined Networking. This is needed for the goal of returning to the operator the same level of autonomous operation that exists within CCAP chassis over distributed and multi-vendor network elements going forward.

In the proposed solution, the combination of orchestration concepts using a product catalog to model a DAA service were explored, application then of these DAA services based on device association were made real. Using SDN based automation and combination with DHCP systems, this specific solution enhanced the network edge for onboarding that presents both an initial step to added edge intelligence and a future option towards 'DAA as a Service' intent-based networks.

With multiple access network options, introduction of virtualization and wireless, we must co-ordinate the systems in the Hub and Headend while deploying any of these remote devices in a Zero-Touch approach whenever possible that maximizes the velocity for the operator overall and minimizes the time spent per field technician installing any of these edge solutions.

Abbreviations

AAA	authentication authorization accounting
API	application programming interface
CIN	converged interconnect network
CCAP	converged cable access platform
DAA	distributed access architecture
DHCP	dynamic host configuration protocol
DOCSIS	data over cable service interface specification
DPoE	DOCSIS Provisioning of EPON
EAP	extensible authorization protocol
EPON	Ethernet PON
EVPN	ethernet virtual private network
FOSC	fiber optic splice closure
GCP	generic control plane
LLDP	link layer discovery protocol
OAM	operations, administration and maintenance
OLT	optical line termination
OSP	outside plant
OSS	operations support system
pCore	principal core

PE	provider edge
PNF	physical network function
PON	passive optical network
QAM	quadrature amplitude modulation
QoS	quality of service
PHY	physical interface
RF	radio frequency
RMD	remote macphy
R-OLT	remote olt
RPD	remote phy
SDN	software defined network
SFP+	small form pluggable optics module; the + indicates 10G capability
s-Leaf	secure leaf; a type of switch
TVL	type length value
vCore	virtual core
VNF	virtual network function
VLAN	virtual local area network
VxLAN	virtual extensible local area network
VTEP	virtual tunnel end point
WDM	wave division multiplexing
xPON	passive optical network