USING SDN AND NFV FOR INCREASING FEATURE VELOCITY IN A MULTI-VENDOR NETWORK

Jeff Finkelstein, Cox Communications
Alon Bernstein, Cisco
Samir Parikh, Gainspeed

Abstract

As shiny objects of the past decade go, few have achieved the shininess of Software Defined Networking (SDN) and Network Functions Virtualization (NFV). You cannot open a trade publication either soft or hard copy, without seeing these TLAs (three letter acronyms). Every conference has topics related to them. In fact, there are even entire conferences dedicated to these technologies with presenters and vendors discussing what it all means to potential customers and the communication industry at large.

This paper is on a similar path, but more specific to cable and how operators (MSOs) may take advantage of these technologies to achieve two key objectives in their quest to improve services. First, MSOs need to more efficiently scale the network on many levels so as to control OPEX and CAPEX, as well as to fit the necessary equipment into existing facilities. Not only does the equipment need to scale, but so too does the support infrastructure. Power, heating, cooling, physical space, cabling and the like all have to reach a scale that was unimaginable just a few years ago.

The second objective MSO's must accomplish relates to the velocity of service enablement, creation, management and provisioning. The operator community is very good at whiteboarding ideas, producing slides showing how these new services will make things better, and talking about them in meetings and conferences. Where they frequently miss is the speed at which they actually create a product and bring it to market so that it adds revenue to their bottom line. If an MSO could get new products and services out more quickly, and at a lower cost, not only would they be more competitive on a technical level, but they could open new ways of getting ahead of our competition.

The challenge is in finding ways to meet both of these objectives — scaling more efficiently and increasing the velocity of service deployments — without increasing complexity. By using SDN and NFV to drive this innovation, operators can actually simplify operations by decomposing network elements and functions and spreading them across multiple devices. Rather than have a large monolithic architecture we distribute the work across smaller pieces of equipment using simple interfaces to communicate between them. As the industry has learned over time, simplification of technologies helps drive cost out.

In our paper, we will present descriptions and use cases of how MSOs may use SDN and NFV in the headend and at the customer premises to achieve these goals.

INTRODUCTION TO SDN AND NFV

A brief introduction is needed to set the stage for the rest of this paper. SDN and NFV can take on different meanings depending on context, so let's start with the basics.

What is SDN?

Historically, the network and the applications have been managed as discrete entities with no direct interaction outside of someone changing network paths or elements as awareness of an application's needs became apparent. Software Defined Networking (SDN) integrally links these two components. Simply put, the key concept behind Software Defined Networking (SDN) is that the network becomes aware of applications and that applications become aware of the network.

There have been a variety of functional definitions of SDN. A couple that have been more commonly used by cable operators and equipment vendors are:

- Separation of control and forwarding functions
- Centralization of control and distribution of processing/forwarding

The importance of SDN was recognized as the use of cloud services grew and their value was better understood, it became clear that the data-center needed to evolve to effectively support these services. The industry realized the need for a method to connect and control the resources (virtual machines, network capacity, storage capacity) making up the cloud in parallel with the applications running on the network. Specifically the network and cloud resources needed to be able to quickly and reflexively react to which applications were being used and how they were being used. Simultaneously, applications needed to be able to communicate what they needed from the network and react to changes in available resources. All of this needed to happen automatically and in real-time SDN had an enormous impact on how we envision and build data-centers in this era of cloud computing.

What is NFV?

Network Functions Virtualization (NFV) was created to leverage standard IT virtualization to consolidate many different network elements into a cloud architecture that uses standards commercial-off-the-shelf (COTS) hardware. This hardware runs many functions which typically have been handled by purpose built platforms. It is important to note that for some functions, particularly those that involve processing packets at line rate, purpose-built platforms remain a best practice. However, placing other functions on a virtual machine (VM) on COTS hardware can accelerate service deployment and optimize equipment spend.

How do SDN and NFV work together?

SDN and NFV are complementary technologies aimed at streamlining the operation of the network and deployment of new services. They can be view as two pillars in the evolution of managed services networks. NFV-enabled applications interact with an SDN-enabled network infrastructure. NFV streamlines the deployment and scaling of new applications in a cloud environment. SDN streamlines the deployment and scaling of traditional network resources. While they can be deployed together to achieve maximum efficiency, they are independent technologies that can exist without each other.

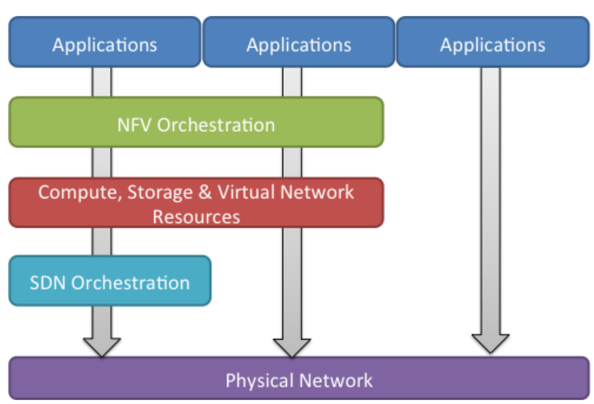

Figure 1: SDN and NFV Layers

Why SDN and NFV?

Graphs showing the increased bandwidth required to meet ever-growing consumer demand are commonplace. In meeting after meeting, MSOs debate why the growth curve cannot possibly continue at the historical CAGR. Yet, it does continue along the existing line and shows no sign of abating.

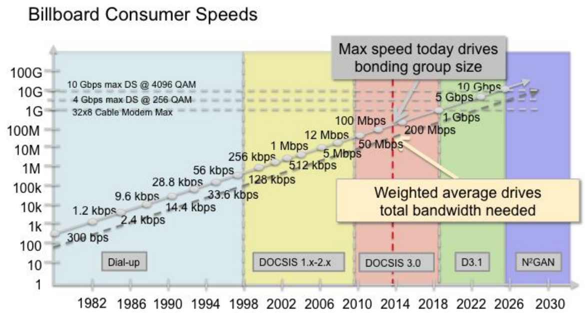

Graph 1: Bandwidth Growth Curve

Accepting that this growth will continue along its present curve, the challenge of keeping pace with bandwidth requirements relative to operators' CAPEX and OPEX spends becomes evident. To keep up with the ever-increasing bandwidth demands, MSOs must replace current hardware with higher capacity equipment and change the plant to support higher order modulations. This in turn allows operators to compete more efficiently and to replace customer premise equipment with next-generation equipment to deliver a new and improved user experience.

In the past, MSOs would have had an opportunity to add capacity and then reap the benefits, realizing a higher profit margin thanks to capitalization of technology assets. However, the useable lifetime of network assets is decreasing rapidly as we approach the asymptote of the capacity growth curve. Capacity requirements now are simply growing at too great a rate for operators to get ahead of the curve in terms of CAPEX and OPEX.

MSOs must find a new way to increase the velocity of service enablement in order to meet or get ahead of their customers' technology curve. At the same time, operators need to manage the scale of new equipment needed to provide those services while reducing equipment density to stay within the power, heating, cooling, and space availability of existing headend and hub facilities.

The advancements of SDN and NFV provide MSOs with a new paradigm for scaling services and support infrastructures through decomposition of functions.

In the following sections, we will discuss how SDN can enable Network Virtualization and how NFV can enable CPE Virtualization.

SDN AND NETWORK VIRTUALIZATION

SDN can take multiple forms in an operator network. This section will focus on how it enables operators to virtualize the headend.

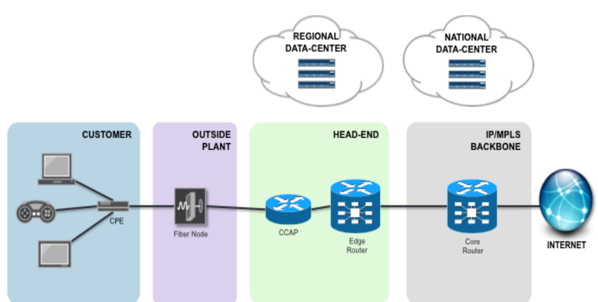

Figure 2: Typical MSO Network Today

As operators scale services, the headend infrastructure (power, cooling, rack space) is coming under increased pressure. The existing architecture of separate video (EQAM) and data (CMTS) platforms, coupled with the combining and splitting networks, occupies a tremendous and growing amount of space and consumes an enormous amount of power. Each new generation of products has provided greater efficiency, but not enough to keep up with demand. Additionally, keeping pace with demand requires operators to churn hardware without maximizing its usable lifespan.

Converged Cable Access Platforms (CCAP), which combines EQAM and CMTS in a single box, are supposed to alleviate this problem. However, evaluating available CCAP products suggests that this is just the next step on the same technology curve.

As previously discussed, MSOs must find a new way to keep up with bandwidth demands on the network, minimize power, cooling and space requirements and prolong equipment's usable life in the network

The key is to build upon CCAP, adding SDN technologies to enhance it, and thereby creating a Virtual CCAP.

This approach allows CCAP to be viewed not as a platform, but rather a collection of base functions that support the services offered by cable operators. These base functions include:

- Cable Control Plane
- IP/MPLS Control and Forwarding Plane
- Subscriber Management
- Video (QAM) processing
- DOCSIS processing
- RF modulation

SDN technologies enable MSOs to move away from the notion that all of these functions must be collocated in a monolithic system centralized in the headend.

In addition, virtualizing these base functions enables operators to leverage best-of-breed products for each function, thereby creating an eco-system which will provide the most feature-rich, reliable and yet lowest cost solution.

SDN technologies allow creativity with regards to where in the network and how these functions are implemented. This paper examines the Virtual CCAP architecture that enables operators to collapse layers/functions without changing their overall network operational model.

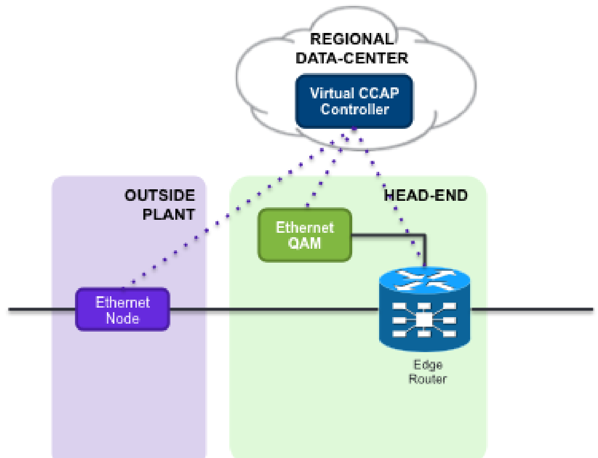

Figure 2: Virtual CCAP Architecture

No changes are required to the CPE in the home, the services delivered to the end customer, or the back-office systems (provisioning, billing, etc.).

<u>Virtual CCAP Components</u>

The Virtual CCAP is composed of the following components:

1. <u>Virtual CCAP Controller</u>: This is a software application in the datacenter that runs the Cable Control Plane and interfaces with existing back-office systems. It is the brains of the Virtual CCAP and is responsible for all control and management of the system. It orchestrates the entire solution, tying all of the pieces together to effectively operate as a single CCAP device.

2. Edge Router: This is the router that is already deployed in the headend. It simply takes on subscriber management functions on top of the functions it already performs (i.e., IP/MPLS control and forwarding plane)

3. Ethernet QAM: This is a new category of edge QAM devices that output a fully groomed MPTS for digital video as a multicast stream on an Ethernet interface as opposed to an RF port.

4. Ethernet Node: This is a new category of fiber nodes that handles all DOCSIS and RF processing. This allows analog transport to be replaced with standard Ethernet.

Benefits of NFV- and SDN-Based Architecture

The benefits of adopting a Virtual CCAP architecture can be broadly categorized as CAPEX and OPEX reduction.

CAPEX Reduction

Operators can achieve significant CAPEX savings with a Virtual CCAP architecture. Savings come from a number of fronts:

1. Replacing legacy analog optical transport with standards based Ethernet transport: 10 Gigabit Ethernet is the lowest cost per bit transport solution available today. It benefits from massive economies of scale from use in data centers and telecommunications infrastructure across wireless, wireline and cable operators.

2. Eliminating the physical CCAP system and leveraging existing headend equipment: Equipment reduction and use of existing gear_lowers the cost of the Virtual CCAP solution.

3. Increasing overall system scalability: The scalability of the solution (number of Ethernet Nodes supported per edge Router, full spectrum agility of the Ethernet Node) provides a much lower cost point as operators add new service group (SG) to the network and scale DOCSIS services on new and existing SGs.

OPEX Reduction

1. Eliminating manual tuning of RF and physical RF combining and splitting: By distributing RF generation to the node and removing it from the headend, spectrum allocation changes can be made quickly and cheaply in software. The manual steps required today to balance RF power levels from the headend towards the transmit laser and in the node between the optical receiver and coax are eliminated. This saves hours and hours of planning and implementation. In addition, physical RF combing and splitting is eliminated and replaced by a "digital" combiner that is controlled entirely by software.

2. Reducing power, space and cooling requirements: By eliminating the physical CCAP, the power and space requirements in the headend are dramatically reduced. This has a follow-on benefit of lowering the cooling capacity required, further increasing the OPEX savings. Virtualizing the physical CCAP also eliminates the need to augment facilities as the network grows.

3. Optimizing use of plant: By pushing DOCSIS (MAC and PHY) processing out of the headend into the node, the need for a tight timing relationship between the headend and node is eliminated. Furthermore, the use of Ethernet transport makes it cost-effective to have longer fiber paths between the headend

and the node. Add to that the reduced footprint that results from virtualization and operators can now consolidate smaller hub facilities into larger, central headends.

4. <u>Leveraging existing management systems</u>: The Virtual CCAP is managed just like a physical CCAP. Control and management flows through the centralized Virtual CCAP controller which now manages up to 400 SGs as opposed to the 64 SGs of a traditional CCAP system.

5. <u>Improving service quality and customer satisfaction</u>: Digital transport to the node increases SNR, improving service performance. This will lower the number of trouble tickets and truck-rolls, resulting in significant savings on an ongoing basis and increased customer satisfaction

Virtual CCAP enables operators to accelerate service velocity and keep up with customer bandwidth demands while concurrently lowering CAPEX and OPEX, all without impacting existing operational systems and processes.

Once this foundation is in place, new and interesting applications can be rapidly delivered to the customer premises by leveraging NFV technologies.

<u>CPE VIRTUALIZATION</u>

In addition to applying SDN and NFV to the headend, an operator can apply these technologies to virtualize the CPE. Virtualizing CPE enables certain capabilities, traditionally implemented in the home to be implemented in the operator's cloud infrastructure.

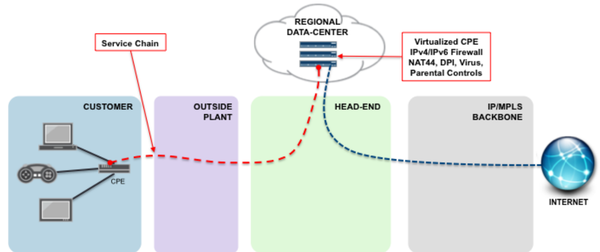

Figure 3: Virtual CPE

<u>Virtualization Candidates</u>

CPE functions can be categorized under the following building blocks:

<u>Networking</u>

1. Layer 1 Physical access media data/control

2. Layer 2/3 networking functions such as routing

3. Upper layer functions such as Network address translation (NAT), firewall, and Deep Packet Inspection (DPI).

It is possible to "virtualize" the control plane or the packet-by-packet handling of the data plane for each of these functions.

<u>Home Appliances</u>

Some of the appliances within the home can be virtualized as well. An example already being deployed in some cable networks is the network video recorder. A household's recorded TV programs would be stored in the cloud instead of on a hard drive in the home. The same way a home backup system can be moved to the cloud. Another example would be the rendering of TV guides, games and other graphic applications in the cloud instead of the home device.

Moving to the Cloud

Which of these functions should be moved away from the home and into a virtualized cloud service must be determined on a case-by-case basis by performing a cost-benefit analysis.

This section will explore the value of CPE virtualization for a couple of specific use cases.

Parental Controls

Parental control refers to the ability to block access to certain URLs. Performing the task of comparing URLs against a large black list is fairly CPU-intensive and thus well-suited for an NFV-type of application rather than standard packet processing.

A virtualized version of parental control would perform the URL filtering in the cloud instead of running it on individual computers or in the home gateway. Updating the list of blocked sites is easier to manage centrally.

Firewalls

Firewalls are network security appliances meant to protect the home network from external security threats. A virtualized firewall would run in the cloud, inspect all packet streams and block those that are a security risk. As with the parental control example, managing the list of security threats centrally is easier and more reliable then having it distributed to a home gateway of consumer devices.

Benefits of Virtualization

The business benefits of CPE virtualization can be categorized as improving time to market, CAPEX and OPEX reduction and system simplification.

Faster Time to Market

1. Accelerating software development: Software development for a virtualized application is less constrained than in a typical embedded system. This not only reduces the time needed to develop apps, but also to add new features as there are more development tools, fewer memory constraints, faster CPUs and a wider pool of engineers capable of programming in a standard environment.

2. Not relying on specific hardware: With virtual applications, there is no need to wait for equipment orders to arrive and hardware to be installed. New services can be launched on the existing server infrastructure.

3. Reducing testing time: Testing of a virtual application can take place in a production environment by starting with a small-scale implementation and then rolling it out as the software stabilizes. In total, this results in shorter qualification cycles.

4. Eliminating dependency on devices in the home: From a logistical standpoint, home devices are not easily replaced. Consequently, new features that depend on an updated home device take years to roll out. With a virtualized environment, this complexity and delay is removed since the dependency on the home is eliminated.

Reduced CAPEX

1. Reuse of resources: The same compute resources can be used for multiple applications, thereby maximizing utilization and reducing the need for additional resources. For example, during high traffic daytime hours, firewall services require a lot of resources.

However, when traffic drops at night, the freed up resources can be used, for back-up services.

2. Statistical multiplexing of resources: Historically, operators could statistically multiplex bandwidth in order to maximize utilization. With NFV operators can also_statistically multiplex the resources needed for a function. Furthermore, the ability to “cloud burst” into other “clouds” when required capacity exceeds available resources allows operators to engineer the network for average usage rather than peak usage.

3. Lowering CPE costs: Because services are enabled and delivered from a virtual machine in the network, they can be enabled for existing customers without replacing operational CPE.

Lower OPEX

1. Reducing manual operations: Because SDN and NFV can facilitate network automation, many operations that previously required manual intervention, including new service deployment, can be automated.

2. Minimizing CPE software upgrades: For every function that runs on the virtual CPE, the cost and complexity of managing CPE software upgrades is eliminated.

3. Improving service quality and reducing support costs: Deploying a unified service that operates independently of the hardware platform deployed at the customer premise leads to better service and less support.

4. Dynamically scaling and optimizing use of resources: CPE consumes power even if it is not forwarding packets. A virtual machine can be completely turned off when a function is not in use and its resources can be reallocated to another function. This enables dynamic scaling of resources, optimizing power, space and cooling.

Operational Simplification

The operational implications are closely related to the OPEX reduction, however it is worthwhile to spell them out:

1. Eliminating the need to manage software versions: CPE functions run as virtualized network functions leverage a single service model. This eliminates the need to deploy and manage software versions across the typically varied CPE footprint and streamlines service roll-out.

2. Speeding error identification and resolution: When software is deployed as a system-wide resource, troubleshooting, isolating and remedying a defect is simpler and quicker.

3. Standardizing infrastructure across multiple applications: Since CPE functions are being run on centralized platforms, operators can standardize the infrastructure (hardware, software, tools), reducing complexity and improving reliability.

4. Increasing customer satisfaction: The home environment is typically a jungle of improperly plugged cables and equipment tucked away without proper air circulation. By moving features/functions into the operator cloud, the complexity of in-home systems is reduced, which will result in fewer customer calls due to improper installs.

Open Issues

The benefits of using SDN and NFV to virtualize CPE functions are considerable. There are, however, a couple of issues which must still be addressed.

1. Requires greater technical skills and knowledge: While the overall operation is simplified by using NFV to perform CPE tasks, a more skilled work force is needed to maintain the system. For example, connecting an Ethernet cable between two physical devices does not require a high degree of skill. However, operating the virtualized network functions requires substantial technical knowledge and a higher skill level.

2. Requires greater security: A cloud environment requires a higher level of security as some of the benefits of NFV could backfire in the event of a security breach.

3. Requires more careful data center planning: A data center that is too heavily oriented around video distribution might need to be upgraded to support virtual CPE applications. Geographic placement of the data center will also come into play for latency reasons. The closer the data center is located relative to the customer, the better response times would be – especially compared to companies that might provide virtual CPE services over the top.

CONCLUSION

In this paper we have explained what SDN and NFV are, how they may be used both in the access network and at the customer premise, what they achieve from a CAPEX and OPEX perspective, and the benefits in terms of resource management.

We can now see how SDN and NFV can work together to help scale the network to levels never before considered. This is achieved by breaking the functional components of devices into software, placing these software elements into a virtual machine environment, and having each use resources only as needed. This helps operators optimize their network spend by using COTS hardware and managing this equipment using an abstraction layer under the control of the orchestration layer and hypervisor.

Real-world examples of how this works are rapidly appearing in labs, proof-of-concept trials, field trials, and in actual deployments. The authors of this paper believe this is the beginning of a new era of cable technologies unlike any we have seen to date. Operators are no longer required to add new hardware except when existing equipment runs out of computing or packet processing resources. We invite you to join the many efforts underway at CableLabs, IEEE, IETF, ETSI, and other standards organizations to help develop the specifications, use cases and models needed to drive these exciting new technologies into daily use so we may all benefit from their adoption.