Delivering Seamless Subscriber Aware Services over Heterogeneous Access Networks using a SDN and NFV framework

Nagesh Nandiraju Ph.D., Yiu Lee and Jorge Salinger
Comcast Cable

Abstract

The cable industry is seeing two major shifts: adoption of new access technologies that are capable of delivering higher data rates and a growing interest towards new advanced services that go beyond traditional service flows. Access network technology is undergoing an evolution. Next generation DOCSIS (DOCSIS 3.1) and Passive Optical Networks (e.g. 10G EPON) enable MSOs to deliver higher data rates to subscribers. On the other hand, Network Function Virtualization (NFV) and Software Defined Networking (SDN) have certainly excited the industry and many innovative network architectures and applications are being proposed using these concepts. MSOs are researching to develop and offer more advanced and seamless subscriber-aware services to users that go beyond the DOCSIS service flows. These services will be abstracted from access network and end-to-end network technology has to be agile and capable to deliver these services.

In this paper, we aim to establish a SDN and NFV framework that enables delivering a uniform set of new services across multiple access networks. We will start by elaborating on what are the new services and key drivers for these services, how operators would offer and manage them. Then we will dive into the changes required in the network that will enable MSOs to quickly deploy new services to subscribers. As a part of this effort, we will present our thoughts on what functions can be virtualized in the access network that will improve the agility of the network.

INTRODUCTON

MSOs are currently seeing several transitions that can help drive their next generation network architectures. First, over the past decade or so, Data has been growing significantly at 40-50% compounded annual growth [1]. The advent of smartphones and cloud-based services has only acted as a catalyst to increase the data consumption. Consumers are increasingly relying on new services and applications that present data to end-user. This translates to more bits delivered per customer but not necessarily increased revenue for service providers. Advances in network equipment and technology have been unable to push down cost per bit at the same rate of bandwidth demand. Many MSOs and telecom operators are increasingly seeing this trend and are looking at ways to enrich the customer experience by adding new value added services to their portfolio in addition to just delivering data.

Second, in the access network, MSOs are continuously improving existing protocols (DOCSIS 3.1) and continuously increase capacity to meet the customer demands for higher speeds. DOCSIS has proven to be one of the most successful broadband access technologies. DOCSIS 3.1 can deliver 1Gbps and will be capable of up to 10 Gbps in an all IP network.

Next, the infrastructure to deliver data in the last mile can use several media – HFC, Fiber, Coax, CAT6. There are considerations of using fiber to the building or units in new green field deployments. Several protocols such as GPON, EPON, RFoG, Ethernet can run over these fiber networks and are capable of delivering 10Gbps or even higher speeds.

On the technology front, Network function virtualization (NFV) and Software Defined Networking (SDN) have clearly been the top buzzwords in the networking industry over the past couple of years. The "talk" has transformed into many interesting developments from vendors. Several vendors are demonstrating general-purpose hardware capable of performing network functions previously only possible over purpose built hardware. These activities are primarily accelerated by the increasing use cases of networking by cloud service providers.

MSOs are actively researching to establish their next generation network architecture to better support these transitions, and how they can offer a wide variety of services with an enriching experience to end customers. As the hype phase of NFV and SDN has passed over, time is ripe for investigating how and where these technology advancements can be applied into the network to monetize and improve operational efficiency. Particularly, MSOs are researching to develop and offer more advanced and subscriber-aware services to users that are seamless and go beyond the DOCSIS service flows. These services will be abstracted from access network and end-to-end network technology has to be agile and capable to deliver these services. The questions that often get asked on NFV are: "Will MSOs embrace NFV and SDN? Can NFV offer performance similar to purpose built hardware and is it cost effective?"

Will MSOs Embrace NFV and SDN?

In the not-so-distant past, many of us were accustomed to carrying a variety of accessories and specialized gadgets such as MP3 players, Cameras, GPS, cell phones, USB sticks, and this list goes on. This trend has suddenly changed with the advent of smartphones. Now, people's traveller kits are much lighter as they only carry smartphones, and all the above-specialized applications are virtualized and underlying hardware components are either natively supported by mobile chipsets or packed into the smartphone. The services provided by these accessories are transformed into applications (apps) running on a generic piece of hardware and an operating system sitting on top of that hardware and interworking with the server functions running in the "cloud" to provide necessary functionality and services to the software applications (API). Applications are almost agnostic of the hardware and somewhat tied up with the operating system. The number of apps in both Apple and Android app stores is well over a million. According to Apple, it paid $13 billion to app developers [1].

NFV and SDN are two concepts to rationalize conventional network design. In traditional model, every day MSOs face two challenges in their network. First, MSOs deployed many purposely-built network equipments such as routers, switches, load-balancers and firewalls. Those equipments are designed to serve specific set of functions. They usually work anomously and require different operational procedures to interoperate them. Network integration and operation is well-known time-consuming and complex. Second, local applications such as home security and VoIP require to implement client applications in the Customer Premise Equipment (CPE). MSOs must develop the applications or contract the CPE vendors to develop the applications. Then, the MSOs must upgrade thousands of CPEs when they are ready to introduce the new service. NFV and SDN attempts to apply software development paradigm into networking world to improve the level of automation and deployment efficiency. At high-level, NFV aims to virtualize conventional network functions into softwares running on Commercial Off-the-Shelf (COTS) and Virtual Machine (VM). SDN aims to orchestrate virtual network functions (VNF) that are running on COTS and VM. The potential benefits of NFV and SDN are enormous. Similar to the

transformation brought by smartphones, NFV and SDN aim at bringing similar transformation for service providers by moving most of the services from running on purpose-built hardware platforms to COTS compute, storage and networking infrastructure. Some benefits are obvious: flexibility to scale for service demand, ability to launch new services with ease, agility in service provisioning, automation in operation, and dynamic scalability. But a fundamental question that is asked by almost everyone in the service provider environment is, how can performance be effective with COTS?

Indeed, similar questions were asked when voice/camera/GPS and other critical apps have been proposed to run over generic smartphone software instead to their purpose-built hardware predecessors. The questions are now history as the industry and smartphone users have overcome the doubts and almost all these apps are running on everyone's smartphone fairly regularly.

Similarly, with NFV, a large variety of service provider applications (e.g., firewall, Intrusion Detection System (IDS)/Intrusion Prevention System (IPS), malware detection) can run over a common infrastructure with generic hardware and virtualization. There are primarily two classes of applications in a service provider environment: data plane and control plane. Control plane applications are relatively easy to run over the COTS platform while data plane applications require hardware accelerated paths and orchestration features that prioritize packets, minimize latency and provide line rate throughput. For many applications not only raw bandwidth, but packet processing in terms of million packets per second (pps) and ultra low latency are vitally important.

CURRENT TRANSITIONS IN MSO ACCESS NETWORKS

MSOs are always in the quest to bring the best and reliable services to their customers. They are constantly augmenting their existing infrastructure or building new infrastructure to support new and upcoming customer demands. Below are some of the transitions that help address the changing demand for current services, diverse geographical serving areas and landscape of new services.

- Multi access architecture
  - DOCSIS has certainly been the most successful access technology for MSOs. With technology advancements and maturity of various fiber based PON solutions, some MSOs are building fiber all the way to new buildings and making use of different PON technologies. Alternate last mile access tehnologies such as Ethernet and WiFi are also seeing some adoption to serve certain end customers. This implies that supporting uniform services across multiple access technologies will be quintessential.
- Remote Architecture for DOCSIS
  - There are several drivers behind the remote architecture in DOCSIS. As a natural evolution, service group sizes in DOCSIS networks are becoming smaller in an attempt to provide higher capacity to subscribers. DOCSIS 3.1 with the help of OFDM enables an efficient use of spectrum by supporting higher order QAM modulations. However existing analog optics and density of RF can be limiting factors to fully reap the benefits. One way to address and benefit

from these changes is by adopting a remote architecture and move from analog to digital optics.

- Remote architecture for PON
  - In an all fiber access network deployments, PON technologies (EPON or GPON) have reach and density limitations. PON technologies have been designed for a 32/64/128 split ratios and typically operate within 20 kilometers. Since MSO networks have traditionally longer reach with their HFC networks and higher densities, the need for a remote architecture in PON deployments is desired.
- Vitualization of Home Network (vHN)
  - MSOs use their robust access network to deliver various services such as voice, video and data to their customers. As we enter into the arena of operating multiple access technologies such as DOCSIS, PON and WiFi, these services should be agnostic to the access technologies and work seamlessly on any type of access networks. However it often isn't the case. Each technology has its network characteristics and its deployment scenarios. The current model couples the service characteristics to the network characteristics. For example: to deliver linear content on DOCSIS network, the current practice requires MSO to statically or dynamically provision a dedicated DOCSIS service flow for the video delivery platform. This process will not seamlessly work in PON or WiFi. The goal of vHN is defining an architecture to decouple the services from the access network technology. Application or service logics are virtualized and could be run inside the network or in the CPE.
- IP CDN
  - With the growing web content, network DVR and IP video, the edge caches will naturally move deeper into the networks, perhaps next to the access network gateways (closer to subscribers).

With these transitions in mind, MSOs are interested in building an architecture that leverages the synergies between these changes. Specifically some of the key value propositions and aspects we are interested in are: How can NFV

a) improve our ability to launch new features and services at a much faster pace reducing the cycle time

b) increase the utilization of common compute, storage and networking resources to save space, power and increase operational efficiencies

c) reduce risk associated with rolling our new services (e.g. lower upfront cost per service per subscriber, enable the elasticity of services).

APPLICATIONS - SERVICE PROVIDER'S APP STORE

Modern MSOs data service is more than just a conduit to deliver IP packets. MSOs are actively working to enhance existing application and introduce new applications to their customers. Some examples are parental control, remote storage, home network health dashboard and home automation. Parental control is one of the key applications that is being virtualized to provide the necessary control.

The number of personal devices per home consuming data is steadily growing and currently ranges from 4-15. This number is only increasing every day. As a result, home networking is getting more complicated.

Customers are expecting more information from their networks with deep data analytics of their data usage and easy access to device policy management. This requires applications at the higher layer to provide accounting with some hooks from the access technology layer.

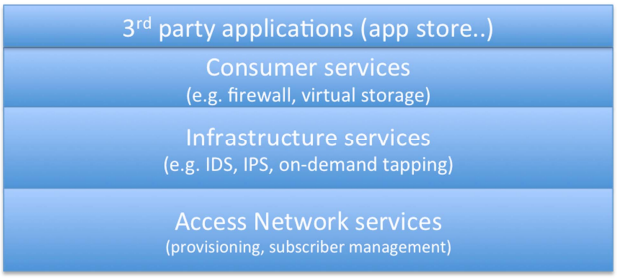

Figure 1: Service Provider Application Stack

Network attack is a real threat not only to big enterprises but also to home users and small businesses. It is not uncommon to hear hackers exploit security vuluerabilies and steal personal data and take control of victim machines to launch attack. Hacking techniques are evolving rapidly, becoming much more sophisticated and harder to detect. Recently, as many high profile data breaches (e.g., Target, Home Depot, Turbo Tax) have been posted in the media, users are more aware and concerned of network security. The traditional client based protections such as anti-virus and personal firewall software require software installation in hosts. Users must constantly upgrade the thread signatures to protect from newly discovered malwares. Besides, client based protections require the software to run on the host, Internet of Things (IoT) such as webcam and smart thermostat may not have the processing power and memory to run the software. The next logical place to protect a user premise is in the ISP network. Customers are expecting their ISPs to offer more protection.

Figure 1 shows a simplistic view of the different layers of services that are likely to run in an MSO network. The bottom layer is the access network services layer which handles applications such as provisioning and subscriber management. The next layer is for applications that are more focused towards network engineers to maintain, manage and protect operator's internal network. Consumer services layer is where the new set of services or applications described above will be located. It is also easy to conceive that having this layered architecture and exposing the right hooks can facilitate development of 3rd party applications that can address a wide variety of application needs.

THE 3 FUNDAMENTAL VARIABLES TO PLAY WITH AND APPLICATION SWEET SPOTS

Compute, storage and networking are the three fundamental blocks that any application or service will require. These are also the three key knobs that vendors and service providers can play with to optimize the performance of applications and delivering the services that drive customer experience. There are wide variety of applications and services that operators are currently providing and/or plan to provide in near future. As shown in figure 2 these applications fall into several classes that can be optimized on one or more variables. For example, consider "policy management" application. This is primarily a control plane application and can be optimized by high CPU power, faster access to memory and persistent storage. Consider another example – VoD service. This services typically uses fixed assets that are preprocessed with multiple bit rates and stored in various storage locations. Delivering this service to end customers relies on substantial storage and network capacity with some need for compute power. The end customer experience (such as smooth access, higher quality video bit rate streams, engaging UI) can be optimized without breaking the bank of service provider (i.e. augmenting network capacity) by playing with the three variables.

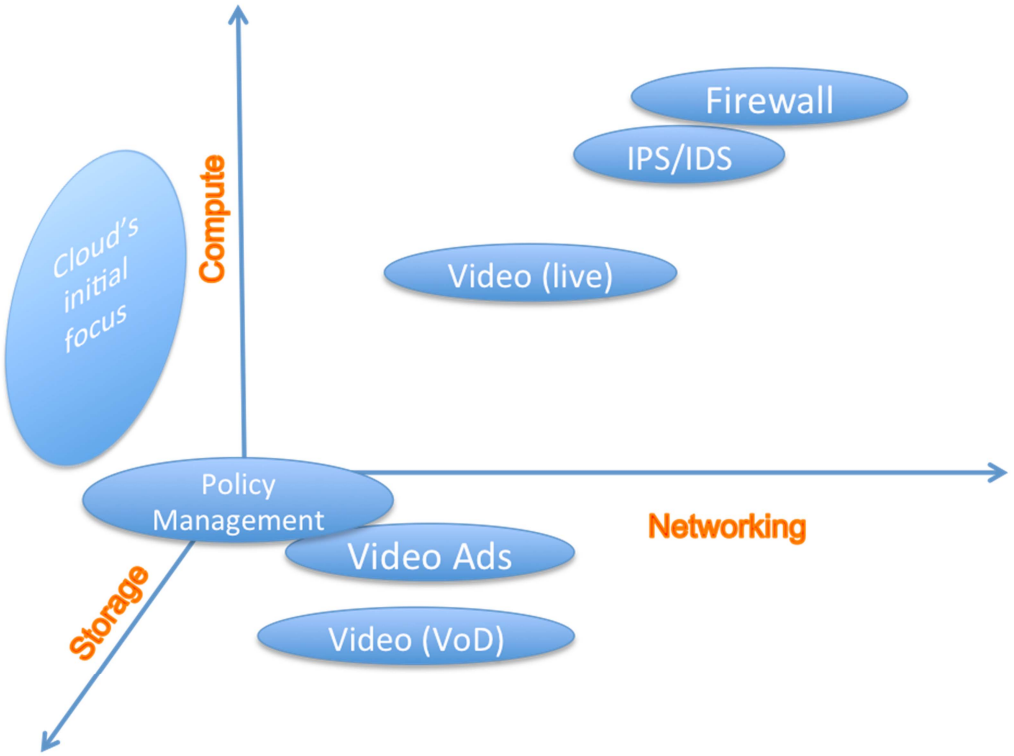

Figure 2: Three Fundamental Variables that can be optimized for applications or vice versa

How does Cloud Computing play a role?

Cloud computing platforms have been optimizing their offerings to enable applications running seamlessly on virtualized compute and followed it by making scalable storage simple to access and reliable. The fact that any fundamental block of operation often involves compute and storage required cloud providers to attach storage to compute and offer storage as a service. The volume of compute and storage in a cloud environment and associated transactions has found short falls in existing networking solutions. This has driven innovation in the networking space both in protocols (e.g. VxLAN, NVGRE, network orchestration) and hardware (dense 10G and 40G switches in 1RU form factor). However the innovation is primarily focused towards a dense, contained data center environment. The innovation is now expanding to address the issues in inter-datacenter traffic (also referred as east-west traffic).

A service provider environment is different when compared to a cloud service provider. In a service provider environment, the primary focus has been on networking to address the challenge of transporting more and more bits to many small and relatively sparse distributed end points.

Optimizing Applications for balancing the tradeoffs between the 3-axis

As show in figure 2, different applications have their own sweet spots in the 3-D space of compute, storage and networking. Some applications can be optimized and architected to change their coordinates in this space. For instance consider two key applications of a

traditional service provider: Video ads and "Video on demand". Both these applications can be hungry for the three resource categories. However if we carefully architect these applications and their delivery mechanisms, we can optimize the resource consumption and move them in the 3-D space. Videos can be preprocessed in one location at a defined time. They can then be distributed to various locations for caching (CDNs) during network's off peak hours. An intelligent controller can monitor requests and analyze trends to prefill video caches with likely movies or ads to locations closer to the requesting end points. Although CDNs have been performing these operations on static web content for many years now, the real application for service provider is ability to transfer these video during off peak hours to maximize the network utilization. This not only increases the network link utilization, but also reduces the need for expensive investment to build new links or upgrading to new technology.

CHALLENGES WITH CURRENT ARCHITECTURES

There are several challenges with current architectures. Network appliances are specifically designed for a service and very slow to change due to limited marketability. New services using network applicances require higher upfront investments and time, so it is difficutly for MSOs to try out a business model without significant capital and operational investment. This contrasts to Web content model where web content providers can try out new applications relatively faster and cheaper.

Silos: Typically an operator's network is a group of silos that can be based on services, technology, vendors, or protocols. Voice, Video and Data are three main silos that have existed for a long time. Voice has slowly transitioned over to Data and the walls between these silo entities are slowly blurring. Next with IPTV, as packet based transport and services become the predominant service, it has become easier than ever to imagine a single platform delivering these services to end customer.

Then next level of silo is based on technology. In the access network, until now DOCSIS has been the dominant technology of choice that runs on a HFC network. The DOCSIS access gateways (CMTS/CCAP) have been vertically integrated with all Layer 1-2, many Layer 3 and above functionalities. Some MSOs are deploying fiber deeper into the network and in some cases all the way to the home. With fiber to the unit or building, Layer 2 based technologies such as EPON or GPON can be used in delivering the data services. In certain cases, particularly MDUs, where Cat6 cabling is present, operators will use of Ethernet in the building to deliver the data services.

As operators are foraying into new services that are subscriber aware, the goal will be offer and maintain the same features across multiple access network technologies. Traditional way of tightly coupling these services to underlying access technology is not desirable.

Clearly with the use of several technologies to deliver the same services, vertically integrated systems tailored for each access technologies can soon become a bottleneck. Abstracting and keeping certain higher layer services independent from the transport technology can help in reducing the time to roll out new features or services both from vendors and MSOs.

Service semantics: Next challenge is coupling "Services" with underlying technologies. QoS is often the most commonly used feature in the access network to differentiate services. Variables such as peak traffic rate, sustained traffic rate, maximum burst size, latency can be controlled and used as distinguishing features in a service. In a DOCSIS

environment, service flows are used to distinguish traffic flows or services. Typically voice would have two service flows, different data tiers will use different service flows. In EPON, use of LLIDs or VLANs can be used as service differentiators. In GPON, combination of Alloc IDs, TCONTs, GEM port IDs can be used to distinguish service treatment.

It remains an open question whether all or a subset of these variables will continue to be the key levers of choice for MSOs to differentiate services in the future. This is especially a relevant question when considering the consumer appetite for feature rich software applications and seamless services, MSOs access network evolution towards smaller serving group sizes coupled with very high peak to average bandwidth ratios in the last mile access. Of these variables, differentiating on latency is likely to be a key lever, considering the use of cloud computing. Today, many web and mobile applications (most mobile applications use HTML 5), are merely thin clients, presenting an immersing user interface but behind the scenes they extensively rely on remote API calls. Minimizing the response times to these API calls will become a critical service.

Although it is hard to predict but given the natural evolution towards smaller serving group sizes and use of high capacity access technologies, it is unlikely that the access network continues to be the choking point or as constrained a pipe as it used to be when DOCSIS was first deployed. Table 1 and 2 shows the Average bandwidth per sub for different access technologies. For instance, consider the average bandwidth per subscriber when we started with DOCSIS 2.0 (0.08 Mbps per sub – see Table 1) and where we are today (see Table 2) - it is a dramatic difference. In the case of DOCSIS, serving group sizes are between 250 and 500 (note: there is a long tail distribuation with larger and smaller serving groups). In case of PON based access network, this serving group ranges between 32 to 128 subscribers. Most MSOs are also considering similar serving group sizes for DOCSIS in the near future. Normalized ratio of capacity to rate tiers offered is nearly 6-12 times higher in just DOCSIS.

Table 1: Average capacity per sub (Mbps) in DS

	500 subs per SG	250 subs per SG
D2.0	0.08	0.16
D3.0 4-ch	0.32	0.64
D3.0 8ch	0.64	1.28
DOCSIS 3.0	1.28	2.56

Table 2: Average capacity per sub (Mbps) in DS

	128 subs per SG	64 subs per SG	32 subs per SG
EPON	7.81	15.63	31.25
GPON	19.53	39.06	78.13
10G EPON	78.13	156.25	312.50
DOCSIS 3.1 (1x96MHz OFDM block +24 ch)	19.00	38.00	76.00
DOCSIS 3.1 (2x192MHz OFDM blocks)	36	72	144

SDN AND NFV FRAMEWORK

In NFV based architecture, the idea is to have a pool of compute, storage and networking resources that can be molded by software to serve multiple services. Combined with virtual networking and SDN where the control functions are managed by the software running in a centralized controller to initiate flows to orchestrate the VNFs running on COTS or VM, we can create a fully agile framework for provider infrastructure the industry has ever had. This will enable operators to:

- Build a service or add any features that offer value to customers and add new revenue streams for operators.
- Try a new service quickly, adapt quickly, deployed at any scale
- Optimize utilization of resources (consider the peak to average ratio of compute)

As shown in figure 3, the access technology can be abstracted from the services layer. By doing this, access technology vendors can

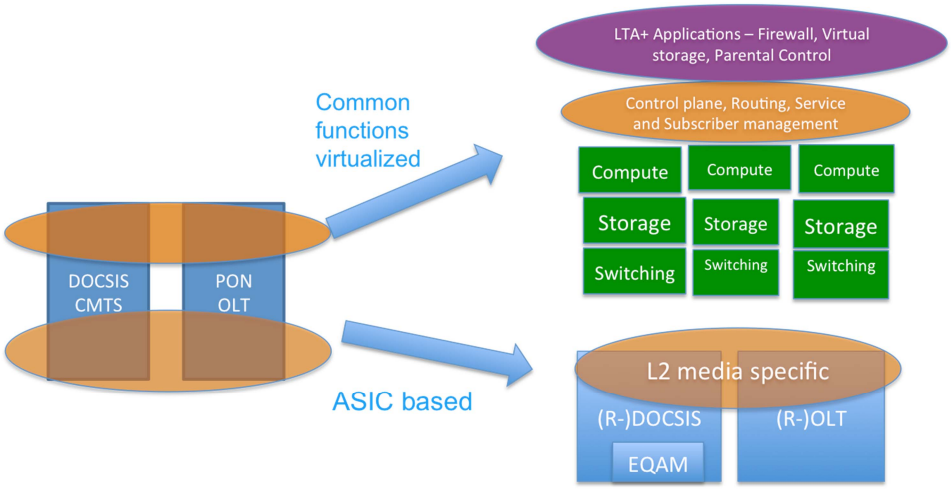

Figure 3: Separating Virtual functions in the Access Network

pack their ASIC based solutions effectively and only migrate their services, control/ management plane to a COTS platform. This will enable both operators and vendors to offer feature parity and reuse of certain components independent of the access technology thereby improving time to market for new services and features.

NFV offers a concept to migrate the hardware appliance based network functions to run on COTS and VM. Recall the example of network security application discussed in an earlier section, MSOs could leverage the NFV technology to run the network security on COTS and VM. The network security application is virtualized in Virtual Network Function (VNF) and deployed inside the network cloud infrastructure. MSOs can use SDN technology to control the flows from the customer to the network security VNF. This enables MSO to flexibly offer network security service in a much faster pace with little upfront investment and minimal physical changes. In comparison, the traditional model requires MSOs to deploy and integrate physical network appliances before serving customers.

<u>High Level Architecture for Virtual Home Network</u>

Figure 4 presents a high level architecture for a Virtual Home Network (VHN). Note, this is a use case study and will likely morph into a different architecture based on our ongoing learning and experiences. For details refer to [4]. This architecture includes:

- a Controller (VHNC) that contains user configurations and policies
- a set of Virtual Network Functions (VNF)
- a Packet Processor (VHNF) that performs functions such as a packet forwarding between multiple VNFs
- a VNF Manager (VNFM) managing the VNFs.

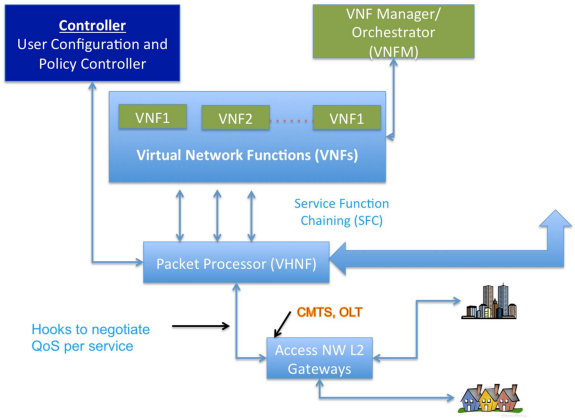

Figure 4: High Level Architecture for Virtual Home Network

There are five key interfaces for control, management and data path:

- The inteface between the Controller and the packet processor is used to exchange configuration and policy. For example, User A's living room TV has Committed Information Rate (CIR) set at 10Mb/s and must be protected by firewall function implemented in VNF.
- The interface between the Controller and VNFs is used to exchange configuration and policy. For example, VHNC could configure the firewall VHF to block any incoming ICMP messages to the User A's living room TV.
- The interface between VNF Manager and VNFs is used to exchange VHN management messages. For example, VNFM could instantiate a new firewall VNF when the current firewall VNF reaches certain capacity.
- The connection between the CPE and VHNF can be a common protocol agreed between CPE and VHNF. It could be Ethernet or any encapsulation technology such as GRE, PMIP or MPLS.
- The interface between the packet processor and VNF is defined by the Service Chaining Function protocol. IETF SFC WG is currently defining the specifications. VNF contains the service definitions and service logic. For example, Virtual Network Function 1 (VNF1) could be a parental control service and manage web filter rules configured by subscriber. Virtual Network Function 2 (VNF2) could be personal firewall that protects a home from botnet and intrusion. MSOs can scale VNFs horizontally to meet user demand. MSOs can also dynamically create VNF per subscriber only when the subscriber wants that service. For example, NSP initiates VNF1 for User X and VNF2 for User Y. In this model, MSO no longer updates CPE for service addition or modification. VHNC stores the user's service subscription. Each user may have different set of home services. For example: User A may have video service. User B may have VoIP service. VHNC contains the user's service subscription and interact with the VNF modules to provide proper services to users. VHNF is usually a device that is optimized for processing packets. It also implements the Service Function Chain function to forward user packets to proper VNFs. CPE is a simple access device that connects to the subscriber's devices at home to the MSO network.

Concerns on Virtual platforms and NFV:

A critically important metric is the packet performance, especially when VNFs are handling a lot of small packets. In a virtualized COTS based environment, there is a virtual switch or router that resides in the host operating system or hypervisor kernel to handle packets in and out of the server and in between Virtual Machines (VMs) in the same server. The virtual switches along with an SDN controller plays a big role in traffic steering and service chaining. Open Virtual Switch (OVS) is a good example that is widely adopted in NFV trials and Proof of Concepts (PoCs). However the packet performance using a purely software based OVS is considered inefficient and incurs high latency. Many vendors are working on a improved virtual switch that improves the efficiency of packet forwarding by up to 10x.

How fast can your physical network elements transmit and receive bit streams?

First, it is the raw speed of the servers and the switches connecting the servers to push packets in and out of the network. In a server, packet I/O is typically handled by a Network Interface Card (NIC). 10Gb/s NICs are very

common but for a high-performance NFV cloud, 40Gb/s or even 100Gb/s will be needed. There are new 25G and 50G standards being developed that provide a wide range of choices [3]. Inside the server, KVM and virtual switching layer play an important role (see question below).

How fast can your system handle packets?

As we currently envision, in a typical aggregation point such as a headend, each serving group in an access network will likely have 10Gbps Ethernet connectivity. Considering this as an example, on a 10G link the theoretical maximum rate of packet performance with 46-Byte packets and 1500-Byte packets is close to 15 million packets per second (pps) and 1 million pps respectively. A critically important metric here is the packet performance, especially when VNFs are handling a lot of small packets. As it turns out, it is a lot harder to achieve lossless small packet handling.

Table 3: Packet Processing Requirements

Packet Data Unit (PDU) in Bytes	Packet Handling Rate on a 10G link (Millions of packets/sec)	Calculations (38 byte IPv4 and Ethernet overhead)
46	14.881	10Gb/s / (84 B * 8 b/B)]
1500	0.813	10Gb/s/(1538 B * 8 b/B)
9200	0.135	10Gb/s/(9238 B * 8 b/B)

How fast can your system handle packets in a virtualized environment?

Even with a good foundation of raw throughput and packet performance, performance can be affected by the components that sit between the VM interface and the server's NIC, OS, and hypervisor. Multiple context switches and memory copies can happen before a received packet is delivered to the destined application (VNF), resulting in sub-optimal packet performance. Several solutions have been developed to over this limitation. With Data Plane Development Kit (DPDK), applications can link to to the DPDK library and call DPDK APIs to keep polling for packets. This eliminates interrupts for packet arrival. This solution can significantly enhance packet performance while still preserving hardware independence. But doing all packet processing in user space and constant polling still poses overhead to the CPU, impacting packet performance, latency, and performance of other applications on the same host.

Another alternative is to use hardware assist to bypass the kernel and use Single Root Input Output Virtualization (SR-IOV) to virtualize a single Ethernet port into multiple lightweight Virtual Functions (VFs) that can be associated with VMs. The communication between VMs and their corresponding VFs are through Direct Memory Access to eliminate lengthy copy operations.

SUMMARY

MSOs are seeing several transitions. We believe NFV and SDN can be used to realize synergies while addressing these transitions. NFV and SDN made a big promise to redefined networking and how MSOs manage and operate their network functions similar to managing and operating software applications. If they execute well to meet their promises, it will rationalize the current networking model, greatly simplify network management and operations. The ultimate success of this architecture will be when MSOs can start launching new applications and services with a short lead time, low upfront cost per user and ability to try out new services similar to the cloud service providers.

NFV and SDN are far from mature, but the potential benefits are enormous if the industry can work together and make it reality. A lot of progress has been made but more work is still required. Vendors are actively proving the feasibility and performance capabilities of running VNFs on COTS platforms. Architecturally, we need to define VNFs orchestration model and Service Function

Chaining among others. Operationally, both vendors and MSOs have to become familiar with the architectures, adopt new processes such as dev/ops to successfully transition.

References:

[1] Growth Architectures: Built to Last, Built to Launch, by Dr. Robert L. Howald. Published in Spring Technical Forum, NCTA 2014

[2] Mobile Apps and Revenue: http://techcrunch.com/2014/06/02/itunes-app-store-now-has-1-2-million-apps-has-seen-75-billion-downloads-to-date/

[3] New 25G and 50G: http://www.bitpipe.com/data/demandEngage.action?resId=1421690810_773

[4] IETF Draft - Problem Statements of Virtual Home Network: http://tools.ietf.org/html/draft-lee-vhs-ps-02, by Y. Lee and R. Ghai.