SCTE ISBE CABLE-TEC EXPO '16
SEPTEMBER 26-29 PHILADELPHIA

SCTE Society of Cable Telecommunications Engineers
ISBE International Society of Broadband Experts

SDN Is Not NFV

A Technical Paper prepared for SCTE/ISBE by

Mark DelSesto
Solution Architect
Hewlett Packard Enterprise
Mark.DelSesto@hpe.com

Table of Contents

Introduction

Today, we are deluged with the hype and promise for the IT'ification of our networks, driven by the lead technology of virtualization. The two pillars of virtualization, mentioned in almost all vendor product sheets, are NFV, network function virtualization, and SDN, software defined networks. While these technologies are almost always discussed together, they are two distinct and separate solutions for different problems. Either technology can be deployed independently, but there are major benefits when both are utilized together.

This paper and workshop will dive into SDN technology, its symbiotic relationship with NFV, and its impact on the virtualization of networks. Starting with a short overview of virtualization and software defined networking and moving on to cover examples of SD-WAN implementations using NFV techniques, SDN techniques, and the combination of the two.

SDN & NFV

Virtualization, from its simple beginnings enabling software to run on different hardware architectures, has now grown into one of the most disruptive forces in the software space. This original software portability combined with the increasing power of the Intel chipset and the decreasing cost of stock servers, fans the flames of disruption. These factors contribute to massive scaling capabilities, while at the same time allowing re-use and oversubscription of hardware resources. And there in lies the major benefits of virtualization – leveraging the ability to leverage COTS hardware for applications that previously required custom bespoke hardware enabling multiple applications to be run simultaneously, while operations teams manage a simplified set of hardware.

As we race towards virtualization and specifically Network Function Virtualization, Software Defined Networking is an integral part. While NFV is very much about virtualizing existing services, SDN is about simplification and the resulting service agility. SDN is geared towards removing logic from core networking elements and shifting this intelligence into applications. In some ways it can be referred to as BYON, bring your own networking.

In the pre-SDN world, applications required networking support for deployment, especially networking intensive applications ranging from access gateways to video delivery components or traffic optimizers. Significant planning and careful deployment processes were required to insert new networking applications, and yet there was always the risk of impacting other applications. The goal of SDN is separation of traffic so that applications can control their own networking requirements. With separation, applications are given the responsibility to manage their network needs without impact to others. One of the key themes in the original OpenFlow paper was the concept of experimentation, the ability to insert new functionality into the network without impacting others.

SDN has a number of characteristics that enable the separation of applications. First is the separation of the control plane from the data plane. In traditional networking there always has been this separation of control and data, however the deployed routers and switches were closed boxes that handled both. Because the box was closed, even though the two planes were separated, the boxes themselves locked

them together. This in turn impacted scaling as both control and data had to be scaled together, and it impacted the ability to manage the network. One of the first steps in the SDN paradigm is the separation of the control and data planes. The data plane, also referred to as the forwarding plane, is simply the set of devices designed for basic packet delivery. The control plane is where the forwarding logic is actually calculated and distributed. The control plane is the key middleware that marries the data plane with applications. On one side, the control plane gets requests from applications and on the other side, it instructs the forwarding plane how to classify and deliver packet flows.

Beyond just the architectural difference and overall openness, the other key tool that SDN uses in practice is overlay tunnels. By leveraging overlay tunnels, SDN can manipulate traffic and impact its destination without altering the actual flow. While overlay tunnels have been used extensively in networking for many years, and are not new, SDN relies on this technology to simplify routing decisions, and enable a light weight forwarding plane.

SD-WAN & Service Function Chaining

While many service providers today look to virtualize or cloudify existing services, the typical approach is initially virtualizing the service as is. The challenge is that for functions that have been deployed for years, or even decades, the shear amount of knowledge contained in the underlying software makes it extremely difficult to modify. Feature improvements, even simple ones, takes months if not years to supply, and even longer to deploy. Beyond software changes, configuration changes themselves can have such wide ranging impacts that they themselves must be carefully considered and well managed for successful deployment.

An alternative approach to simple virtualization is decomposition of the service. This is where functionality is decomposed into separate smaller chunks, similar to the use of micro-services in the web world. One “box” that is ripe for this type of decomposition is the PE router. The PE router in some ways has provided network virtualization for many years. The PE router participates in each enterprises WAN but separates all the traffic and information base for each enterprise. This separation is exactly the same type of technique used within an SDN approach. The challenge for the PE router is the amount of information that must be learned per network, and then stored, and the complexity of changes on the PE router when it contains this much state information. The criticality of the PE router in each of the enterprises it services, as well as the all the state information it contains, leads to very long lead times for changes, either in software, or just configuration due to stringent service level agreements (SLAs) declaring strict maintenance windows and maintenance notice requirements. Any change that impacts the PE device is certain to require planning and attention to detail. It is exactly this type of challenge that the combination of SDN and NFV techniques can excel. Focusing on a use case related to the PE router will highlight the approaches of each technique and hopefully give a good example of how to apply these techniques to additional use cases.

The SD-WAN use case generates significant excitement in the service provider realm due to a few key aspects:

1) SD-WAN generates revenue

Providing the ability to interconnect branch offices together, as well as connectivity to a central data center remains a use case with strong demand for medium and large size businesses. While more and more corporate applications are transitioned to cloud based services, there are still significant drivers for enterprises to remain directly connected with layer 2 or layer 3 VPN interconnections. This type of service is above and beyond basic internet connectivity and drives a significantly higher premium, approximately on the order of 25x more. Corporations pay this amount for the perceived benefits of QoS and stringent SLAs.

2) WAN deployment times are excessive

Unfortunately for operators there is a weak link for this use case and that is the time to deploy. Generally VPN service requires the configuration of MPLS circuits for interconnecting the various sites of an enterprise. The MPLS configuration requires changes on the Provider Edge (PE) router, and this is a heavy weight instance with significant history. Configuration changes are intensive processes that require highly specialized knowledge and strict adherence to processes to avoid unintentional outages. All together 30-60 day time periods for the connection of a site to a VPN are average. Even customers that pay express setup fees are still subjected to lengthy wait times on the order of weeks. As you can see, the combination of a high price differential over and above ordinary internet service, combined with a significantly delayed installation time makes this a use case ripe for an overhaul, as well as unfortunately attracting attention from 3rd parties.

3) Over-the-Top VPN solutions are feasible

Clearly 3rd parties have turned their sites to offering enterprise VPN services. There is a wide difference in price between basic internet connectivity and VPN service, implying there is significant room for alternate providers to offer an in between level of service. In addition, these alternate providers can leverage newer technology as well as a nimbleness that service providers configuring MPLS tunnels do not have. Therefore the competitive threat of an over-the-top provider grabbing the VPN business of customers within a service provider's footprint is significant. This in turn drives a significant focus by operators to make their VPN services operationally more efficient.

1. Virtualization Solution

In this environment of NFV, the conventional wisdom is that the first step to enhancing operational efficiency is to virtual the PE router. This is the primary source of configuration delays and migrating this entity to more of a cloud-based instance with scale-in/scale-out capabilities is very appealing. In addition, PE routers have been deployed for many years, and simply virtualizing the existing hardware-based design is a solid way to incrementally move the technology forward but maintain the years of knowledge and experience in the product. Unfortunately, the real challenge that must be solved for the VPN use case is to ensure advances in configurability leading to faster VPN deployment operations. The configuration issues generally stem from all the state management of the PE router, and it's ability to keep this information correctly separated. One possible step in the right direction is to assign specific PE

instances to specific instances. Whereas in the past one large monolithic PE had to support all customers within a geographic area, in the virtual world, the number of PE instances can be increased in theory to one-to-one PEs to enterprises. This type of multiple instances model solves the information and configuration challenge as modifying a PE instance dedicated to a single enterprise has no impact upon others. Signing up new customers only requires the creation of a new PE instance.

While a PE instance per customer potentially solves the configuration timing and complexity, it does introduce other challenges. First there is the management of potentially hundreds or thousands of PE router instances. This requires a new level of orchestration that muse be implemented beyond how the PE is managed today. Another significant challenge is compute power. PE routers are generally deployed close to the edge in POP sites, typically with space and power limitations. Preparing many sites to receive an increase in the number of virtual PE instances requires enough COTS hardware. But distributing this hardware to many sites is expensive, and most of the hardware is unused until new customers are activated. If enough hardware is not deployed, then when a new customer is signed, the POP location needs additional hardware, which defeats the agility goal of virtualizing the PE router to begin with. To truly make this transition work, the virtual PE router function has to be pulled to a more central location – this enables early deployment of hardware in advance of demand, and aids OpEx costs of maintaining the hardware.

This in turns leads to another challenge. Traffic for the PE router is designed to flow to the distributed POP sites. To bring the traffic from the POP locations back to the centralized cloud locations, an additional technique must be leveraged - overlay tunnels. Here is the first intersection of SDN with NFV. Using SDN techniques, tunnels from the CE devices to the vPE can be orchestrated and connecting the customer premise directly to the operators cloud data center. At this stage, the PE router has been virtualized and migrated to a more centralized data center. CE traffic leverages tunnels to get routed from the enterprise to the centralized PE instances. The PE router still needs to maintain the same amount of state as previously done, however the use of PE instances per enterprise can potentially minimize this aspect.

2. SDN Solution

Given that in the end to virtualize the PE, we needed SDN overlay techniques to make it possible, let us look at what a deployment that leads with SDN technology looks like instead. Instead of virtualizing the PE, we can instead leverage SDN to remove the VPN application from the PE and instantiate it as part of a networking application.

Again one of the primary challenges with the PE router itself is simply the amount of state information it must manage about all the tenets and the tenets sites. This information model makes it difficult to perform configuration changes as it has the potential to impact many enterprises. Regardless of approach, the solution to faster SD-WAN deployments is the simplification and separation of the information model, so that configuration changes can be isolated with impacts only to the appropriate tenet.

For an SDN deployment, an SDN controller is needed, along with a SD-WAN (VPN) application requesting service from the SDN controller. The key is to minimize the centralized storage of information that makes future configuration changes expensive due to potential impacts to multiple enterprises.

To start with an SDN approach, the CE first needs to get IP connectivity. The solution will use tunnels to route traffic to the appropriate destination, but at a minimum the CE needs to begin with an IP address and we can continue to leverage the PE router for this function. Note that this will be the only piece of information that the PE router needs to maintain for a site within an enterprise. There is no need for MAC learning of entities behind the CE, or participation in BGP traffic to learn enterprise subnets. Going even further, there is no need to separate enterprises from each other as the PE router is only maintaining the underlay network. The SDN controller and the overlay tunnels will enforce enterprise separation, and assist with MAC/subnet learning, which obviously greatly simplifies the functionality of the PE router (actually given these changes, it is not really a PE router anymore).

With basic connectivity in place, the SDN controller and the VPN application on top of it can now begin the process of inter-site connectivity. From a service provisioning perspective there has to be some type of service orchestrator that is used to create the WAN service to begin with (i.e. layer 2 service, layer 3 service, point-to-point, point-to-multipoint, etc..). The orchestrator interfaces to the VPN application which in turn signals to the controller. The application converts service requests into tunnel creation requests, which are then delivered eventually into the deployed CE devices. The CE devices create overlay tunnels, much in the same way as the tunnels were created for a centralized virtual PE router. The major difference is the CE devices create CE-to-CE direct tunnels. For example if an enterprise needs connectivity for 5 sites, then a full mesh of tunnels is created with each CE supporting a direct tunnel to the other 4 neighboring sites.

The final piece is the networking information required per enterprise. Here each CE must learn the MAC addresses associated within the enterprise network. This is done with a standard MAC learning approach. For any packet arriving at the CE, the CE must classify and determine which port the packet should egress on (upstream, or downstream in the simplest case), and for upstream traffic, which tunnel or tunnels should it use. If the CE does not yet have a rule for the packet, then it is sent to the controller. For MAC learning, the lack of a rule indicate the CE has yet to see this MAC address before, and sending the packet to the controller just enables the controller to create a future rule for handling that MAC address again. Subnet learning is slightly more involved. One approach here is to use a vRouter, either centralized, or potentially distributed to the CEs, which participates in the enterprise BGP traffic. As subnets are advertised this information is shared with the controller which distributes the information to the various enterprise CEs. The use of the virtual router is solely for the purpose of learning enterprise subnets, and it is not actually used for routing per se.

Summarizing, the key change from the monolithic PE approach is that all the state information to enable an enterprise WAN is distributed to the edge and into the CE devices. The CE device itself stores all the rules necessary to enable the WAN function, and the CE leverages an SDN controller when new rules are required. The SDN controller can have many instances, potentially up to 1 instance per CE device, and all can leverage a shared distributed database. While the information stored within the SDN controller and the distributed database is similar to all the data the monolithic PE stores, the database storage has been separated out from the actual forwarding plane. Changes can occur on the database itself, or more likely specific instances running the database, with no impact to the underlying forwarding plane with the data is actually used. In addition another benefit is that enterprise traffic flows directly site to site without having to traverse through the operators data center. While core networks generally have sufficient bandwidth, this also represents a significant savings over the virtualized PE approach.

This distribution of the networking information all the way to the edge is the largest paradigm shift with the SDN approach. It is this change which in the end simplifies the service providers network, and enables the agility required to deploy services faster.

3. Moving Beyond Just SD-WAN

With the SDN tunnel approach in place, it then becomes relatively straightforward to introduce multiple connectivity links for redundancy purposes. Specialized CEs with alternate broadband capabilities from DSL to 4G connectivity can then be leveraged. Separate tunnels across the different connectivity options are created, and the CE can then implement an appropriate selection algorithm that either balances the traffic across the different tunnels, or uses them in a more traditional active/standby fashion. While the routing complexity of this arrangement on a PE is significant, with an SDN approach the heavy lifting is shifted to the CE itself minimizing any complexity the network itself must deal with.

Beyond additional links, additional network functions are also high value add-ons to a VPN service. Generally the number of boxes sitting between the customer's internal network and the providers WAN network is actually expanding. Functions ranging from WAN optimization, intrusion detection and prevention, malware detection, as well as content filtering are all functions with high value to businesses today. Before virtualization, these instances were almost always implemented with yet another box. Boxes are expensive not just due to their physicality, but also the truck roll and down time required to install. Clearly this is an area where virtualization and NFV can excel. Layered along with SDN so that virtual functions can be inserted seamlessly simply by adding new overlay tunnels for the interconnection.

Conclusion

NFV and SDN are both important new technologies that are leading to many disruptions in the approach service providers can take when expanding their networks, improving the agility, and reducing the CapEx /OpEx. However there are specific realms that each technology excels at and the right use cases need to be matched with the right technology. Virtualization tends to excel when paired with use cases requiring increased utilization. Virtualization enables concepts such as scale-in/scale-out, increasing the number of instances, and potentially altering redundancy schemes to unlock additional capacity. SDN on the other side is very much about streamlining networking changes when deploying new applications. SDN is about service agility and shifts networking requirements to applications instead of applications imposing requirements on the network. This paradigm shift forces applications to take on networking control, but at the same time hands applications the keys for faster rollouts and rapid changes in the future. When paired together these two techniques can create decomposed applications that are much simpler and more highly focused that what has existed to this point. The key is to leverage the strengths of both of these technologies and to work on creating the right decomposition of services going forward. Simply moving software that previously ran on custom hardware to now run in a VM on a COTS server is a great first step, but it is not a final step. Likewise, simply leveraging an overlay tunnel to enable traffic to terminate inside a data center instead of a distribution center is a piece of an SDN approach but not the complete picture of what is possible. In summary, service providers have a number of new disruptive tools at their disposal. It is important to see that these tools can be used to create radically new models of existing

services but that requires services to be decomposed into simpler components and then rebuilt. Just migrating legacy services into the SDN or NFV world may not get all the benefits that are expected.

Abbreviations

CE Device	Customer Edge Device
COTS	Commercial Off-The-Shelf
NFV	Network Function Virtualization
PE Router	Provider Edge Router
SDN	Software Defined Networking
VM	Virtual Machine
vPE	Virtualized PE router

Bibliography & References

D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky and S. Uhlig, "*Software-Defined Networking: A Comprehensive Survey*," in *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14-76, Jan. 2015.

ONUG SD-WAN Working Group. "ONUG Software-Defined WAN Use Case." Open Network User Group., Oct. 2014.

McKeown, Nick, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. "OpenFlow: Enabling Innovation in Campus Networks." Open Networking Foundation, 14 Mar. 2008.

