Leveraging SDN and NFV in the WAN

Introduction

Software Defined Networking (SDN) and Network Functions Virtualization (NFV) are two of the key components of the overall movement towards software defined IT[1]. Other components include virtualized servers and software defined storage as well as the management and security functionality that is required in a software focused IT environment.

A key enabler of the movement to software defined IT is that in a growing number of instances, commercial off the shelf (COTS) hardware platforms have enough compute power that proprietary hardware-based solutions are no longer necessary. Some of the drivers of the movement include the desire to:

- Have programmatic interfaces that enables applications and orchestration systems to program the network and request services from it;
- Move network functionality off of a lengthy hardware life cycle and onto a shorter software lifecycle;
- Increase agility by being able to dynamically create and move IT resources and by automating as much functionality as possible;
- Move away from proprietary solutions.

One goal of this white paper is to provide insight into SDN and NFV: What are they? What's driving them? How are they related to each other? Another goal of this white paper is to describe the value that SDN and NFV bring to the WAN.

Software Defined Networking

What is SDN?

The Open Networking Foundation (ONF)[2] is the group that is most associated with the development and standardization of SDN. According to the ONF, "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions."

According to the ONF, a SDN is:

- **Directly programmable**: Network control is directly programmable because it is decoupled from forwarding functions.
- **Agile**: Abstracting control from forwarding lets administrators dynamically adjust network-wide traffic flow to meet changing needs.

[1] Software Defined Everything (SDX) is another phrase that is commonly used to describe the broad based movement away from hardware-based solutions and towards software-based solutions.
[2] https://www.opennetworking.org/

- **Centrally managed**: Network intelligence is (logically) centralized in software-based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch.
- **Programmatically configured**: SDN lets network managers configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs, which they can write themselves because the programs do not depend on proprietary software.
- **Open standards-based and vendor-neutral**: When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.

What's driving the interest in SDN?

The *2015 Guide to SDN and NFV*[3] (The Guide) reported on the results of a survey that was taken by 246 IT professionals in late 2014. The survey respondents asked to indicate their company's interest in SDN. Their answers to that question indicates that the interest that enterprise IT organizations have in SDN has grown considerably over the last year and is likely to grow considerably over the next year.

In order to understand what is driving the interest in SDN, the survey respondents were asked to indicate which challenges and opportunities they thought SDN could help them to respond to. Their responses indicate that IT organizations are optimistic that SDN can help them respond to a wide range of opportunities and challenges. At the top of the list are:

- Better utilize network resources;
- Perform traffic engineering with an end-to-end view of the network;
- Ease the administrative burden of configuration and provisioning.

Because SDN introduces new functionality into the network, at the very bottom of the list is reducing complexity.

Where will SDN be deployed?

While the initial discussion of SDN focused primarily on implementing it in the data center, in the last year or two there has been a lot of discussion about implementing SDN in campus networks as well as in Wide Area Networks (WANs). In order to understand where SDN will likely be implemented, the survey respondents were asked to indicate how broadly they expected their campus, WAN and data centers networks would be based on SDN three years from now. Their responses (Table 1) show that IT organizations believe three years from now that SDN deployment in data centers will be highly pervasive and that there will also be significant SDN deployment both in the WAN and in campus networks.

Table 1: Planned SDN Deployment

	Campus Networks	**WAN**	**Data Center Networks**
BASED EXCLUSIVELY ON SDN	1%	2%	6%
Mostly SDN	10%	6%	20%

[3]http://www.ashtonmetzler.com/

Hybrid, with SDN and traditional coexisting about equally	34%	36%	50%
Mostly traditional	29%	31%	10%
Exclusively traditional	13%	13%	4%
Don't know	12%	12%	10%

What is a Software Defined (SD) WAN?

A group that is working to advance the application of SDN in the WAN is the Open Network User Group's (ONUG's) SD-WAN Working Group. That group wrote a white paper entitled the *ONUG Software-Defined WAN Use Case* which is available from their web site upon registration[4]. That white paper identified 10 business requirements that a SD-WAN must meet. Those requirements are listed in the appendix to this white paper. Some viable requirements for a SD-WAN that are not included in the ONUG white paper include the:

- Inclusion of optimization functionality;
- Ability to route traffic to the appropriate cloud services provider based on characteristics of those providers;
- Portal based control of QoS and bandwidth allocation;
- Sophisticated analytics into the network, the network functions and the applications;
- Automation of the branch office functionality that is enabled by NFV.

As is the case with any software defined network, a SD-WAN centralizes the control function into a controller which abstracts the network specific policies and exposes application-based business policies to the user. Leveraging the underlying WAN platforms, which may include physical or virtual routers, the controller sets up virtual overlays that are both transport and technology agnostic. Under the direction of the controller, the WAN platforms implement functionality such as quality of service, path selection, optimization and security, often using dynamic multi-pathing over multiple WAN links. The WAN platforms also provide deep visibility to application performance and export information to the controller or other collectors.

What is the value of a SD WAN?

As was correctly pointed out by the survey respondents, the centralized processing associated with a SDN means that optimum routes can be calculated for each flow of WAN traffic by understanding the required service levels, having knowledge of the performance of each WAN link and by leveraging a complete model of the end-to-end topology of the network. Given the programmatic interfaces that are a key component of a SDN, it is also possible to include the existence and severity of security events into the calculation of optimum routes.

As was also correctly pointed out by the survey respondents, a SDN enables better utilization of network resources. This fact was demonstrated by Google's G-Scale WAN, which is the WAN that links Google's various data centers and which was the first use of SDN in the WAN to

[4] http://opennetworkingusergroup.com/archive/fall-2014/sd-wan-working-group/

gather a lot of attention. Google has identified a number of benefits associated with its G-Scale WAN, including that Google can run the network at utilization levels up to 95%[5].

There are many other benefits of leveraging SDN in the WAN. This includes the ability to:

- Provide bandwidth on demand;
- Compute in advance a set of fail-over routes for each possible link or node failure;
- Implement new services faster;
- Provide users with the ability to dynamically control their network.

How will SD WANs be implemented?

An enterprise network organization can implement a SD-WAN as a private network[6] or acquire WAN services from a service provider who operates a SD-WAN. One of the factors that network organizations typically use to decide between implementing a private network and using services from a carrier is the level of complexity. An example of that choice is MPLS. Due to the complexity of MPLS, the vast majority of network organizations have chosen to avoid implementing a private MPLS network and to acquire MPLS services from a carrier. Given the complexity that the survey respondents associated with SDN, it is highly likely that many, if not most enterprises will acquire SD-WAN based services from a carrier.

Network Functions Virtualization

The group that is most closely associated with the development of NFV is the Industry Specifications Group for Network Functions Virtualization that was formed under the auspices of the European Telecommunications Standards Institute[7] (ETSI NFV ISG). Some of the key characteristics of the ETSI vision for NFV include[8]:

- Achieving high performance virtualized network appliances which are portable between different hardware vendors and across different hypervisors.
- Achieving co-existence with hardware based network platforms.
- Managing and orchestrating many virtual network appliances while ensuring security from attack and misconfiguration.
- Implementing automation to enable the scalability of the solutions.
- Ensuring the appropriate level of resilience to hardware and software failures.

As defined by ETSI, NFV is applicable to all data plane packet processing and control plane functions in both fixed and mobile networks. In a NFV environment, a Virtual Network Function (VNF) is responsible for handling specific network functions that run on one or more virtual machines (VMs) on top of a COTS-based infrastructure. A classification of the virtual network functions (VNFs) that are associated with NFV is shown in Figure 1.

[5] https://www.opennetworking.org/images/stories/downloads/sdn-resources/customer-case-studies/cs-googlesdn.pdf
[6] In this context, a private network is one in which the enterprise network organization does the bulk of the associated planning, designing, implementing and ongoing management.
[7] http://www.etsi.org/news-events/news/644-2013-01-isg-nfv-created
[8] https://portal.etsi.org/NFV/NFV_White_Paper.pdf

Network Element	Function
Switching elements	Broadband network gateways, carrier-grade network address translation, routers
Mobile network nodes	Home location register/home subscriber server, gateway, GPRS support node, radio network controller, various node B functions
Customer premises equipment	Home routers, set-top boxes
Tunneling gateway elements	IPSec/SSL virtual private network gateways
Traffic analysis	Deep packet inspection, quality of experience measurement
Assurance	Service assurance, service level agreement monitoring, testing and diagnostics
Signaling	Session border controllers, IP Multimedia Subsystem components
Control plane/access functions	Authentication, authorization and accounting servers, policy control and charging platforms
Application optimization	Content delivery networks, cache servers, load balancers, accelerators
Security	Firewalls, virus scanners, intrusion detection systems, spam protection

Figure 1: Classification of VNFs

It will be a while before all of the functionality listed in Figure 1 is available as a VNF. However, some of that functionality, such as routers and firewalls, is already available as a VNF.

What is the relationship between SDN and NFV?

The conventional wisdom in the IT industry in general, and on the part of the ONF and the ETSI NFV ISG in particular, had been that SDN and NFV were separate topics and didn't need to be formally coordinated. That conventional wisdom officially changed in March 2014 when the ONF and the ETSI NFV ISG announced the signing of a Memorandum of Understanding (MOU). As part of the announcing the MOU[9], the ONF and ETSI said that "Together the organizations will explore the application of SDN configuration and control protocols as the base for the network infrastructure supporting NFV, and conversely the possibilities that NFV opens for virtualizing the forwarding plane functions."

In a recent white paper[10] ETSI made the following comments about the relationship between SDN and NFV. "NFV creates a very dynamic network environment, driven by customers needing on-demand services and operators needing to manage utilization and performance of services. Tenant networks will come and go, and VNFs and their connectivity will change frequently to

[9] http://www.rethink-wireless.com/2014/03/19/etsi-nfv-group-closer-operator-sdn.htm
[10] http://portal.etsi.org/NFV/NFV_White_Paper3.pdf

balance load across the infrastructure. The capability to programmatically control network resources (through a centralized or distributed controller) is important in an era of continuous change."

What are the primary NFV use cases?

The ESTI NFV ISG has identified 9 potential use cases for NFV. A thorough description of the use cases is available on the ETSI web site[11].

It will be a while, if ever, before all of the 9 use cases are broadly implemented. However, one of the use cases, Virtual Network Functions as a Service[12], is already being implemented. This use case is targeting the fact that many enterprises are deploying numerous network service appliances at their branch offices. Network services commonly installed at the branch can include access routers, WAN optimization controllers, stateful firewalls, intrusion detection systems, and DPI analysis devices. If a number of these functions are implemented on dedicated physical appliance platform, the result can often be a complex, expensive, and difficult-to-manage branch office network.

An alternative solution for enterprise branch office networks is to subscribe to VNFs that are hosted on servers either in the network service provider's access network PoP or which are provided on the customer's premise. VNFs delivered as a Service are analogous to cloud networking SaaS applications where the subscriber pays only for access to the service and not the infrastructure that hosts the service. The subscriber is also freed of any responsibility for the planning, designing, implementing or ongoing management of the functionality.

Next Steps

Perhaps the most important step for network organizations is to realize that after a long period with little if any fundamental innovation, the WAN is now the focus of considerable innovation. As a result, for the first time in a decade network organizations have an opportunity to make a significant upgrade to their WAN.

The next step is for network organizations to either begin or to extend the analysis that they have done of SD WAN solutions, whether those solutions are private or carrier based. Initially at least, this analysis should be done in a passive way – such as talking to vendors and carriers. However, in order to accurately assess the impact of a SD-WAN, network organizations will need to adopt a more aggressive approach to the analysis and implement a trial of a SD-WAN. The goal of the trial is to determine if the advantages of implementing a SD-WAN that were described in this white paper are achievable and if there are any operational considerations that would prevent a broader implementation of a SD-WAN.

Appendix

The following are the 10 business requirements that the ONUG's SD-WAN Working Group identified that a SD-WAN must meet.

[11] http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf
[12] This use case is also commonly referred to a virtual CPE.

1. Ability for remote site/branch to leverage public and private WANs in an active-active fashion for business applications.
2. Ability to deploy CPE in a physical or virtual form factor on commodity hardware.
3. A secure hybrid WAN architecture that allows for dynamic traffic engineering capability across private and public WAN paths as specified by application policy, prevailing network WAN availability and/or degradation at transport or application layer performance.
4. Visibility, prioritization and steering of business critical and real-time applications as per security and corporate governance and compliance policies.
5. A highly available and resilient hybrid WAN environment for optimal client and application experience.
6. Layer 2 and 3 interoperability with directly connected switch and/or router.
7. Site, Application and VPN performance level dashboard reporting.
8. Open north-bound API for controller access and management, ability to forward specific log events to network event co-relation manager and/or Security Incident & Event Manager (SIEM).
9. Capability to effect zero touch deployment at branch site with minimal to no configuration changes on directly connected infrastructure, ensuring agility in provisioning and deployment.
10. FIPS 140-2 validation certification for cryptography modules/encryption with automated certificate life cycle management and reporting.