THE RIGHT SDN IS RIGHT FOR NFV

STRATEGIC WHITE PAPER | NFV INSIGHTS SERIES

To maximize the potential of Network Functions Virtualization (NFV) and open up new opportunities, such as service chaining, the network must be made as dynamic and programmable as the virtualized network functions. SDN is an essential and complementary technology that enhances the advantages of NFV, but the SDN implemented must be tailored to an operator's specific NFV requirements. The degree of integration between SDN and NFV can vary, bringing different degrees of reward to the service provider. This paper explores the value a carefully-chosen SDN can bring to NFV. It also explores the four stages of integration, which provide increasing levels of operational optimization.

About the NFV Insights Series

NFV represents a major shift in the telecommunications and networking industry. NFV applies virtualization and cloud principles to the telecommunications domain, something that appeared to be impossible until recently due to the stringent performance, availability, reliability, and security requirements in communication networks. Many service providers are now keen to implement NFV to help them gain an advantage through automation and responsiveness to deliver an enhanced customer experience and reduce operational costs. This series of white papers addresses some of the key technical and business challenges on the road to NFV.

TABLE OF CONTENTS

WHY SDN FOR NFV?

In October 2012, a group representing leading network operators around the world prepared the seminal paper that defined the benefits, enablers and challenges of Network Functions Virtualization (NFV) [1]. As the authors wrote, NFV and Software-Defined Networks (SDN) are different and complementary. SDN can be deployed without NFV, and NFV can be deployed without SDN. Nevertheless, SDN will be a critical component in the majority of NFV deployments. The purpose of this paper is to understand how NFV and SDN can be combined to provide optimum value.

In the past, the network was semi-static. Any changes to it had to be made manually through command line interfaces or a variety of management systems. Manual changes were error prone and only specially trained and dedicated staff was authorized to make carefully documented changes. For example, carriers use pre-defined fixed addressing according to defined policies. The design of a detailed network and IP addressing architecture alone can take weeks. Deploying a new network element in this environment involves following stringent methods of procedure and rigid rules to avoid conflicts with existing installations.

This is where NFV and SDN can bring significant advantages to enable new services, improve service request response times and react faster to changing services. Currently such changes can take weeks or months. SDN virtualizes networks in a way similar to the server virtualization for compute and storage. SDN provides easy-to-use network abstractions with open northbound APIs — such as OpenStack® Neutron — to allow a wider group of people as well as automated systems to provision and configure networks. This cuts down response times to minutes or seconds.

However, the introduction of NFV and SDN is not only a change in technology; it also requires a change in mindset and procedures. Service providers will need to accept a much higher level of autonomous behavior in their systems making it possible to introduce network-wide changes within minutes. Such changes can be the deployment of virtualized network functions (vNFs), software upgrades to enable new features, service scaling, realigning of network resources, and others.

The degree of integration between SDN and NFV can vary, which will in turn bring different degrees of reward. An advanced SDN and NFV solution offers a high level of integration, with a declarative model of network configuration for the vNFs. Network configurations will automatically follow the changing needs of vNFs and authorized persons will be able to track the owner of a service and the reasons why certain configurations exist.

WHAT IS SDN?

Software-Defined Networking (SDN) is a new approach to designing, building and managing networks. The basic concept is that SDN separates the network's control and forwarding planes to make it easier to optimize each. A controller acts as the brains, providing an abstract, centralized view of the overall network. Through the controller, network administrators can quickly and easily make and push out decisions on how the underlying systems (switches and routers) of the forwarding plane will handle the traffic.

WHAT IS NFV?

With NFV, network functions are virtualized and liberated from the need for dedicated hardware. Virtualized network functions run on common off-the-shelf servers, storage and network infrastructure. The infrastructure in effect becomes a horizontal layer shared by many network functions. The benefits of NFV are a cost-effective high-volume hardware infrastructure, reduced operational costs due to automation, and high service agility.

CRITICAL NETWORK REQUIREMENTS FOR NFV

NFV introduces cloud practice into service provider networks. Network functions become virtualized and automated to run on a shared server infrastructure that provides the necessary compute, storage and network resources. However, network functions are more demanding than most IT applications. Figure 1 illustrates the key areas of NFV-specific requirements.

Figure 1. NFV requirements

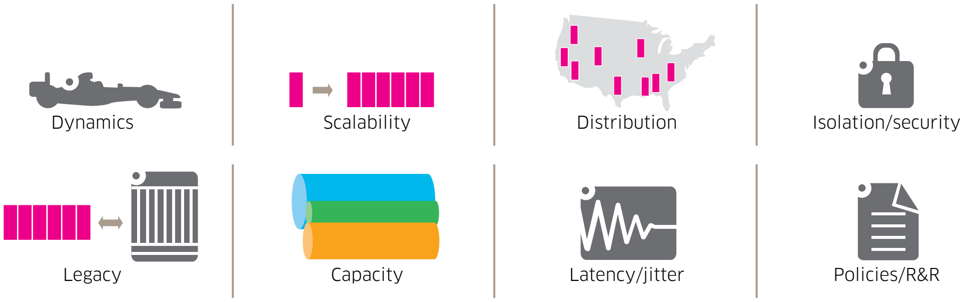

Dynamics and scalability. NFV infrastructures need to be dynamic. They need to support highly scalable applications that can respond to changing service uptake. When a vNF is scaled out or moved to a different location, the networks need to follow without manual intervention.

Connectivity in a distributed environment. The role of the network in NFV — SDN or not — is first of all to provide connectivity between vNF components (vNFCs). Most NFV applications require Layer 2 and Layer 3 connectivity. For some applications, Layer 1 and Layer 0 network control may also be needed (transport SDN). SDN networks provide static or dynamic IP addressing, floating IP addresses (public IP addresses dynamically assigned to one or more ports), multicast/broadcast/anycast, as well as middlebox services (load balancing as a service (LBaaS), firewall as a service (FWaaS) and virtual private networks as a service (VPNaaS)).

While IT clouds strive to centralize and consolidate data centers, NFV nodes need to be carefully distributed throughout a geographic coverage area to guarantee performance and high availability, and to avoid unnecessary backhaul of traffic to centralized data centers [2].

Security. The network needs to provide a level of security, for example, to restrict connectivity to those elements that are supposed to talk to each other, and to allow only legitimate data traffic (through firewalls and security groups). vNFs need to be sufficiently isolated from any “noisy” neighbors for performance and security reasons. Security from any type of external and internal attacker is a fundamental requirement of any carrier infrastructure [3].

Legacy interworking. The introduction of NFV to the network will be a gradual process, taken in a stepwise manner. Interworking with legacy networks will therefore be critical to ensure uninterrupted services throughout the evolution to a fully NFV-based infrastructure.

WHAT IS AN NFV PLATFORM?

An NFV platform provides everything needed to run and manage virtual network functions efficiently. This includes server, storage and network resources distributed across the network geography. These resources are managed as a cloud – a pool of resources that can be flexibly allocated to network functions. It also includes tools to automate the management of the hardware resources as well as the lifecycle of the virtualized network functions including onboarding, deployment, scaling, and healing.

Capacity and reliable performance. vNFs often support high performance data and media traffic meaning that sufficient bandwidth and packet throughput needs to be available both across the wide area network and between the server network interface cards (NICs), hypervisors and virtual switches. Latency and jitter sensitivity is also a major consideration for many network functions with real-time performance requirements. The network needs to guarantee service availability in case of failures and force majeur disasters.

Policies and changing roles and responsibilities. Today, many services are supported by service silos, that is, each service comes with its own hardware, software and operational team. With NFV, service providers deploy a more horizontal model where the NFV platform with its compute, storage and networking resources is a common layer that doesn't need to be duplicated for each service. This changes the roles and responsibilities of the operational teams. NFV and SDN will be much more policy-driven to ensure a coherent operational model and better automation.

USE CASES

SDN can play several different roles in an NFV network. Two notable examples are the virtual backplane and service function chaining (SFC).

The virtual backplane

In an NFV environment, the network needs to take on an additional task that was solved with special hardware for physical network functions. The larger physical network elements consist of a number of processing blades and interface cards connected via a backplane with a switching fabric. As these network elements become virtualized (vNFs), blades and cards are mapped to components (vNFCs) on virtual machines (VMs) running on the same or different servers or even distributed across different data centers (see Figure 2).

Figure 2. Backplane connectivity is replaced with network connectivity

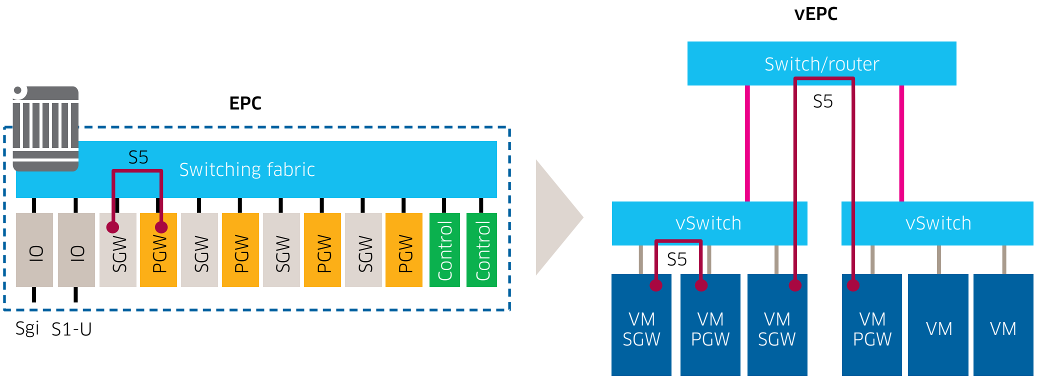

As shown in Figure 2, in a virtual Evolved Packet Core (vEPC) functions that are traditionally housed inside a router platform may now be virtualized into separate VMs running on different servers. These functions include Serving Gateway (SGW), Packet Data Network Gateway (PGW), control, and termination. Internal communication paths now need to be mapped to (virtual) networks between servers or even across the wide area from one data center to another.

For physical network functions, packet performance is guaranteed and optimized by tuning the hardware and software together. Within a virtualized environment, the management of traffic flows between functions is a new task, which will be discussed below.

Service function chaining and virtual CPE

SFC has been widely discussed in the industry as an important application of NFV and SDN. It allows service providers to create and sell packages of value-added services by adding them dynamically to the customer's data path (see Figure 3). A service chain is a sequence of vNFs spliced into the data path between a traffic source and a traffic sink. For example, the blue service chain includes a firewall and a video optimizer. The red chain also includes the firewall but adds anti-virus and parental control functions. Service chains can be applied to fixed broadband networks and mobile networks. Virtualized customer premises equipment (vCPE) can use SFC to steer traffic through a sequence of vCPE functional components in the network.

Figure 3. Service function chaining

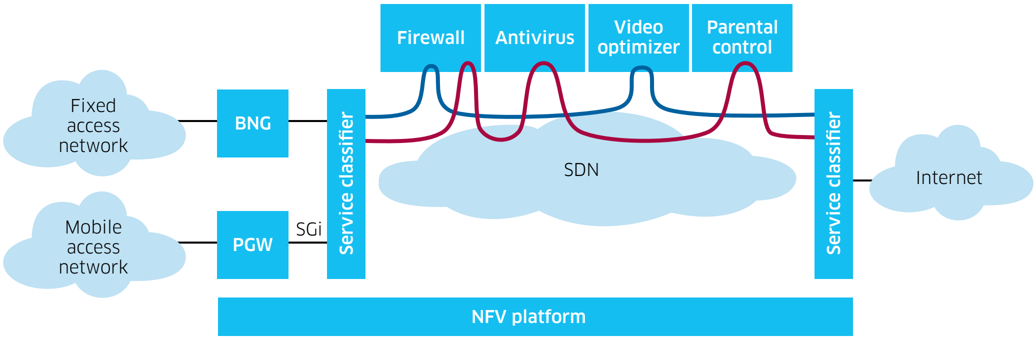

With NFV and SDN, service providers can assemble service function chains through software configuration without physical installation of appliances and without rewiring or manually reconfiguring network connectivity. They can use NFV to dynamically deploy virtual appliances and scale out these appliances to match traffic demand. And SDN enables them to steer subscriber traffic according to the service chain configurations through different sequences of service functions. Service providers can modify and enhance service chains on the fly, for example, by replacing one service function with a newer version or with a similar function from a different vendor, and subscribers can add or remove functions to their personal chains via a self-service portal.

EXAMPLE SERVICE CHAIN FUNCTIONS

- Network address and port translation (NAPT)
- HTTP header enrichment
- IPv4 - IPv6 translation/ transition
- Video optimization
- Firewall
- WAN optimization
- Intrusion detection/prevention system (IDS/IPS)
- Application delivery controller
- Secure tunnel gateway
- Caching/content delivery network (CDN)
- Anti-virus
- Lawful intercept
- Parental control
- Traffic mirroring
- Load balancer
- Traffic shaping/throttling
- TCP optimization
- Charging
- Web optimization
- Network probe

SDN NETWORK MODELS FOR NFV

The network for NFV needs to span both the data center and the wide area between data centers, access networks and customer premises. Service providers, such as Deutsche Telekom with TeraStream and AT&T® with its User-Defined Network Cloud, are reconsidering their network architectures in light of NFV and SDN technologies. The new architectures must be designed to respond to the frequent and automated changes driven by NFV. They must also offer simplified networks based on a common IP/optical layer with all other services implemented on top of this foundation.

There are many ways to design the network for NFV (see Figure 4). In the classical network, integrated switches and routers provide both the forwarding plane with Layer 2 and Layer 3 capabilities, and the control plane where the routing protocols are executed. The original idea of SDN was to break open this connection between forwarding and control planes to enable independent programmability and control over the network behavior. This is shown in Figure 4 as the Infrastructure SDN. To achieve this programmability, operators configure the SDN controller and implement network-wide changes without having to access each switch manually and across multiple vendor platforms. In addition, this process should bring automation. Network applications access the SDN controller via open northbound APIs.

More recently, an overlay SDN model has been introduced. In this model, the SDN controller controls virtual switches in data center servers or selected physical switches, but not the majority of physical switches, which form an underlay network and provide the basic connectivity between the SDN elements.

Figure 4. SDN network models

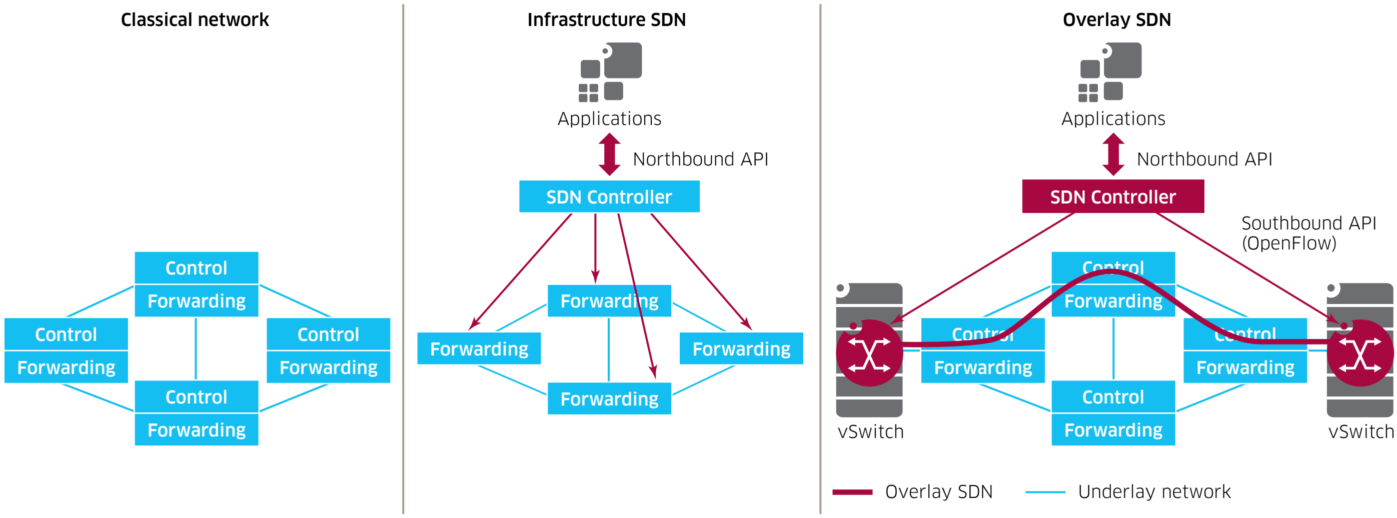

The overlay is formed from a mesh of virtual end-to-end tunnels, based, for example, on VXLAN or GRE. This way, the overlay SDN can create virtual private networks (VPNs) connecting the vNFs without needing any control over the underlay switches. This makes the SDN solution compatible with any existing underlay and obviates the replacement of the switches and routers that are already there.

The overlay/underlay model is effective for creating the custom connectivity needed between vNFs and their components. The tunnel-based VPNs create a measure of security by restricting access to vNFs to legitimate endpoints only.

Quality of service in an overlay SDN model

Assuring the quality of service (QoS) needs of the vNFs is not quite as easy to do. A simple solution for QoS is to provision the network with enough resources to cover even the highest demand. This eliminates bandwidth bottlenecks, packet loss and much of the jitter. Some service providers intend to go this route at least for large parts of their networks.

Another approach is to define a set of universal service classes and the static configuration of these service classes into the underlay network (diffserv model). For mobile network operators, the QoS classes with QoS class identifiers (QCI) 1 through QCI 9 defined in 3GPP TS23.203 may be a good choice (see Table 1). vNFs can then take advantage of the configured QoS classes by marking their outgoing packets with the appropriate QCIs. Switches in the underlay network apply the configured queuing policies to the packets.

Table 1. LTE QoS classes

QCI	Resource type	Priority	Packet delay budget	Packet error loss rate	Example services
1	GBR	2	100 ms	10^{-2}	Conversational voice
2	GBR	4	150 ms	10^{-3}	Conversational video (live streaming)
3	GBR	3	50 ms	10^{-3}	Real-time gaming
4	GBR	5	300 ms	10^{-6}	Non-conversational video (buffered streaming)
5	Non-GBR	1	100 ms	10^{-6}	IMS signaling
6	Non-GBR	6	300 ms	10^{-6}	Video (buffered streaming) TCP-based (e.g. www, email, chat, ftp, point-to-point file sharing, progressive video)
7	Non-GBR	7	100 ms	10^{-3}	Voice Video (live streaming) Interactive gaming
8	Non-GBR	8	300 ms	10^{-6}	Video (buffered streaming) TCP-based (e.g. www, email, chat, ftp, point-to-point file sharing, progressive video)
9	Non-GBR	9	300 ms	10^{-6}	Video (buffered streaming) TCP-based (e.g. www, email, chat, ftp, point-to-point file sharing, progressive video)

The diffserv model works well with a carefully dimensioned network and a careful allocation of vNFs. Already today, highly optimized data-plane-intensive vNFs can churn out several gigabits per second of data traffic per CPU core. Therefore, placing too many vNFs into a single server or into a single rack could exhaust available NIC capacity or the capacity of top-of-rack switch uplinks. This means that the resource (cloud) management system must have smart placement algorithms that have a view of resource availability both in the server and in the network. The algorithms must only place network functions where sufficient resources exist to continue to meet the SLA.

EXAMPLE – A VIRTUAL AAA SERVER

A virtual authentication, authorization and accounting (vAAA) server will illustrate the role of SDN in NFV (see Figure 5). This particular vAAA server has been virtualized into four components running on four VMs:

- Load generator for testing purposes. In real deployments, it would be replaced by a broadband network gateway (BNG), for example.
- RADIUS Routing Agent (RRA). The RRA is essentially a load balancer, sending incoming AAA requests to one of a group of RADIUS Front Ends.
- RADIUS Front Ends. The front ends do the real work in this application. When service demand increases, the vNF manager automatically allocates additional front end VMs, a process called scale-out. Conversely, when the service demand decreases, RADIUS Front Ends are removed (scale-in).
- Universal State Server (USS) as a backend database function. The USS stores the session state.

Figure 5. A vAAA server network function

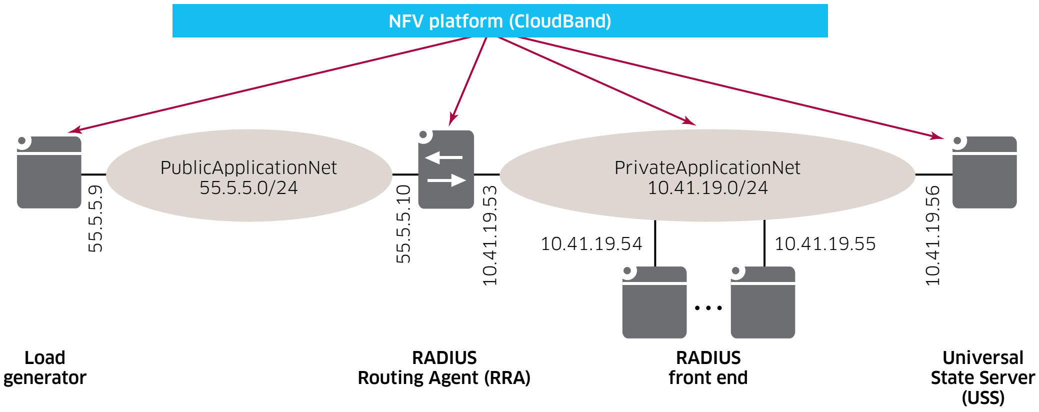

The vAAA network function uses two networks: PublicApplicationNet and PrivateApplicationNet. The private network serves as the virtualized backplane of the vAAA and is designed to hide and protect the inner structure of the vNF. The RADIUS Front Ends and the USS are connected to the private network. The RRA is the gateway between the public network and the private network. It is connected to both of them with two vNICs addressed by two IP addresses.

When the vAAA is first deployed — manually or automatically under control of a vNF manager — the four different types of VMs are connected to the public and the private networks and they are assigned IP addresses from the associated address ranges of the two networks. When the network function is scaled out, additional RADIUS Front End VMs are created and connected to the private network with newly assigned IP addresses. Figure 6 shows a view of the VMs and virtual networks and how they are mapped to two physical servers. Each server contains a virtual switch to relay data traffic among VMs and the physical NIC. In keeping with the overlay SDN model described above, the SDN controller controls the virtual switches but not the physical switches. The PublicApplicationNet and PrivateApplicationNet are implemented using tunnels based on VLANs, VXLAN or GRE.

Figure 6. Mapping of VMs and networks to physical servers and switches

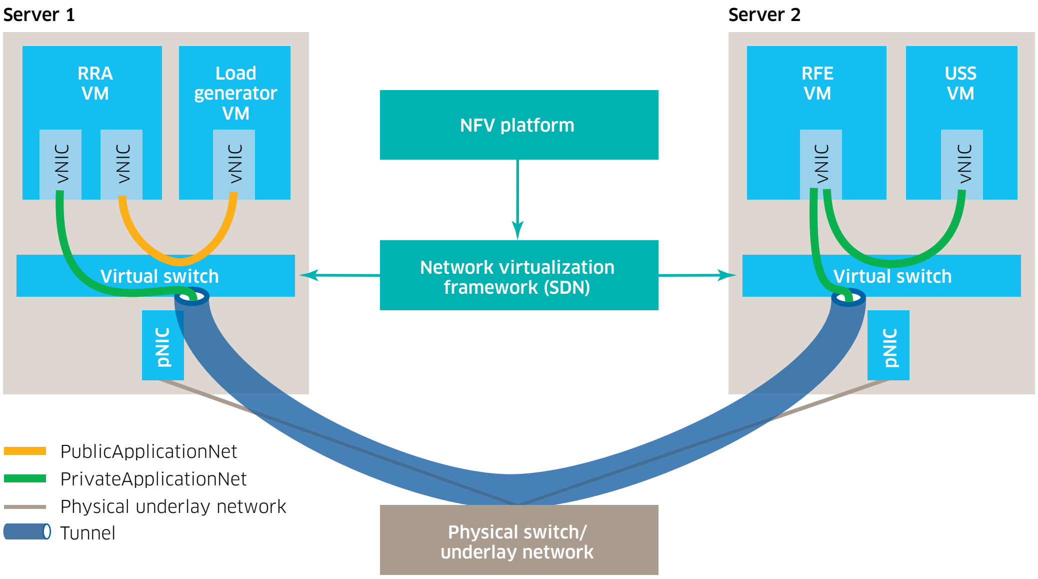

INTEGRATING SDN INTO AN NFV PLATFORM

The NFV platform (Figure 7) provides all the compute, storage and network resources needed to run vNFs (NFV infrastructure). In addition, it provides management and orchestration capabilities to manage the infrastructure as a shared resource as well as capabilities to automate the lifecycle management of vNFs via vNF managers. The network requirements of vNFs are captured in a database of vNF descriptors, for example expressed as Topology and Orchestration Specification for Cloud Automation (TOSCA) templates, an OASIS standard.

Figure 7. NFV platform

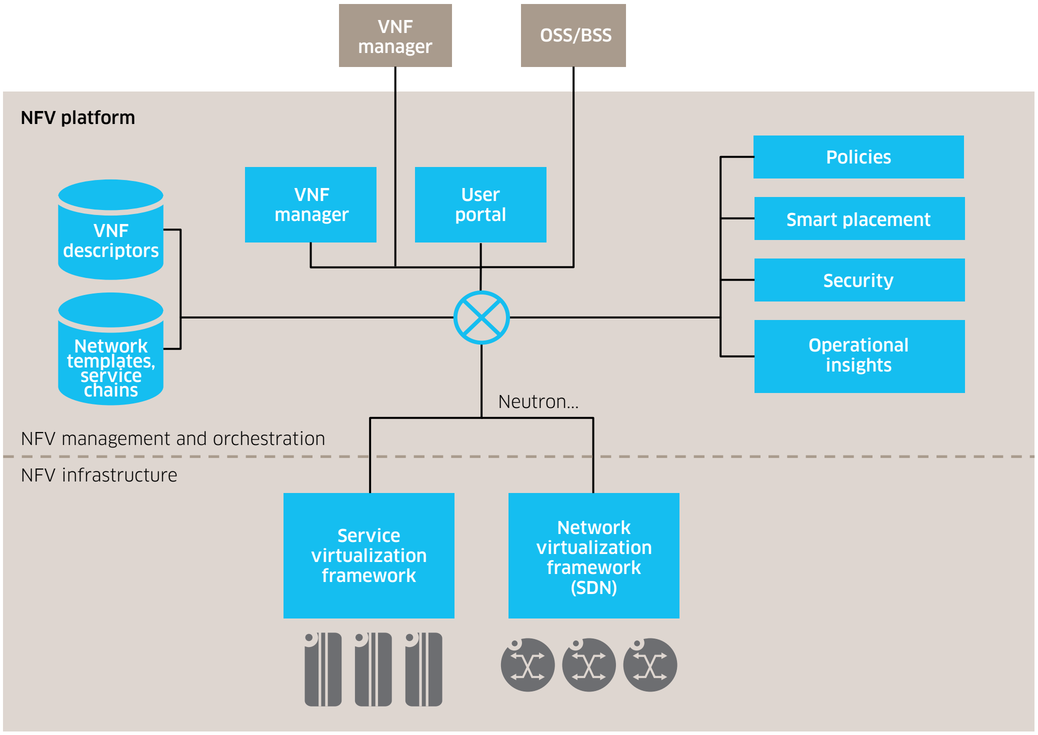

A key part of the NFV platform is a networking framework that is automatically able to adapt to vNF placements. It is widely accepted that SDN meets this need due to the capabilities discussed in the previous section.

Operators create and manage network entities either manually, through a user portal, or in an automated way via a vNF manager. The vNF manager, user portal and OSS access the SDN framework through an API, most likely, based on OpenStack Neutron. (The SDN framework may however contain plugins to adapt to different network implementations). One example may be a plugin that supports an SDN controller, that is, the plugin is essentially an adapter from the Neutron API to the native API of the SDN controller.

Most vNF lifecycle actions will impact the network. When a vNF is to be deployed on one or more VMs, the placement algorithm will match the networking policies and requirements of the vNF with the available resources. These requirements may include throughput, latency and jitter values. vNFs can also indicate the kind of network drivers they support for NICs with optimized data plane handling, such as with Single Root I/O Virtualization.
When the vNF is deployed, the virtual network interface cards (vNICs) will be connected to virtual network ports, and the virtual switches and vNICs from the overlay network need to be configured to apply the desired network policy: traffic handling, marking and encapsulation.

FOUR STAGES TOWARD FULL NETWORK AUTOMATION IN NFV

NFV network control can be automated in different ways. The target for the integration of SDN into an NFV solution is full automation of all network configuration processes. For many network operators used to current methods of procedure, this is a large step to take. Some operators will want to maintain the option to give the final go/no-go decision for any significant changes.

Achieving the goal of full automation can be broken down into four stages.

Stage 0 (no SDN). Any network change is sent as a work order to the network department. The network department executes the work order manually through a command line interface to the involved network switches and router or through a network management system. This is the current state in many networks.

Stage 1 (SDN, non-integrated). The work order is executed manually but with an SDN. Operators have two separate screens: one for the NFV platform and one for the SDN framework. Using the SDN framework already simplifies change management as there is only a single place that needs to be configured: the SDN controller. The individual switching elements will be configured automatically based on SDN protocols, such as OpenFlow™.

Stage 2 (SDN, integrated). This is a model-driven process. The NFV platform includes a repository of network templates and service chains that can be instantiated using the SDN framework. The network models are automatically instantiated and synchronized with the SDN (see Figure 8). Whenever a new VM is created on a particular server and in a particular data center, the NFV platform makes sure that the associated networks reach that data center and server. For example, if necessary, the NFV platform will set up additional overlay tunnels. Explicit work orders are no longer necessary. The NFV platform may also have direct access to configuration data and key performance indicators from the SDN framework. This information can be used for smart placement algorithms, root cause analysis, capacity management and accounting.

Stage 3 (SDN, lifecycle automation). In Stage 3, network setup and configuration becomes part of the vNF lifecycle automation. When vNFs are deployed, scaled, healed, upgraded or terminated, the necessary networks will be automatically created, adapted or removed. During vNF onboarding, the networks needed by the vNF are described in network templates associated with the vNF descriptor (see Figure 9). When a vNF is deployed, the network template will be automatically instantiated without any need of a human operator. Similarly, when a vNF is scaled out, the NFV platform with the SDN framework will extend the necessary networks to the location of the new VM whether it is in the same or a different data center.

Figure 8. Stage 2 – Automatic synchronization between NFV and SDN

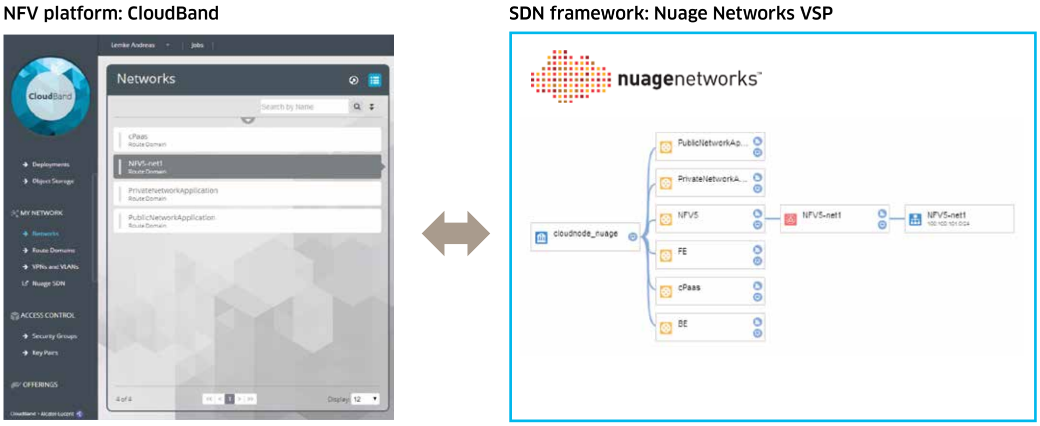

NFV platform: CloudBand

SDN framework: Nuage Networks VSP

Figure 9. Stage 3 – A TOSCA-based vNF descriptor of the vAAA application

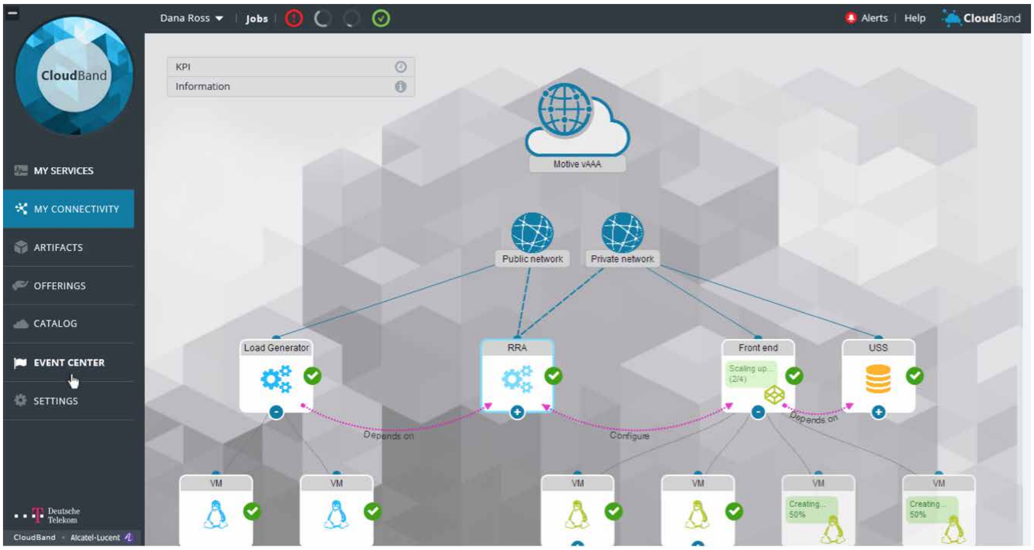

SDN also helps service providers monitor network behavior and detect failures and performance issues (service assurance). Any SDN switching point will be able to deliver a variety of data and measurements about the network. For more extensive monitoring as well as to fulfill regulatory requirements, data traffic can be mirrored to specialized inspection functions. SDN is an important element of an NFV analytics component. With detailed network knowledge, an NFV platform can not only quickly restore service in case of failure but also help to detect the root cause of a failure.

ALCATEL-LUCENT CLOUDBAND AND NUAGE NETWORKS

Service providers' most valuable resource is their network, known for its high reliability, availability and performance. But they need to find ways to leverage this asset in an automated, efficient way. Alcatel-Lucent CloudBand incorporates network control for hybrid SDN/classical networks as an integral part of the NFV platform. It leverages advanced SDN technology from Nuage Networks to provide the network abstractions critical for NFV applications and to automate network provisioning (see Figure 10). CloudBand delivers NFV networking that extends from the data center across the WAN using route domains and network templates to enable flexible aggregation of networks. It offers integrated network automation from Stages 0 to 3 (as defined above) as well as full lifecycle management. CloudBand is capable of real-time monitoring of key network performance indicators for root cause analysis and other purposes.

In addition to having integrated SDN, CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins.

Figure 10. CloudBand and Nuage Networks VSP

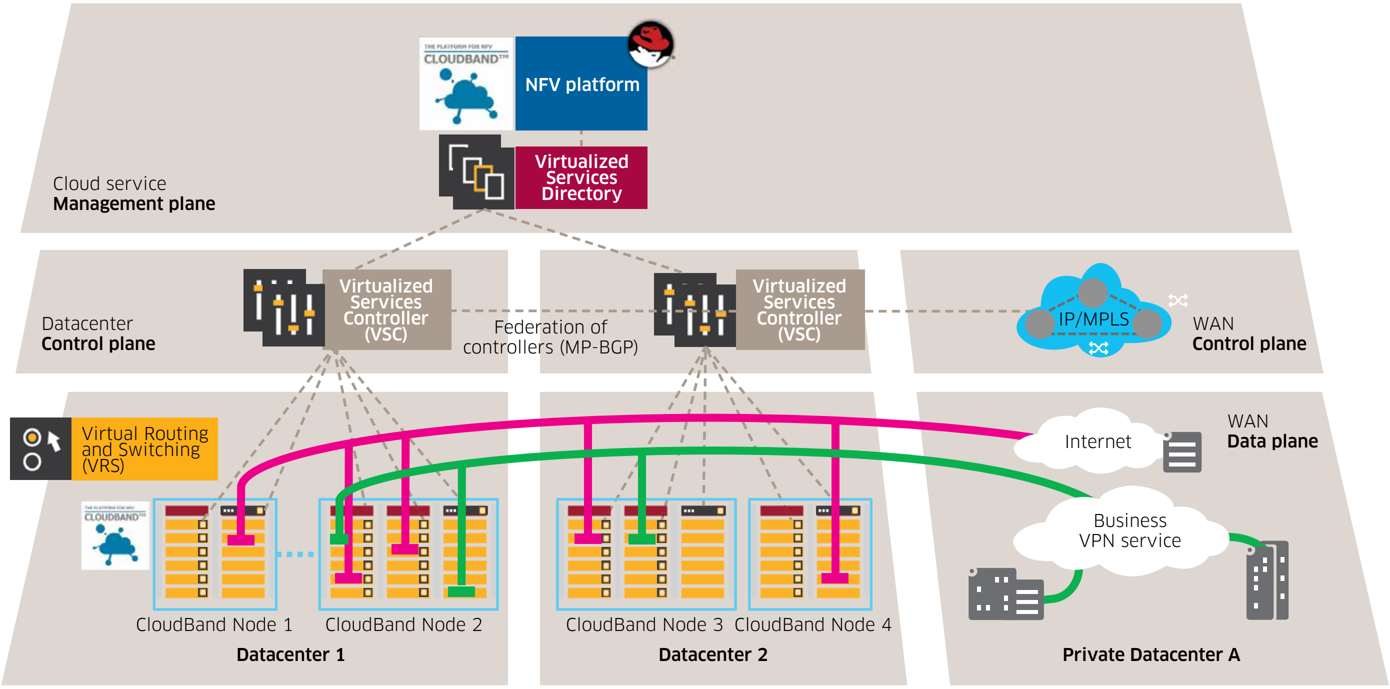

Nuage Networks Virtualized Services Platform (VSP) is a second generation SDN solution. It leverages SR OS, the software base of the Alcatel-Lucent 7750 Service Router family. SR OS supports a wide variety of networking protocols and capabilities, such as VXLAN, GRE and MPLS. Nuage Networks VSP enables the kind of distributed NFV infrastructure needed for carrier applications. With Nuage Networks VSP, networks extend transparently across multiple data centers, each with its own SDN controller (the Virtual Services Controller or VSC). These SDN controllers communicate using standard MP-BGP.

CloudBand and Nuage Networks VSP are policy controlled. CloudBand delivers the network policy required by each virtual network function via a Neutron plugin and additional REST interfaces to the Virtual Services Directory (VSD). As vNFs are created, Nuage Networks VSP detects the new virtual machines, pulls the network policy from the VSD and automatically creates the necessary tunnels across servers and data centers, and programs the virtual switches in the servers. Nuage Networks VSP also provides connectivity to MPLS VPNs or other existing network technologies.

Nuage Networks supports policy-based routing that can be defined via a graphical user interface and web APIs. With this facility, groups of vNFs can be easily combined into service chains. Service chains can also be modified after they are created. For example, in the course of a software upgrade process, a new version of a firewall virtual appliance can be spliced into a service chain to replace an old version without affecting the other functions in the chain.

CONCLUSION

SDN is ideally suited to delivering a networking solution for NFV. SDN's programmability and easy configuration make it a perfect match for the networking needs of rapidly changing NFV applications. However, NFV needs more than a data center SDN. An NFV infrastructure needs to support network functions across many geographical locations. It also needs to be a highly reliable, high performance solution that can interwork easily with legacy networks. The benefits of SDN for NFV depend on which one of four stages of integration a service provider can realize. The Alcatel-Lucent CloudBand NFV platform, integrated with Nuage Networks Virtualized Services Platform, can help services providers on this path by capturing the networking requirements of virtual network functions and automating network control and assurance.

REFERENCES

[1] NFV Whitepaper, http://portal.etsi.org/NFV/NFV_White_Paper.pdf

[2] NFV Insights Series: Why Distribution matters for NFV, a joint Telefonica/Alcatel-Lucent white paper, http://resources.alcatel-lucent.com/?cid=180110

[3] NFV Insights Series: Providing security in NFV - Challenges and opportunities, http://resources.alcatel-lucent.com/?cid=178552

For further reading:

[4] NFV Insights Series: CloudBand with OpenStack as NFV Platform, a joint Red Hat/ Alcatel-Lucent white paper, http://resources.alcatel-lucent.com/asset/180265

[5] NFV Insights Series: Business case for moving DNS to the cloud, http://resources.alcatel-lucent.com/?cid=178476

[6] Why service providers need an NFV platform, http://resources.alcatel-lucent.com/?cid=170811

ACRONYMS

API	Application programming interface
BNG	Broadband Network Gateway
CDN	Content Delivery Network
EPC	Evolved Packet Core
IMS	IP Multimedia Subsystem
MP-BGP	Multiprotocol Border Gateway Protocol
NFV	Network Functions Virtualization
NIC	Network interface card
OASIS	Organization for the Advancement of Structured Information Standards
OSS	Operations support system
PGW	Packet Data Network Gateway
pNIC	Physical network interface card
QCI	Quality of service class identifier
QoS	Quality of service
R&R	Roles and responsibilities
RRA	RADIUS Routing Agent
SDN	Software-Defined Network(ing)
SFC	Service function chaining
SGW	Serving Gateway
TOSCA	Topology and Orchestration Specification for Cloud Applications
USS	Universal State Server
vAAA	Virtual authentication, authorization and accounting
vCPE	Virtualized customer premises equipment
vEPC	Virtualized Evolved Packet Core
VM	Virtual machine
vNF	Virtualized network function
vNFC	Virtualized network function component
vNIC	Virtual network interface card
VPN	Virtual private network
VRS	Nuage Networks Virtual Routing and Switching
VSC	Nuage Networks Virtualized Services Controller
VSD	Nuage Networks Virtualized Services Directory
VSG	Nuage Networks Virtualized Services Gateway
VSP	Nuage Networks Virtualized Services Platform

Alcatel·Lucent