NOKIA

The telco cloud and SDN

Comparing Nuage Networks SDN to OpenStack Neutron

White paper

As we enter the 5G era, telcos are in the midst of a large-scale architectural shift towards the cloud, network virtualization and software-defined networks. In order for telcos to fully reap the benefits of an NFV infrastructure (NFVi), the chosen NFVi must include among other things network automation, as well as scalability, efficiency and reliability. The Nuage Networks SDN solution enables a smooth transition from a legacy central office style core-edge network to a more data-center-style, leaf-spine network and virtual networking. In this white paper, we examine the requirements of the evolving telco network. We also compare the Nuage Networks SDN solution to the OpenStack Neutron solution arguing that the latter falls short on some key performance characteristics.

NOKIA

Contents

NOKIA

Introduction

The emergence of new technologies and applications such as internet of things (IoT), edge computing and 5G are providing tremendous opportunities for providers of communications services. However, in order to capture these opportunities, they need to evolve their infrastructure, even though it is the foundation on which they've built the differentiated services and applications that has been their traditional business.

To be traditional managed service providers (MSP) or digital service providers (DSP)? That is the choice facing telecommunications organizations today. To address the challenges of increasingly demanding customers and fierce competition from over-the-top (OTT) operators such as Google and Amazon, they will have to re-invent themselves.

MSPs that don't make the digital leap are going to fall away. Those that do must radically shift to adopt cloud-native architectures. While MSPs have traditionally focused on hardware, DSPs must maximize the benefits of software. With the right BSS/OSS, supported by cloud-native, software-defined networks (SDN) running on virtualized infrastructure (network function virtualization or NFV), they can unlock new dimensions of possibility. Virtual, modular, easily scalable and responsive: these are the benefits of the virtualized software approach.

From an operational perspective, development, integration, deployment and maintenance need to be automated. High availability and "continuous everything" should be the watchwords. Software should be stateless, operations distributed, and processes should be decomposed. Horizontal and vertical modularization will facilitate scaling. Containerization will ensure that software packaging and delivery are consistent and easily automated.

With cloud-based solutions, DSPs can configure everything quickly and easily according to customer requirements. They can build, test and launch services in fail-fast mode, innovate more and automate practically everything. However, legacy silos will need to be broken down, and future infrastructure will need to expand to deliver business rules and service models associated with today's over the top (OTT) players.

NOKIA

The path to digital transformation

Worldwide, MSPs are undergoing multi-pronged digital transformations in order to increase service and business agility. The aim of these transformations is to enable them to better compete in the digital economy with a flexible and programmable network powered by NFV and SDN. This transformed, digital network will provide a foundation for the delivery of both on-demand services, such as software-defined wide area networking (SD-WAN), and an improved and superior customer experience.

The telecoms industry has chosen NFV/SDN-based networks to power 5G. The 5G core will be cloud-native and virtualized. There will be new transport network architectures with disaggregated fronthaul, mid-haul and backhaul, which will allow the service provider to distribute core functions depending on the needs of the user application. For instance, low latency services or time-sensitive networking (TSN) in 5G is in part addressed by moving functions previously hosted in the core to the edge of the network. This multi-edge computing (MEC) will be critical to support applications such as automation and other machine-to-machine communications.

5G will also enable network slicing, which will allow the service provider to provide a dedicated slice of the virtualized network services to a customer. As an example, a slice could be dedicated to a single enterprise's distributed operations or to a vertically oriented group, such as a health care authority and all its members. In private networks, one slice might be dedicated to office communications, while another, reserved for factory operations.

Operational automation for service agility

These are the sorts of shifts that will require MSPs to increase business and service agility using elastic, cloud-based digital networks, rapidly creating and decommissioning services to reflect the fast-changing demands of consumers and enterprises. They will need to have highly automated and data-driven operations that are delivered at a fraction of the cost of their current operations. These changes will be essential for them to scale their networks to match the traffic growth from new services and deliver superior customer experience by delivering instantaneous and personalized digital experiences across all channels. These changes will enable them to evolve to be proper DSPs.

As we have seen, virtual network functions (VNFs) will play a key role in achieving this dynamic, automatable network fabric. The software equivalent of appliance-based networking components, VNFs as cloud applications or software instances can be dynamically created and modified on the fly. They form the building blocks to compose, implement and switch on new services on demand.

For instance, 5G slices are multiple logically independent networks that are organized by the SDN controller. An orchestrator assembles the necessary NFVs to create the virtualized network services that compose the slices, and it maps virtual functions to the underlying transport layer to assure the appropriate service levels.

In a 5G network, both the distribution of core functions to the edge, in order to support TSN, and the creation and decomposition of slices are intended to occur on demand, depending on the service policy being automatically triggered by a given customer application. Thus, the traditional operations model, which is primarily based on manual and reactive approaches, is not only uneconomical, it simply cannot operate quickly enough to meet the highly variable demands of 5G or any similar virtualized network.

Next-generation enterprise services will also be offered using universal CPE platforms. In this case, the CPE is simply a standard compute and storage server that can support any kind of VNF to provide everything from simple broadband access services, to firewalls or even enterprise-wide slices that might extend from the branch to a specific server in the enterprise cloud.

MSP's envision a future in which customers use a self-service digital interface to request new services, and expect these services to be supplied and working in a matter of minutes. In this future operational model, the digital network, the digital channels and the operational systems must work in perfect harmony to deliver and assure services with minimal manual intervention. MSPs need a new operational model that is based on the principles of automation.

The role of SDN in virtual networks

Traditional network designs deployed in carrier networks used network elements to implement control, data and management planes. They also implemented services on the same hardware, such as connectivity-related services like layer 2 and layer 3 (L2/L3) virtual private networks (VPNs). Some chassis-based systems also had service blades that implemented stateful network functions such as firewalls, load-balancers and NAT.

Figure 1. Traditional network designs used network elements to implement services

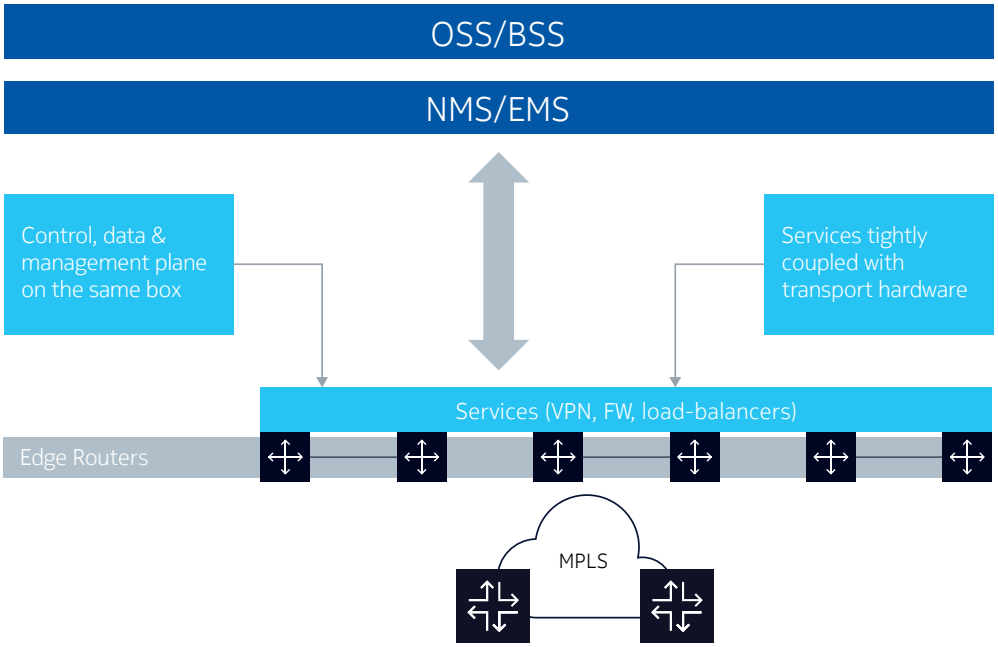

With the evolution to cloud, there has been a paradigm shift in how services are built and delivered. The transition from legacy designs to the newer telco-cloud architecture is based on four key principles:

1. The use of commercial off-the-shelf (COTS) hardware for servers
2. Harnessing the agility of the cloud
3. Implementing software-defined networking (SDN)
4. Implementing network function virtualization (NFV)

Figure 2. The four principles of the telco-cloud architecture

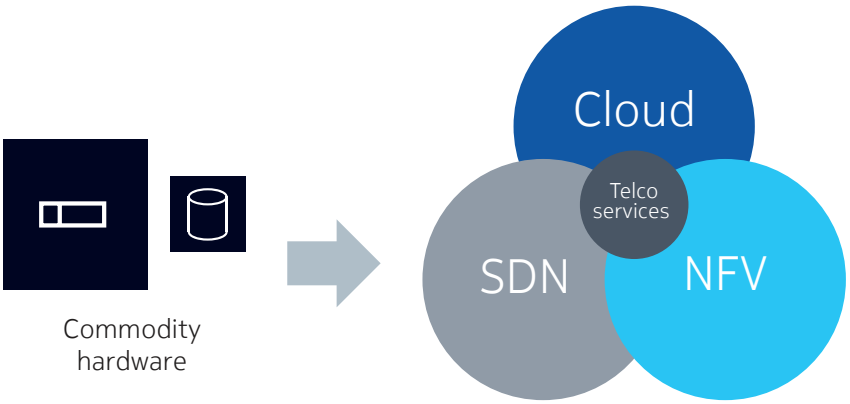

COTS server hardware

This was one of the key paradigm changes in the evolution to the cloud; by using commercial Intel-based servers, cloud operators were able to scale-out compute, network and storage with clustering technologies. The server thus became the new networking platform. Another key change was to implement software-defined distributed storage by pooling small computer system interface (SCSI) and non-volatile memory express (NVMe) storage disks within individual servers to create much bigger logical storage.

Harnessing the agility of the cloud

Software platforms running on commodity server hardware enables the rapid creation of new virtual services. Telcos initially used virtual machine (VM) software platforms, but are increasingly looking beyond VM-based VNFs, to cloud-native architectures based on containers and micro-services. Cloud-native software will provide the full benefits of the cloud by providing better hardware utilization, dynamic horizontal scaling and quicker software deployment. Unlike VNFs, decomposed, cloud-native network functions will not require specialized VNF managers and complex post-boot initialization procedures.

Implementing software-defined networking (SDN)

SDN has a lot of definitions in the market. Specifically in the telco domain, SDN is the separation of control, data and management planes so that each layer can scale independently. It is also the ability to have more flexible change management, enabling, for instance, the separate upgrading of the data, control and management planes. One of the key goals is to separate services like L2/L3 VPNs and other network functions from the transport/underlay hardware elements. This allows them to have a cleaner separation of duties, thereby simplifying their overall network design and operational model. While separating the services from the transport, another key architectural shift is to move the network edge from the edge router hardware to the edge cloud server.

Implementing network function virtualization (NFV)

Because network functions virtualization (NFV) is based on industry-standard COTS servers, it mostly avoids problems involved in using proprietary network hardware equipment. The necessity to install network-specific equipment is reduced, depending upon the use case requirements and economic benefits. The ETSI[1] Industry Specification Group for Network Functions Virtualization (ETSI ISG NFV) sets the requirements, reference architecture, and the infrastructure specifications necessary to support virtualized functions.

1 The European Telecommunications Standards Institute (ETSI) is an independent standardization group that develops standards for information and communications technologies (ICT) in Europe.

In general, the NFV architecture, as specified by ETSI, has the following components:

Virtualized Network Functions (VNFs): the software implementation of network functions, such as routers, firewalls, load balancers, broadband gateways, mobile packet processors, servicing nodes, signaling and location services.

NFV Infrastructure (NFVi): the physical resources (compute, storage, network) and the virtualization layer that make up the infrastructure. The network includes the datapath for forwarding packets between virtual machines and across hosts. This allows you to install VNFs without being concerned about the details of the underlying hardware. NFVi forms the foundation of the NFV stack. NFVi supports multi-tenancy and is managed by the Virtual Infrastructure Manager (VIM).

NFV Management and Orchestration (MANO): manages all the tasks required throughout the lifecycle of the VNF. The main functions of MANO are service definition, automation, error-correlation, monitoring and lifecycle management of the network functions offered by the operator to its customers, decoupled from the physical infrastructure.

Figure 3. ETSI NFV reference architectural framework

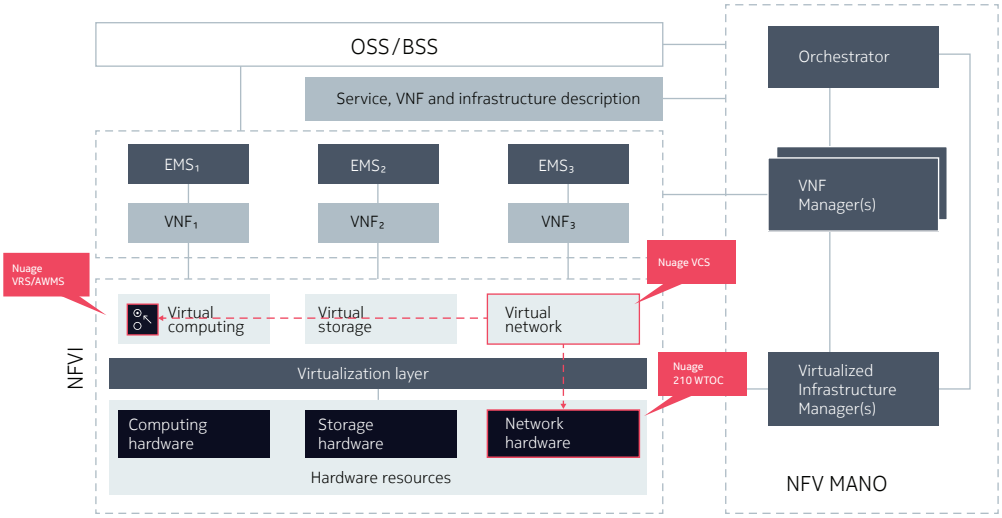

This decoupling requires an additional layer of management, provided by the virtual network function manager (VNFM). The VNFM manages the lifecycle of the virtual machines and VNFs by either interacting directly with them or through the element management system (EMS) provided by the VNF vendor.

The other important component defined by MANO is the NFV orchestrator (NFVO). The NFVO interfaces with various databases and systems including operations and business support systems (OSS/BSS) and with the VNFM. If the NFVO wants to create a new service for a customer, it asks the VNFM to trigger the instantiation of a VNF, which may result in multiple virtual machines.

SDN in the wide area network

With the virtualization of the network edge and the adoption of SDN, it became possible to deliver per tenant virtual network services using overlays like VXLAN. The services can be deployed and managed on a per tenant virtual network basis. It also allows for the virtualization of existing physical network services, such as firewalls, load-balancers and NAT. SD-WAN is a good example of a value-added telco service that can be offered as a managed service, especially catering to global enterprises.

With SD-WAN, the enterprise customer is able to achieve the separation of per tenant connectivity and value-added services from the transport/underlay. This is shown below in figure 4.

Figure 4. A layered telco data center view of an SD-WAN architecture

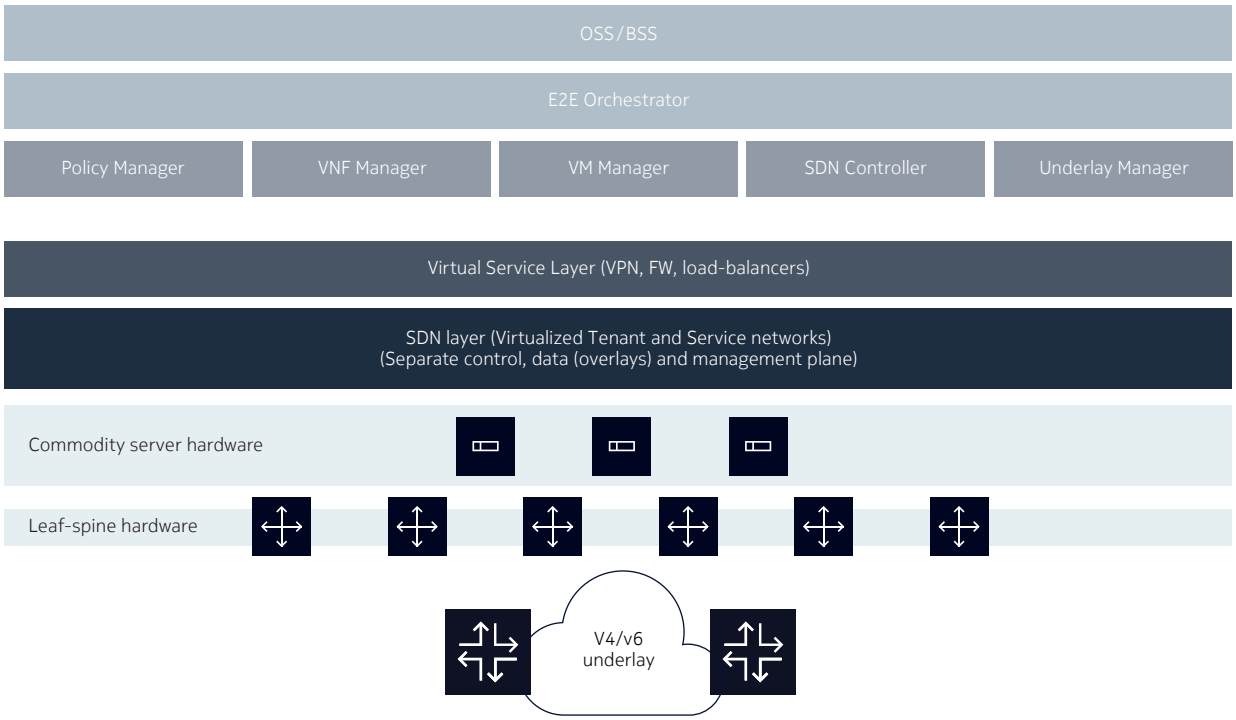

The performance of the NFVi depends on the implementation of the physical and virtual layers for compute, network and storage. For an end-to-end solution, however, the critical aspect is management simplicity throughout the lifecycle of the VNF or set of VNFs. Some degree of automation is essential to make these virtual services cost-effective. Ideally, policies can be defined that automatically set up the virtual service and make a particular deployment use-case simple to manage during the VNFs' lifecycle.

The Nuage Networks SDN solution

The Nuage Networks SDN solution is an example of how to deploy virtual services in a telco cloud environment. It enables a smooth transition from a legacy central office style core-edge network to a more data center style, leaf-spine network and virtual networking. Figure 5, below, highlights the alignment of the Nuage Networks SDN platform with the telco cloud architecture as we've described it above. It is capable of supporting multiple hypervisors and clouds, including private cloud and software-defined branches using SD-WAN technologies.

Figure 5. The alignment of the Nuage Networks SDN solution in a typical telco cloud.

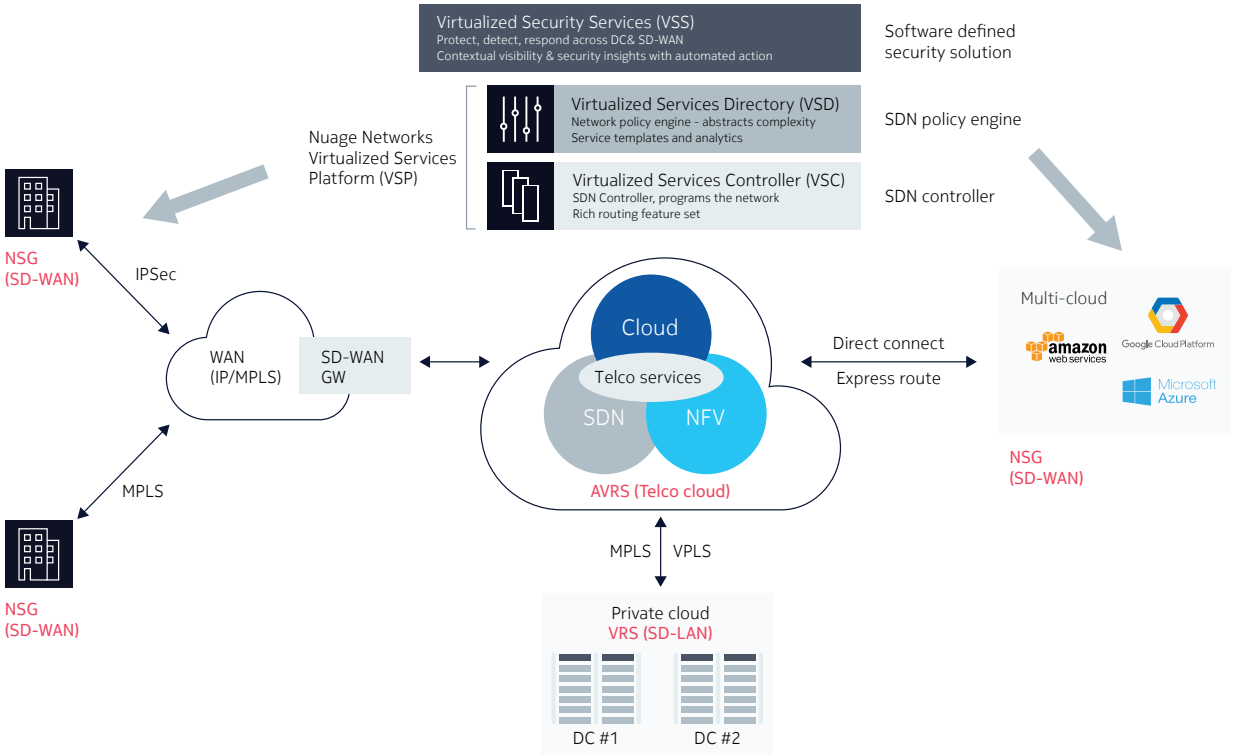

In order to reap the benefits from deploying an NFVi, as we discussed above, it is important to consider network automation, as well as the scalability, efficiency and reliability of the NFVi. The Nuage Networks SDN solution delivers a distributed L2/L3 forwarding plane to provide traffic isolation and security between overlays. Each overlay can either be based on a template that can be used many times over by VNF applications or customized to each VNF application. It scales to 64,000 L2 overlays and 16,000 L3 overlays.

To forward tenant traffic, it uses distributed routing and switching for the overlays in conjunction with an IP fabric underlay, which utilizes equal-cost multi-path (ECMP) for efficient traffic forwarding. Every leaf is the same number of hops away from every other leaf for efficient traffic distribution. The Nuage Networks SDN solution uses a distributed routing concept for service overlays, meaning that traffic between VMs will always take the shortest, most efficient path through the fabric.

We will look at the contrast between the Nuage Networks approach and OpenStack in the next major part of the paper. However, for now it is interesting to note that as a point of contrast, the non-SDN use-case with OpenStack networking provides shortest path forwarding only for L2 traffic. All L3 traffic requires virtual LANs (VLANs) to provide inter-network routing, which is usually the function of an external router or the data center gateway (DC-GW). This leads to tromboning of traffic, which is inefficient in terms of both capacity and latency, but as well, since it involves stitching VLANs on physical switches, it also complicates provisioning due to:

- The need to track VLAN-IDs
- Reduction in flexibility: moving a VM could involve renumbering VLANs
- Reduction in scalability (VLAN scalability)

- Reliability could be reduced: to counter the scalability downsides of VLANs one could build larger L2 segments but this will impact reliability due to media access control (MAC) scaling issues, MAC learning and broadcast, unicast and multicast (BUM) flooding.

Nuage Networks SDN also provides control plane reliability by choosing to use Ethernet VPN (EVPN) as a distributed control plane, which in turn programs the distributed L2/L3 forwarding plane. EVPN is based on multi-protocol border gateway protocol (MP-BGP), so it permits the Nuage Networks SDN to leverage the highly scalable capabilities of MP-BGP. MP-BGP with EVPN is also the control plane of choice when looking beyond the DC to an end-to-end (e2e) unifying control plane.

Telco VNF additional requirements

As described in the previous section, the Nuage Networks SDN provides an efficient, scalable and reliable SDN solution. However, today's telco VNF applications have additional requirements that go far beyond what enterprise applications require, such as routing protocol support, interface liveness, high capacity connectivity to the fabric and traffic mirroring. On top of the above benefits of an SDN layer, there are additional capabilities that the Nuage Networks SDN brings:

- **Routing protocol support for VNFs** requires that the VNF can peer, usually using border gateway protocol (BGP), with the overlay control plane to distribute tenant routes (e.g., UE address pools). Nuage Networks SDN supports both BGPv4 and BGPv6.
- **Interface liveness detection**, usually in the form of bi-directional forwarding detection (BFD), can be critical for some telco applications in detecting when an application needs to failover from one interface to another. Nuage SDN supports BFD for liveliness detection.
- **High-capacity connectivity** is not only about the raw speed of network interface cards (NICs), but also whether VNFs provide accelerated throughput. There are several different approaches to this, ranging from single root input/output virtualization (SR-IOV) VNFs, data plane development kit (DPDK) VNFs or utilizing the open virtual switch (OVS) offload capabilities within the current and next-generation of compute NICs. Nuage Networks supports all of these options allowing any type of VNF to be deployed on top of a Nuage Networks SDN solution.
- **Monitoring and visualization** - Operators also need visualization and near real-time monitoring capabilities to find out what is going on within their network. This could be for troubleshooting or for input into customer experience management (CEM) systems which require knowledge of customer and network information. Mirroring can be used to provide this type of input information. In the Nuage Networks SDN solution mirroring on virtual switches or fabric switches (for SR-IOV traffic) can be automated via API so thata an external system can turn on and off the mirroring as required.

Why Nuage SDN vs OpenStack Neutron

OpenStack is often used in conjunction with NFV technology in data centers to deploy cloud services. On paper, OpenStack holds a lot of promise. The goal of OpenStack in the NFV world was to build large, scalable services very efficiently and economically. The nature of using open technologies as the foundation for certain SDN and NFV frameworks, however, means that although they are driven by community standards, those standards are often the lowest common denominators agreed to by the committee.

After committing to OpenStack, some service providers have experienced the limitations of OpenStack's compromises. The biggest problem with OpenStack scaling in NFV is the networking piece, Neutron. Many operators have reported that Neutron can tap out with less than 200 requests. Due to its strong networking DNA, the Nuage Networks SDN has overcome the Neutron deficiencies and provides network plugins so that the networking scales to the 'carrier-grade' requirements. OpenStack's limitations are described in the subsections below.

Data plane complexity

The Neutron data plane was complex when designed initially and it has remained so. In comparison, Nuage Networks SDN has simplified the data plane with distributed switching, routing and firewalling, as shown in figure 6.

Figure 6. The complexity of Neutron compared to Nuage Networks simplicity.

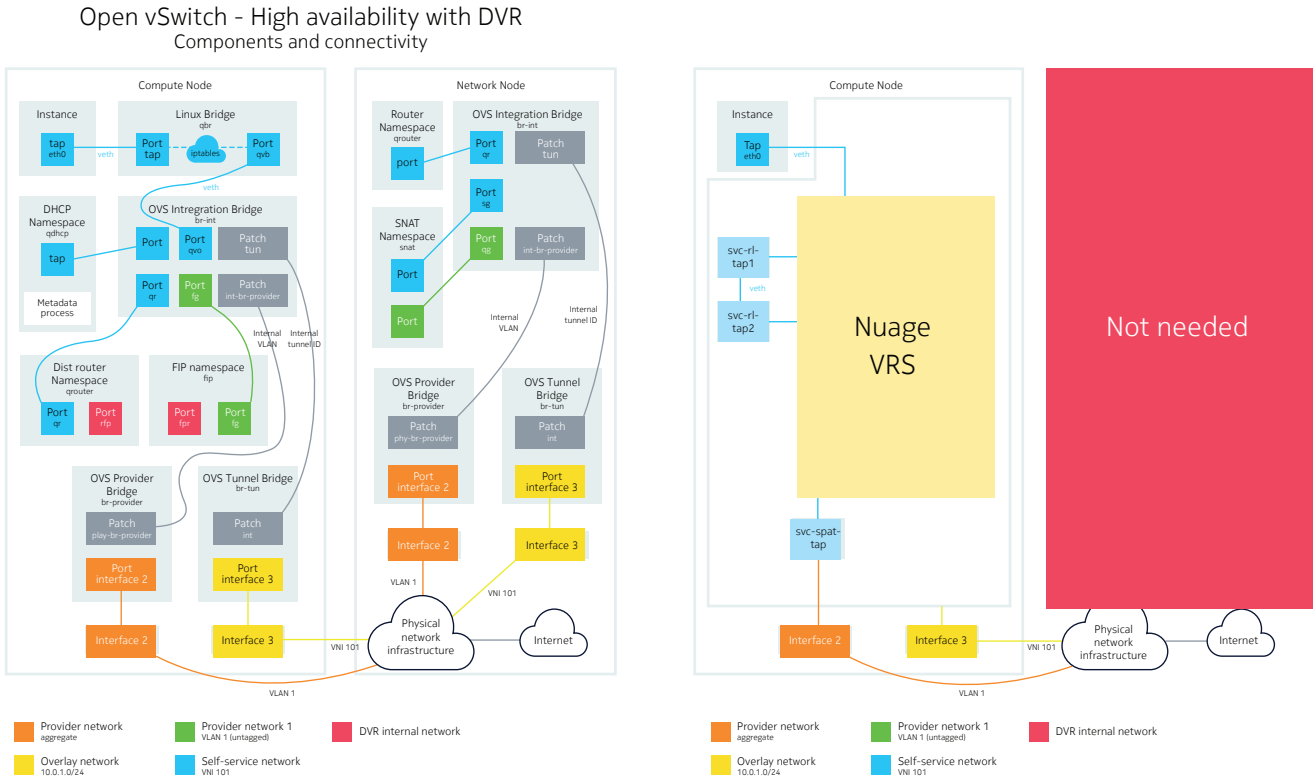

Nuage Networks SDN achieves the simplicity because the VRS is a single OVS bridge that is flow-based, performs firewalling, switching, routing and NAT, as well as processing ARP and DHCP locally. There is no dedicated network node either in the non-DVR case for routing, DNAT, SNAT and DHCP, or in the DVR case, for SNAT and DHCP.

Control plane scalability

OpenStack has scaling issues for a number of reasons, outlined in figure 7, including bottlenecks at high numbers of nodes or high rates of change.

Figure 7. Control plane scaling concerns with OpenStack vs. Nuage Networks

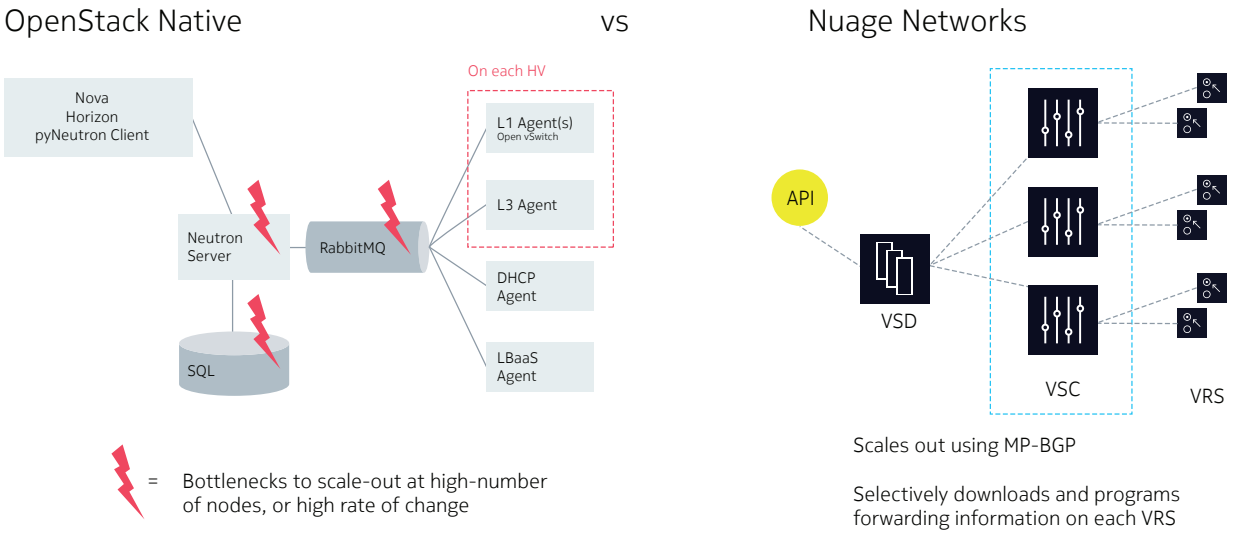

The following table provides a summary of the differences between OpenStack and Nuage Networks.

Table 1. Native OpenStack vs. Nuage Networks SDN

Telco cloud characteristic	Native OpenStack	Nuage Networks
BGP	Software only	Software/hardware VTEP support
QoS	Port rate limit	Port + FIP + BUM rate limit
QoS	DSCP tagging	DSCP tagging and remapping
SNAT	Centralized	Distributed
Cross-tenant routing	no	yes
Route leaking to shared hub	no	yes
BareMetal mapping to tenant subnet	no	yes
Optimized L3 multicast send/receive	no	yes
Direct routing to underlay for direct bare metal access Exit without network node	no	yes
Port-address-translation	no	yes
DC-interconnect	no	yes
Detailed network analytics	no	yes

Telco cloud characteristic	Native OpenStack	Nuage Networks
Connected to MPLS networks	Separate (software) gateway is required to terminate datacenter overlay	1. Provisioning of Nokia 7750PE devices is fully automated 2. Routing reachability is done using MP-BGP (EVPN) – fully interoperable with Nokia 7750PE
Geographically dispersed	Networks are isolated and have no interaction / knowledge of each other	1. Provides single pane of glass for networks 2. Offers fully distributed switching / routing implementation across locations for IPv4/IPv6. 3. Federates its control plane for scale and local robustness
VM overlays need to be robust and scalable	L2 fabrics are used for simplicity, but they are difficult to troubleshoot and very prone to loops, broadcast storms, and have low link utilization	SDN overlays deployed on a BGP-based leaf/spine-based fabric allow for easy horizontal expansion as well as the ability to leverage IP-based ECMP across all links. ISIS or OSPF can also be used for fabric control plane.
VNFs with SRIOV	No automated connection of VNF-VLANs to overlay domains, hence needing explicit fabric provisioning	Automated SRIOV provisioning to allow SRIOV VLANs to be automatically mapped to overlay networks
VNFs that like to leverage DPDK	No DPDK module is available that works for L3	DPDK enabled acceleration for tenant VMs for both L2 and L3
SmartNICs (eSwitch) for hardware-offload of overlays	No strategic collaboration	Close partnerships with SmartNIC vendors to deliver accelerated throughput for tenant VMs
Support for dynamic routing protocols BGP and BFD	Not available	Available both in TOR and in (A)VRS
Mirroring/tapping	Not available	Not available
Security enforcement	Limited segregation and logging options available	Extensive L2-L4 stateful firewalling are available with logging and network threat-detection capabilities

Conclusion

As we enter the 5G era, telcos are in the midst of a large-scale architectural shift towards the cloud, network virtualization and software-defined networks. In order for telcos to fully reap the benefits of an NFV infrastructure, the chosen NFVi must include, among other things, network automation, as well as scalability, efficiency and reliability. The Nuage Networks SDN solution enables a smooth transition from a legacy central office style core-edge network to a more data-center-style, leaf-spine network and virtual networking.

The Nuage Networks SDN solution meets the needs for automation with VNF templating, as well as a distributed L2/L3 forwarding plane to provide traffic isolation and security between overlays. It also uses a distributed routing concept (ECMP) for service overlays, meaning that traffic between VMs will always take the shortest, most efficient path through the fabric. For improved reliability, it uses EVPN as a distributed control plane, which enables the Nuage Networks SDN to leverage the highly scalable capabilities of MP-BGP.

Figure 8. The Nuage Networks SDN simplified architecture

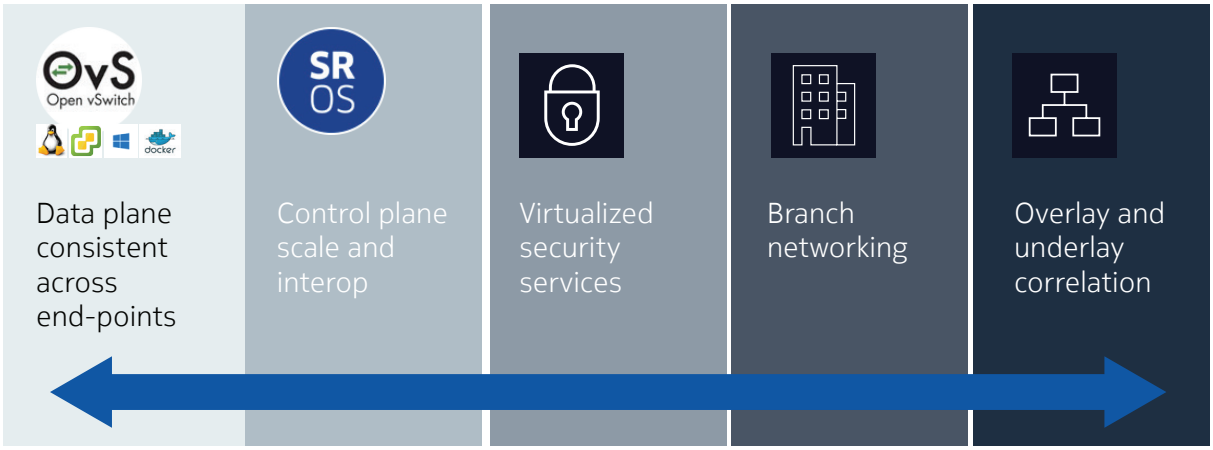

When compared to OpenStack Neutron, we are reminded of the limitations of decisions by committee. Neutron has many architectural shortcomings and with each release it goes through many changes. This makes for a lack of stability and reduced agility as there are many more checks and pre-requisites to deploy new versions and services.

The superior, highly scalable architecture of Nuage Networks SDN supports a wider range of end-point types in bare metal and docker. It provides a feature-rich networking solution in an OpenStack Cloud environment that combines underlay and overlay networking information for improved troubleshooting. Due to its strong networking DNA, the Nuage Networks SDN has overcome the Neutron deficiencies and is more than capable of meeting 'carrier-grade' requirements.

NOKIA

Acronyms

API	Application program interface
ARP	Address resolution protocol
AVRS	Accelerated virtual routing and switching
BFD	Bi-directional forwarding detection
BGP	Border gateway protocol
BSS	Business support system
BUM	Broadcast, unicast and multicast
CEM	Customer experience management
COTS	Commercial off-the-shelf (hardware servers)
CPE	Customer premises equipment
DC	Data center
DC-GW	Data center gateway
DHCP	Dynamic host configuration protocol
DNAT	Destination NAT
DPDK	Data plane development kit
DSCP	Differentiated services control point
DSP	Digital service providers
DVR	Distance vector routing
E2E	End to end
ECMP	Equal-cost multipath
EMS	Element management system
ETSI	European telecommunications standards institute
ETSI ISG NFV	ETSI industry specification group for network functions virtualization
EVPN	Ethernet VPN
FCoE	Fiber channel over Ethernet
FIP	FCoE initialization protocol
FW	Firewall
GW	Gateway
HW	Hardware
ICT	Information and communications technology
IP	Internet protocol
IP/MPLS	IP multiprotocol label switching

ISIS	Intermediate system to intermediate system
L2/L3/L4	Layer 2, 3 or 4
LBaaS	Load balancer as a service
MAC	Media access control
MANO	Management and orchestration
MEC	Multi-edge computing
MP-BGP	Multi-protocol border gateway protocol
MQ	Message queue
MSP	Managed service providers
NAT	Network address translation
NFV	Network function virtualization
NFVi	NFV infrastructure
NFVO	NFV orchestrator
NIC	Network interface card
NMS	Network management system
NSG	Network security group
NVMe	Non-volatile memory express
OSPF	Open shortest path first
OSS	Operations support system
OTT	Over the top
OVS	Open virtual switch
QoS	Quality of service
SCSI	Small computer system interface
SDN	Software-defined networking
SD-LAN	Software-defined local area network
SD-WAN	Software-defined wide area network
SIP	Session initiation protocol
SQL	Structured query language
SNAT	Source NAT
SR-IOV	Single root input/output virtualization
TOR	The onion router
TSN	Time-sensitive networking
UE	User equipment

VLAN	Virtual local area network
VM	Virtual machine
VNF	Virtual network function
VNFM	Virtual network function manager
VPLS	Virtual private LAN service
VPN	Virtual private network
VRS	Virtual reference station
VSC	Virtual services controller
VSD	Virtual services directory
VSP	Virtual services platform (Nuage Networks product)
VSS	Virtual security service
VXLAN	Virtual extensible local area network

About Nokia

We create the technology to connect the world. We develop and deliver the industry's only end-to-end portfolio of network equipment, software, services and licensing that is available globally. Our customers include communications service providers whose combined networks support 6.1 billion subscriptions, as well as enterprises in the private and public sector that use our network portfolio to increase productivity and enrich lives.

Through our research teams, including the world-renowned Nokia Bell Labs, we are leading the world to adopt end-to-end 5G networks that are faster, more secure and capable of revolutionizing lives, economies and societies. Nokia adheres to the highest ethical business standards as we create technology with social purpose, quality and integrity.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: SR1912040532EN (January) CID207060