NOKIA

Accelerate Telco Clouds with Nuage Networks Virtualized Cloud Services

Application note

NOKIA

In this paper the new challenges that 5G and IoT bring to the Telco Cloud will be outlined along with how Nokia's pre-engineered NFVI blueprint solution addresses them. The paper then focuses on Nuage Networks from Nokia's Telco Cloud SDN offering called Virtualized Cloud Services (VCS). VCS' overall deployment architecture is discussed showcasing openness, deployment flexibility, and scalability. Both intra- and inter- DC network automation is addressed. The paper also discusses various application acceleration methods that can be deployed to improve packet processing efficiency.

NOKIA

Contents

The challenges facing today's Telco Cloud	4
Nokia's NFVI Telco Cloud blueprint solution	6
The Data Center solution	6
The NFV Management and Orchestration (MANO)	7
Nokia NFVI Telco Cloud Services	7
Nokia NFVI Telco Cloud blueprint benefits	7
Nuage Networks VCS	8
VCS network abstraction	8
Nuage Networks VCS Telco Cloud architecture	10
Virtualized Services Platform	10
Architectural Component of Virtualize Cloud Services	11
Virtualized Services Controller	11
Virtual Routing and Switching	12
HW VTEPs: 210 WBX or third-party equivalents	13
Packet acceleration methods	14
VRS offloading	14
AVRS - OVS offload using DPDK	14
OVRS - OVS offload using SmartNIC technology	15
VLAN direct to HW VTEP	15
VCS Telco Cloud in action	15
Intra-DC Telco Cloud network automation	15
VXLAN tunnel termination on border leaf nodes	17
VXLAN tunnel termination on DC GWs	18
Network Functions Interconnect to address 5G and network slicing	18
Conclusion	19

NOKIA

The challenges facing today's Telco Cloud

The virtualization of the Telco Cloud infrastructure is rapidly progressing, and many Communication Service Providers (CSPs) are well underway in virtualizing their data center infrastructure. Proprietary and purpose-build hardware platforms are being replaced with their virtualized counterparts creating an open ecosystem where commodity general-purpose x86-based hardware is deployed at a much lower price. These general-purpose x86 platforms are used to host Virtual Machines (VMs) and containers that provide the compute resources for new Virtualized Network Functions (VNFs) that enable both consumer services (e.g. 3G, LTE, 5G, Internet broadband, etc.) and enterprise services (e.g. managed SD-WAN, IP/MPLS, security, IoT, etc.). With this virtualization infrastructure in place (via ETSI's NFV Management and Orchestration (MANO)), manual operational tasks are being done automatically with massive scale reducing operational costs and time to affect changes.

The biggest challenge with this virtualized compute environment is it creates an operational and scalability issue with the statically configured IT network. Software-defined networking (SDN) was created to solve this problem by providing the automated networking framework that dynamically inter-connects these virtualized resources. It does this by separating the forwarding plane from the control plane while enabling centralized policy control to program overlay VPNs across the existing "underlay" network. This allows IT operators to program the network centrally and use templates to replicate configuration tasks in a much more scalable manner.

These overlay VPNs or "tunnels" are created in the forwarding plane by adding a further layer of encapsulation to native Ethernet frames making it routable yet isolated from all other neighboring traffic. In addition, SDN solutions inherently have application level visibility and control allowing for the IT operator to program the network to optimize and serve the networking needs for each application. This allows for a highly programmable and intelligent networking fabric. With this infrastructure in place applications and their underlying packet flows will be automatically delivered to their designated VMs or containers with efficiency and scale while at the same time being able to dynamically adapt to the dynamic virtualization infrastructure. With the predicted deluge of 5G- and IoT- empowered applications, a complete SDN programmable infrastructure in place is no longer a luxury but a requirement.

With 5G and IoT technologies, Communication Service Providers (CSP), have a great opportunity to transform their business creating a boon of new services and technologies. CSPs need faster networks and the ability to rapidly adapt to changing demands and new opportunities. In addition, CSPs face their biggest expenditure in decades so keeping costs down and making immediate profits will be a top priority. CSPs can't afford any service delays or issues deploying their new infrastructure. Any delays will result in significant loss of revenue and market share.

Today's modern Telco Cloud will have to adapt to a few key challenges brought on by cloud architectures and exacerbated by 5G and IoT:

- Meet unpredictable data growth with on demand scalability – growth will be huge, but it will also be unpredictable and spikey requiring an adaptable and elastic framework that can handle the extreme demands that the next generation of applications will place on it.
- Deliver an improved and measurable customer experience – the next generation of applications will require measurable performance, and many will have stringent latency requirements. Automated vehicles, robotic surgery, drone navigation, mobile gaming, etc., will all place severe requirements on the next generation Telco Cloud infrastructure.

- Achieve a lower Total Cost of Ownership (TCO) – lowering the TCO is also an expectation most CSPs share. With the programmability, abstraction and simplification that the next generation of cloud solutions promise, economies and efficiencies of scale and programmability are built into the business case.
- Accelerate innovation cycles – the industry is already witnessing accelerated innovation cycles due to cloud-based automation through DevOps and this innovation is expected to continue with the next generation of Telco Cloud services. Service agility and faster time to market are considered table stakes.

As depicted in Figure 1, it is at the confluence of cloud computing, NVF technology, and SDN where CSPs need to look to address these transformational challenges. SDN's ability to intelligently automate the network is paramount in transitioning CSP networks to achieving their goals.

Figure 1. Telco Cloud key challenges

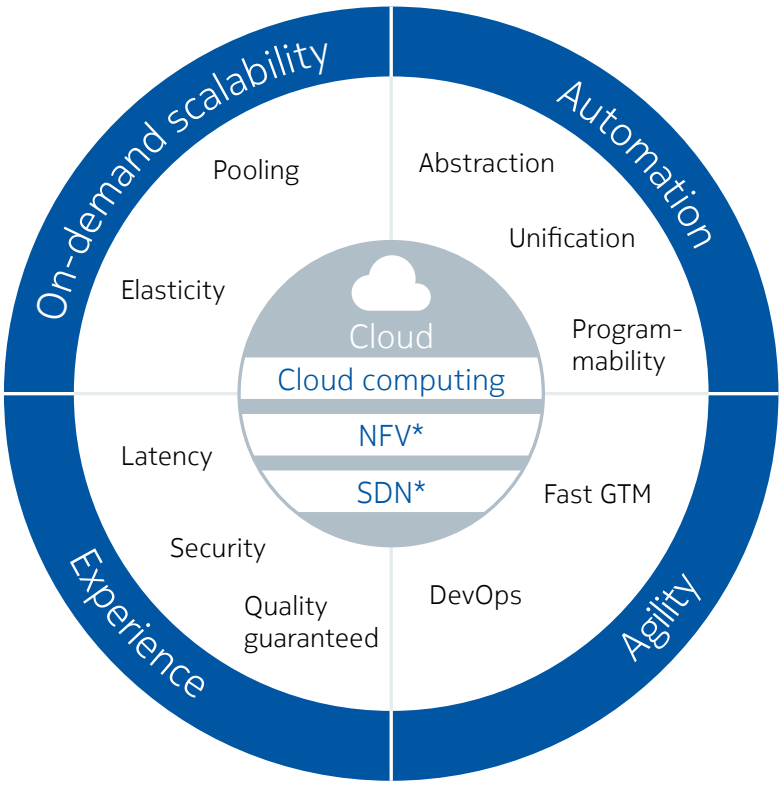

Nokia's NFVI Telco Cloud blueprint solution

Today, nearly all CSP growth opportunities will leverage their Telco Cloud infrastructure. Cloud technology has matured from initial inflated expectations to real life deployments. Early adopters have overcome the initial hurdles and are beginning to experience tangible benefits and are starting cloud expansion and the harmonization of initial multi-cloud silo deployments. CSP buying behavior is shifting rapidly towards acquiring pre-integrated and validated Network Function Virtualization Infrastructure (NFVI) systems for the transformation of their Telco Cloud.

Nokia NFVI Telco Cloud offers a pre-engineered, pre-validated Telco Cloud infrastructure ensuring a solid foundation on which to transform or renovate CSP's existing infrastructure. The Nokia NFVI Telco Cloud provides compute, storage, and networking infrastructure to be shared by VM-based and containerized network functions and applications.

As shown in Figure 2, Nokia NFVI Telco Cloud blueprint consists of a few pre-integrated areas:

The Data Center solution

Virtualized Cloud Services (VCS) is Nuage Network's SDN solution allowing for intelligent and programmable networking to connect applications to VMs, containers, and bare metal appliances. This solution offers ultimate flexibility as it operates in any data center environment, in any cloud management system, across any workload. It is also a highly scalable solution built from a proven high-performance routing stack.

Nokia AirFrame OpenRack Hardware represents the hardware components of the NFVI Telco Cloud blueprint. Nokia AirFrame offers extremely low TCO with industry leading density per rack, low power consumption and lean operations.

Figure 2. The elements of Nokia's NFVI Telco Cloud blueprint solution

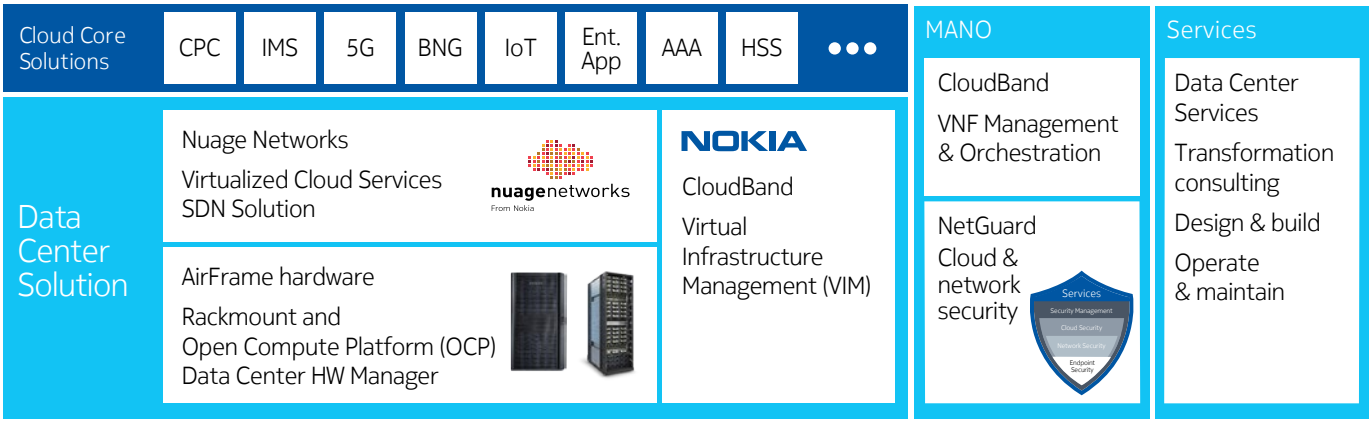

The NFV Management and Orchestration (MANO)

Nokia Cloudband provides the complete MANO architecture with the following major components:

- CloudBand Infrastructure Software (CBIS) - represents Virtualized Infrastructure Manager (VIM) from the MANO architecture. CBIS is a ready to use, open source based virtual infrastructure that provides serviceability, operability and universal applicability for any type of workload in an NFVI domain.
- CloudBand Application Manager (CBAM) - represents a ready-to-use Generic Virtualized Network Function Manager (G-VNFM) from the MANO architecture. It automates VNF lifecycle management and cloud resource management.
- CloudBand Network Director (CBND) - represents the NFV Orchestrator (VNFO) from the MANO architecture. CBND automates network services delivery and operation in a distributed, multi-tenant, multi-vendor environment while optimizing and governing the usage of the infrastructure resources.

Nokia NFVI Telco Cloud Services

Nokia's NFVI Telco Cloud services team can be leveraged to provide transformation consulting services, design & build services, and operation & maintenance services. Nokia NFVI Telco Cloud services have offices around the world that include cloud design centers, delivery centers, and cloud labs. Staffed by the best cloud engineers in the world, CSP customers have witnessed an increase in quality and speed of delivery while lowering cost and mitigating risk through these offerings.

Nokia NFVI Telco Cloud blueprint benefits

By adopting Nokia NFVI Telco Cloud blueprint, CSPs will realize the following benefits:

- Fast deployment (time to value) - fast track deployment from initial PO through acceptance to commercial use
- Stay up-to-date with latest capabilities - stay current with latest OpenStack, SDN and x86-based hardware releases & capabilities
- Fast and proven system upgrade path - system life cycle management with an assured forward-path based on certified systems. Optimization and pre-certification of upgrades for shortest upgrade duration and outage
- Off-load integration burden, cost and risk - safe-harbor certification of product release combinations covering functional, operational, performance & robustness testing

NOKIA

Nuage Networks VCS

VCS for Telco Cloud in 7 points:
Provides support for all major cloud management systems, hypervisors, and network gear. VCS leverages VMs on any x86-based hardware
Uses programmable business logic and policies to fully automate and simplify network service creation
Provides support for L2-L4 services including programmable security services
Optimizes and scales Telco Cloud connectivity and is deployable on heterogeneous networks
Offers unrestricted placement of VM, container or bare metal workloads to maximize efficiency of server resources
Includes extensive data analytics and performance monitoring capabilities
Integrates public, private and hybrid cloud applications into managed VPNs

Nuage Networks Virtualized Cloud Services (VCS) is a Software-Defined Networking (SDN) solution that provides network virtualization and advanced automation across any Telco Cloud data center infrastructure and automatically establishes connectivity between virtualized compute resources whether virtual machines, containers, or legacy bare metal servers upon their creation ensuring application traffic is served by all functions reliably and efficiently. Leveraging programmable business logic and a powerful policy engine, VCS provides an open and highly responsive solution that scales to meet the stringent needs of massive multi-tenant Telco Clouds. VCS is a software solution that can be deployed over any existing datacenter network environment.

VCS network abstraction

Nuage Networks VCS allows enterprise administrators to define their networking requirements in application terms, without being burdened or slowed down by network implementation details. CSPs can express security requirements (e.g. firewall and ACL policies), SLA requirements, load balancing, user access-rights, and more in an intuitive and abstracted way that is translated to network policies which are further translated to network configurable parameters.

Network behavior is also governed on an as-needed basis. By using an event-driven model with a policy pull approach, VCS reserves network resources as they are required, triggered by compute instance creation, change, migration or deletion. This ensures that the demands of cloud-based applications and services can be met across thousands of users in an efficient and timely manner. Refer to Figure 3 depicting VCS' abstraction model. In this case a secure virtual network is setup to provide network connectivity between appliances, VMs, and containers across a hybrid virtualized environment that includes bare metal servers (appliances).

Figure 3. VCS abstraction model

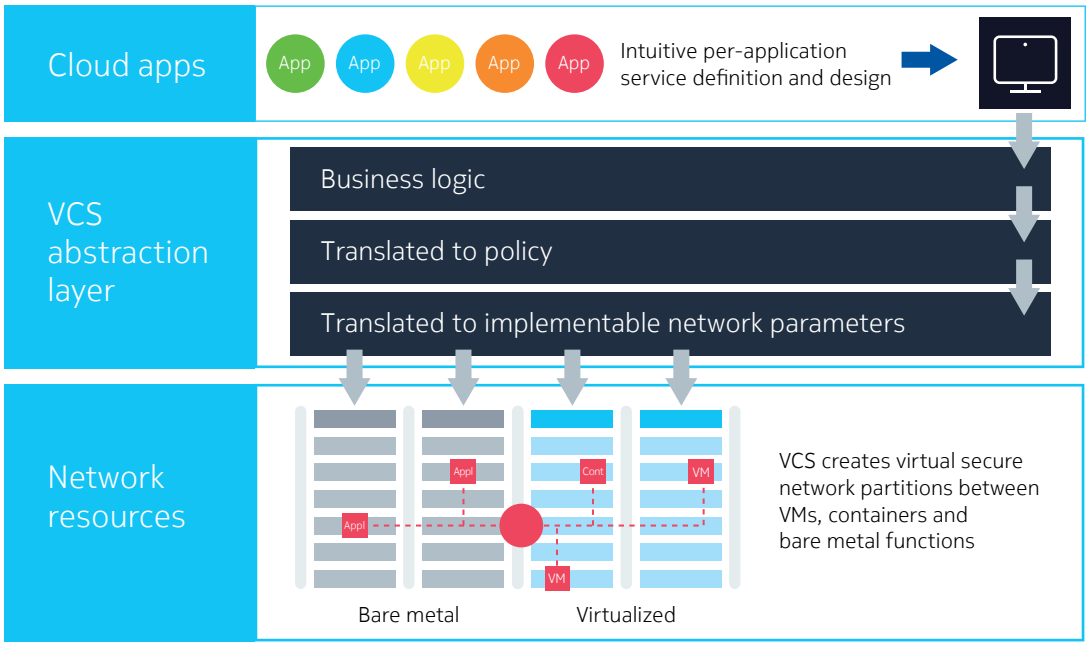

Cloud deployment of complex applications within a Telco Cloud requires more than automated L2/L3 connectivity. To meet these needs, VCS deploys the full range of L2-L4 networking services on a per-tenant and per-application basis. This ensures each application receives the services required. VCS enables seamless interoperability across administrative domains and with existing VPN services. It does this by leveraging Multi-Protocol BGP (MP-BGP) technologies to federate across multiple domains.

One of the key aspects of VCS is its flexibility. It can be deployed in any Telco Cloud environment independent of hypervisor, or cloud management systems deployed. VCS significantly improves Telco Cloud resource utilization by allowing VMs, containers and bare metal workloads to be freely placed wherever compute resources are available, within or across Telco Cloud Data Centers.

VCS offers CSPs the ultimate in Telco Cloud security with programmable security capabilities such as micro-segmentation to protect network resources from each application, per-application network analytics used to detect suspicious network activity while allowing for programmable remedial actions capable of reacting before the network is infected.

Nuage Networks VCS Telco Cloud architecture

Virtualized Services Platform

Nuage Networks VCS Telco Cloud architecture is based on Nuage Networks Virtualized Services Platform (VSP). VSP is a network automation platform enabling a complete range of SDN, SD-WAN, and cloud solutions. VSP provides advanced network automation across networks and clouds of all sizes and architectures from Telco Cloud data center private clouds to large enterprise wide area networks (WANs) to some of the largest public clouds in the world.

As shown in Figure 4 there are three network automation services that are all offered from VSP:

- Virtualized Cloud Service (VCS) – provides advanced network automation for Telco Clouds
- Virtualized Network Services (VNS) – provides advanced end-to-end SD-WAN services connecting branches, data centers, and clouds together with a single governance model
- Virtualized Security Services (VSS) – provides a set of automated security capabilities that protect the network, automatically detect threats and network anomalies, while providing an automated framework to program responses to ward off threats

Figure 4. The services that the VSP enables

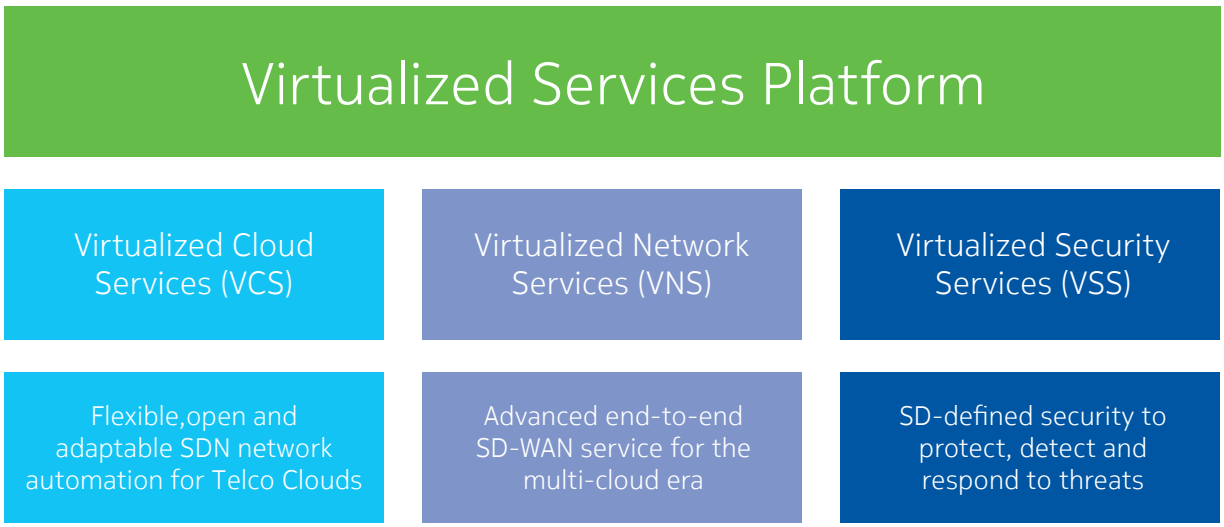

Architectural Component of Virtualize Cloud Services

Figure 5 shows a drawing depicting the architectural components of VCS.

Figure 5. The VCS Telco Cloud architecture

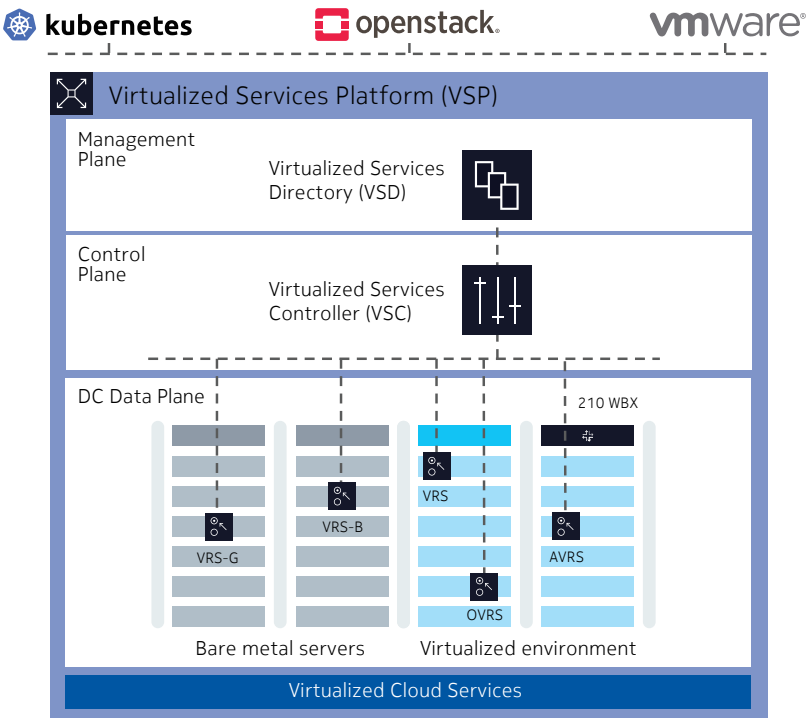

Virtualized Services Directory

The Virtualized Services Directory (VSD) is a programmable policy and analytics engine. It provides a flexible and hierarchical network policy framework that enables IT administrators to define and enforce network resource policies in an intuitive and user-friendly manner.

VSD contains a multi-tenanted service directory that supports role-based administration of users, compute and network resources. For service assurance, VSD allows the definition of sophisticated statistics rules such as collection frequencies, rolling averages and samples, as well as Threshold Crossing Alerts (TCAs). When a TCA occurs, it will trigger an event that can be exported to external systems. Statistics are aggregated over hours, days and months and stored to facilitate data mining and performance reporting. VSD offers all functionality via Restful API allowing third party management and orchestration systems full secure access. VSD is deployed in a cluster solution to drive resiliency.

Virtualized Services Controller

The Virtualized Services Controller (VSC) is a multi-tenant SDN controller that functions as the network control plane for data centers, maintaining a full view of each tenant's network and service topologies. Through VSC, virtual routing and switching constructs are established to program the network forwarding plane using the OpenFlow™ protocol. Multiple VSC instances can be federated within and across datacenters by leveraging MP-BGP — a proven and highly scalable network technology.

Virtual Routing and Switching

Virtual Routing and Switching (VRS) is a software module that is installed in the hypervisor layer for VMs or as part of a container structure in virtualized server environments. It creates and manages the virtual endpoints (i.e. VXLAN Tunnel Endpoint (VTEP)) that are used for the virtual "overlay" tunnels between the workloads (e.g. VMs and containers) within a prescribed overlay VPN. These tunnels are created by adding a VXLAN encapsulation to the original Ethernet frame from designated traffic allowing it to be directly routable to other VMs or containers within that overlay VPN.

As shown in Figure 6, VRS is based on the Open vSwitch (OVS) which is an open-source implementation of a distributed virtual multilayer switch which also provides the VTEP function. In addition to providing the virtual switching and forwarding plane for VMs and containers, VRS also offers L2-L4 capabilities such as offering a distributed L4 ACL-based firewall.

Figure 6. The Anatomy of the VRS

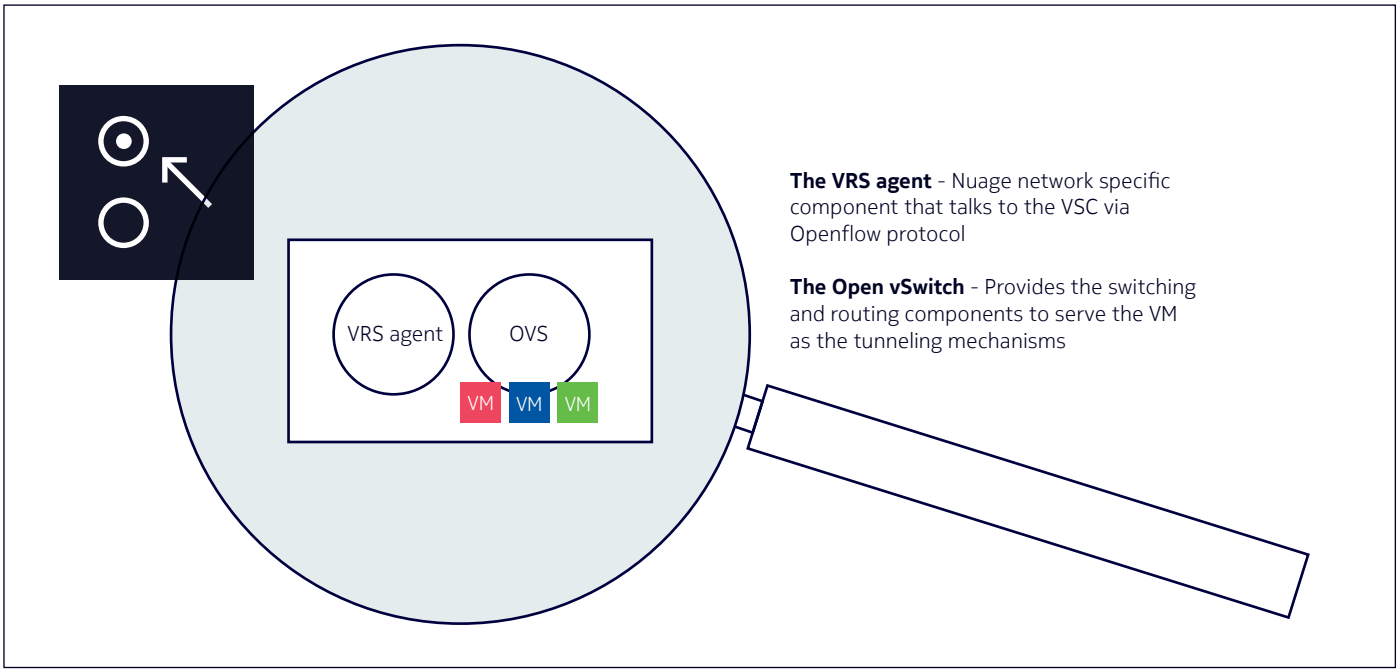

The other component of VRS is the “VRS agent” which is responsible for communication to the VSC through the Openflow protocol. Through this channel, changes in the compute environment (e.g. change in VM location, a new VM deployed) are immediately detected by the agent triggering instantaneous policy-based responses in network connectivity to ensure application connectivity and performance are not compromised.

There are two other variants of the VRS. The Accelerated VRS (AVRS) and the Offload VRS (OVRS)). The AVRS leverages Data Plane Development Kit (DPDK) technology to help improve packet performance by bypassing the Linux kernel while accessing dedicated processing cores for SDN packet processing. The OVRS leverages SmartNIC technology to offload SDN packet processing from the Linux kernel to the SmartNIC’s own processing resources.

Each of these VRS variants can be configured and deployed on bare metal servers to support a couple of different modes of operation: the gateway mode (G) and the bare metal mode (B). VRS-G is the software gateway that allows the aggregation of bare metal functions through their VLAN IDs. VRS-G provides full VTEP capabilities thus allowing bare metal functions to be part of an SD-WAN overlay VPN concurrently with virtualized functions running on VMs and containers. VRS-B allows a bare metal server and its network function (e.g. WAN optimization, L7 firewall, load balancer, etc.) to participate in a given SD-WAN overlay VPN.

HW VTEPs: 210 WBX or third-party equivalents

210 WBX is a purpose-built switch designed to meet the demands of next generation data centers and cloud services. It is deployed as a scale-out underlay leaf or spine node in the Clos-based data center architecture and is equipped with advanced IP routing and L2-L4 capabilities. It serves as a powerful SDN VTEP to create Layer 2 and Layer 3 Virtual Private Network (VPN) services in the overlay. It also provides data plane aggregation of bare metal servers and appliances while providing a VTEP functionality between these legacy assets and other virtual endpoints in the network. This allows functions hosted on legacy bare metal resources to be part of the network automation with no restrictions.

Although 210 WBX is a pre-integrated component of Nokia’s NFVI Telco Cloud blueprint other HW VTEPs can also be integrated into the overall solution with a limited amount of interoperability testing.

NOKIA

Packet acceleration methods

As an overlay SDN solution, VCS uses tunneling protocols such as VXLAN to further encapsulate the traffic. Processing encapsulated packets can cause significant network I/O performance degradation and can create large CPU overhead, especially as server I/O speeds increase and packet sizes decrease.

To overcome these inherent performance challenges, packet acceleration techniques based on offloading processing burden from the Linux kernel (i.e. OVS offload) are needed. VCS offers various OVS offload techniques to accelerate packet processing that are supported in both VM and container environments. Figure 7 summarizes the various offload methods to accelerate packet processing.

VRS offloading

In a VRS environment OVS packet processing is done in the Linux kernel assisted by stateless offloading in the server NIC for certain functions such as segmentation and checksum offloads. Post processed packets are then sent back through the OVS kernel to the relevant VM. The advantage of this method is that since nearly all data center NICs support stateless offloads, this approach is hardware independent offering a very flexible approach. This approach speeds up packet processing and is efficient for larger packet sizes but is not optimized for application flows with smaller IP packets which require much more processing. This approach is well suited for management traffic or enterprise applications that require less stringent processing.

AVRS - OVS offload using DPDK

In an AVRS environment DPDK technology is leveraged to accelerate the processing of VXLAN encapsulated packets by bypassing the Linux OVS kernel completely resulting in an accelerated “fast path” for packet processing. Performance is improved dramatically as a fixed number of dedicated cores running in user space are programmed to focus only on packet processing. This approach is also hardware independent provided the NIC it will use has a DPDK Poll Mode Driver (PMD) installed. The disadvantage of this approach is that these dedicated CPU cores can only be used for the packet processing. In addition, during periods of inactivity, these expensive CPU cores can spin in loops, while waiting for packets to arrive. This approach is well suited for applications with more stringent processing requirements and with smaller packet sizes.

Figure 7. Summary of Packet Acceleration Methods

	VRS	AVRS	OVRS w/SR-IOV	WBX w/SR-IOV
Description	A module in the hypervisor performing OVS switching and VTEP	Accelerates packet processing using a DPDK-based ‘fast path’ process running in the Linux user space	Offloads packet processing to the eSwitch on SmartNIC	VM traffic bypasses the hypervisor and VRS and is terminated at the WBX
Best suited traffic	Management traffic, large packets	Packet flows with smaller packet sizes and VNFs with moderate processing requirements	Processing intensive VNFs like 5G/LTE Data Path, BNG data path, 5G xHaul data path	Processing intensive VNFs like 5G/LTE Data Path, BNG data path, 5G xHaul data path
Strengths	Flexible as it is HW independent, deployed in any environment	DPDK accelerates packet processing over VRS with little HW dependence	High performance	High performance
Limitation	Performance for intensive VNFs and packet processing req.	Compute cycles could be stranded	HW dependency	Least control over traffic

OVRS - OVS offload using SmartNIC technology

In an OVRS environment, Nuage Networks has partnered with SmartNIC vendors to leverage their resident switching and packet processing capabilities for OVS offloading. This solution combines the performance and efficiency of networking hardware on the SmartNIC with the flexibility of VRS all while leveraging the Single-Root Input/Output Virtualization (SR-IOV) standard. This offload approach scales performance linearly with the number of NICs to achieve some of the highest packet performance results. This approach is ideally suited for the most stringent and demanding applications such as the infrastructure components of the 5G packet core.

VLAN direct to HW VTEP

In this case packets avoid the hypervisor and xVRS completely and are terminated on an external HW VTEP such as the Nuage Networks 210 WBX or other third party HW VTEP. This approach uses SR-IOV to leverage the external HW VTEP resources to provide packet processing and VM communication. VCS provides full programmability to leverage the HW VTEP processing to support SR-IOV workloads in an OpenStack environment This method provides high performance processing at the cost of control.

VCS Telco Cloud in action

Intra-DC Telco Cloud network automation

Most large modern Telco Cloud data center networks are built with a spine/leaf architecture where every leaf switch is connected to every spine switch. This approach optimizes server to server traffic within the Telco Cloud. These are also referred to as Clos networks and this is the model that will be used in the rest of this paper to demonstrate VCS capabilities within the Telco Cloud environment.

In Figure 8 there are four leaf nodes connected to two spine nodes. Each spine has connectivity to every leaf node enabling easy network access across each rack of servers. In this drawing there are three sets of VMs that are associated with three different tenants supporting a unique application. Secure network connectivity is required between each set of VMs. The green VMs represent tenant #1's cloud packet core workloads, the red VMs represent tenant #2's IoT workloads, and the blue VMs represent tenant #3's BNG workloads. Assume that there is VRS software installed in the hypervisor layer of each of the servers and is used as a VTEP for each of the VMs.

VCS will allow the IT administrator to program VXLAN L2 tunnels (or L3 tunnels) from each of the VTEPs to provide network connectivity between each VM as shown in Figure 9. The red lines in the drawing represent a single overlay domain created between the VTEPs creating connectivity between all the red VMs. As the environment changes the network connectivity will adapt. For example, if one of the VMs is deleted and re-established on another server, the VRS will detect this change and send a notification up to the VCS whereupon a new tunnel will be created to connect this VM from its new location to all other VMs in this domain.

Figure 8. The virtualized environment showing VMs from three tenants

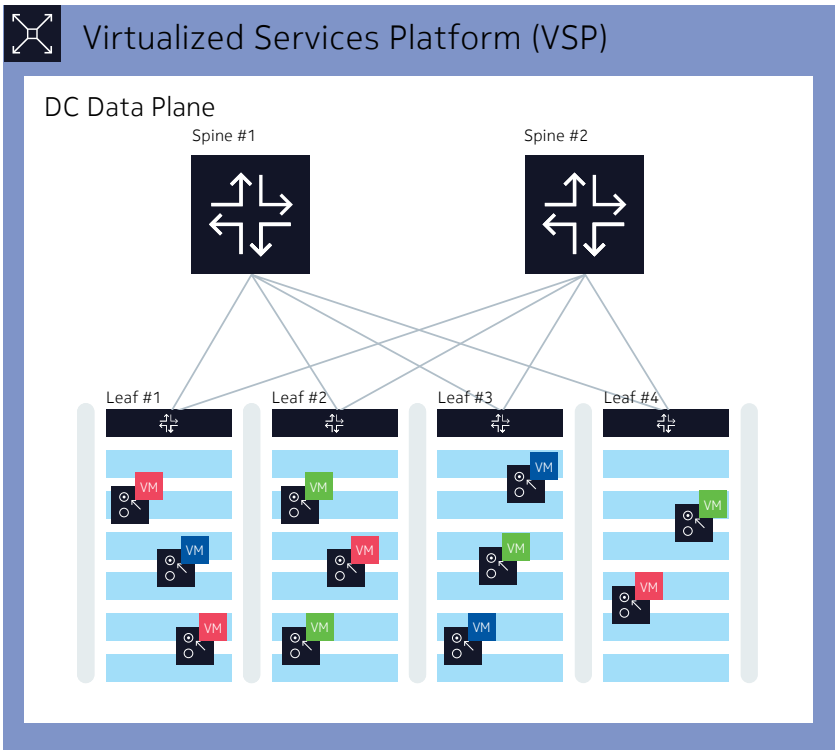

Figure 9. Overlay VPN connecting the red VMs

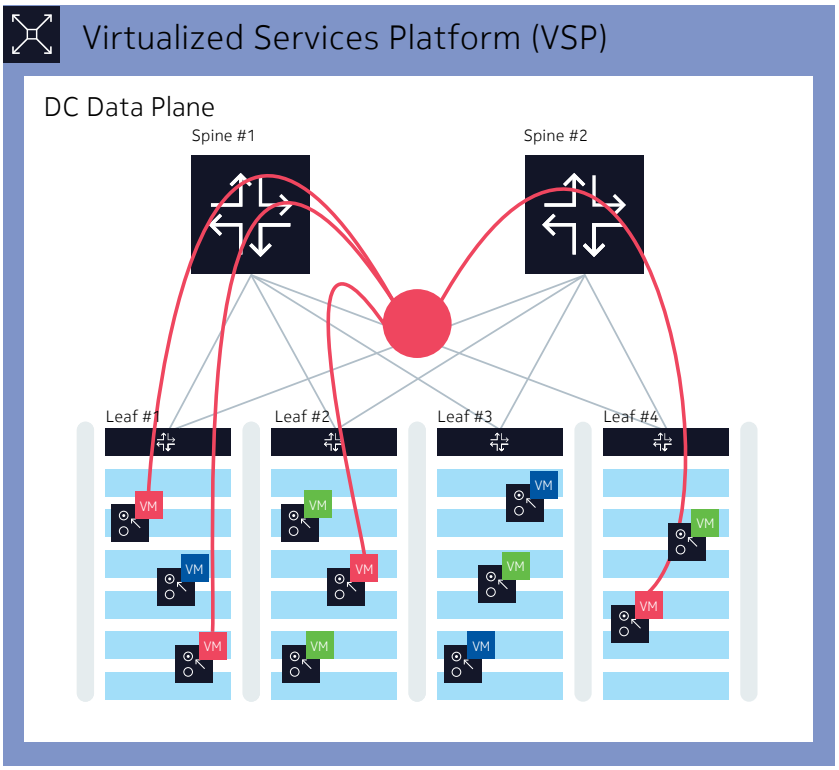

VCS provides full support for applications that run on legacy bare metal servers or appliances. Figure 10 shows how the 210 WBX can be used to aggregate applications run on these legacy servers so they can also participate seamlessly in the overlay VPN. The dotted red lines represent traffic from the legacy appliances. The 210 WBX aggregates traffic from an application on an appliance by identifying the L2 flow using the VLAN ID. It then creates a VTEP to connect application packet flows from these appliances with the rest of the overlay VPN.

Figure 10. VCS in action aggregating functions running on bare metal

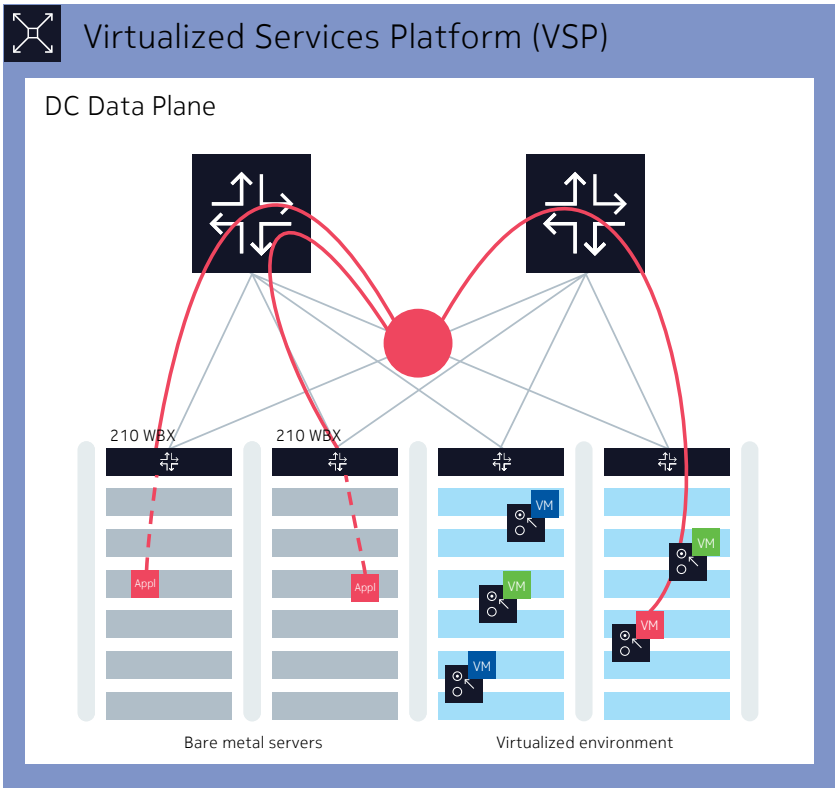

Inter-DC Telco Cloud network automation

VCS offers several flexible deployment models to provide inter-data center network automation.

Figure 11. VCS in action for Telco Clouds that span data centers

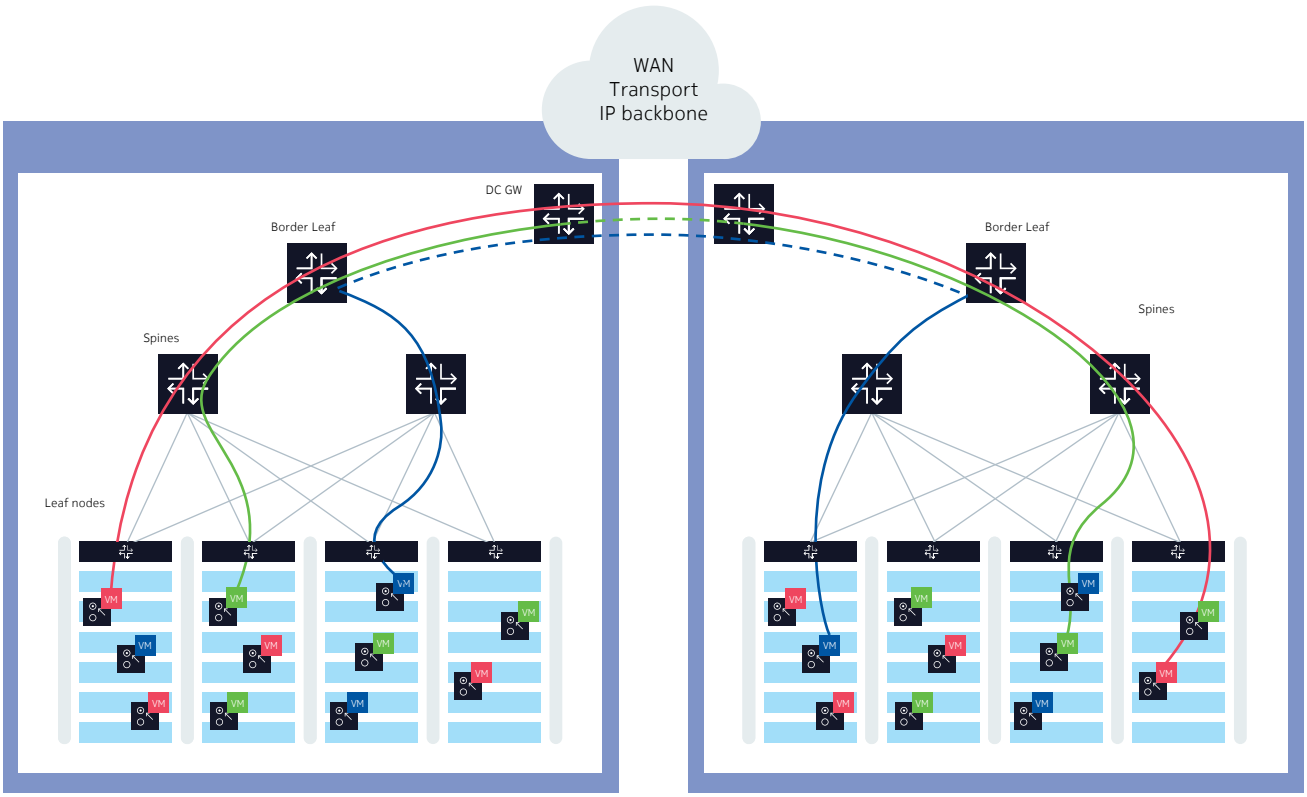

Inter-DC end-to-end VXLAN tunnels

The traditional approach to connect workloads (e.g. VMs, containers, or appliances) from one data center to the next is to extend the VXLAN tunnel across the WAN to the virtual endpoint or VTEP in the other data center. This model creates one large seamless virtual data center where workloads can be connected through VXLAN tunnels that span DCs. For this model to work the network must be running Ethernet VPN (EVPN) routing across the WAN so the destination VTEP can be discovered and the VXLAN tunnel can be created to extend across the network from VTEP to VTEP. In Figure 11 this model is depicted with the contiguous red tunnel.

VXLAN tunnel termination on border leaf nodes

The other approach is to terminate the VXLAN tunnels within the edge of each data center and leverage existing WAN routing protocols to transport the traffic through a dedicated path. An MPLS path or a segment route can be used. This requires some additional “stitching” or configuration. In Figure 11 this model is depicted with the blue tunnel and the WAN segment is depicted by a dotted blue line.

It is important to note that terminating VXLAN tunnels will necessitate that all service state for the upstream EVPN or VPRN addresses from the WAN will exist on the Border Leaf as it is effectively acting like a Provider Edge (PE) Router for the data center. In a Telco Cloud environment services are often exposed to

large routing domains driving up scaling requirements. Because of this scaling requirement the Border Leaf may not be ideal for Telco Clouds that have large service or control plane scaling requirements.

In this model SDN service provisioning ends on the Border Leaf node. This means that VLAN management for the Border Leaf PE-Customer Edge (CE) connection into the Data Centre Gateway (DC-GW) and provisioning of the DC-GW are responsibilities outside of the SDN controller. These tasks would typically become the responsibility of a Cloud Management System or service orchestrator.

VXLAN tunnel termination on DC GWs

The more scalable solution for large Telco Cloud deployments is to terminate the VXLAN tunnels in the Data Centre Gateway (DC-GW) directly. This will also require some "stitching" or configuration work to interwork the tunnels into the MPLS control and data-plane. However, a DC-GW is usually equipped with the technology and resources to be a PE router offering massive service scale. By using this model an MPLS VPN or segment route can be used to securely deliver tenant tunnel information across the WAN. In Figure 11 this model is depicted with the green tunnel and the WAN segment is depicted by a dotted green line.

Another advantage of this model is the SDN service provisioning model. VCS SDN can provision services directly onto the DC-GW through its Netconf interface.

As part of the Nokia's NFVI Telco Cloud blueprint solution the Nokia 7750 SR was modeled as the DC-GW. However, other third-party DC-GWs can also be managed by Nuage VCS.

Network Functions Interconnect to address 5G and network slicing

5G network slicing is a network architecture that enables the multiplexing of independent logical networks on the same physical network infrastructure. As network slicing is an end-to-end concept, it also impacts the transport network and how slices are extended into the Telco Cloud data center environment.

The previously discussed models where the VXLAN tunnels are terminated on the Border Leaf or DC GW, represent a termination point for services. Since this is a service stitching point and service state needs to be maintained, true end-to-end path control is not possible in this model. Since the stitching points are always intermediate points in the end-to-end path, implementing end-to-end SLAs is compromised and the model may not scale as efficiently as with a true end-to-end implementation.

Nokia offers a solution that addresses this called Network Functions Interconnect (NF-IX). In this architecture, the core IP/MPLS network WAN leverages Segment Routing (SR) which is extended into the Telco Cloud data center through Segment Routing over UDP (SRoUDP). In this architecture the DC-GW translates encapsulations (SRoMPLS to/from SRoUDP) but does not terminate services directly therefore full end-to-end path visibility and control are maintained.

This model also introduces the Segment Router Interconnect Controller (SRIC), which is a path computation element (PCE) that is updated constantly from the BGP (EVPN) control plane. When a VM or container is instantiated as part of a new virtualized service construct the associated SLA requirements for network function interconnections are advertised via BGP. The BGP (EVPN) control plane propagates these updates to SRIC, which computes the optimal route to support the service's SLA.

Conclusion

Telco Clouds are the service engine for modern CSPs and network automation is essential to accommodate the future service and scale requirements that cloud-based architectures have unleased. With 5G technologies and IoT ramping up Telco Clouds need to add network scale and flexibility to their network infrastructure. In addition, CSPs will need to deploy new services will agility and reliability. With Nuage Networks VCS, CSPs will be able to deploy an SDN solution that fits their specific environment while offering dependable scale with no performance degradation. It is what is needed for the next generation of Telco Clouds.

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

Document code: SR1912040318EN (January) CID207030