NETWORK FUNCTIONS VIRTUALIZATION – CHALLENGES AND SOLUTIONS

STRATEGIC WHITE PAPER

EXECUTIVE SUMMARY

Network functions virtualization (NFV) will revolutionize the way telecommunications networks are being built and operated. NFV promises benefits in two main areas:

- Operational and capital expenditure (OPEX and CAPEX) savings obtained from using general purpose hardware;
- Increased automation, resulting in operations simplification, business agility, and faster time to market.

Furthermore, NFV is expected to create an environment that fosters innovation and new services, thus becoming an important tool for future network monetization.

To fully benefit from NFV, several challenges need to be overcome. Specifically:

- Operators need to deploy a carrier-grade cloud platform, leveraging de-facto industry cloud management standards such as OpenStack application programming interfaces (APIs);
- The scale and fluidity of virtualized compute and storage configurations need to be matched by the ability to control and dynamically change network behavior. Software defined networking (SDN) is an important technology to meet this need;
- Operators need to deploy industry-standard tools for automated on-boarding, deployment and scaling of virtualized network functions;
- NFV will introduce new security concerns, but automation through tools and solutions related to and enabled by NFV will enable service providers to address these concerns effectively and to even improve security;
- The nature of operations and business management will change. Some functionality currently supported by operations and business support systems (OSS/BSS) will move to new, real-time management solutions that handle network-wide resource and service management.

Virtualizing certain telecommunications functions is straightforward, but many network functions have stringent real-time and reliability requirements. Virtualizing those network functions requires fine calibration between tasks supported by the cloud platform and tasks under the control of the network functions themselves.

Alcatel-Lucent is a leader in NFV. The CloudBand offering is an industry-leading cloud infrastructure and management solution. Alcatel-Lucent venture Nuage Networks provides a solution for zero-touch, SDN-based network control. Several network functions offered by Alcatel-Lucent are already available in a virtualized form, and for many others virtualization is a near-term roadmap item. Alcatel-Lucent has also deeply analyzed the impact of virtualization on reliability, availability and operations, and is building a consultancy practice to assist operators transition to NFV.

TABLE OF CONTENTS

INTRODUCTION

Network Functions Virtualization

The term network functions virtualization (NFV) was introduced in a whitepaper co-authored by several large telecommunications service providers [1]. The whitepaper and the subsequent creation of an NFV Industry Specification Group (ISG) in the European Telecommunications Standards Institute (ETSI) have put NFV at the center of industry attention, together with complementary initiatives on software defined networking (SDN).

NFV promises many benefits and will revolutionize the way telecommunications networks are built and operated. To fully benefit from NFV, several challenges need to be overcome. This whitepaper identifies the challenges and proposes solutions.

NFV BENEFITS

The original operators' whitepaper [1] described the benefits of NFV in detail. These benefits are primarily conveyed in two areas:

1. OPEX and CAPEX savings, due to, among other things, the use of lower-priced general purpose hardware; the ability to share computing resources between functions due to hypervisor technology; reduced energy consumption and associated carbon emissions; the ability to leverage a more widely available skills base for operating cloud infrastructures; and, efficiencies in sparing.
2. Automation gain. Service providers can use the introduction of virtualized networking and cloud technologies to adopt tools similar to those used by the information technology (IT) industry to automate many aspects of operations and management. This can lead to significant reductions in operational expenses, but more importantly, it will enable service providers to meet the needs of the telecommunications market through faster service introduction, automated scaling of resources up and down to meet changing demands, and the ability to continuously optimize resource allocation based on the results of sophisticated analytics-based algorithms.

Furthermore, NFV is expected to create an environment that enables new business models and services, increased innovation, and prompts new vendors to enter the telecommunications market place. This will result in NFV creating new ways to monetize telecommunications infrastructure.

To achieve these benefits, service providers need to evolve their infrastructures as well as their operations and business management practices. Figure 1 depicts several evolutionary stages service providers may go through as they deploy virtualized network functions.

Figure 1. NFV Deployment Evolution

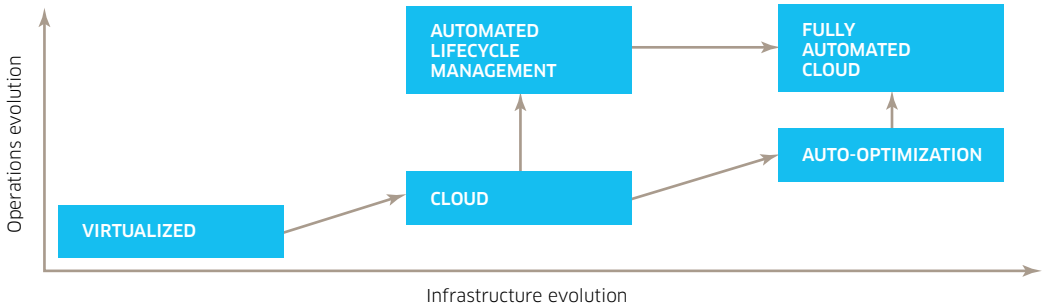

Virtualized is a primary NFV evolutionary stage in which network functions run on hypervisors and general purpose hardware, but may require dedicated physical resources, selected hypervisors and customized configuration solutions.

Cloud is the stage in which virtualization is based on standard interfaces and where virtualized network functions (vNFs) are interoperable with – and portable between – hardware platforms, hypervisors, and cloud resource control[1] systems of different vendors.

Automated Lifecycle Management is the stage in which service providers use tools – similar to the ones deployed in the IT world but adapted to carrier requirements – to manage the vNF lifecycle from onboarding to phase out.

Auto-optimization is the stage in which vNFs are able to dynamically and automatically adapt to changes in service demand (e.g., by scaling up and down) and in which vNF placement and configuration can be automatically modified to optimize resource use or improve customer experience, for example through reduced latency.

A service provider could – and often will – be in different deployment stages simultaneously. The provider may, for example, deploy several virtualized network functions that adhere to industry-standard interfaces and that can be managed through automated lifecycle tools, while other functions may only be available in a virtualized form.

Once a service provider reaches the **Fully Automated Cloud** stage for all virtualizable network functions, all potential benefits of NFV can be realized.

The remainder of the paper will analyze various aspects of these evolution stages.

THE NEED FOR CARRIER-GRADE INFRASTRUCTURE

The infrastructure required for a virtualized telecommunications network differs from the cloud as it is currently used in the IT world in some or all of the following points:

1. Reliability and availability requirements are often much more stringent.
2. Latency is a critical attribute for many telecommunications services. The placement of – and connectivity between – virtual machines (VMs) that constitute a vNF may need to take quality of service (QoS) requirements into account.
3. Functions in a telecommunication network are interconnected in complex arrangements. This requires additional networking considerations not generally found in the IT world.
4. The infrastructure will likely be much more distributed for two reasons: 1) telco service providers have many geographically distributed central offices they can use to host parts of the virtualized infrastructure; and, 2) telco service providers offer services that benefit from being located as close to the end-user as possible.
5. NFV needs to be massively scalable to support millions of subscribers while meeting the above requirements.

1 Cloud resource control is the function responsible for the allocation of cloud compute and storage resources by administering the hypervisors' settings. The open-source software created by the OpenStack community is an example of this function. Cloud resource control is sometimes referred to as *Cloud OS*.

This whitepaper discusses how service providers can meet reliability, availability and security requirements in a cloud context. Also addressed is service providers' need for a cloud management and orchestration function that facilitates and simplifies management of a large, distributed infrastructure and a networking solution that supports complex networking needs.

WHAT TO VIRTUALIZE?

Figure 2 shows an assessment of the benefits of virtualizing various network functions. Network appliances refers to functions such as firewalls and load balancers. Network apps refers to functions such as IP Multimedia Subsystem (IMS) telephony application servers or presence servers.

Figure 2. Virtualization benefits

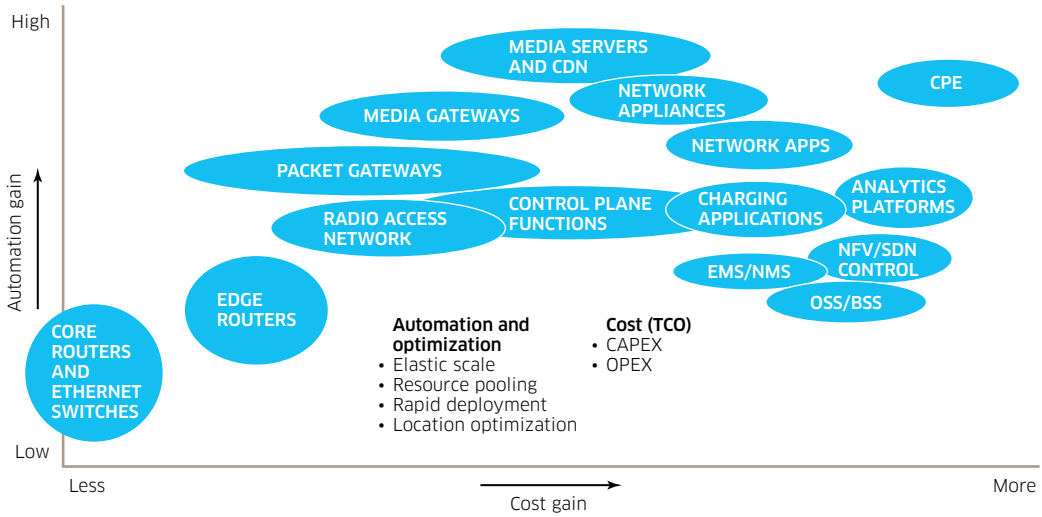

In a majority of cases, service providers will benefit from virtualizing network functions, even functions that perform real-time packet processing and whose current implementation uses network processors or application-specific integrated circuits (ASICs). Implementing the same functions on general purpose processors (GPP) may increase CAPEX due to the higher number of server blades required, but this increase would be offset by OPEX savings and automation gain.

There are a few exceptions. For example, it is unlikely virtualizing high-performance routers or Ethernet switches will be cost effective, and several other network functions are in a "grey" zone. Virtualizing products primarily focused on packet forwarding could result in cost savings in certain deployment scenarios but not in others, depending, for example, on the ratio of control traffic versus data plane traffic. But even when there are no cost savings, virtualization could sometimes be justified by automation gain arguments, such as the ability to roll out new services faster.

RELIABILITY AND SERVICE LEVEL AGREEMENTS

Many network functions within telecommunications infrastructure have stringent requirements for latency, reliability and availability. Virtualizing these functions poses challenges generally not encountered in the IT world and which are unique to NFV. For instance, temporary unavailability of a sales support system or a video-on-demand service may be extremely frustrating; temporary unavailability of telephony service could mean the difference between life and death.

Today, vendors design products that combine software, hardware and networking (i.e., the backplane) in a single box. They design the product to meet the specific performance, reliability and availability targets directly associated with service level agreements (SLAs). In a virtualized environment, application software is decoupled from hardware and hypervisor software. Furthermore, distribution of VMs across different equipment racks (or even different datacenters) introduces new potential points of failure. The evolution to NFV will, therefore, have profound implications on the way reliability and SLAs are managed. In particular:

- Focus will shift from reliability and availability per network element to end-to-end service availability. This will require new systems for monitoring, analyzing and managing the end-to-end infrastructure.
- The industry will need to define new frameworks for expressing the reliability and availability metrics of individual components – hypervisors, VMs, vNFs and networking domains – in a way that enables operators to specify requirements on those components and predict the availability of an end-to-end service.
- Service providers will likely have different operations teams for the physical infrastructure and for the virtualized software layer, which will have a significant impact on the way they operate.

Real-time behavior is another important attribute of telecommunications networks. Table 1 depicts different types of network functions in today's telecommunications solutions and the varying timing demands those functions require to maintain end-to-end QoS. Management plane functions (e.g., OSS, off-line charging and network element managers) have been designed to have high tolerance to long response times and delays, with little-to-no impact on the QoS and end user experience. Network functions in the applications and control and signaling planes (e.g., telephony application server, session control and subscriber databases) can tolerate small timing delays (up to seconds). Media services, packet processing, and networking infrastructure functions (e.g., media gateways, session border controllers and radio access network controllers) have a very low tolerance for long response times or timing delays.

Table 1. Timing requirements of various network functions

CATEGORY	REQUIREMENTS	EXAMPLES
Management	Non-real time	Element management, device management and off-line charging systems
Application	Near real time	Telephony application server; presence server; analytics solutions
Control and signaling	Near real time	Session control; subscriber data bases; policy management systems
Media services and packet processing	Real time	Media gateways; session border controllers; content delivery network distribution nodes; deep packet inspection
Networking infrastructure	Real time	Routers; switches

Virtualization impacts the ability of vNFs to meet their timing requirements. If multiple network functions share a single microprocessor core, each entity gets a time slice. This means a certain amount of additional latency may be incurred before a network function is activated. Alcatel-Lucent is working with processor and hypervisor suppliers to evolve capabilities that enable low-latency processing in a virtualized infrastructure. In the meantime, service providers may need to dedicate multi-processor cores to a single function for certain vNFs to avoid latencies due to time slicing.

SECURITY

Evolving to NFV will introduce several new security challenges:

- Due to the dispersion of VMs that belong to a vNF across racks and datacenters, and due to migration of VMs for optimization or maintenance purposes, the physical perimeters of network functions become blurred and fluid, making it practically impossible to manually define and manage security zones.
- The introduction of hypervisors creates new attack surfaces that could result in, among other things, compromised isolation between VMs.
- Hardware, hypervisors, vNFs and cloud resource control solutions may be provided by different vendors, increasing the risk of security holes due to mismatched assumptions and expectations.

Service providers already employ a comprehensive set of tools and best practices that can be used to address existing challenges and new ones that result from evolving to NFV. There are, however, three additional ways service providers can address the new challenges and potentially achieve even higher levels of security:

- While traditional networks firewalls are often large appliances that protect entire security zones, virtualizing firewall functionality enables service providers to instantiate smaller, dedicated virtual firewalls tailored to protecting specific network functions or domains. The same is true for load balancers, which play an important role in the protection against denial-of-service attacks.
- Automating the placement and associated provisioning of virtual firewalls enables service providers to create, manage and adjust security zones in concert with the configuration and placement of VMs.
- Increased reliance on monitoring and analytics for reliability and SLA management will also apply to security. For example, malicious activities by intruders or infected devices and systems can be detected through comprehensive, continuous monitoring before they cause wide-spread harm.

THE ADAPTABLE NETWORK

NFV will require the scalability and flexibility of networking infrastructure to be increased, as follows:

1. Each individual VM needs to be addressable.
2. NFV is more fluid than today's network due to situations such as VM migration and scaling. The fluidity of NFV must be matched by a similar ability to adapt the network on the fly.
3. NFV introduces a level of networking complexity far beyond what is required in current networks.

To understand the third point, consider the following example. A vNF, P, may have an internal configuration of sub-components (VMs), as shown in Figure 3. Load balancing (LB) is required to distribute load across C1, C2 and C3. For reliability purposes, VMs A and B have redundant components A' and B'. For security, the firewall functions FW and FW' are being deployed.

From an end-to-end networking perspective, vNF P needs to connect to other network functions, Q, R and S.

Figure 3. Abstract view of networking needs of an example vNF

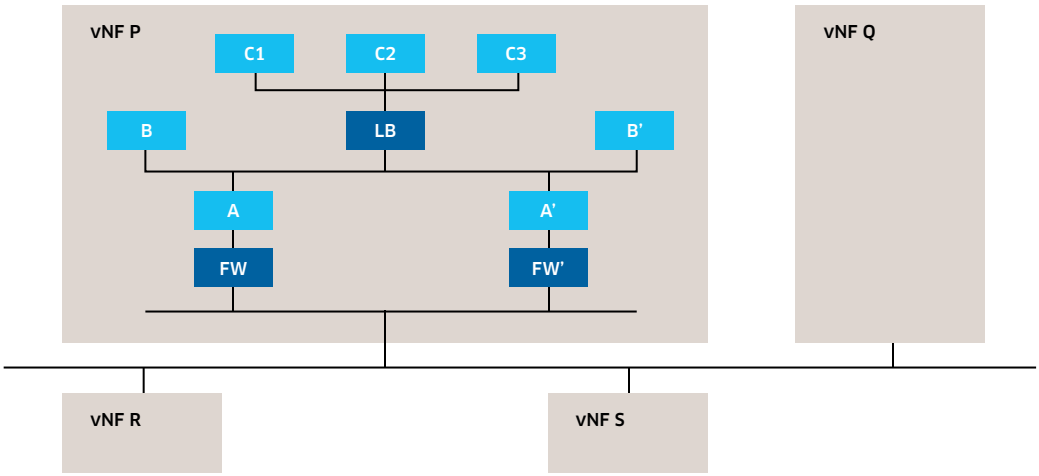

A specification of the desired configuration must address two considerations:

1. Rules that determine relative positioning of VMs, the so-called *affinity* and *anti-affinity* rules. For instance, A and A' need to be on different servers or in different racks to ensure they are not brought down by a single failure. Conversely, other components may need to be placed close enough to each other to minimize internal communication latencies and minimize the number of potential failure points between them.
2. Network design rules that capture the desired connections and forwarding rules, as well as requirements such as the maximum latency or other performance requirements between components.

First, the interconnections between VMs that constitute a single vNF form a virtual private network (VPN). The interconnections between distinct vNFs may also form a separate VPN. Clearly, the entire collection of VMs for all vNFs creates a complex network topology that may extend across multiple racks and even multiple datacenters. Managing such a network efficiently requires a solution that allows operators to specify networking requirements at an abstract level in terms of service requirements and policies, hiding all low-level details.

Second, large numbers of VPNs are required at Layer 2, Layer 3, and Layer 4. Supporting multiple VPNs is difficult with traditional techniques (e.g., using VLANs on traditional Ethernet switches). In general, NFV will require datacenter infrastructure designed specifically to support massively scalable VPNs.

Third, whenever changes are made to the cloud infrastructure – for example, when creating new VMs for capacity growth or migrating VMs from one location to another for optimization or maintenance purposes – network configuration needs to be adapted immediately. This requires new forwarding rules to be programmed in different parts of the network. As discussed in the previous section, this includes enforcing security zones (i.e., making sure that traffic flows through the appropriate firewalls and load balancers).

SDN provides the solution to meet these requirements. While multiple definitions of SDN are used in the industry, the two fundamental characteristics that appear in all definitions are *abstraction* and *programmability*. As explained above, those characteristics are needed to realize NFV.

Alcatel-Lucent regards SDN as the yin to NFV's yang. To achieve the full benefits of NFV, a carrier-grade, scalable and responsive networking solution is critical, first and foremost within datacenters, but also in the telecommunications network at large.

AUTOMATED LIFECYCLE MANAGEMENT

In the IT world, tools such as Cloudify, Puppet and Chef[2] are widely used to automate server configuration. With the virtualization of the telecommunications industry, service providers have the opportunity to introduce similar tools as part of their operational arsenal.

Figure 4. Automated vNF lifecycle management

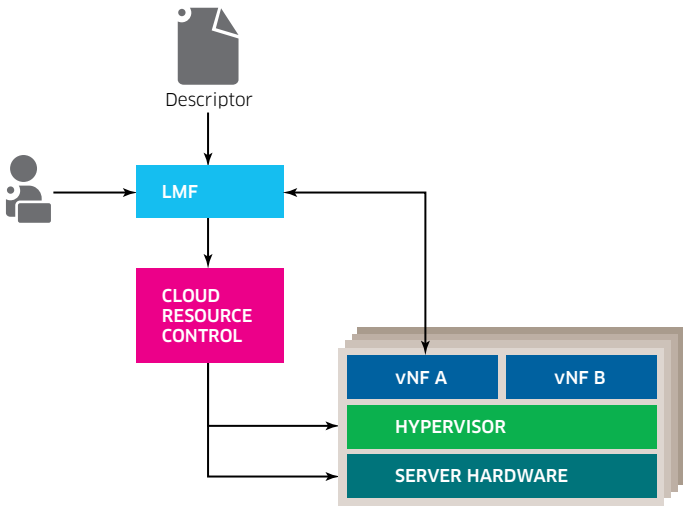

Figure 4 illustrates the role of a vNF lifecycle management function (LMF). It uses a descriptor that is generally provided by the vNF vendor. This descriptor defines the structure of the vNF (as it may consist of various sub-functions where each needs to run as an independent VM) and deployment and operational aspects, such as computation, storage and networking requirements. These descriptors are mapped to requests to the cloud resource control to create VMs and to identify software images to be downloaded and initiated on those VMs. Once the VMs are up and running, the LMF configures parameters of the vNF components based on instructions in the descriptor.

2 Cloudify, Puppet and Chef are open source tools created by GigaSpaces Technologies, Puppet Labs and Opscode, respectively.

Lifecycle management also plays a role in elasticity – scaling the vNF up and down– adding, removing or resizing VMs as needed. Triggered by either an operator decision, a vNF event (e.g., a traffic overload alert), or an infrastructure event (e.g., VM load crossing a defined threshold), the LMF executes the required actions.

In principle, every virtualized network function could come with its own proprietary lifecycle management, but service providers would obviously benefit from a consistent approach across all vNFs. The use of tools for automated lifecycle management and convergence on one or a small set of compatible tools will provide significant benefits to both, service providers and vendors.

IAAS OR PAAS?

Will vNFs rely only on infrastructure as a service (IaaS) or will they take advantage of platform as a service (PaaS)? With the former, vNFs run on a managed infrastructure of compute, storage and network resources, but are otherwise autonomous. In the latter, the PaaS services provide hooks and tools that make it easier to develop and deploy vNFs. One such set of hooks and tools concerns vNF lifecycle management, as described in the previous section.

Alcatel-Lucent believes both models will apply in parallel. Some network functions will benefit from hooks and tools provided by the platform; other functions are, by necessity, self-contained and will just use the physical resources according to the IaaS model. Ideally, however, all vNFs should rely on the automated installation and initialization functionality offered as part of an NFV PaaS.

An example that helps to understand the difference between the two approaches is elasticity. PaaS services could include monitoring processor loads, memory utilization or other parameters, and automatically allocating additional resources whenever certain thresholds are exceeded (e.g., initiating new VMs). vNFs, for which processing power is the primary constraint, could benefit from such PaaS scaling services and could be developed without having to consider elasticity control themselves.

Conversely, there are many vNFs for which the load of signaling or control messages is the primary constraint. Automatic elasticity of the vNF is not necessarily the right answer to a high signaling load, since it may trigger overload situations in other systems. A key responsibility of such vNFs is to deal intelligently with overload situations so the end-to-end integrity of the service is maintained. Such vNFs cannot, therefore, rely entirely on a PaaS scaling service but need to implement their own specific scaling functionality to achieve the required service delivery constraints.

FROM VERTICAL TO HORIZONTAL – THE NEED FOR NEW OSS

In the traditional model, OSS are specific to a network technology and form the funnel through which all operations and maintenance actions flow. This model will need to change, and the previous sections highlight two reasons why, namely:

- The need for automated lifecycle management requires new, more dynamic forms of resource management, ideally cutting across all vNFs.
- Decoupling hardware and software will result in an operational model in which the physical layer is managed independently from the virtualized software layer.

In addition, the ability to move workloads dynamically (e.g., VM migration for optimization purposes) will further change the nature of vNF operations. Whereas operations is currently a fairly static, top-down process, in the future it will require a paradigm similar to that used in mobile networks. Rather than a push model in which a central, omniscient OSS configures the network, the resource management of the NFV infrastructure will be more dynamic and autonomous and will pull policies from (logically) centralized policy systems to guide local decisions. Furthermore, to maintain a consolidated view of the cloud infrastructure, information will flow from the bottom up and the task of future OSS will be to aggregate and present dynamic state information in near real time for use in operational decisions.

The increasing use of big data analytics in telecommunications networks represents another change in the operations philosophy. Rather than implementing performance management for individual network elements, interfaces or network segments, service providers can now gather data from many sources and use sophisticated algorithms to both determine the state of the network and to better manage the service quality offered to end-users.

Current management solutions can be seen as a set of vertical silos predominantly focused on network elements (NE) or groups of NEs that make up a service. An NFV infrastructure will be much more horizontally focused, including: operations across compute and storage resources and across the end-to-end network; automated lifecycle management across all vNFs; and, analytics to determine end-to-end service quality. The future OSS will reflect these changes.

The changes outlined in this section are not unique to NFV. To some extent they are happening anyway, but are accelerated and amplified by the evolution to NFV.

Interestingly, the service providers' NFV whitepaper advocated introducing virtualization without changes to OSS and BSS systems. Ultimately such changes are inevitable, but the desire not to change current systems reflects a fundamental truth about NFV: evolution will take place gradually over several years.

ALCATEL-LUCENT AND NFV

Alcatel-Lucent provides a market-leading NFV solution that addresses the issues identified in the previous sections regarding compute and storage, and the networking aspects of the NFV infrastructure:

- Alcatel-Lucent CloudBand™ provides a carrier-grade cloud management and orchestration function that supports lifecycle management and other NFV PaaS services, among other things. It is the stepping stone towards a new, cloud-focused OSS, as discussed in the previous section.
- Nuage Networks provides an SDN-based networking solution. It is an open, software-based overlay that works with any existing vendor equipment and provides the functionality introduced in the section on the adaptable network.

These solutions can be used independently or in combination, for even greater benefit. This section discusses these two solutions in more detail. In addition, it discusses virtualizing other network functions produced by Alcatel-Lucent.

Alcatel-Lucent CloudBand

CloudBand consists of two distinct elements: the CloudBand Node and the CloudBand Management System.

The CloudBand Node is a pre-integrated, deployment-ready cloud infrastructure consisting of compute, storage and switching hardware, hypervisors, and cloud resource control software built for large-scale deployments. The installation of CloudBand Nodes is highly automated and takes less than three hours to go from bare metal to being fully functional.

The **CloudBand Management System** provides the following capabilities for managing cloud resources, virtualized network functions and applications:

- Cloud Management and Orchestration – This function manages and orchestrates resources across the end-to-end infrastructure. It leverages distributed cloud concepts and aligns with main cloud computing functions, including on-demand self-service, broad network access and resource pooling. It maintains a global cloud resource status view across all applications. The CloudBand Management System interfaces with the Nuage Networks Virtual Service Platform (see below) or other network controllers.

  The CloudBand Management System exposes its functionality to other entities via open APIs, and it interfaces via open – and commonly used proprietary – APIs to the underlying NFV infrastructure (e.g., CloudBand Nodes, or other vendors' infrastructures).
- Carrier PaaS – This function is responsible for vNF lifecycle management. It supports:
  - ¬ Automated application deployment: Installation, configuration and upgrade
  - ¬ Application monitoring and automated self healing
  - ¬ Automated scaling up and down according to circumstances, such as workload

  Carrier PaaS is hardware/platform agnostic. It can be used to manage the vNF lifecycle on any infrastructure, ranging from bare metal to a variety of cloud infrastructures. The Carrier PaaS function relies on the Cloud Management and Orchestration function for virtual resource allocation.

- Cloud Optimization – In a distributed carrier cloud environment, optimally locating virtual machines and virtual storage volumes is not a simple task. The CloudBand cloud optimization function, based on placement algorithms developed by Bell Labs, calculates optimal resource locations, taking into account vNF-specific policies, including availability requirements, (anti)affinity rules, inter-node latencies, capacity requirements, and regulatory requirements. These placement algorithms help ensure consistent service quality, enabling, for example, network distances to be minimized or utilization bottlenecks to be circumvented.
- Cloud Analytics – This function collects cloud events at the infrastructure and vNF/application level and provides input to the global cloud resource status view on an on-going basis. It also provides historical and real-time data processing and analysis, including event correlation, anomaly detection, automatic baseline, event prediction and root-cause analysis.

Nuage Networks Virtual Service Platform

The Nuage Networks **Virtualized Services Platform** (VSP) fully virtualizes and automates any datacenter network infrastructure, transforming it into a dynamic environment that rapidly establishes the network services required to deliver cloud services for thousands of applications in a policy-driven manner.

The VSP is a second-generation SDN solution for datacenter networks. The solution lays the foundation for an open and dynamically controlled datacenter network fabric to accelerate application programmability, facilitate unconstrained VM mobility, and maximize compute efficiency for cloud service providers, web-scale operators and leading tech enterprises.

The VSP eliminates the constraints that have traditionally held back datacenter network responsiveness and efficiency by:

- Making the datacenter network as dynamic and consumable as compute infrastructure through automated instantiation of network services;
- Eliminating the need to hair-pin and trombone intra-datacenter traffic;
- Eliminating the cumbersome configuration processes for datacenter networking that are today's norm;
- Separating and simplifying the definition of network service requirements and policies from the manner in which network services are established;
- Virtualizing any existing datacenter network infrastructure (Layer 2 through Layer 4); and,
- Scaling to meet the demands of thousands of applications, each with their own unique networking requirements.

The VSP consists of three elements:
Virtualized Services Controller (VSC) serves as the robust control plane of the datacenter network, maintaining an abstracted, full view of network and service topologies. Through network APIs and protocols such as OpenFlow, the VSC configures the datacenter network independent of datacenter networking hardware.

Virtualized Services Directory (VSD) serves as a policy, business logic, and analytics engine for the abstract definition of network services. Through RESTful APIs to the VSD, administrators can define and refine service designs and manage associated policies.

Virtualized Routing & Switching (VRS) software for any hypervisor platform extends the seamless control of network interfaces – across both virtualized and non-virtualized services and appliances – as extensions of the datacenter fabric.

Virtualized network functions

Several network functions offered by Alcatel-Lucent are already available in a virtualized form and virtualization is a near-term roadmap item for many others.

Alcatel-Lucent development of virtualized network functions with stringent real-time requirements is well advanced, using techniques such as appropriate CPU allocation mechanisms and solutions to reduce sensitivity to temporal constraints. For example, the Wireless Cloud Element is a fully virtualized, carrier-grade, high-performance solution specifically addressing radio access network (RAN) controller functions for WCDMA and LTE, such as the radio network controller (RNC) and the LTE evolved multimedia broadcast multicast service (eMBMS).

Figure 5 illustrates the scheduled availability of various Alcatel-Lucent products.

Figure 5. Virtualization roll-out sequence

Exact availability dates of specific virtualized network functions are determined by product management for the individual products. Contact an Alcatel-Lucent sales representative for further information.

SUMMARY

Network Function Virtualization can provide service providers with significant gains in automation and reductions in costs. With its engagement in industry initiatives, such as the ETSI NFV ISG, innovative solutions like CloudBand and Nuage Networks available now, and with numerous virtualized network functions available or under development, Alcatel-Lucent is committed to making NFV a success.

GLOSSARY

API	Application programming interface
BSS	Business support system
CAPEX	Capital expenditure
GPP	General purpose processing
IaaS	Infrastructure as a service
ISG	Industry Specification Group (working group in ETSI on a specific topic, such as NFV)
IT	Information technology
LMF	Lifecycle management function
LTE	Long Term Evolution (of 3GPP 4G technology)
NFV	Network functions virtualization
OPEX	Operational expenditure
OS	Operating system
OSS	Operations support system
PaaS	Platform as a service
PGW	Packet data network (PDN) gateway (function defined in LTE networks)
RAN	Radio access network
RNC	Radio network controller
SDN	Software defined networking
VM	Virtual machine
vNF	Virtualized network function
VPN	Virtual private network
WCDMA	Wideband Code Division Multiple Access

REFERENCES

[1] Network Functions Virtualization - An Introduction, Benefits, Enablers, Challenges & Call for Action; October 22 2012; Contributing organizations: AT&T, BT, CenturyLink, China Mobile, Colt, Deutsche Telekom, KDDI, NTT, Orange, Telecom Italia, Telefonica, Telstra, Verizon

CONTACTS

For further information please contact:
Peter Busschbach, Core Networks & Platforms Business Strategy
peter.busschbach@alcatel-lucent.com

Alcatel·Lucent

Alcatel·Lucent

Alcatel·Lucent

Alcatel·Lucent