WHY SERVICE PROVIDERS NEED AN NFV PLATFORM

STRATEGIC WHITE PAPER

Network Functions Virtualization (NFV) brings proven cloud computing and IT technologies into the networking domain to help service providers reduce equipment and operational costs, power consumption, and time-to-market for new services and functionality. But IT technologies alone are not enough. Service provider applications are more demanding than most IT applications. To meet these requirements, service providers need an NFV platform that merges the best attributes of IT — agility and low costs — with capabilities needed for carrier applications. This paper focuses on the key characteristics and requirements for NFV platforms and highlights the central role NFV plays in service providers' drive toward zero-touch operational models.

TABLE OF CONTENTS

1. THE TREND TOWARD NFV

Network Functions Virtualization (NFV) is a new trend, and it is taking the telecommunications industry by storm. Service providers have teamed up to convince their vendors to move away from special-purpose equipment and toward software-only solutions that run on industry-standard commodity servers, switches, and storage equipment. Some of the frequently discussed use cases for NFV include:

- virtual IP Multimedia Subsystem (vIMS)
- virtual evolved packet core (vEPC)
- virtual content delivery networks (vCDN)

By embracing NFV technology, service providers expect to significantly reduce equipment costs, power consumption, and time-to-market for new services and functionality. Due to its software-driven approach, NFV also allows service providers to achieve a much higher degree of operational automation and to simplify operational processes such as capacity planning.

By many accounts, NFV will massively change the networking landscape and is leading the industry's network transformation to a virtualized, cloud-based approach that is characterized by:

- An independent generic hardware layer
- Automated network operations
- An agile application development and deployment model

NFV brings proven cloud computing and IT technologies, many of which are embodied in open source projects such as OpenStack, into the networking domain. But these technologies alone are not enough. Service provider applications are more demanding than most IT applications; they must deliver real-time services with high availability. To meet these requirements, NFV solutions must merge the best attributes of IT — agility and low costs — with capabilities needed for carrier applications.

The NFV community, including almost all of the major service providers and networking vendors, has formed an industry specification group under the auspices of the European Telecommunications Standards Institute (ETSI). This specification group is defining the requirements and architecture for NFV. The group recognized early on that a common infrastructure with compute and networking resources — referred to as NFV Infrastructure (NFVI) — along with a common management and orchestration layer are essential for successful NFV implementation. Together, these two components comprise the NFV platform.

2. VIRTUAL NETWORK FUNCTIONS ARE DEMANDING

Virtual network functions must fulfill the same strict latency requirements as traditional network functions. They need to provide five-nines availability and offer the same quality of service and security levels as telecommunications networks. Service providers will not accept the idea that network functions become unavailable if a cloud data center or an entire region loses service as has happened with public cloud providers.

Traditionally, network functions have been delivered on purpose-built and telecom-specific hardware, such as Advanced Telecommunications Computing Architecture (ATCA) platforms. Service providers have depended on the reliability and long mean times between failures (MTBF) these hardware elements provide to achieve the required availability levels. In contrast, following models developed by web-scale companies, such as Google® and Facebook®, NFV uses cost-effective commodity servers and hardware elements that may not be as reliable as those service providers have been using. Therefore, the perspective of availability needs to shift from individual elements to the system level, where availability is delivered through many cooperating elements in conjunction with the NFV platform.

Enterprise IT departments strive to consolidate datacenters and centralize applications in just a few datacenters to reduce operational costs. This model is generally not suitable for service providers. To deliver carrier-grade performance and availability, virtual network applications need to run on a distributed cloud and be embedded into core, metro, edge, and sometimes even access networks. Instead of hauling network traffic into one or a few centralized datacenters, which requires massive network capacity, traffic is processed in distributed datacenters. This approach minimizes network load by shortening the distance to the subscriber and by offloading traffic as early as possible. Because multinational service providers must meet regulatory requirements that restrict placement of personally identifiable data within national boundaries, their datacenters must be distributed across countries.

3. SERVICE PROVIDERS NEED OPERATIONAL EFFICIENCY

Putting network functions onto virtual machines is not enough to give service providers the cost efficiencies and agility they are looking for. However, NFV also gives service providers a new level of operational efficiency.

The software-based approach of NFV enables a much higher degree of automation, a key theme we will see throughout the migration toward NFV. In an NFV environment, operational processes such as service deployment, on-demand resource allocation, failure detection and on-time recovery, and software upgrades, can be programmed and executed with little or no human intervention (zero-touch). This provides an opportunity to reduce process intervals from months and years to days and weeks. Operations staff no longer needs to be sent out to remote locations to deliver, install, or repair network elements. Now these tasks can be completed remotely as long as staff can access sufficient compute, storage, and networking resources.

Under the banner of DevOps, advanced Chief Information Officers (CIOs) have established a new culture of innovation based on agile development methodologies. DevOps integrates and extends agile development methodologies with operations and quality assurance. With NFV, service providers can apply DevOps methods to the networking and communications service domain to accelerate innovation with a high rate of production releases. They can enhance:

- Product lifecycle management
- Capacity and inventory management
- Configuration management
- Service assurance
- Security management

For example, in capacity management, planning for individual types of network elements is no longer the main focus. Instead, the capacity of the shared compute, storage, and networking resources must be planned at an aggregate level taking into account that capacity shortages in one location can often be compensated with available capacity in other locations.

4. NFV PLATFORM REQUIREMENTS

To reap the potential of NFV, service providers and application developers need an NFV platform that allows them to take full advantage of virtualization. Figure 1 illustrates the conceptual architecture for such a platform.

Figure 1. The NFV platform must allow service providers to make the most of virtualization

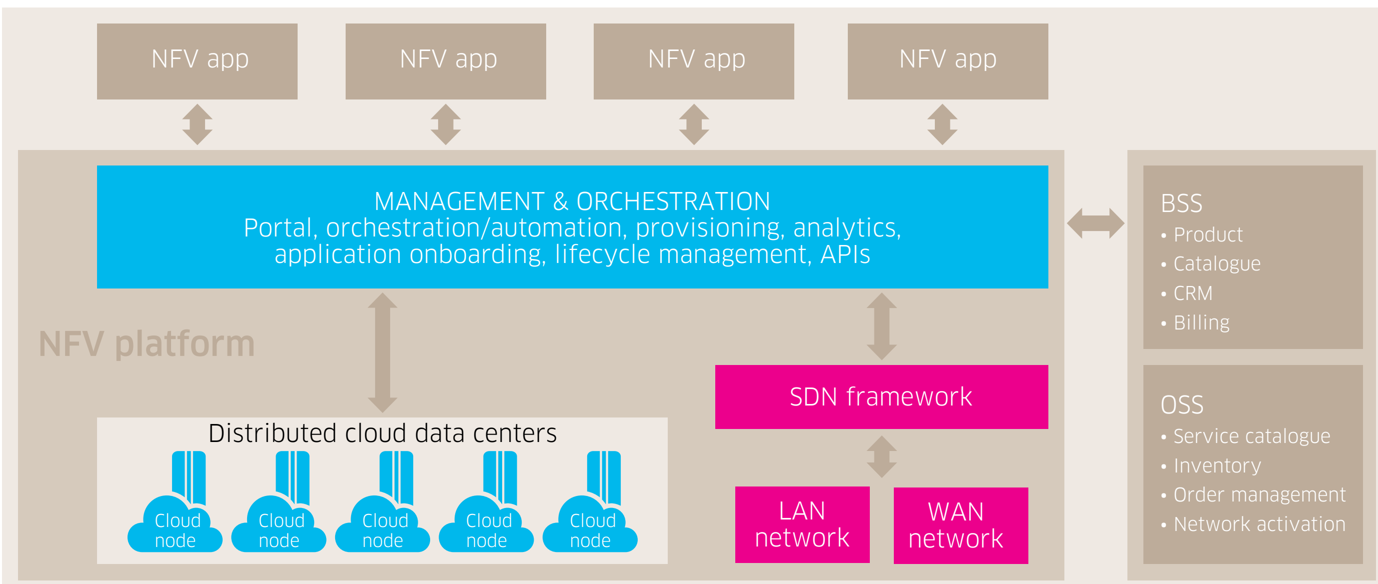

The NFV platform should provide a management and orchestration layer, including an NFV Platform as a Service (PaaS) capability that simplifies and accelerates development and deployment of virtual network functions. In addition, the NFV platform should provide an NFV infrastructure consisting of cloud nodes with compute and storage resources as well as a software-defined networking (SDN) framework for rapid provisioning of the virtual network structures needed for NFV applications. Together, these elements provide operational tools that will give service providers new horizontal and highly efficient operational models.

The NFV platform must combine virtualization and cloud technologies to fulfill the following requirements:

- Distributed cloud infrastructure
- Automated cloud nodes
- Automated application lifecycle management
- Network automation
- An open and multivendor environment

4.1 Distributed cloud infrastructure

In the IT world, the focus is on datacenter consolidation and centralization of applications in one or a few datacenters. In the telco world, latency and availability requirements are very demanding, requiring a more flexible, distributed architecture with datacenters and points of presence located throughout a coverage area. For this reason, network functions need to be embedded into the network and deployed in the network core, the metro, the edge, or the access tier.

An NFV platform must support a distributed architecture that:

- Gives application developers fine-grained control over the placement of network functions and sub-functions
- Automatically finds the optimal workload locations based on service provider policies and resource availability

Figure 2 illustrates these concepts.

Figure 2. The NFV platform must intelligently place workloads based on service provider policies and resource availability

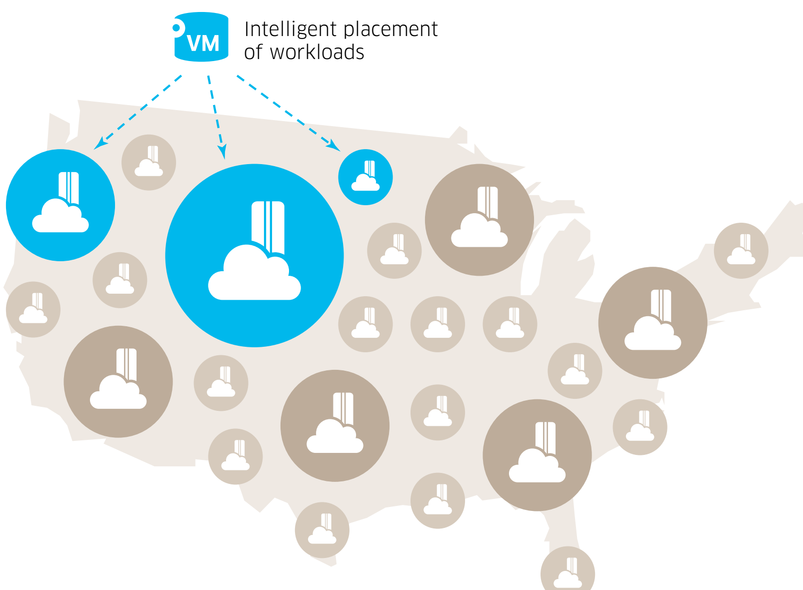

In addition, service providers need to manage and orchestrate the distributed datacenters and networks as a single virtual cloud. With this ability, operational staff can analyze and monitor the entire cloud platform in real time using tools to correlate events and metrics from various inputs for more effective decision making.

4.2 Automated cloud nodes

In a distributed NFV infrastructure, datacenters and points of presence must be added and changed more frequently than in a centralized cloud with centralized datacenters. There will be tens, hundreds, or even thousands of cloud nodes across the coverage area, each of them providing compute, storage, and network resources.

To attain the capital expenditure (CAPEX) advantages of NFV, service providers are looking to deploy cost-effective, mass-produced hardware. They also want the freedom to choose their preferred vendors without being locked into a particular vendor that also delivers the management and orchestration system or particular virtual network functions. This approach enables a maintenance model where individual components are no longer repaired, but are simply taken out of service until the cloud node is replaced. In this kind of environment a highly automated and pre-configured type of cloud node is important to keep operational costs low. While it can take a week or more to install a regular cloud node, an NFV cloud node can typically be bootstrapped and configured in a matter of hours and with most of the work completed remotely.

4.3 Automated application lifecycle management

The pure software nature of NFV applications enables more efficient lifecycle management processes. In the best case scenario, a new service can be deployed at required locations with the push of a button, without having to procure and install new equipment. To enable this automation, an application is first onboarded onto the NFV platform. This includes describing the application components and specifying how to execute the lifecycle phases shown in Figure 3.

Figure 3. The NFV platform automates all lifecycle phases

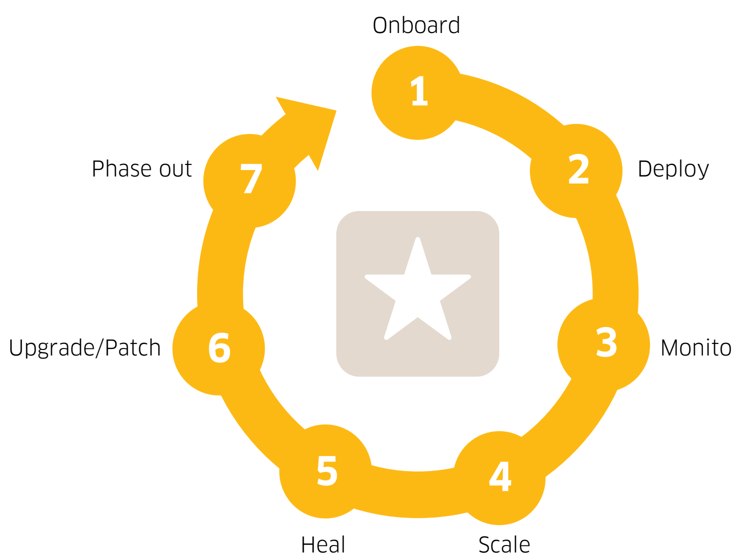

Service provider applications often consist of multiple virtual machines, storage volumes, and network configurations. That means service providers need to see how these components group to form applications, what applications are running where, and how many resources they use in addition to seeing the individual virtual machines.

To accomplish this goal, the cloud management system must see the application as a set of virtual network functions. For example, a DNS application consists of master nodes and slave nodes. An IMS application may consist of 10 or more different types of virtual network functions, each with its own deployment, scaling, and maintenance rules.

Today, deploying network elements for a new service and configuring the required network connectivity are still separate processes executed by separate teams. In an NFV environment, these processes are further complicated by the large number of datacenters and points of presence along with the frequent changes required when locations are added or removed. This complexity makes traditional manual deployment processes virtually impossible in an NFV environment.

The ideal NFV platform makes deploying a network application then allocating and configuring the necessary compute, storage, and network resources an integrated and automated process that takes minutes instead of weeks. The deployment process takes into account application policies and requirements described in the application descriptors. In most cases, deployment processes can be executed remotely, reducing the need to send staff to various locations.

4.4 Network automation

NFV applications must be embedded in the network — at the LAN and WAN levels — to ensure the desired performance and availability levels. To accomplish this, NFV applications such as virtual IMS and virtual evolved packet core often need sophisticated network structures with different handling of media, signaling, and management traffic.

Virtual network functions are also much more dynamic than their physical counterparts. They are deployed and removed on demand and scaled to match changing traffic. In addition, communications paths between blades that are internal in a physical network function are mapped to communications paths between virtual machines in the case of NFV. These communications paths must be provided by the NFV infrastructure.

Achieving all of this requires rapidly configurable and flexible network abstractions, which classical router and switch networks cannot easily provide. The NFV platform should have access to SDNs which provide enhanced programmability. With combined NFV and SDN, the network needs of virtual network functions can be captured and network capabilities can be created automatically as needed. This approach helps service providers address one of the key challenges in today's networks — labor-intensive and error-prone configuration processes. Figure 4 illustrates the conceptual relationship between NFV and SDN.

Figure 4. The NFV platform should combine NFV with software-defined networking

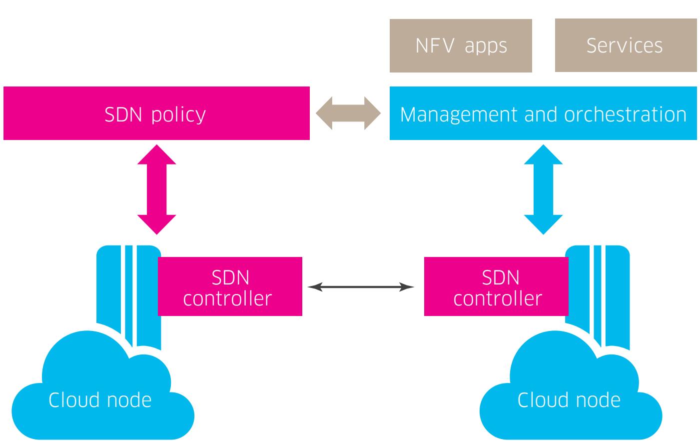

4.5 An open and multivendor environment

The NFV platform needs to be an open, shared environment capable of running applications from different vendors. Vendor-specific NFV silos with vendor-specific hardware and vendor-specific platform capabilities defeat the purpose of NFV. Service providers must be free to make their own hardware selection decisions, change hardware vendors, and deal with heterogeneous hardware. The NFV platform should be multivendor, and it should shield virtual network functions from the specifics of the underlying infrastructure.

At the NFV application level, service providers want the freedom to choose best-of-breed applications, no matter who the vendor is. Therefore, it is important that the NFV platform supports industry-standard application programming interfaces (APIs). OpenStack, for example, is an open source software project that provides a number of APIs for managing compute, storage, and network resources. NFV applications can access these APIs on the NFV platform and, for example, request additional compute and network resources when they are needed.

5. NFV ENABLES NEW OPERATIONAL MODELS

With an NFV platform, service providers can evolve their operational processes away from the siloed models of today and toward a more horizontal approach. The starting point for the horizontal approach is the shared infrastructure where all virtual network functions are running. Service providers need a single entity that manages the compute, storage, and network resources for all NFV applications deployed on it. Because the resources are no longer dedicated to specific network functions, there must be a common inventory, resource, and capacity management function supporting all of the applications running on top of the NFV infrastructure.

Operations Support Systems and Business Support Systems (OSSs and BSSs), such as fulfillment, assurance, and billing systems, can take advantage of the NFV platform to retrieve data about infrastructure and applications and to trigger configuration and management actions.

The impact of NFV on OSS and BSS architectures is just starting to be understood. Service providers have voiced their desire to evolve their operational models to:

- Become more real-time, zero-touch, and horizontal in approach
- Save costs
- Increase their flexibility to adapt to customer needs
- Become more innovative

6. CLOUDBAND: A PLATFORM FOR NFV

Alcatel-Lucent CloudBand™ is an NFV platform that is designed to meet the needs of service provider applications. CloudBand consists of two major components:

- The CloudBand Node, which provides the ETSI NFV Infrastructure
- The CloudBand Management System, which provides the ETSI NFV management and orchestration framework

The CloudBand Node provides the compute, storage, and local networking infrastructure to be deployed at the distributed NFV points-of-presence. The core technology is an advanced set of infrastructure management, monitoring, and virtualization software for efficient remote operations. The node software comes pre-integrated with industry standard x86 servers and is open to servers from multiple vendors.

The CloudBand Management System provides the following functions for managing virtualized network functions and applications:

- Cloud management and orchestration: This function manages and orchestrates resources across the end-to-end NFV infrastructure. Service providers gain a horizontal view of resource and service usage across the distributed locations.
- North and southbound APIs: CloudBand Management System functionality is exposed to applications and cloud nodes through industry-standard open APIs, such as OpenStack and CloudStack.
- Carrier PaaS: This function is responsible for the lifecycle management of NFV applications.
- Cloud optimization: In a distributed carrier cloud environment, optimal placement of virtual machines and virtual storage volumes is not a simple task. The CloudBand cloud optimization function runs placement algorithms developed by Alcatel-Lucent Bell Labs.
- Cloud analytics: This function collects cloud system and application events and provides near real-time input to the global cloud resource status view.

CloudBand integrates with the Nuage Networks programmable SDN solution as well as other SDN frameworks. With the Nuage Networks SDN framework, application developers and service providers can specify the networking policies required by the application. The Nuage Networks framework then automatically sets up the associated network structures as virtual machines are created throughout the NFV infrastructure.

7. THE NFV PLATFORM IS A POWERFUL TOOL TO ACHIEVE NFV GOALS

Service providers expect Network Functions Virtualization to deliver significant cost savings and service agility improvements that will help them address their pressing need to remain competitive.

The NFV platform is a powerful new element that helps service providers achieve these goals. To meet the demands of service provider applications and to enable more efficient operational models, the NFV platform must:

- Support a distributed cloud architecture
- Support a high level of automation in the cloud nodes and in managing the lifecycle of virtual applications and other management processes
- Be integrated with a software-defined network
- Be open and multivendor

Alcatel-Lucent recognized early on the potential impact of NFV technology. With CloudBand, Alcatel-Lucent provides an NFV platform that helps service providers reap the benefits of NFV technology while increasing operational efficiency.

8. ABBREVIATIONS

ETSI	European Telecommunications Standards Institute
NFV	Network Functions Virtualization
vIMS	virtual IP Multimedia Subsystem
vEPC	virtual evolved packet core
vCDN	virtual content delivery network
NFVI	Network Functions Virtualization Infrastructure
ATCA	Advanced Telecommunications Computing Architecture
MTBF	mean time between failures
CIO	Chief Information Officer
PaaS	Platform as a Service
SDN	software-defined networking
CAPEX	capital expenditure
API	application programming interface
OSS	operations support system
BSS	business support system

9. MORE INFORMATION

For more information, see:

Cloud opportunities: http://www.alcatel-lucent.com/solutions/cloud

CloudBand solution: http://www.alcatel-lucent.com/solutions/cloudband

CloudBand ecosystem program: http://www.alcatel-lucent.com/partner/cloudband-ecosystem

Alcatel·Lucent