CHARTING THE COURSE TO A VIRTUALIZED CARRIER NETWORK

STRATEGIC WHITE PAPER

Network Functions Virtualization (NFV) and cloud computing technologies create an opportunity for network operators to streamline and re-engineer their operational procedures and to capture new business benefits. However, despite its many promises, NFV also brings its share of challenges. As a result, network operators must carefully assess and evaluate the benefits and motivations for moving to NFV and develop a structured plan for introducing NFV into their existing environment. To help network operators chart their course to a virtualized carrier network, this paper discusses the impact of NFV on the various aspects of the carrier's operations. It also proposes a roadmap for transitioning to an NFV model.

TABLE OF CONTENTS

INTRODUCTION

Network Functions Virtualization (NFV) is the abstraction of network applications and services from dedicated hardware to run on general purpose hardware and hypervisors, storage and switching hardware (NFV White paper, 2012). To capture full benefits of NFV and cloud as well as successfully manage network transition challenges, operators must first understand common key characteristics of an NFV environment: common hardware, multi-tenancy, scalability, elasticity, simplicity, mobility, pro-activeness, and automation. A network operator's ability to turn those characteristics into business value will depend on how they are implemented and how its operations are re-engineered to support this new environment.

The opportunities and costs of evolving to a virtualized infrastructure are significant. To maximize the business benefits, network operators will need to develop a structured plan for introducing NFV into their existing environment. This plan should address the impact to stakeholders across the organization, and should identify the opportunities and challenges associated with each stakeholder group. The purpose of this paper is to describe some of the factors that network operators will need to consider in their development of an NFV evolution plan, and to identify the key milestones that should be addressed.

Despite its many promises, NFV is not without its share of challenges (Alcatel-Lucent, White paper). Network operators must assess and evaluate the benefits and motivations of introducing NFV into their unique business environments as part of the challenge. While this paper discusses the "how" of operational transformation, network operators must start their journey by understanding "what" function(s) to virtualize and "why" to virtualize.

The three key dimensions to address in identifying the early use case candidates for virtualization are:

- **Priority.** The significance of the business problem being addressed.
- **Impact.** The sophistication of the operations environment supporting the use case.
- **Risk.** The impact to service offerings in case of a failure or outage.

NFV is still in the early stages of development, and it will take some time before legacy services are available in a cloud environment. As early NFV deployments happen, network operators will also have to learn to manage co-existence, or hybrids, within their legacy network functions. It is likely that the co-existence of NFV and legacy network functions will continue for some years. Indeed, technical, strategic, and tactical reasons can make it impractical to reach a 100 percent cloud environment.

As illustrated in Figure 1, network operators may go through several evolution stages in their deployment of virtualized network functions.

Figure 1. NFV deployment evolution

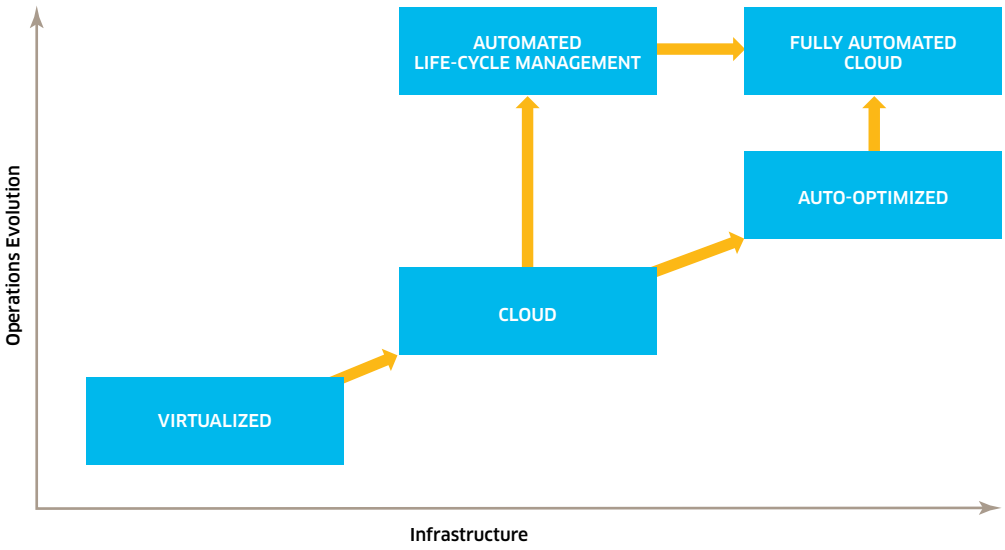

It should be noted that a network operator can, and often will, be in different deployment stages simultaneously. The operator can, for example, deploy several virtualized network functions that adhere to industry-standard interfaces and that can be managed through automated lifecycle tools, while other functions may only be available in a virtualized form. Once a network operator reaches the fully automated cloud stage for all network functions that can be virtualized, all of the potential benefits of NFV can be realized (Alcatel-Lucent, White paper).

The introduction of NFV and cloud computing technologies creates an opportunity for network operators to streamline and re-engineer their operational procedures and to capture fuller benefits of this new environment. Roles and responsibilities may be changing with the introduction of NFV, but the value-added activities and functions in an operator's organization will still have to be done, perhaps in a different way.

In this paper, we will use the term NFV in conjunction with cloud computing as cloud computing is one of the enablers of NFV as discussed in (NFV White paper, 2012). This paper is divided into two sections:

- **Section I** discusses the impact to the various aspects of the carrier's operations environment.
- **Section II** proposes a roadmap for transitioning to an NFV model.

Throughout this document, Alcatel-Lucent will express its views on how a carrier's operation will potentially be impacted by the adoption of NFV and these impacts will clearly be carrier-specific.

SECTION I: IMPACT OF NFV ON THE CARRIER'S OPERATIONS

This section discusses the impact of NFV on the carrier's operations environment with an emphasis on business and service management and operations.

- Business management and operations owns the business, is in charge of revenue, margins, and service evolution.
- Service management and operations is responsible for the implementation and management of network operator services.
- Resource management and operations is in charge of physical network assets.

Business and service management and operations are the consumers of network and datacenter resources, whereas resource management and operations are the providers of such resources. As such, there must be tight coordination and timely flow of information between these groups.

Business management and operations

Service lifecycle management

In today's environment, the service lifecycle — planning, sourcing, deployment, operations and support, and end-of-life termination — is executed over many months or years. In addition, considering the challenge of managing many legacy services and increasing demand for new services, network operators are facing the daunting task of managing the lifecycle of the services they offer.

In a cloud environment, designing for services, deploying and delivering them as per demand, detecting failures and enabling on-time recovery, and upgrading services provide an opportunity to reduce intervals from months and years to days and weeks. Because the rate of innovation and service introduction in a virtualized environment is expected to be much faster than today's dedicated hardware environment, a tighter collaboration between teams will be needed.

As networks become more horizontal, that is, general purpose hardware is pooled to deliver compute, storage, and networking resources, organizational barriers also need to come down in order to match the new paradigm. One important way to facilitate this adoption is by introducing lifecycle management as a service to teams performing various tasks at different stages of the lifecycle.

Some examples of how cloud computing technologies can be used to help manage the lifecycle include:

- Automating lifecycle management through the blueprint, deploy, monitor, scale, heal, upgrade, and tear-down lifecycle stages.
- Reducing application complexity by enabling the application to utilize the infrastructure it requires without requiring the application to understand the details of the infrastructure or how it is provisioned.
- Using application blueprints to define scaling behavior, performance, and resource requirements. These blueprints can be used repeatedly and dynamically for various deployment scenarios.
- Scaling applications dynamically to meet demand.

To maximize business impact, these capabilities should be presented to stakeholders in accordance with their respective roles in managing the lifecycle of services. Though it is too early to say whether there will be one tool to manage the entire lifecycle of all virtual network functions (vNFs), or if each vendor will provide their own set of tools, network operators should consider converging toward a small set of tools. The use of a common automated tool for lifecycle management should provide significant benefits to network operators.[1] It will likely limit incremental service management costs as an obstacle to new service launch and operation.

In today's environment, service planners introduce new services that typically require deployment of hardware, which adds to the existing challenge of lifecycle management: for example, shortening hardware lifecycles increases the costs of service delivery. In the cloud model, the ability to introduce new services will not be contingent upon hardware implementation lifecycle. As a result, the focus of business operations will shift from the logistics of deployment in the traditional sense to service development and deployment on an existing set of resources. Network operators will be able to introduce new services faster and have better ability to introduce new functionality and changes.[2]

Business planning and budgeting may need to be re-evaluated for their fit with the new cloud model, which is likely to reduce individual business unit capital expenditure (CAPEX) projections but increase operating expenditures (OPEX). Such CAPEX allocations will likely be transformed into a utility consumption model from a shared pool of compute, storage and networking resources. Business teams may have less clarity in understanding the cost structure being used to charge, as they will be using slices of datacenter assets. Especially for early services to be deployed in the NFV environment, the cost of charging is likely to be unfavorable to service profitability as network operators will have the burden of covering more of the fixed costs. Thus, business and resource teams need to agree on a well-designed and equitable cost model to charge for services as they scale.

Federated cloud

Cloud federation is the ability to bring two or more independently managed distributed clouds under one umbrella and represent them to the cloud consumer as a single cloud. Federation will become an important capability for carrier because it is unlikely that every network operator will have the scale and ability to achieve cost structures that are comparable to the best-in-class cost models achievable through the use of warehouse-scale computers (WSC). This will become an issue for early adopter business teams because they will likely bear the burden of an unfair share of network and datacenter startup costs.

1 The Alcatel-Lucent CloudBand Management System provides a hardware-agnostic carrier Platform as a Service (cPaaS) function that is responsible for vNF lifecycle management. It supports automation of application deployment, application monitoring and self healing, and automated up- and down-scaling.

2 There are related discussions to the NFV topic suggesting that a next-generation operations support system (OSS) may be required to support the additional capabilities of the NFV model. We are urging some caution here because updating OSS and business support systems (BSSs) to take new services and applications into account can involve substantial customization effort before services can be monitored and sold as a product.

WSC differ significantly from traditional datacenters: they belong to a single organization, use a relatively homogeneous hardware and system software platform, and share a common systems management layer. The software running on these systems, such as Gmail or Web search services, execute at a scale far beyond a single machine or a single rack; they run on no smaller a unit than clusters of hundreds to thousands of individual servers (Barroso & Hölzle, 2009). The major cost advantage of WSC comes from the economies of scale. Estimates from 2011 show that Google has 900,000 servers and Amazon Web Services has 454,400 servers, of which 40,000 are dedicated to running Amazon Web Services' EC2 servers in 7 datacenter hubs around the globe.

This is the kind of scale that is not attainable by every network operator. By tapping into the public cloud, network operators can consume compute and storage resources on demand and at low cost. Federated cloud may offer a viable business case for network operators in service development, whether the network operator is conducting in-house development or working in collaboration with third-party software developers.

Service management and operations

Planning and design

In this section we discuss the planning and design considerations for the service domain. Teams responsible for traditional service-level network planning and design will need to focus on two discrete domains: the physical domain and the service domain. While the physical domain involves design and planning of physical assets under resource management and operations, the service domain centers on the logical design upon which the service will be introduced.

The role of planner within the context of service operations and management is expected to evolve to being an application planner, who will build upon the virtual compute, storage, networking resources supplied by resource management and operations. Unlike today's static designs, planning teams will be able to design more dynamic network environments for their respective services. Because NFV enables the abstraction of network functions from underlying assets and matches those assets to service consumption on an ongoing basis, planners can create variants, or scenarios, of their design based on different traffic patterns and network policies. These design scenarios can then be translated into sets of rules that trigger network elasticity and service rescaling. As a result, network operators can introduce services more gradually using minimal network resources and allow the cloud management system (CMS)[3] to orchestrate its growth based on the set of rules defined by the planners.

Once a service is introduced, planners can then focus their attention on service characteristics, asset consumption, traffic profiles and other indicators that will help them better understand the service behavior under real usage conditions. This monitoring allows designers to refine their design rules that trigger events, such as starting up a new virtual machine (VM), moving a VM to a different zone, and scaling up or down as per pre-agreed service policies.

3 In this white paper, cloud management system refers to a generic management entity that facilitates the supply and demand between compute, store, and network assets downstream and services upstream.

Another important job for planners becomes optimizing the utilization of resources — the service costs in accordance with demand for the service and the most cost-effective utilization of assets. This close matching of supply to consumption potentially exposes business and service operations teams to the risk of under-serving their customers. As illustrated in Figure 2, there may be spikes on demand that surpass the amount of resources available. This could be a one-time event which normally would not warrant investing in additional resources.

Figure 2. Spikes in demand can exceed available resources

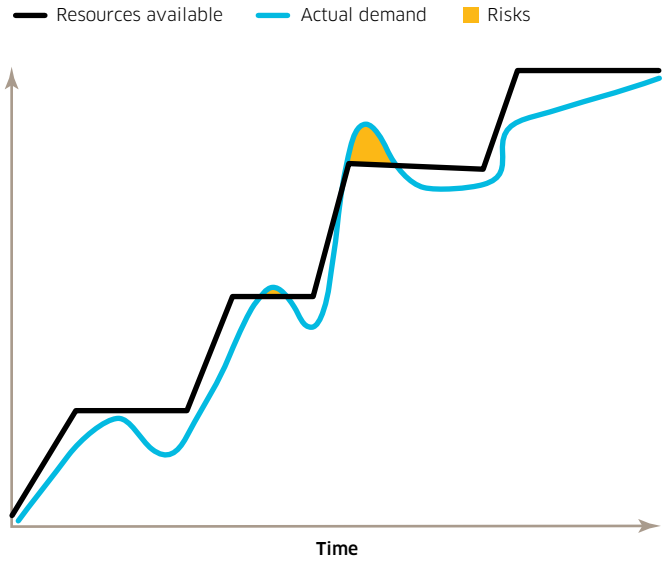

To mitigate the risks of under-serving customers, planners may choose to include a federated cloud solution in their environment. Rather than resource management and operations procuring more equipment, planners may choose to address spikes in service demand using cloud federation to tap into third-party cloud services on-demand or for reserve. However, planners must be familiar with cloud design principles beyond that of their own organization, and furthermore, they must validate that the third-party cloud meets the requirements of their environment. For each network operator, the case for using a federated cloud will be different as there are other highly critical considerations beyond supply-demand and margins. These include control, security, performance, legal, and regulatory concerns.

Deployment and integration

In a cloud environment, deployment and integration teams are concerned with deploying the software-based vNFs in the carrier cloud and making necessary provisions for enabling the associated services. As a result, there are no recurrent delays due to ordering and receiving custom-built hardware, conducting site surveys and site preparation for it, and its engineering and installation. Furthermore, there will be less dependency on the field engineering teams for deployment and integration. This will make program coordination and communication easier while improving efficiency and speed. Network operators can even perform this activity by using their own resources remotely using CMS with perhaps some level of support from their vendors.

However, network operators must have certain capabilities in their arsenal to make deployment and integration as simple, seamless, and automated as possible.

First, all knowledge that is required to install and configure an application or vNF must be encapsulated. This encapsulation can be offered to the service owner through the automated lifecycle management solution as discussed earlier. Such a solution allows the service owner to deploy the service functions by using descriptors.[4] Examples of the information in descriptors includes:

- Infrastructure as a Service (IaaS) requirements
- Service definition
- Healing rules
- Scaling rules
- Key performance indicators (KPIs) to monitor
- Deployment rules
- Number of initial instances and scaling
- Geo-locations and constraints
- Affinity rules
- Network requirements

Figure 3 illustrates the relationship between the descriptor and the service functions.

Figure 3. Descriptors can be used to deploy service functions

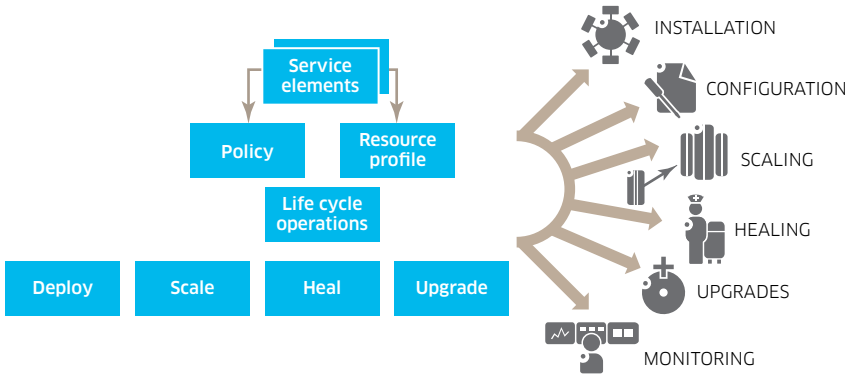

Unlike traditional deployment and integration processes, this automated process enables network operators to execute centrally and is less error prone at a much faster speed. Furthermore, deployment schedules can be more flexible as the need to depend on underlying equipment and field resources is less than traditional deployment scenarios.

Testing

Today, testing can add an additional 6 to 18 months of delay to deployment or upgrade of a solution after it is delivered by the vendor. Furthermore, network operators incur additional operating costs for lab infrastructure setup and maintenance. Moreover, once a solution or upgrade passes lab testing, it is typically tested in the field, which brings field engineering teams into the picture for support. However, this triggers additional complexity in program communications and coordination, and will likely delay other daily field activities.

4 The descriptor defines the structure of the vNF as well as deployment and operational aspects, such as computation, storage, and networking requirements.

Cloud computing is not the only answer to better testing processes and reduced test cycles. It does, however, foster a new operations environment where testing activities can be performed in a more effective fashion. This environment, however, is not new; it has been widely accepted and utilized by the software development community, and many tools and processes have already been developed to support it. The telecommunications industry has a chance to adapt some of these practices and tools to maximize the benefits offered by NFV.

Figure 4. In a cloud environment, different environments can easily be created and presented to users based on their respective roles.

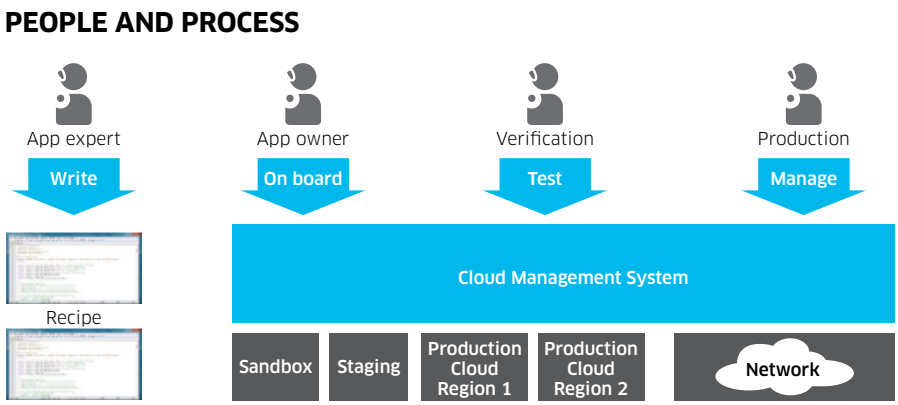

Teams in charge of testing can create their own test environment in the cloud through the use of automated lifecycle management solutions. Such environments can be created fully in the cloud without any dedicated equipment, or within a sandbox in the datacenter infrastructure. The sandbox can be set up by CMS requesting dedicated physical equipment downstream from the datacenter. The test environment set up is completed in a matter of hours, once test engineers complete their test plans and supply specifications to CMS for test environment creation.

In addition to reduced timelines for test staging and environment creation, test teams will also have increased ability to test incremental releases as opposed to waiting for every major release from vendors. Furthermore, test teams can run multiple instances of test environments to simultaneously test different vendors and releases. However, to fully achieve these efficiencies, network operators need to invest in testing automation as it is crucial to cope with the dynamic environment that NFV brings and to stay ahead of the competition.

Service quality management

Virtualization technology, along with resource sharing and additional aspects of NFV, can introduce additional impairments, such as potential points of failure, to the compute, memory, storage, and networking services that support cloud-based applications. Figure 5 highlights this risk. Service quality management tools and processes must be in place to isolate impairments to service qualities such as availability, reliability, or latency to the true root cause, be it the application software or the underlying compute, memory, storage, or networking infrastructure so that the true root cause can be rapidly identified and proper corrective action taken.

Figure 5. Virtualized Infrastructure Impairments to Cloud based Applications

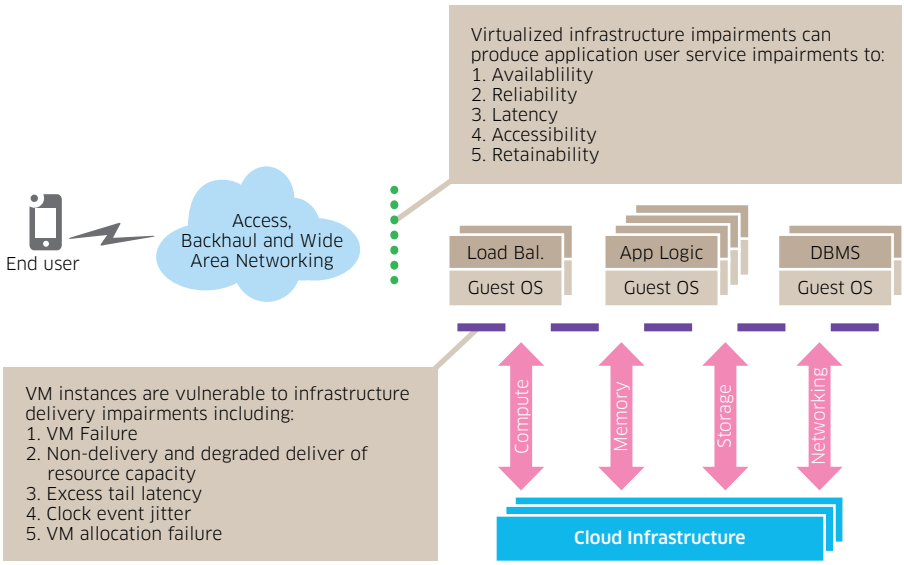

As network functions are decoupled from hardware, the lines of ownership and accountability will blur. While business and service operations teams own the service and are accountable for service performance, they will not have control over underlying hardware and network assets. As consumers of such assets, they will expect resource management and operations teams to make such resources available to them. This outsourcing[5] model necessitates new SLAs between the consumers (business and service operations) and the providers of cloud infrastructure (resource operations). These SLAs should define the boundary between the application and the infrastructure. Once the boundary is defined, application KPIs (e.g., service latency, reliability, availability), and infrastructure KPIs (e.g., compute, memory, storage and networking), can be defined and monitored accordingly (Bauer & Adams).

Inventory management

Within the scope of service management, inventory is related to software management. Many software products, each with possibly multiple versions of itself, can be running concurrently in the NFV environment. To make matters even more daunting, as network operators venture into the do-it-yourself (DIY) domain or collaborate with third parties for innovation, they will have to manage many software packages in the development environment with several versions being collaborated upon by different teams.

As network operators venture toward NFV, service management and operations teams must have the ability to manage the software inventory.

5 Cloud consumer (business and service management and operations) outsourcing the ownership and operation of the compute, memory, storage and networking to resource management and operations.

Network operators need:

- A software repository where software packages, including license management, may be retrieved and deployed
- The ability to provide automated deployment of applications across all environments
- A robust distributed version control system (DVCS) to support DIY or collaboration with third-party developers
- The ability to produce snapshots of the network for backup, restore, and auditing

Resource management and operations

In the NFV environment, resource management and operations can be expected to maintain and operate datacenter infrastructures, providing cloud computing resources as IaaS, or even PaaS, for consumption by the vNFs.[6]

Resource management and operations teams will continue to have physical layer requirements such as capacity, availability, redundancy, resiliency, and connectivity to name a few. They will need to consider system-wide requirements, spanning datacenters and central offices, regardless of where cloud resources are located. They will be providing slices of compute and storage to service management and operations for consumption by vNFs. This will ensure that, at the service layers, there will be a hardware-agnostic view southbound. From a workforce standpoint, teams managing the application layer of legacy technology will need to evolve to an IT datacenter profile. With the likelihood of operating an infrastructure that supports hybrid services — continued operations of some legacy functions plus vNF — operators will need to consider a support model to accommodate this new eventuality.

SECTION II: ROADMAP FOR TRANSITION TO NFV

This section discusses key milestones to introduce NFV into the operations environment. The model presented can be repeated at each evolution stage — virtualization, cloud, automated lifecycle management, auto-optimization, and fully automated cloud — with relevant tasks. The teams executing this three-step model should be composed of experts representing the organizations that will be most affected by the setup, delivery, and operations of the NFV environment.

Figure 6. A three-step model helps operators transition to NFV

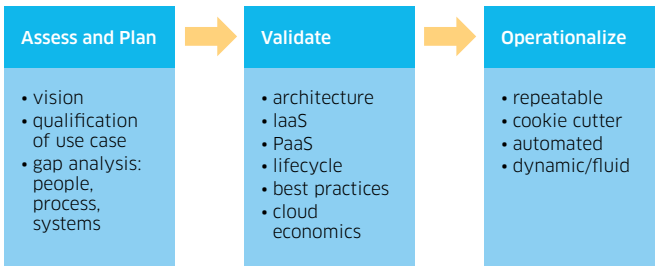

6 The Alcatel-Lucent CloudBand Management System manages and orchestrates resources across the end-to-end infrastructure. It leverages distributed cloud concepts and aligns with main cloud computing functions such as on-demand self-service, broad network access, and resource pooling. It maintains a global cloud resource status view across all applications. The CloudBand Management System exposes its functionality to other entities through open application programming interfaces (APIs), and it interfaces through open, and commonly used, proprietary APIs to the underlying NFV infrastructure.

Assess and plan

In this stage, network operators define their target future mode of operation (FMO) and draw a roadmap of how they want to get there. A key action in this process is identifying a targeted, prioritized use case such as virtualizing the IP multimedia subsystem (IMS), the mobile core network, or the content delivery network (CDN) as the first stepping stone into NFV.

The first question to ask when identifying the use case is "what business problem is being solved". Is it about new revenue opportunities, optimizing spending, bringing innovation into the network or something else? The use case(s) may have an impact and potentially add risk to current operations.

It is also reasonable to assume that current operations will continue while the NFV environment is developed and operated in parallel, with incremental costs associated with these activities. The more sophisticated or integrated the operations supporting the use case, the greater the potential impact. Most operators will encounter a learning curve as they transform to an NFV model. There may be new or different risks associated with services operations, organizational tuning to support the lifecycle, and the previously mentioned business case modifications that all need to be taken into consideration. Operators may want to consider the first engagement as a learning experience, but they must be tolerant of risks and failures and plan for both. Is the impact to customer satisfaction major upon a service failure? How fast can the network operator recover from fallout? For example, virtualizing the mobile core network may have a much bigger positive business impact than virtualizing customer premises equipment (CPE) in the home environment. However, the fallout from an implementation failure will be less harmful in the case of virtual CPE.

In addition, there are a number of security considerations that must be accounted for when moving a function to a virtualized environment. Network operators expect that the virtual appliances — the set of network functions — have the same level of security as the physical appliance, meaning that both hypervisors and virtual appliances have to be security certified.

Other factors associated with NFV use case planning include the legal and regulatory requirements imposed on network operators in the regions in which they operate. These laws can be related to how well the data is protected from a confidentiality aspect as well as to how this data is stored, transferred or accessed. A good example is the case of virtualizing the home environment; copyright rules may preclude network operators from storing end-user recordings in the cloud through a virtual personal video recorder (PVR) service. Not only do network operators have to comply with these regulations, they must also make sure that all the legal and regulatory requirements are fulfilled by their vendors and business partners.

NFV gives network operators an important opportunity to transform their business model, infrastructure, and operations processes. Current best practices can be carried into this stage and included in the plan. Such practices will then need to be fine-tuned based on the results of the validation stage that comes next. Roles and responsibilities must be identified in relation to different customer and market segments in order to clarify which operations functions and processes are impacted by virtualizing a specific function. Once these components are understood, a gap analysis is performed focusing

on the differences between the present mode of operation (PMO) and FMO of the function in question. This analysis should capture differences in staff skills and expertise, operational processes, and support systems to provide estimates of the magnitude of change.

Network operators will need the right tools to tightly link operations, technology, services, and impacts together in an integrated and streamlined business case. Cloud modeling and planning tools, such as Cloud Modeling Tool developed by Alcatel-Lucent Bell Labs, can help network operators address such need. By looking into an end-to-end array of engineering and economic modeling options representative of cloud infrastructure solutions across various deployment scenarios, the Cloud Modeling Tool enables network operators to select the solution that best suits their specific needs.

Validate

During validation, the output of the planning and assessment stages is executed. The validation itself is iterative because early cloud orchestration and automation solutions to support NFV will come in incremental steps. As a result, one approach, shown in Figure 7, is to have sequential and parallel work streams to iteratively evaluate and validate new features released by vendors.

Figure 7. Sequential and parallel work streams are needed for validation

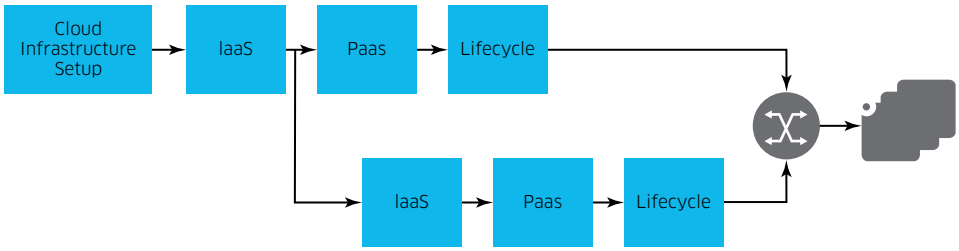

Network operators can consider entering the validation stage by deploying a self-contained "pseudo" cloud, which does not require much upfront investment and is fast to deploy. With this approach, operators can start testing their use cases in a cloud environment without having to wait for large datacenter deployments. Furthermore, over time, this self-contained pseudo-cloud infrastructure can possibly be rolled into the actual virtual network environment, thereby protecting the initial investment. Many network operators have already recognized the need to set up a dedicated carrier cloud environment to test NFV. These test beds should reflect the operating requirements necessary to support NFV.[7]

With the infrastructure in place, the operator can start validation of the use case by testing against IaaS capabilities such as networking, compute, storage, southbound APIs, federation, performance, availability, change log, and backup.

With the operator's chosen use case deployed, IaaS testing is followed by testing of PaaS capabilities such as on-boarding, monitoring, self-healing, scaling, automation, deployment scenarios, security, multi-tenancy, northbound APIs, migration if it is an existing service or cloud federation, for example with Amazon.

7 The Alcatel-Lucent 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 from bare metal installation to fully functional.

The final step in this sequential process is lifecycle management validation. This is where the NFV environment is tested as a whole, going through the various stages of lifecycle management activities and arriving at a set of best practices and workflows, including:

- Planning for capacity
- Sourcing
- Testing
- Environment creation, including development, testing, and production
- Deployment
- Installation
- Integration
- Patch management
- Monitoring
- Availability
- Backup
- Service upgrade
- Service retirement

These steps must be applied in accordance with the NFV evolution stage that the network operator is part of: virtualized, cloud, automated lifecycle management, auto-optimization, and fully automated cloud. The output of the validation stage should yield:

- Best practices
- Engineering guidelines
- Fine-tuning of the cloud economic modeler
- Workflow and process documents for lifecycle management
- Service quality management (SQM) metrics
- Service level agreement (SLA) templates
- Establishment of ownership, responsibility for resources, and accountability for results

Operationalize

While the plan, assess and validate steps lead to fine-tuning of expectations from NFV for the selected use case scenario, they also set the ground for operationalizing the knowledge. The process from planning to validation is iterative every time a service is introduced or network operator moves from one stage to another in the NFV environment. As network operators gain new knowledge, they must transform this knowledge to a set of well-defined, repeatable, automated and measurable tasks by re-designing their best practices.

Network operators may choose to follow different transformation frameworks that fit the best for their situation. For example, transformation may start from the segment of the network layer that is non-real-time but has high virtualization potential, and move toward real-time with fewer potential segments; for example, from BSS/OSS to elements of packet core or even access, or CPEs and set-top boxes. Or, transformation could be a bottom-up approach where virtualization starts from the infrastructure and moves northbound to services and business processes.

Regardless of the path chosen, there are additional practices and capabilities that network operators should consider as they gain more experience with NFV.

Service characterization is an essential practice for evaluating the fitness for virtualization in a network operator's environment. Simply seeing compute and storage as an economic choice as opposed to buying dedicated hardware misses the true potential of NFV. Consider, for example, how a cyclic application uses resources and, as a result, leaves resources heavily underutilized. A high-growth application, on the other hand, requires the ability to scale. Is the service compute, network, or database-access-intensive? What are the throughput and latency requirements? Is it coupled with specific hardware for encryption?

Service characterization helps to determine how services are delivered over the cloud, and whether it makes business sense for a network operator to do so. This characterization is repeated over many services, and the evaluation must capture essential characteristics such as complexity, rate of change (from a few releases per year to dozens of releases per year), rate of use (from millions to hundreds of millions of users per month), likelihood of exceeding expected parameters, and likelihood of competing for resources with other services during peak period, to name a few. As a result, network operators will have a better view of how much, if any, economic value-add each service will contribute, and whether that value meets their expectations or not.

One size does not fit all. As networks become more horizontal, network operators also have the choice to opt for a DIY strategy with the help of software engineering practices to maintain control and oversight over this new environment. Whether it is to enhance productization, quality, productivity, maintainability, or reusability, network operators can further improve their ability to tackle business and operational issues in an NFV environment by introducing homegrown solutions.

Not only can a homegrown solution complement application descriptors, but it can also bridge the gap between the network and the operator's unique operations environment. As services talk to each other and the network through standard APIs, a homegrown solution can enhance the quality of operations by introducing additional capabilities into the operations environment. The operator can, for example, create its own dashboard, or correlate KPIs and alarms to alert field teams in advance and present demand forecast to datacenter operations. Such an approach can also be used in building a test library over time that results in only incremental test case creation as new functionality is introduced into the network. It can also be used to manage the network operator's APIs, exposing them to partners for collaboration or productizing them. In summary, a homegrown solution can give the network operator greater flexibility in managing its operations activities within a variety of domains as well as across domains.

Unhindered and timely flow of information across all relevant organizations becomes even more critical as network operators adopt NFV. Each network operator must assess its business and operations workflows. They must also identify the stakeholders and channels that those workflows touch. Furthermore, the stakeholder needs and how best to communicate with them must be defined and captured in communications plans. In doing so, network operators must utilize technology to automate this information flow and sharing across departments and geographies. As networks become more adaptable, dynamic and programmable, such information flow must be relevant, timely, and actionable.

CONCLUSION

The activities and functions of the three key stakeholders discussed in this white paper will not be disappearing with the introduction of NFV. Rather, there will be a skill shift and an organizational realignment from business-level operations that are responsible for the service to low-level operations that manage the infrastructure to adapt to the NFV operating environment.

Business management and operations responsibilities will include:

- Institution of automated lifecycle management as a service
- Transition from a logistics program management style to innovation-driven service introduction by leveraging cloud computing technologies
- The ability to de-risk service profitability by advocating equitable cost accounting for use of infrastructure
- Consideration of new alternative service delivery options such as federated cloud, where feasible

Service management and operations responsibilities will include:

- Setup and delivery of services in accordance with costs
- The ability to leverage IaaS and PaaS to deploy and integrate new services and updates while further minimizing downtime
- Facilitation of collaborative innovation with third parties by creating development, testing, and production environments
- Shortening of lifecycles for deployment, integration, testing, and production
- The ability to work with resource management to define, monitor and enforce SLAs upon which compute, storage, and networking resources are provided
- Management of the software repository, version control systems, and backups

Resource management and operations responsibilities will include:

- Accountability for physical assets such as access, connectivity, datacenters, and field resources
- The ability to adapt to new requirements driven by NFV
- The ability to manage the skill shift to cloud computing from traditional telecommunications and IT operations

The journey will be unique for each operator, and perhaps it will include multiple steps for many. Consequently, operators must continually adapt their operations environment and best practices as they gain new knowledge.

ABBREVIATIONS

API	application programming interface
CAPEX	capital expenditure
CDN	content delivery network
CMS	cloud management system
CPE	customer premises equipment
DVCS	distributed version control system
DIY	do it yourself
FMO	future mode of operation
IaaS	Infrastructure as a Service
IMS	IP multimedia subsystem
KPI	key performance indicator
NFV	Network Functions Virtualization
OPEX	operating expenditure
PaaS	Platform as a Service
PMO	present mode of operation
PVR	personal video recorder
SLA	service level agreement
SQM	service quality management
VM	virtual machine
vNF	virtual network function
WSC	warehouse-scale computers

WORKS CITED

Alcatel-Lucent white paper. (n.d.). Network Functions Virtualization – Challenges and Solutions.

Barroso, L. A., & Hölzle, U. (2009). The Datacenter as a Computer: – An Introduction to the Design of Warehouse-Scale Machines. Synthesis Lectures on Computer.

Bauer, E., & Adams, R. (n.d.). Service Quality of Cloud Based Applications. ISBN 9781118763292. Wiley-IEEE Press (Tentative Publication December 2013).

NFV white paper. (2012). Network Functions Virtualisation – An Introduction, Benefits, Enablers, Challenges & Call for Action – Contributing organizations: AT&T, BT, CenturyLink, China Mobile, Colt, Deutsche Telekom, KDDI, Orange, Telecom Italia, Telefonica, Telstra, Verizon. SDN and OpenFlow World Congress. Darmstadt – Germany.

Contacts

For further information please contact your Alcatel-Lucent sales representative.

Contributors

Okan Tanrıkulu, Morgan Stern, Jonathan W Cohen, Peter Busschbach, Tanja De Groot, Andras Menyhei, Mark Windy, Enrique Hernandez-Valencia, Daniel R Johnson, Michael Kohring, Nape Kona, Christopher Larcher, Chris Helliwell

Alcatel·Lucent