Phil Tilley

THE ROLE OF NFV AND SDN IN FUTURE COMMUNICATIONS

The concept of virtualisation – and the ability to run multiple functions on a single machine or server – is at the heart of today's notion of cloud computing. Our modern version of cloud computing started in the late 1990s and continues to evolve. But it could be argued that the move to mainframe computing was the start of cloud computing back in the 1970s. And even earlier, Herbert Grosch (1918-2010) predicted that the world would operate on dumb terminals run from around 15 data centres worldwide.

The benefits of today's virtualisation-based approach are clear: a shared common infrastructure partitioned to support multiple users provides significant operational cost benefits. Separating control and switching is not unknown in the telecommunications sector as this is the way we have been running networks since the introduction of digital exchanges and TDM (Time Division Multiplexing) switching. Intelligent Networks also relied on this model to provide enhanced digital services and greater flexibility in call routing and pricing plans.

In the last few years, the market has shifted from voice-centric networks to data-centric networking, dramatically changing the telecom playing field. Today, telecom professionals need to focus on handling massive data volumes while continuing to protect the quality and reliability of both voice and data services. This focus has led to integrated hardware and software, and software optimised around proprietary hardware appliances. The introduction of an Advanced Telecom Computing Architecture (ATCA) chassis was meant to be the point at which we started developing software modules to run on a common set of hardware. However, demand for ATCA chassis never really reached critical mass. In contrast, demand for commercial compute servers accelerated dramatically and server pricing became very competitive.

Price point alone would make the business case for running telecom applications on commercial server hardware very compelling. But the rewards and operational benefits become irresistible when hardware can be virtualised and telecom applications can be distributed across multiple standard virtual machines. Not only do you get flexibility in hardware utilisation, you gain the advantage of being able to distribute traffic loads. This ability to balance traffic brings resiliency, another critical requirement for telecom applications.

Figure 1

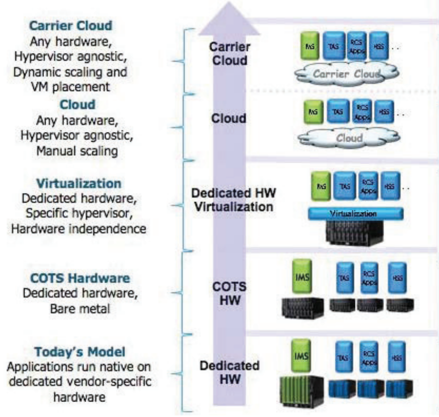

TELCO CLOUD ATTRIBUTES

Hosting	Multi-tenancy	Software only delivery	Scaling and mobility	Full support of industry std cloud API's
x	x	x	x	x
x	x	x	Manual	
x	x			
x				
x				

Figure 1: Evolution and attributes defining the ultimate NFV operational model

This demand for increased capacity, flexibility and resiliency at an affordable price point is driving the telecommunications community's interest and motivation for what has been termed Network Functions Virtualisation (NFV).

NETWORK FUNCTIONS VIRTUALISATION

NFV is an approach that is being documented and driven by the world's largest telecommunications providers, who formed a working group within ETSI in 2012. The ETSI working group has gained significant momentum very quickly, producing over five definitional specifications and growing to over 100 participating members.

NFV is at the start of a multi-year journey, a journey that is being made possible as a result of many technical advances coming together simultaneously. The desired end result will be new user-centric services where the network adapts to users' demands rather than the situation today where the network defines what the user can do. These new services promise to be much more flexible and offer a better user experience. However, for service providers to remain viable businesses, it is critical for the success of the service provider that these new services do not cost more to deliver than today's services.

When looking at the evolution towards an NFV operational mode, we can consider three broad strands of development all interlinked but also having their own evolutionary path:

1. Virtualisation of telecom functions (voice, video and mobile core)
2. Orchestration and management of virtual machines to enable telecom applications and functions
3. Integration with the underlying network to provide connectivity

ADAPTING EXISTING FUNCTIONS FOR VIRTUALISATION

While it is inevitable – and highly desirable – that new applications and functions be developed, telecom service providers already have significant investments in the applications running today's networks. If we are to continue offering the services people already receive, these applications must be brought into the new NFV environment. This adaptation will take place via a number of steps to ensure continuity and integrity.

Figure 1 demonstrates one path that may be adopted. Under the existing mode of operation, the software is tightly embedded with the hardware. For NFV, software must evolve towards a cloud operational model that is fully scalable in line with demand. In the cloud model, software should be completely decoupled from hardware. This approach gives the service provider greater flexibility in choosing suppliers as well as increased flexibility in operations.

Traditionally, service providers have built complex operational support systems to manage the appliances within their network. In the new world in which applications run on virtual machines, much greater emphasis must be placed on the management and orchestration of the virtual machines that will host the applications. While cloud management systems exist today and are running some very large scale cloud operations, the telecom environment brings additional and unique challenges. If these are not addressed correctly, the end user experience will suffer. For example, the networking of voice, video and mobile applications requires careful engineering. Any delay or delay variation introduced into the network may be instantly detectable by the user – as a period of silence, an echo, loss of picture or in the case of mobile, a dropped call.

So as we start looking at the introduction of NFV to run our networks, we are challenged to ensure that we can scale out virtualised functions in a way that maintains an acceptable performance level. Performance will be determined by the manner in which the expected data and signalling load are processed, as well as by how well packet throughput across the network is maintained under all load conditions.

The challenge of performance can of course be overcome in a virtualised computing world by throwing more resources at the problem. However, this brute force approach leads to an increased cost and will undermine the expected payback from virtualisation.

LOOKING FOR THE REWARDS OF VIRTUALISATION

The need to balance performance and cost reductions forces service providers to make critical decisions about how to prioritise their roadmap for virtualisation. Some functions will achieve significant cost gains using today's available commercial off-the-shelf (COTS) hardware while other functions will require further hardware development. Alcatel-Lucent Bell Labs has studied the return on investment for moving different applications and functions into the cloud and assessed the balance between cost improvement and improved dynamic control.

Service providers rely on a significant number of management and support applications which are prime candidates for moving to virtualised cloud hardware. Ultimately this will lead to significant cost gains due to improved resource utilisation but less benefit from dynamic control (as shown in Figure 2).

Where user demand is highly variable, either because the user base is growing or because of subscriber mobility, there are significant gains to be made by enabling greater dynamics since NFV will allow, the compute resources required to be scaled according to user demands, so capacity matches demand much better than is possible today. Media servers and packet gateways are examples of applications where demand is flexible and the greater dynamics will lead to significant operational cost savings.

Whilst there are significant cost benefits in running a media or packet gateway on COTS hardware it introduces some new operational challenges. The quality of any voice service is very susceptible to delay and whilst a dedicated appliance is used, the delay can be predicted and engineered around. However in a COTS-based world, quality can be impacted by the performance of the server and more specifically by its packet I/O. Potentially, different functions may need to traverse across multiple virtual machines and these may be on different servers geographically non-adjacent. These challenges are overcome by ensuring that the orchestrator understands the virtual machine distribution rules.

Figure 2

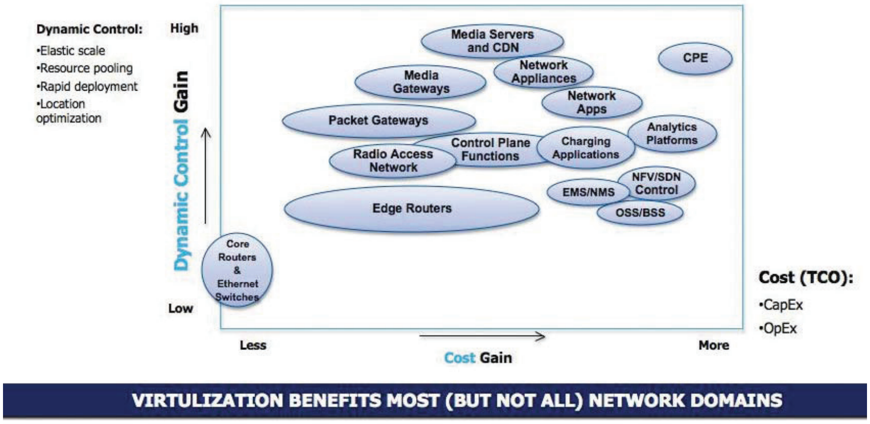

Figure 2: Benefit analysis of virtualising different telco functions

ORCHESTRATING AND MANAGING AN APPLICATION'S HOST ENVIRONMENT

Whether the service provider is offering a mobile app or real-time voice within a Web app (WebRTC) there will be many software modules all interconnected and sharing data across Application Programming Interfaces (APIs). Each of these modules is uploaded onto a virtual machine image within a server. As a result, the telecom domain requires many thousands of virtual machines which may be widely distributed. This scenario is behind the development of one of the key components of NFV: the manager/of virtual machines.

The orchestrator automates the process of preparing and tracking virtual machines within the service provider's network. Each telecom function requires a different virtual machine set-up and configuration and it is the role of the orchestrator to know what configuration is required to support each application. When a service provider decides it wants to launch a new function within the network, an available virtual machine will be located and made available with the correct configuration automatically. The orchestrator tracks that virtual machine and function during its life-cycle and scales resources in line with subscriber demand. Figure 3 highlights the application life-cycle for which the orchestration layer is responsible.

THE ROLE OF SDN: NETWORKING VIRTUAL MACHINES AND HOST APPLICATIONS

The third key aspect of telecom function virtualisation sets it apart from other cloud applications. Our users are not always close to the software functions they use. To avoid latencies and back-hauling a lot of data over long and potentially expensive transport connections, we need to consider how we distribute the applications. Consequently the distribution of virtual network functions becomes another critical success factor.

While we scale-up NFV we must look at how we manage the underlying network infrastructure. To do this cost effectively, we need to automate the network to ensure it is in step with application demand. This is where Software-Defined Networking (SDN) comes in. SDN has emerged as another

Figure 3

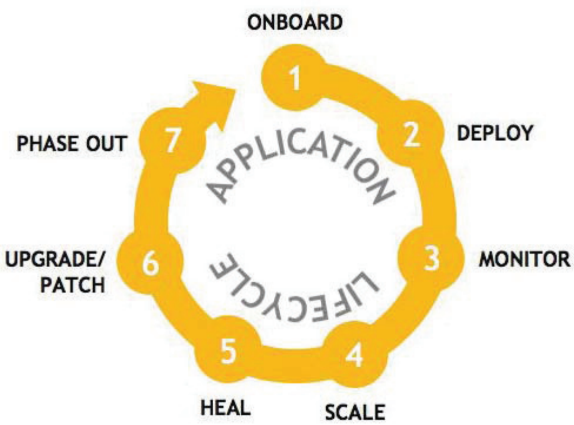

Figure 3: Application life-cycle management

Figure 4

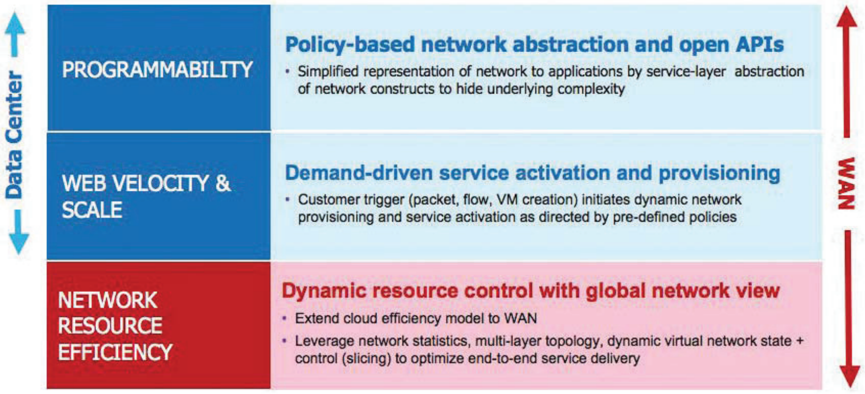

Figure 4: Key attributes of SDN

hot topic over the past few years and there are many interpretations of its definition. Fundamentally, it is the ability of the network to adapt automatically based on software commands or machine-to-machine communication.

An SDN infrastructure can be defined as one that is automatically programmable through a northbound API; it provides an abstracted view that can be understood by a higher level controller. The abstracted view is presented as a model of the underlying network elements.

SDN is being deployed as an enhancement both within the data centre and the Wide Area Network as shown in Figure 4. SDN was originally conceived as a university project, as a means to allow independent development of the control logic and the underlying data fabric. The Open Networking Foundation has championed its development ever since. The Foundation defined OpenFlow™ as a protocol to enable efficient communication between the layers so that a controller could manage a range of different switches.

Since inception, SDN has been widely embraced and a number of different implementations are taking shape. The first deployments are in data centres where an overlay control layer is proving critical in addressing the networking demands of the rapidly rising number of virtual machines driven by cloud computing. In these deployments, SDN is expected to help service providers ensure that network connections can be made as fast as the virtual machines within a server are created.

The adoption of cloud computing brings a much shorter service life-cycle combined with increased application mobility. The location of the host for a service can move at short notice and hence the environment is much more dynamic than has previously been experienced.

Adoption of SDN within the Wide Area Network has been identified as a way to improve the resource and capacity utilisation of the network by automating adjustments based on real-time usage. If the network is not as agile as the cloud services it supports, the cost of service delivery will rise dramatically as the service provider will need to over-provision the network to maintain the quality of experience.

AUTHOR'S CONCLUSIONS

The current era of cloud computing, smartphones and user choice is placing a lot of pressure on the traditional telecom networking paradigm and some radical changes are required. The Telecommunications sector must learn from the IT world and adopt cloud computing for its own network functions. Telecom functions introduce new and unique challenges in the area of latency and function distribution and these are being evaluated and tested as part of ongoing NFV activity.

Associated with NFV is the need to make the network as agile as the virtual machines used to support its functions, and SDN holds much promise in this area. While the two terms SDN and NFV are frequently spoken about together they are in fact separate but related topics.

The Telecommunications sector is at the start of a transformation journey as significant as the introduction of digital switching of the past but with expectations of accelerated time-scales and much promise for the future.

ABBREVIATIONS

API	Application Programming Interface
ATCA	Advanced Telecom Computing Architecture
COTS	Commercial Off-The-Shelf
NFV	Network Functions Virtualisation
SDN	Software-Defined Networking

ABOUT THE AUTHOR

PHIL TILLEY

Phil Tilley is Marketing Director Cloud solutions and strategy and is responsible for the positioning and promotion of Alcatel-Lucent Cloud Shift and NFV solutions. He has a BSc in Engineering from University College of North Wales and has over 20 years engineering and marketing experience and has always been an evangelist for the latest telecoms technologies.