Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING
WHITE PAPER

The Integration Challenges of Software-Defined & Virtualized Enterprise Networking

A Heavy Reading white paper produced for Ekinops

AUTHOR: JAMES CRAWSHAW, SENIOR ANALYST, HEAVY READING

INTRODUCTION

Traditionally, network deployments within the enterprise domain have been fairly static, with far less programmability than is typically seen in the data center. Networks are typically provisioned via command line interface or with simple scripting tools. The network estate is comprised of proprietary hardware running tightly coupled, proprietary software. All this proprietary hardware leaves an enterprise with a complex and inflexible network that is expensive to procure and manage.

SDN and NFV are set to bring significant benefits to enterprise networking in terms of automation, agility and lower total cost of ownership. NFV has matured sufficiently such that most enterprise network functions (routing, firewall, etc.) can now be hosted on Intel architecture-based servers.

As such, instead of deploying different dedicated hardware devices at the customer premises for each function, operators are looking to deploy universal customer premises equipment (uCPE) that can deliver a variety of branch office networking functions – such as routing, switching, network address translation (NAT), access control, quality of service (QoS), load balancing and performance monitoring – on a single appliance. Moreover, this uCPE has the capacity for additional virtual network functions (VNFs) that enable communications service providers (CSPs) to deploy value-added services, such as security or software-defined wide-area networking (SD-WAN).

CSPs are rushing to deploy SD-WAN services to their enterprise customers, driven by competitive forces. In the U.S., cable operators see SD-WAN as an opportunity to provide enterprise services outside of their existing footprint. Incumbent enterprise service providers are having to respond with their own SD-WAN offerings to remain competitive. In other cases, enterprises are looking to deploy SD-WAN on a do-it-yourself basis. Managed service providers are keen to avoiding losing business to the DIY approach, and hence have added SD-WAN to their existing MPLS and VPN offers.

Until recently, most enterprises were content to deploy another dedicated appliance for SD-WAN. Increasingly, they are looking to deploy SD-WAN as an application running, alongside others, on a universal CPE device, as shown in **Figure 1**.

Figure 1: Universal CPE Architecture

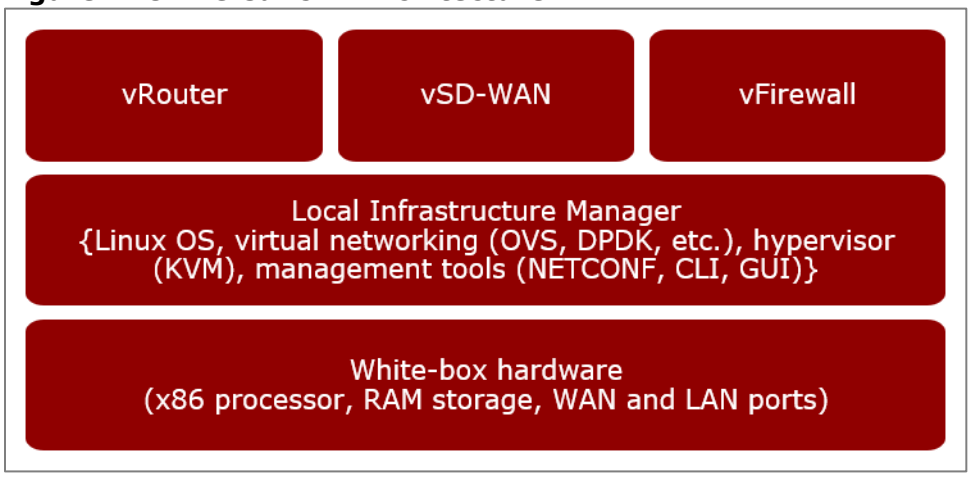

Source: Heavy Reading

SD-WAN is helping businesses reduce their connectivity costs and increase service agility. As the survey results in **Figure 2** show, lower cost, faster deployments and the ability to integrate multiple functions into one device are seen as some of the key attractions of SD-WAN solutions. Operators see the ability to offer additional VNFs as a major differentiator for their SD-WAN services. These VNFs can be best-of-breed or price-performance solutions selected either by the operator or its enterprise customer.

Figure 2: SD-WAN Differentiating Features

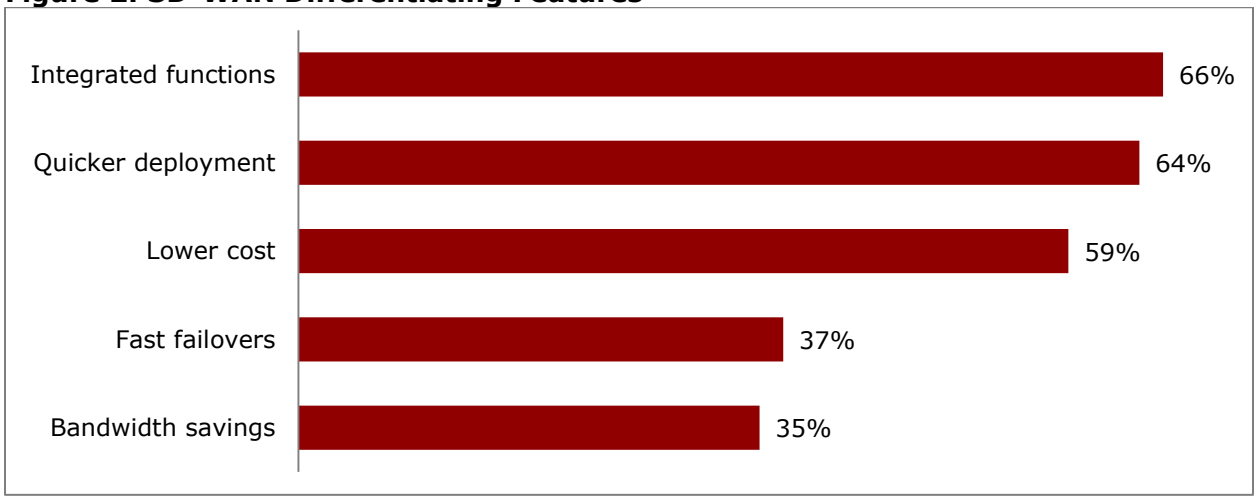

Source: Heavy Reading

Although NFV is not as topical with enterprise customers as SD-WAN, larger enterprises are interested in its potential, especially when combined with an SD-WAN deployment such that one piece of customer premises equipment can support multiple functions (e.g., firewall, IDS/IPS, load balancing, WAN acceleration) in addition to the core SD-WAN capability. This is particularly attractive for smaller branches, where the enterprise might not have much local IT support (or even space for multiple boxes).

INTEROPERABILITY & INTEGRATION CHALLENGES

To serve the nascent SD-WAN and NFV opportunities, service providers are generally going to market with a couple of different universal CPE platforms, a couple of different SD-WAN partners, and a couple of different security (mainly firewall) partners. However, even when restricting the number of options, CSPs have found it quite challenging to integrate their chosen SD-WAN solutions, firewalls and other VNFs with their chosen uCPE. NFV introduces additional layers of operational complexity that put the onus on the operator to integrate technologies that were traditionally integrated within a single appliance by their vendors.

VNF interoperability is a major challenge that comes from the disaggregation of network management systems, software, operating systems (OSs) and the underlying hardware. With physical infrastructure, this all came pre-integrated; with NFV, the operator has the fun job of gluing it all together themselves.

This task is quite complicated, as not only are operators looking to work with several different VNF providers, but each supplier also has multiple VNF versions and license options.

This can make it very hard to create a service chain that spans multiple VNFs from different vendors and works in the way that was intended.

As operators look to increase the number of virtualized network functions in their portfolios they face increasing integration challenges. For every VNF, service providers may need to make significant tooling investments to integrate with the VNF vendor's systems to provide a single pane of glass for their enterprise customers that covers monitoring, ticketing and performance/capacity management. The quick fix is to take a turnkey solution from one vendor, however, this leads to vendor lock-in, undermining a key part of the business case behind virtualization.

Instead operators that pursue a multi-vendor strategy will need to conduct thorough testing to ensure correct VNF instantiation, configuration and service chaining. In addition, they will need to test the integration of VNFs with northbound management systems to ensure automated deployment can be achieved.

TECHNICAL CHALLENGES OF NFV

In their haste to offer virtualized network functions, many vendors have simply ported their existing software from its traditional hardware environment (ASICs, FPGAs, network processors) to a commercial off-the-shelf (COTS) server platform (x86-based). They haven't had the time (and perhaps not the motivation) to re-architect their software to take advantage of the new NFV environment and as a result their software does not fully capture the benefits of virtualization.

Consequently, operators are finding that many VNFs require significant workarounds to get them to perform adequately on their chosen universal CPE platform. A lot of integration work is required to match the underlying hardware capability with the VNFs that are running on top. That might require adaptation of the hardware platform or tuning of the VNFs themselves.

A common issue encountered with VNFs is that they lack a feature-complete application programming interface (API) for a full integration with northbound management systems. Sometimes the management user interface of a VNF requires a Windows application to be used, instead of a simple web browser. VNFs often suffer from poor documentation about how to deploy. More detailed and meaningful log and alarm information is often required to troubleshoot installation and operational problems.

As you get into the practical details of VNF operation, operators are finding other pain points. For example, VNF serial numbers might be set using a shell command that works if executed from a console but is buggy if executed over Secure Shell (SSH). The devil, as they say, is in the details.

Virtual Networking & Hardware Acceleration Compatibility

One area that is proving troublesome for many VNFs is the use of techniques to improve the throughput of x86-based devices. As enterprises have moved from TDM-based connectivity to Ethernet-based connectivity in the WAN, the CPE devices at the demarcation between the WAN and the LAN are increasingly similar to the x86-based server platforms used in the data center. X86-based server technology has been able to run network functions such as

routing and firewall security for some time, although until recently the performance lagged behind what could be achieved with more specialized devices.

For example, in tests run by AT&T (see Building the Network of the Future, Donovan et al., 2017), the packet processing of virtual routers was shown to be 20 percent to 35 percent that of physical devices using the same amount of processing power. Service providers such as AT&T have worked with their suppliers to take advantage of new technologies such as Data Plane Development Kit (DPDK) and Single Root I/O Virtualization (SR-IOV), in addition to general software performance tuning focused on I/O and CPU optimization, in order to achieve higher packet processing on COTS (i.e., x86-based) hardware.

Linux Bridge

The basic method for virtual networking in Linux is called a bridge. This sits between the host OS (Linux kernel) and the hypervisor (e.g., KVM). As an Ethernet frame arrives at the physical network interface card (NIC) on the machine, it is copied by the host OS onto the Linux bridge. The frame is then switched to the virtual NIC of the VNF at the corresponding destination MAC address, as shown in **Figure 3**.

Figure 3: Virtual Networking – Linux Bridge

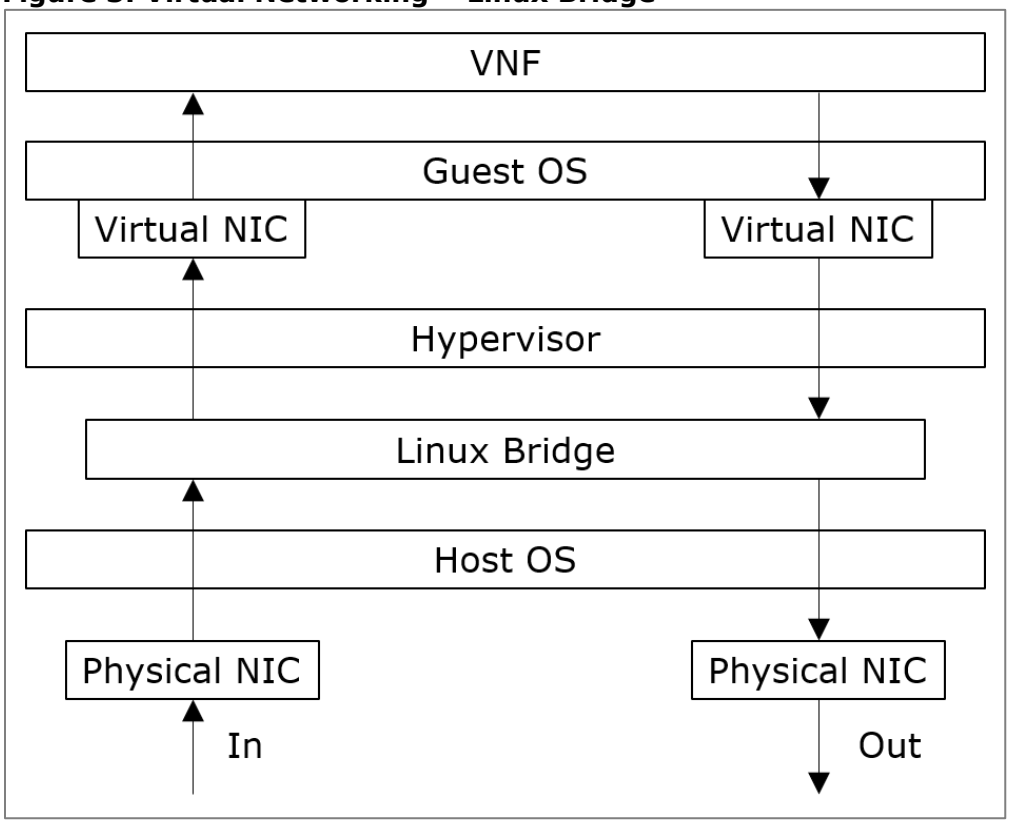

Source: Heavy Reading

Linux bridging can switch based on MAC address or VLAN tags but doesn't enable more granular traffic policy controls. Additionally, when creating a service chain of multiple VNFs running on the same machine, multiple Linux bridges are required in order to keep the traffic segregated.

Open vSwitch

Open vSwitch (OVS) was developed as an alternative switch for hardware virtualization environments. It offers more functionality than the basic Linux bridge such as support for VXLAN, OpenFlow, GRE, QoS, NetFlow and granular traffic policy. However, this comes at the cost of lower performance due to the increased processing required for each frame switched.

In order to mitigate this, Intel developed the Data Plane Development Kit, which uses special libraries designed to accelerate packet processing on x86 hardware. To enable it, the VNF must implement a virtual driver to make use of DPDK-enabled virtual switching. The VNF is then able to copy a packet directly from the physical NIC of the machine to its own memory space (and vice versa) without having to make an intermediate copy in the kernel space (the host OS). While DPDK was initially introduced by Intel it is now an open source project and is supported on other CPU architectures including ARM and PowerPC.

SR-IOV

SR-IOV enables even greater virtual switching performance by enabling Ethernet frames to pass straight from the physical NIC of the machine to the virtual NIC of the VNF without needing to be processed by the hypervisor or an intermediate virtual switch or bridge (as is the case with OVS or Linux bridge). With SR-IOV each VNF is granted direct access to the physical NIC. As the frame arrives at the physical NIC it is directly placed into the hardware queue associated with each VNF's MAC address or VLAN tag.

The downside with SR-IOV is that the VNF cannot be easily moved to another server as it is tied directly to a particular hardware resource. SR-IOV is not suitable for a dynamic NFV environment where VNFs are frequently moving to different servers (e.g., in different geographic locations). However, for a use case such as on an enterprise CPE (where the VNF must always run on the same hardware), this is not a problem.

Figure 4: Comparison of Virtual Switching Technologies

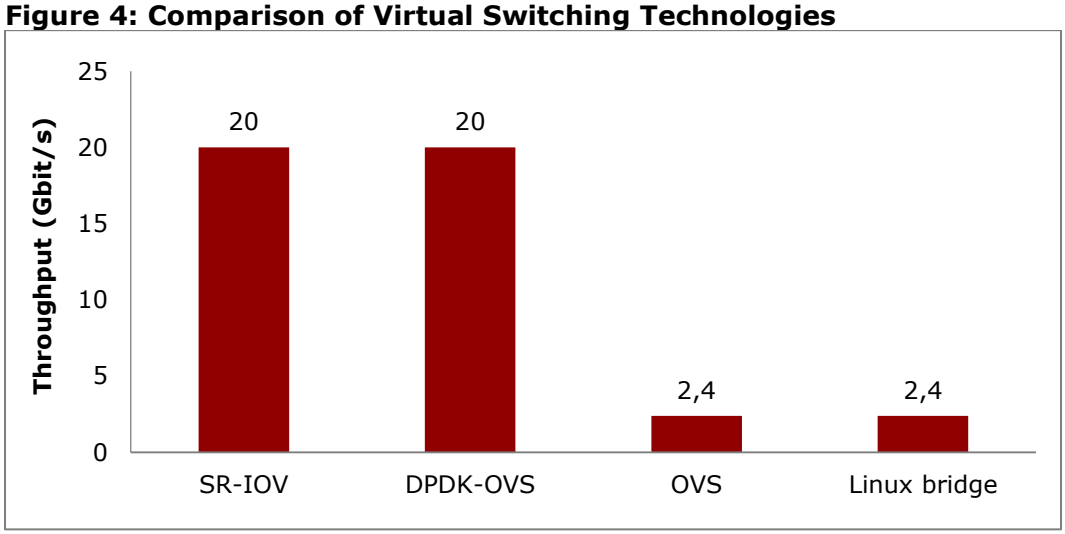

Source: AT&T; packet size = 384 bytes

While the benefits of SR-IOV over basic OVS seem clear, operators have found that not all VNFs fully support SR-IOV. This can mean that the throughput of the VNF is significantly lower than anticipated.

OS & CPU Dependency

It is not just networking functions that can struggle to work well on uCPE. One example we learned of was a cybersecurity solution that was designed to run on a Windows server and became very resource-hungry when it was run using the KVM hypervisor on Linux. Another issue is CPU dependency: Some VNF components might require CPUs that support particular features (e.g., Advanced Vector Extensions) that are only available on certain Intel CPU families. As a result, these VNFs may not perform as expected on a uCPE that uses a processor type that lacks these features.

THE NEED FOR VNF TESTING

Given all the technical challenges of NFV, operators will need to create a repository of VNFs that they have tested on the different uCPE hardware versions they plan to deploy. The validation process should cover manageability, scaleability, functionality, chainability, upgradeability, data-plane performance and performance on different x86-based hardware. Testing is not limited to the VNFs in isolation; operators must also test how VNFs perform as part of a service chain.

Manageability

Manageability includes licensing validation and day zero (initial) configuration: creating a service model, defining an SLA, instantiating a service and configuring the service. Operators may manage the day 1 configuration (service monitoring) by sending a specific workflow API to the VNF. This will require that the VNF has a northbound API, so it can be configured from a third-party management system. Even if a VNF provider uses NETCONF (though most today use REST) they use different naming conventions and parameters. Each VNF has a unique information model which requires mapping to the VNF manager's information model.

Most VNF vendors come from an appliance background, where they are used to having their own management tools and controllers. VNF vendors all have their own specific workflow to instantiate, manage and deploy their VNFs. Although NETCONF offers a standardized protocol for configuration management, each VNF vendor is at a different level of maturity in implementing or even assessing NETCONF. Supporting NETCONF/YANG requires a significant rewrite of software. As industry-standard YANG models come out, vendors will gradually migrate toward them. However, the current state is that most VNFs operate like islands, without taking into account the broader circumstances when they need to service chain with other vendors' VNFs.

Scaleability, Functionality, Chainability & Upgradeability

Operators need to test how VNFs handle vertical scaling (increased processing power drawn from the same machine) for small, medium and large VNF instances (number of CPU and amount of RAM varies). The specific functionality of the VNF needs testing. For example, for

an SD-WAN, does it support Forward Error Correction (FEC)? For WAN optimization, how much gain can be achieved?

Operators should evaluate what technology the VNF supports in order to be integrated into a service chain. They also need to consider if the VNF manages its own software upgrades or if this has to be triggered by a third-party management system.

Data-Plane Performance

Each vendor provides its own unique workflow for configuring the VNF. The VNF has to be aware of the underlying hardware to take advantage of hardware acceleration features such as SR-IOV or the availability of encryption acceleration capability. Using a traffic generator, operators should measure at what point the VNF starts dropping packets. It should be tested using different hardware acceleration technologies:

- OVS with DPDK (single and multi-queue mode)
- Virtio, a virtualization standard for network and disk device drivers where just the guest's device driver "knows" it is running in a virtual environment and cooperates with the hypervisor. This enables guests to get high-performance network and disk operations.
- The Fast Data Project (FD.io) is a packet processing engine that can run on a Linux-based host and uses vector packet processing (VPP) mechanisms to achieve high performance for packet processing, routing, switching and NAT. FD.io provides an abstract environment for building virtual routers, switches and packet processors. FD.io can use DPDK to communicate directly with NIC cards. DPDK provides hardware acceleration to FD.io. However, hardware acceleration and usage of DPDK is optional – an FD.io-based application can still run without DPDK.
- SR-IOV enables switching packets between VNFs without consuming any CPU cycles. Its availability depends on the network interface model. For example, you might test VNFs using the Intel I350 (Ethernet Server Adapter) and X553, X772 (backplanes), depending on the specifics of your chosen uCPE.
- Peripheral Component Interconnect (PCI) passthrough – allows guests to have exclusive access to PCI devices (any piece of computer hardware that plugs directly into a PCI slot on a computer's motherboard) for a range of tasks. PCI passthrough allows PCI devices to appear and behave as if they were physically attached to the guest OS.

Performance on Different Processors

There are a variety of x86-based processors that can be used for uCPE, e.g., Xeon (D, W, E, etc.) and Atom (C, E, X, etc.). Not all service chaining mechanisms may be available with some processor families. Some VNFs require specific CPU instructions and if this is not available within the hardware CPU on the machine the hypervisor path cannot emulate the instructions and hence the VNF will not work.

Today's CPUs can offer very sophisticated instructions, but they can differ significantly between architectures (e.g., you may have different instructions for moving a byte from one register to another). VNFs that have been designed for a specific CPU version may malfunction when run on a different CPU type that lacks a specific instruction.

Service Chaining

In addition to a VNF test repository, operators should create a repository of tested service chains (e.g., SD-WAN version 123 plus firewall version xyz) that have been defined with a Network Service Descriptor (NSD). The NSD is a chaining of VNFs with interfaces and embedded services, ready to be deployed on a particular uCPE. Once ready, the NSD is exported to an orchestrator for deployment. However, creating an NSD is difficult and requires a lot of expertise and analysis of VNF documentation.
To chain two VNFs running on separate virtual machines, you need a virtual network. There are many options for creating this virtual network (e.g., OVS, SR-IOV, DPDK, PCI passthrough), though as discussed above not all VNFs will support them. Hence the need to design the NSD taking into account the idiosyncrasies of each VNF.

For example, the virtual network could rely on OVS. But with OVS, there are at least two options: OVS with Virt.io interfaces and OVS with DPDK interfaces. DPDK offers better performance but consumes more processing power than Virt.io. Furthermore, there are two sub-options with DPDK: single queue per interface or multi-queue for faster processing. Whichever permutation you decide on for your service chain has to be supported by all the VNFs in the chain. Given the immaturity of NFV, VNFs and NSDs are constantly changing (e.g., new features or API extensions are being added). This requires retesting of the VNFs and NSDs each time a change is made.

An overview of what a validation setup might look like is shown in **Figure 5**. In the tests, customers A and B each have a private office with its own uCPE, and both share a third uCPE in a service provider datacenter. All three uCPEs are connected via both an Internet connection and a dedicated MPLS circuit. Each uCPE runs an SD-WAN instance that is controlled by an SD-WAN orchestrator.

Figure 5: Simplified View of uCPE & VNF Validation Setup

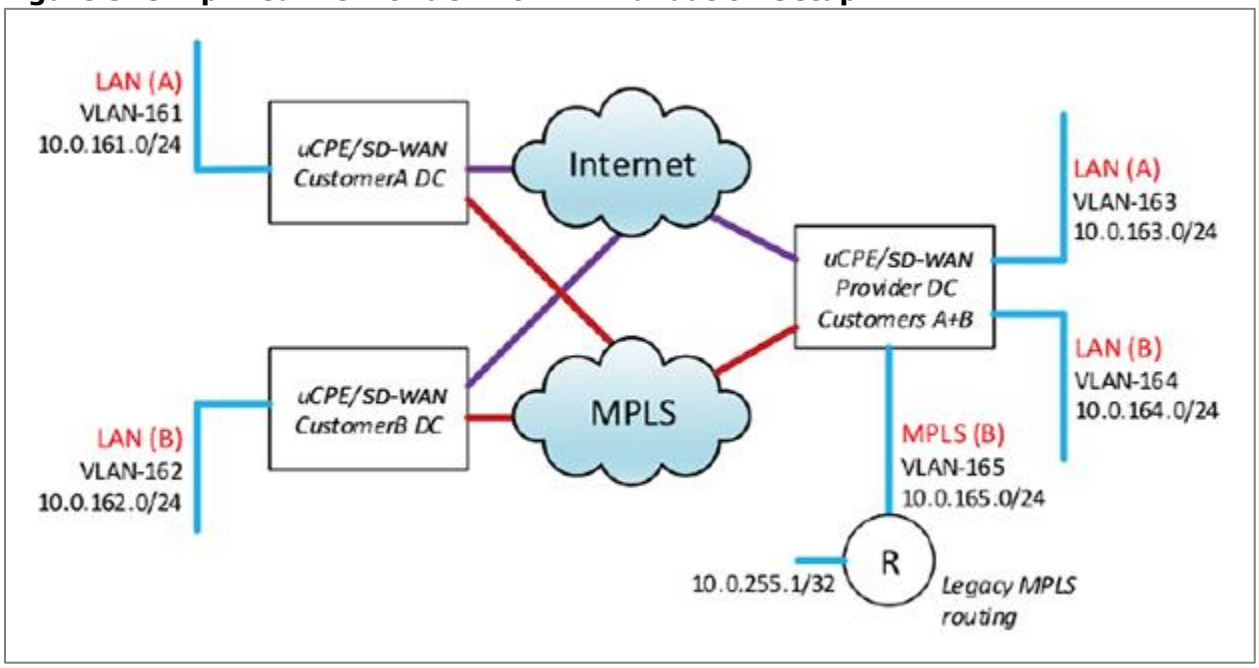

Source: Advantech

ORGANIZATIONAL CHALLENGES OF NFV

As well as the technical challenges of NFV, this new paradigm also introduces several organization challenges. Operators used to only need to concern themselves with networking technology. With SD-WAN and NFV they need new skills in virtualization, software management and even writing YANG models to describe their service offering.
When problems arise with VNFs running on uCPE, the operator may face difficulty in identifying who is to blame: the VNF vendor, the uCPE vendor, other VNFs or PNFs in the same service chain, or some combination of these parties. With multiple VNFs running on the same hardware, operators may struggle to find the root cause of a problem. This can make it challenging to ensure resiliency to support the SLAs expected by an application or end customer.

To avoid such problems, the temptation is to overprovision the uCPE, but this can make it unnecessarily expensive. Overprovisioning could potentially make a uCPE solution (plus VNFs) even more expensive than the collection of physical appliances that it is meant to replace.

Another challenge in managing VNFs is that it can be unclear where the management demarcation lies between the VNF Manager and the higher layer resource orchestrator when it comes to scaling up and down of VNFs. In general, CSPs are not familiar with orchestration – many are still using static scripting to configure and manage devices/functions.

Figure 6 shows a generic architecture for NFV. The NFV Infrastructure (NFVI) equates to the uCPE, which comes with its own Virtual Infrastructure Manager (VIM) – essentially a Linux OS plus middleware. The VNFs have their own managers, which are directed by multiple layers of orchestration above. The demarcation of responsibility between these layers is unclear, and expectations will differ across VNF suppliers.

Figure 6: Orchestration Layers

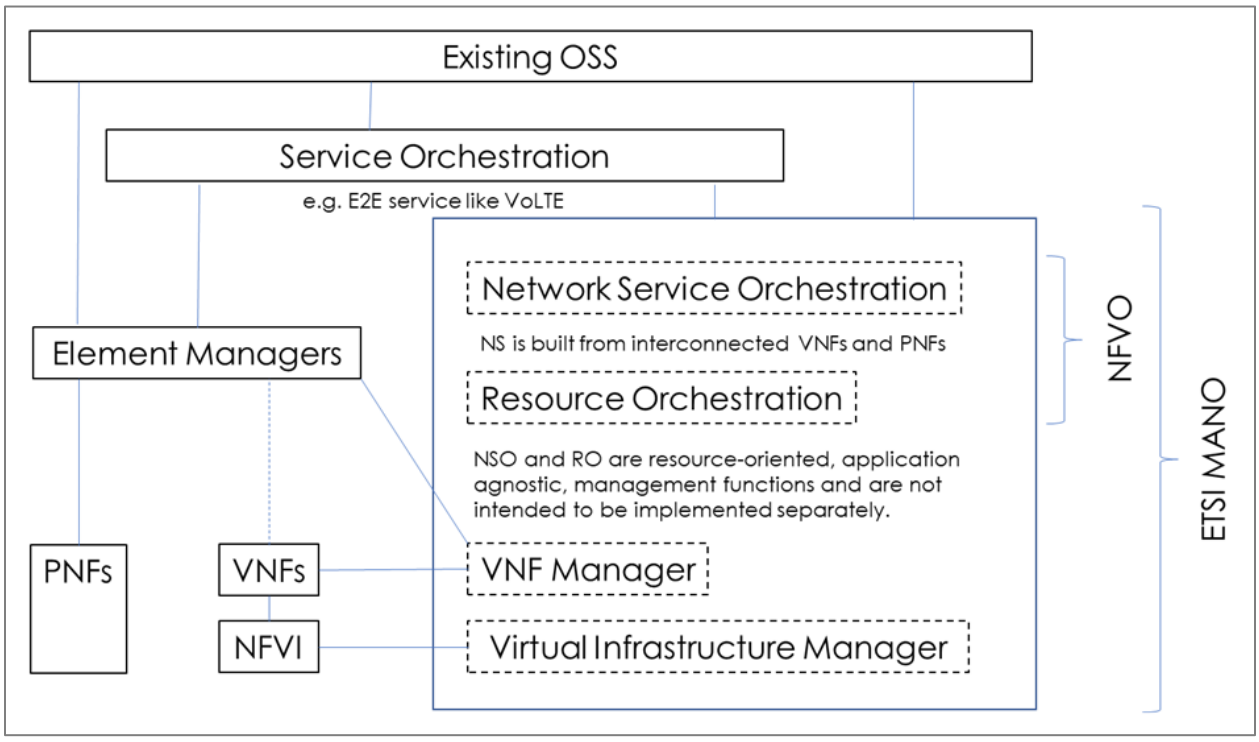

Source: Heavy Reading

Another challenge is the degree of control that operators are prepared to give their enterprise customers in the management of the uCPE and the VNFs that run on it. For example, the enterprise admin may want to make changes to firewall settings but may have to request that the CSP implement these. Large enterprises want to keep control of their security systems and are unlikely to see much benefit in virtualization if it leads to a loss of control.

CONCLUSIONS

Over the last 30 years, the provision of telecom technology moved from turnkey solutions, where one company would supply all the infrastructure and integrate it, to a fragmented supply chain where multiple, best-of-breed suppliers are selected for each element in order to avoid supply risks and ensure competitive pricing. With NFV, the industry is moving to a new level of fragmentation, with the supply of basic hardware being separated from the software that runs on top. This software itself is further split into separate operating systems, hypervisors, VNFs and management systems (orchestrators and controllers). All these elements are interdependent, which brings more flexibility in system design.

However, operators are finding there are too many options with NFV, and it is very time-consuming to understand them all. It becomes even more complicated in an environment with mixed versions of VNFs, and each VNF uses its own management APIs. Each time a new version of a VNF is released, an operator might need to go through hundreds of pages of documentation to get the necessary API information to manage it. Frequent releases of VNF software can lead to uncertainty about which is the best version to pick. CSPs have to spend considerable effort to discover the configuration options that each VNF provides. For example, an SD-WAN VNF may have options for WAN optimization and FEC. There may be options for deployment on one or multiple vCPUs and different amounts of RAM.

Given that hardware is now split from software, VNF providers have built generic applications, not customized for the particular hardware an operator has chosen for their uCPE. These VNFs might be affected by how the uCPE is configured. This can lead to a lot of finger-pointing when things go wrong. Is it the VNFs fault, the uCPE, or a management-layer issue? Often there is no clear accountability.

What operators need is for someone to take ownership of the NFV deployment – either internal project managers, external systems integrators, or suppliers that can work as a lead implementor, resolving issues within the ecosystem of NFV suppliers. These lead implementors can help with the VNF interoperability testing and onboarding process, ensuring that VNFs will work in the specific NFV environment an operator has chosen (i.e., its specific combination of computing infrastructure, VIM, VNFM, NFVO, etc.).

For successful NFV deployments, operators need to partner with vendors that can provide systems integration capabilities and train more of their own staff in key technologies such as Linux, virtualization, OpenStack and DevOps tooling. To make enterprise deployments as painless as possible, operators should select uCPE vendors that support a broad range of third-party VNFs. Agreements with uCPE vendors should include terms for the support of new VNFs (and new versions of existing VNFs), such as testing and evaluation of the impact on service chains.

To learn more about the integration challenges of software-defined and virtualized enterprise networking, please watch the webinar sponsored by Ekinops, "The Integration Challenges of Software Defined and Virtualized Enterprise Networking," first shown on December 18, 2018.

ABOUT EKINOPS

Ekinops is a leading provider of open and fully interoperable Layer 1, 2 and 3 solutions to service providers around the world. Our programmable and highly scalable solutions enable the fast, flexible and cost-effective deployment of new services for both high-speed, high-capacity optical transport as well as virtualization-enabled managed enterprise services.

Our product portfolio consists of two highly complementary product sets. One, marketed under the Ekinops 360 brand name, provides a single, fully integrated platform for metro, regional and long-haul applications. The other, marketed under the OneAccess brand name, provides a wide choice of physical and virtualized deployment options for Layer 2 and Layer 3 network functions.

As service providers embrace SDN and NFV deployment models, Ekinops' solutions enable them to deploy today in the knowledge that they can seamlessly migrate to an open virtualized delivery model at a time of their choosing. Drawing on its experience of integrating with different hardware platforms, third-party VNFs and different orchestration systems and controllers, Ekinops complements its virtualized product portfolio with an extensive range of integration services to act as a single point of contact, extend its customers in-house skills and speed up time-to-market for managed virtualized services.

A global organization with operations in four continents, Ekinops (EKI) – a public company traded on the Euronext Paris exchange – is headquartered in Lannion, France, and Ekinops Corp., a wholly owned subsidiary, is incorporated in the U.S.