IDG

WHITE PAPER

# NETWORK VIRTUALIZATION STRIPPED DOWN

## Network Function Virtualization Revolutionizes Network Economics, But It Doesn't Fit Every Scenario

Software-defined infrastructure has finally reached the network layer, and that is welcome news for IT organizations. Traditionally the domain of high-priced proprietary gear, networks are now enjoying the same cost dynamics that have enabled IT organizations to replace many of their legacy servers and storage devices with low-cost "white box" equipment while implementing value-added functions in software.

Although this trend is good news for both technology budgets and user choice, it is not without its dangers. Technologies like software-defined networking (SDN) and network function virtualization (NFV) have created elevated expectations that software can displace specialty equipment at every level of the network stack. Reality has not yet caught up with that promise, however. The transition will take some time, and even then some layers of the network stack may always favor the use of special-purpose hardware.

Gartner's July 2015 *Network Hype Cycle* placed both NFV and SDN in the "Trough of Disillusionment." Translation: It's a much-hyped new technology that so far fails to live up to its billing. Every successful new technology must navigate this stage before realizing its potential, so now is a good time to look at the myths and realities of network virtualization.

PRESENTED BY

Copper River Information Technology

JUNIPER NETWORKS

## DEFINITIONS

Let's start by defining the terms.

NFV decouples network functions such as network address translation, firewalling, intrusion prevention systems, VPN and caching from proprietary hardware appliances, and places those functions in virtual machines, similar to the way server virtualization uses software to simulate hardware environments. Thanks to Moore's Law and the rapid evolution of off-the-shelf X86-based hardware, compute performance rivals some specialty appliances that have populated network racks for years. This allows some modest-capacity tasks to be virtualized with NFV.

SDN is sometimes confused with NFV, but it is actually quite different. SDN provides a high level of abstraction that virtualizes the entire network infrastructure, enabling administrators to create independent virtual subnetworks and to manage them in software. While NFV focuses on virtualizing appliance-specific functions that traditionally have been physically placed in the network communications path, SDN virtualizes the network as a whole. The two technologies are different but complementary, and they are intended to work in harmony.

## NFV BENEFITS AND USE CASES

The principal benefit of NFV is to provide a low-cost alternative to specialized hardware that can be rapidly deployed and scaled out to address changing demands. Organizations can reduce vendor lock-in, decrease time to market and gain a level of flexibility that enables them to easily redeploy software and services as needed. By automating routine operations, NFV can save on operational costs, simplify complexity and increase delivery speed.

NFV permits many functions to run on commodity servers with accompanying savings in setup, cabling and management costs. It provides centralized orchestration and management, and removes manual operations, thereby reducing configuration errors. Troubleshooting can be as simple as swapping in a new virtual machine and then configuring the software—rather than the costly and time-consuming process of ordering and installing new hardware.

A good use case scenario for NFV would be a managed service provider or enterprise looking to launch into a new geographic region with predictable network growth expectations. Many small and midsize businesses find NFV to be an attractive option when handling functions at the edge of the network such as virtual routers, virtual firewalls and virtual WAN acceleration. This allows these routine but essential operations to be centralized in order to simplify local operations. With flexibility to scale as needed, NFV can effectively replace many of today's specialized hardware appliances.

## NFV TRADE-OFFS

Considering the compelling benefits of NFV, why isn't the entire market shifting toward virtual solutions? The answer to this question mirrors the reason that engineering companies buy $10,000 CAD/CAM software instead of $99 consumer-grade alternatives. While the low-end products offer outstanding price-performance, there are certain specialized functions that can only be handled by high-end tools.

Hardware scalability and economic feasibility are the two most significant limitations of current NFV solutions.

There are two kinds of scalability: scale out and scale up. Scale-out functionality excels in cases in which additional devices can be added or subtracted to accommodate variable demand. In this scenario, traffic can typically be split into multiple streams and processed in parallel without a significant impact on performance. Think of the analogy of a large toll plaza on a highway, which fans out traffic into many streams and then reconsolidates those streams. Web servers are an excellent example of scale-out functionality, because spikes in demand can usually be handled by adding servers that duplicate the functionality of others on the network.

Scale-up functionality is required when large amounts of traffic must pass through a single point or network function, as is the case with high-volume routing and firewalls. Parallelizing such traffic is impractical, since each packet must be handled in sequence to avoid corrupting data or introducing vulnerabilities. To use the high-

way example, an on-ramp could be considered a scale-up scenario. It would be impractical and dangerous to take the scale-out approach of adding more lanes. The only way to increase volume is by widening the road onto which traffic flows.

In its current incarnation, virtualization is best suited for scale-out functions, and physical servers better suited for scale-up. Software makers have focused their attention on these markets because that is where the volume is greatest. Highly scalable NFV software may still be years away.

Economic feasibility is a function of total cost of ownership. It takes into account all the costs associated with buying, maintaining and administering equipment, including "soft" costs like salaries as well as the business risk of equipment failure. As we shall see, inexpensive hardware and software do not necessarily equate to low cost of ownership.

## WHEN SPECIAL PURPOSE IS BEST

In scale-up scenarios, dedicated, purpose-built hardware that supports open standards and extensible network operating systems is usually the best choice. This is both an architectural and an economic decision. Consider these three scale-up scenarios in which an integrated hardware/software combination makes sense.

Juniper Networks' SRX Series of Gateways provide integrated threat intelligence delivered on a resilient platform. They can be scaled up to 100 gigabits of Ethernet throughput and up to 2 terabits per second of performance in the data center. This is a capacity that no virtualized alternative can currently match.

Juniper's MX Series of 3D Universal Edge Routers are the Swiss Army knives of network edge platforms. Functioning as a giant virtual cross-connect, the MX is at once a switch, a router, a quality-of-service manager, an SDN gateway and a translator. Built on open standards, it can translate nearly any protocol into any other protocol. Products like the MX Series must sequentially translate very large amounts of data in order to do their job. This demands scale-up functionality, which is not a strength of virtualization technology.

Juniper's QFX10000 Spine Switches contain a subset of the functionality of the MX Series, but in a highly scalable package designed specifically for data center applications such as core switching or data center interconnect. Using custom application-specific integrated circuits (ASICs), the versatile QFX10000 switches scale up to an incredible 96 terabits per second of throughput. Such speeds are not possible with current virtualized-based equipment.

## MAKING SMART CHOICES

Over time, many of the functions of the high-end routers and switches just described may become available in virtualized form as computing power improves, but there is no guarantee of when and whether that may occur. Currently, commodity hardware cannot approach the speed of custom ASICs.

In fact, it may be nearly impossible to ever duplicate some scale-up functionality in a scale-out architecture, regardless of hardware power. Some scenarios run up against the physical limitations of bandwidth; ASICs will always be faster than software. The compromises required to parse network streams for processing by multiple boxes may also create inherent performance or data integrity problems that simply can't be solved with a scale-out approach. Rearranging traffic for parallel processing carries intrinsic overhead penalties (such as latency and out-of-order packets).

Substituting multiple scale-out devices for a single scale-up device introduces additional complexity and points of failure. For example, imagine the administrative burden of managing 20 NFV servers instead of one. Given the staff costs, many IT organizations opt for a simpler infrastructure—even at a higher cost—because the total cost of ownership is lower, and there is less risk of failure.

Distributed NFV devices are also more time-consuming to manage than specialized equipment. For example, software upgrades must be performed on each device individually, which takes up time and resources while increasing the number of potential failure points. Automation is addressing some of these shortcomings, but there is no question that multiple devices require more administrative oversight than fewer devices. Also, automation cannot overcome some of these challenges.

One of the most compelling arguments for NFV technology is that it increases customer choice and reduces the risk of lock-in by separating software functionality from hardware. However, many of the same flexibility benefits can be achieved by adopting technology based on open standards. For example, Juniper Networks' Junos operating system supports traditional open standards-based routing and switching protocols as well as the SDN based OpenFlow communications protocol, which provides access to the critical forwarding plane of a network switch or router to a centralized controller. This gives IT organizations the flexibility to integrate with other OpenFlow-based equipment without extensive reprogramming or customized support development.

## BOTTOM LINE

NFV is a powerful new option for IT managers to consider when evaluating equipment purchases. Like many novel technologies, however, hype sometimes gets ahead of reality. Consider carefully your network growth expectations as well as all total-cost-of-ownership factors in order to make a decision that doesn't limit your business growth in the future. ■

## RECOMMENDATIONS

One truism about network capacity is that demand always catches up with supply. For example, many organizations now deliver streaming video and videoconferencing to employee desktops. These high-bandwidth, latency-intolerant applications were once the domain of specialized equipment, but now are accommodated in off-the-shelf hardware and software. Here are some factors to consider:

- What new demands will the mobile workforce, big data analytics or real-time streaming video place on networks three years down the road? Overprovisioning for future growth, although not very efficient, has traditionally not been a bad strategy when it comes to network capacity.
- When deploying hybrid virtual/physical network infrastructure, use the NFV management and organization (MANO) framework as a guide to provide consistent management and orchestration of all resources.
- Use comparable evaluation criteria when comparing capabilities between physical and virtual options.
- Likewise, management elements should also be similar in features to allow for a smooth transition between physical and virtual environments. Limit the number of network equipment vendors, or ensure that all adhere to open standards.
- To ensure choice and prevent vendor lock-in, look for scale-up solutions that use open protocols while providing excellent bandwidth scalability and efficient power management. Choose solutions that can easily interconnect with the network and later support NFV functions if needed. Vendors should be willing to provide a roadmap that ensures any investment will be protected, even as needs change.
- Consider the costs and complexity of scale-out architectures to solve specific bottlenecks in the network. While they make sense in applications for which high-capacity bandwidth is not an issue, scale-up environments demand different approaches.
- When calculating cost, consider carefully the administrative overhead, data center footprint, software licenses and failure risks at each key point in the data's path through the network. If an NFV solution introduces a potentially performance-crippling vulnerability at any point, it may be better to invest in a flexible scale-up alternative.