White Paper

Platforms for Accelerating the Virtual Infrastructure

Prepared by

Simon Stanley
Contributing Analyst, *Heavy Reading*
www.heavyreading.com

on behalf of

www.advantech.com

September 2014

Executive Summary

Network infrastructures are shifting from physical systems to virtual functions, and this requires a new class of network appliance that provides high-performance processing, balanced input/output (I/O) and hardware or software acceleration. Software-defined networking (SDN) and network functions virtualization (NFV) are changing the way networks and services are provisioned. This new virtualized infrastructure requires a combination of standard server technology and modular systems that can be configured to support line-rate performance with network interfaces up to 100 Gbit/s.

The increasing use of smart devices, such as mobile phones and tablets, and cloud-based services, such as remote storage and video on demand (VoD), is driving data bandwidth and requiring a much more flexible network. Mobile networks require a complex infrastructure that includes systems to handle data and voice connectivity, quality of service (QoS) and subscriber management. Data centers need high-speed connectivity and access to storage and other resources based on the services being provided. Conventional networks require significant investment and can take days or weeks to provision.

SDN and NFV enable a virtualized infrastructure where functions and resources can be provisioned and reallocated to meet short-term requirements. This gives service providers significant flexibility in deploying expensive hardware resources to meet customer demands. Service providers are expecting to drive new revenues and dramatically increase return on investment (ROI) by using standard server platforms that support virtual functions.

The challenge for anyone deploying SDN and NFV is delivering these benefits while still maintaining line-rate performance. Many network ports that are already running at 10 Gbit/s will quickly move to 40 Gbit/s and 100 Gbit/s as data rates continue to increase. Virtual environments extend this challenge by increasing the East-West traffic between virtual functions running on different hardware platforms. The key to meeting this challenge is to deploy hardware platforms that can support the SDN- and NFV-based virtual infrastructure and have integrated hardware and software to support high-speed network interfaces and the acceleration of critical functions, such as security processing and load balancing.

The purpose of this white paper is to examine these issues. The paper explores the requirements for delivering line-speed performance in a virtual infrastructure environment and reviews an exciting solution to this challenge that is based on a 2U rack-mount chassis with four Intel Xeon E5-4600 v2 series processors, with up to 12 cores per processor, and integrated support for up to 640 Gbit/s of I/O bandwidth. This highly-integrated solution provides the flexibility to implement security acceleration up to 400 Gbit/s and stateful load balancing across many virtual servers and networking I/O. The platform supports multiple high-speed network interfaces, including 100 Gigabit Ethernet (100GE). The paper also describes an off-the-shelf software solution for supporting NFV and other virtual environments on this platform with line-rate performance.

SDN, NFV & the Virtual Infrastructure

Service providers need a virtual infrastructure where resources and applications can be provisioned automatically to meet customer needs. This allows service providers to deliver flexible services and maximize ROI. SDN and NFV are two complimentary approaches to delivering this virtual infrastructure. **Figure 1** shows the key elements of each approach.

Figure 1: SDN & NFV

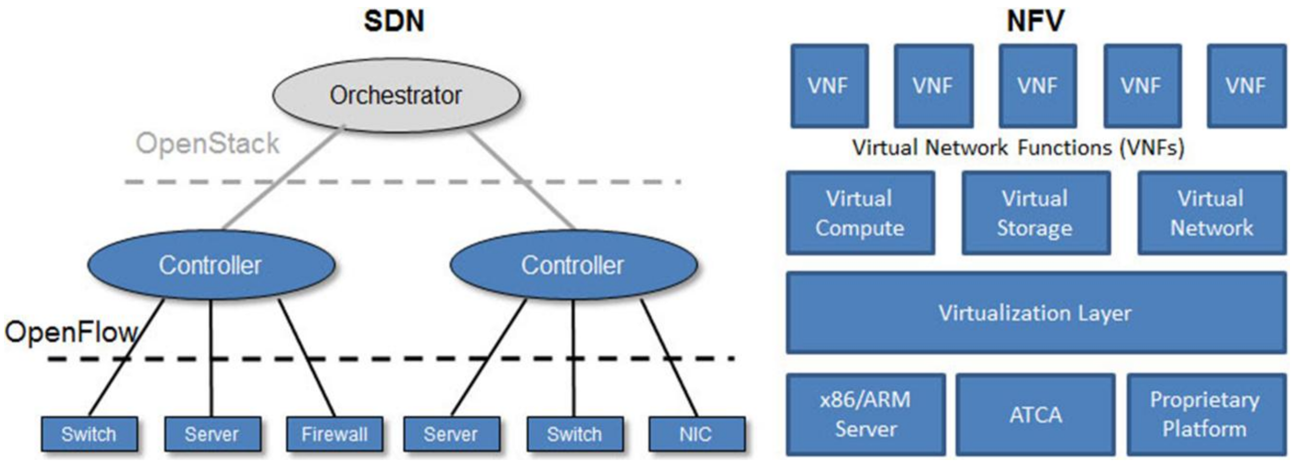

Source: Earlswood Marketing

SDN replaces the physical network with a mix of virtual and physical elements connected through open interfaces, such as OpenFlow and OpenStack. At the lowest level are network elements, including switches, servers, firewalls and network interface cards (NICs); these can be physical systems or virtual implementations running on standard server hardware. These elements are configured and monitored by controllers that are virtual elements running on standard server hardware. At the top level is an orchestrator that ensures the correct virtual and physical elements are in place to deliver the services required.

NFV replaces the physical network functions with virtual network functions (VNFs) that implement the same capability. The VNFs have access to virtual compute, storage and network resources that are mapped through a virtualization layer to physical systems. These physical systems may be x86- or ARM-based servers, ATCA platforms or any proprietary platform with the necessary resources.

There is much effort going into the development of virtual infrastructure implementations that take advantage of both SDN and NFV and make maximum use of standard server technology and handle data traffic at line-rate speeds.

Key to delivering this virtual infrastructure are high-performance server platforms with balanced I/O, hardware acceleration for key functions, such as security processing and load balancing, and accelerated software that can deliver line rate performance in a virtualized environment.

Standard Server Platforms

Most standard server platforms are based on Intel processors. Server systems for SDN and NFV do not need to use the highest performance processors but, to realize the full cost benefits of virtualized systems, must deliver high performance at moderate power consumption.

The Intel Xeon E5-4600 v2 series processors integrate hardware support for virtualization and offer an optimal mix of performance, I/O bandwidth and power consumption for SDN and NFV applications. The Intel Xeon E5-4600 v2 series processors have 4, 6, 8, 10 or 12 cores, 10-20M cache and integrate two high speed (up to 8 GT/s) QuickPath Interconnect (QPI) interfaces supporting cache coherent, quad processor implementations. A single server can, therefore, support up to 24 cores with two processors or 48 cores with four processors. The maximum thermal design power (TDP) consumption per processor ranges from 95W to 130W.

Figure 2 shows an Intel Xeon E5-4600 v2 server platform with four Intel E5-4600 v2 series processors and an Intel C600 Series chipset. In this particular implementation, each processor has four DDR3 memory channels and 40 PCI Express (PCIe) Gen 3 lanes, providing huge I/O bandwidth that matches the processing power of the four processors and eliminates possible QPI bottlenecks. The PCIe lanes on each processor can be configured to support two x16 PCIe Gen 3 interfaces and a single x8 PCIe Gen 3 interface or five x8 PCIe Gen 3 interfaces. Standard PC interfaces, including SATA, SAS and USB, are supported through the Intel C600 Series chipset.

Figure 2: Intel Xeon E5-4600 v2 Server Platform

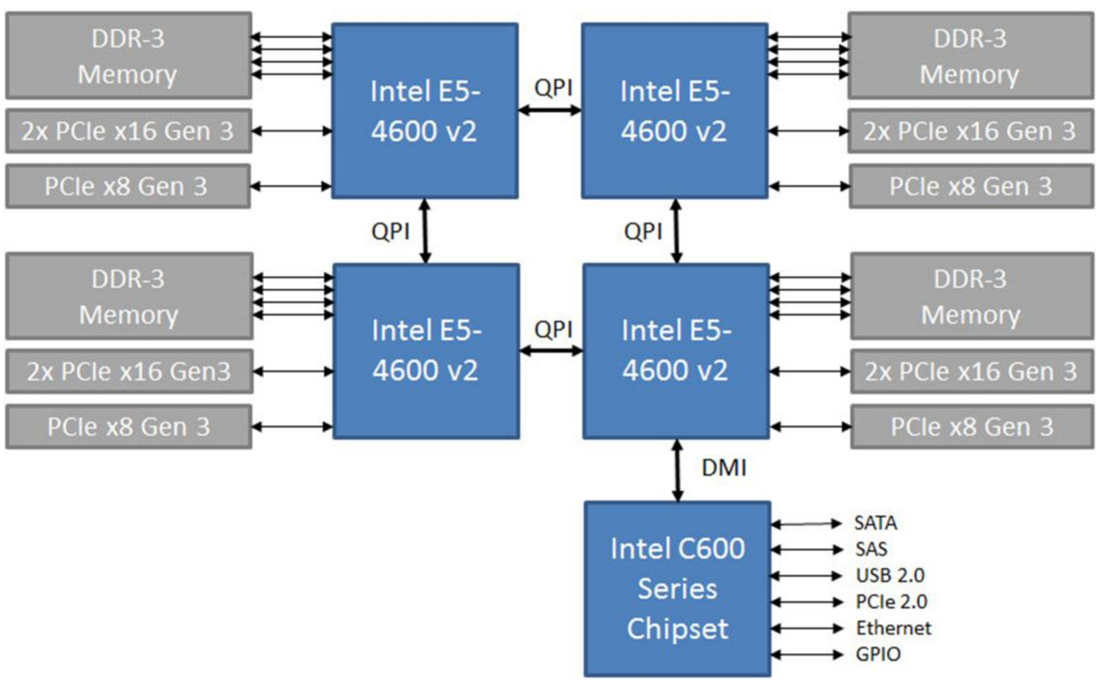

Source: Heavy Reading, Intel

2U Network Application Platform

The Advantech FWA-6512 (**Figure 3**) is a 2U network application platform that integrates four Intel Xeon E5-4600 v2 processors and an Intel C604 chipset. It supports processors with up to 130W maximum TDP, has two dual-processor motherboards and supports up to 1TB DDR3 memory and two hot – swappable 2.5" SATA HDDs or SDDs.

Figure 3: Advantech FWA-6512 Network Application Platform

Source: Advantech

The 160 PCIe lanes on the processors support four internal PCIe slots and up to eight network mezzanine cards (NMCs). Each of the eight NMCs can support two 40GE interfaces, enabling a total I/O bandwidth of 640 Gbit/s. The chassis has six hot-swappable cooling fans and redundant 1,400W AC or 1,500W DC power supplies.

Figure 4: Advantech Network Mezzanine Cards (NMCs)

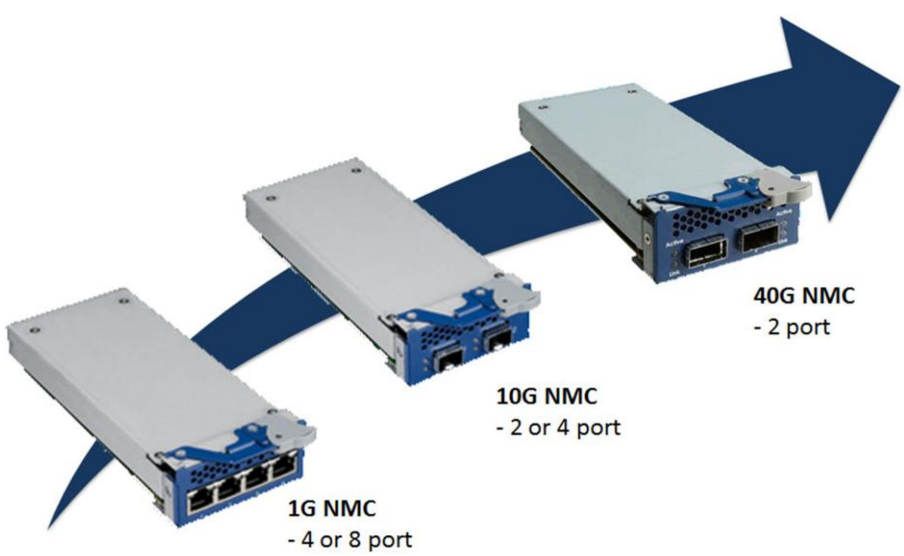

Source: Advantech

NMCs are available with four- or eight-port GE interfaces, two- or four-port 10GE interfaces and two-port 40GE interfaces (**Figure 4**). Security processing and compression on the platform can be accelerated by using up to four PCIe adapters with two Intel Communications Chipset 8955 devices. Each adapter will support IPsec up to 86 Gbit/s, SSL up to 100 Gbit/s and up to 48 Gbit/s compression. With four adapters, the system will support security processing up to 400 Gbit/s.

Accelerated Software

The 6WIND NFV Infrastructure (NFVI) software provides a high-performance open networking platform for COTS-based systems that supports NFV. The NFVI uses the 6WINDGate packet processing software with 6WIND's Accelerated Open vSwitch (OVS) and Data Plane Development Kit (DPDK). The platform supports VNFs running on multiple OSs and delivers over 10x the performance of a pure Linux implementation, transparently without any change to the OS, OVS, hypervisor or management.

Figure 5 shows the 6WIND NFVI software running on the Advantech FWA-6512 Network Application Platform. 6WIND's Accelerated OVS software uses DPDK to bypass the Linux kernel stack and provides its own stack with a fast-path architecture, including additional Layer 2-4 protocols, such as VLAN, VXLAN, LAG, GRE, Filtering, NAT, IPsec and more, in addition to basic switching features (using OVS or Ethernet bridging).

Figure 5: COTS Hardware With Accelerated Software for Virtual Infrastructure

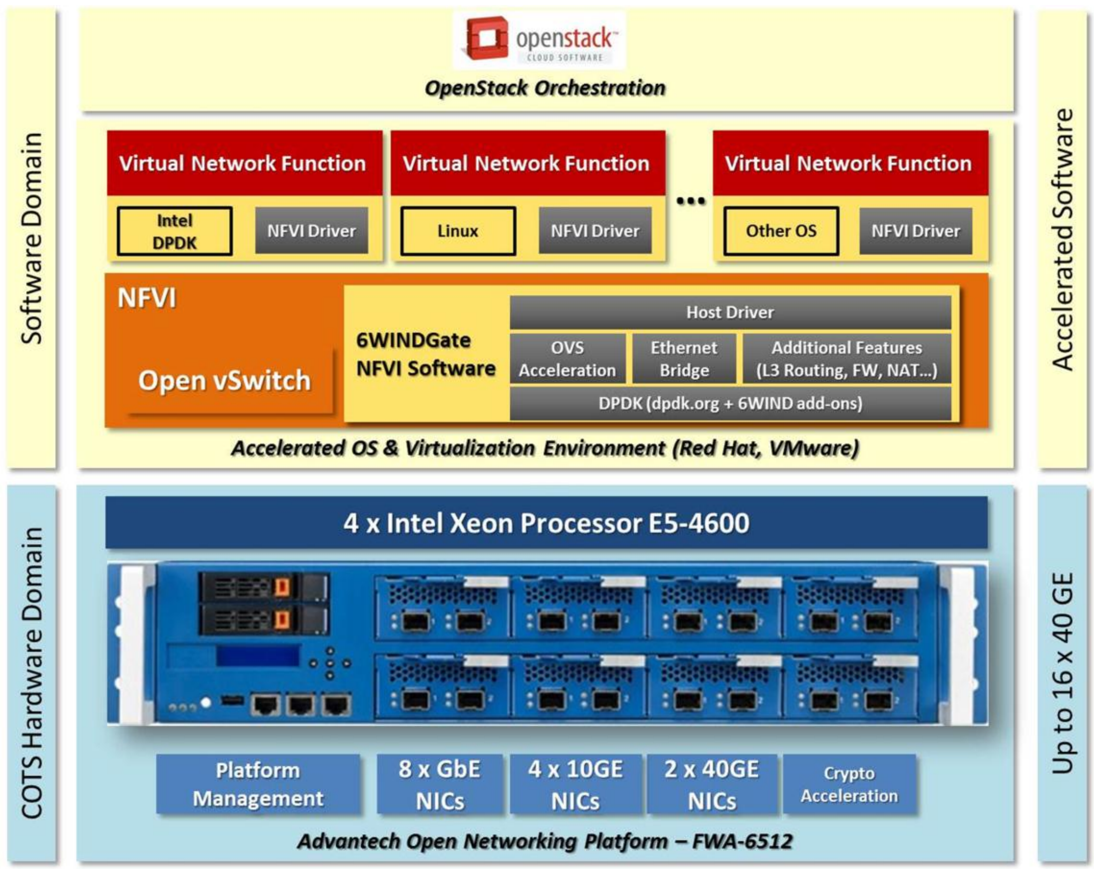

Source: 6WIND, Advantech

The 6WIND software works with standard orchestration solutions that use OpenStack for NFV and SDN implementations. It has also been demonstrated with OpenFlow controllers and other SDN software. By delivering performance without impact on the environment, this solution provides a smooth, incremental path to new architectures.

Advanced Server Acceleration

The Advantech FWA-6512 described above is a 2U server platform with four processor sockets and up to 48 x86 cores for dense compute and networking applications. The internal PCIe slots and flexible NMC modules support a range of network interfaces from GE to 40GE with optional acceleration for security and compression. The performance of the network appliance can be further enhanced by the use of NICs based on high-performance network processors with network interfaces supporting up to 100 Gbit/s.

Figure 6 shows the Advantech FWA-6512C Network Application Platform. The FWA-6512C incorporates a 100G FlowNIC board in place of the NMC slots. The FlowNIC features a Netronome NFP-6xxx network processor that can be used to offload packet processing functions, including packet classification and network overlay processing. Unlike conventional networking interfaces that use the x86 QPI to route traffic between cores, the FlowNIC is connected to the Intel Xeon processors using four separate PCIe x8 Gen 3 connections, which enables network traffic to be load balanced to the destination core(s).

Figure 6: Advantech FWA-6512C Network Application Platform

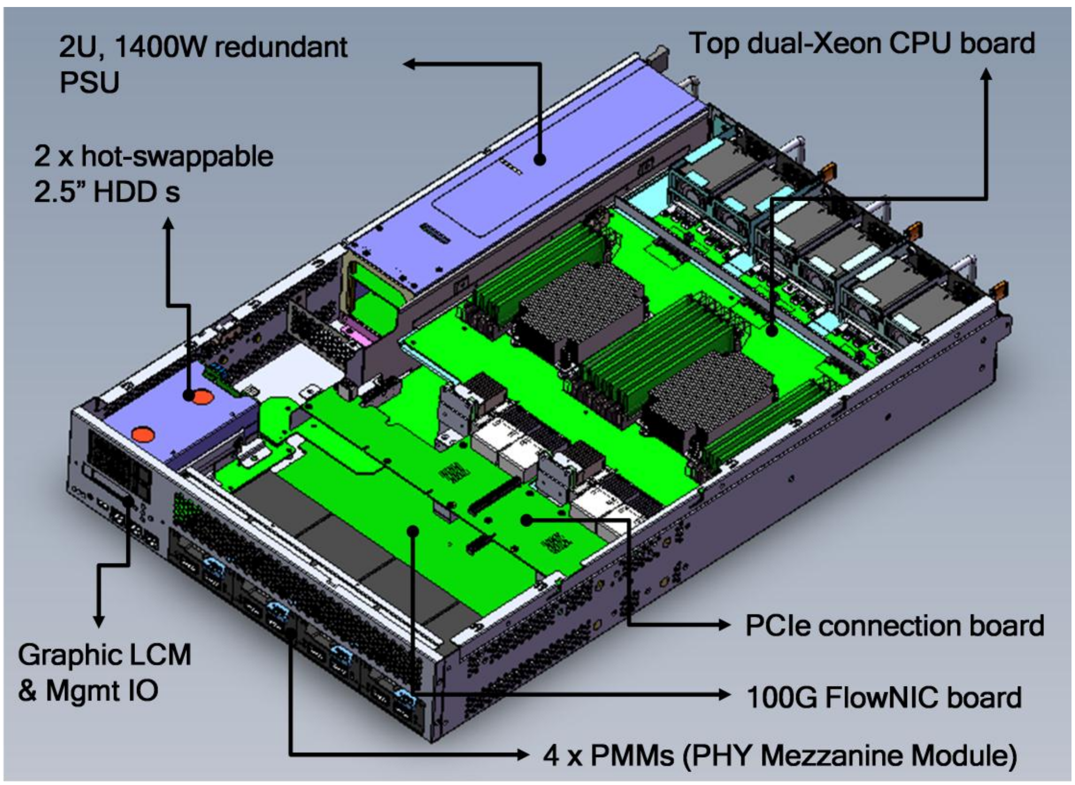

Source: Advantech

The FlowNIC board integrates support for four PHY Mezzanine Modules (PMM). Each PMM has 10 SerDes lanes connected to the FlowNIC running at 10 Gbit/s or 25 Gbit/s. The initial PMM's will support a single 100GE port, two 40GE ports or eight 10GE ports. The FlowNIC is treated by the system software as a standard intelligent NIC with DPDK driver and control application programming interface (API) for flow management. There is no requirement for custom network processor software.

Applications running in a virtualized environment use virtual ports and virtual switches, which are implemented in software. This software switching can be offloaded to the FlowNIC, freeing up x86 processor cycles for running applications. The implantation can support large flow tables with tens of millions of entries.

Figure 7 shows the mapping of virtual port queues from the virtual switch, implemented on the Netronome NFP-6xxx network processor, to the x86 application queues. Traffic from the physical ports is forwarded through the virtual switch to the applications running on the server system.

Figure 7: Virtual Port-to-Queue Mapping

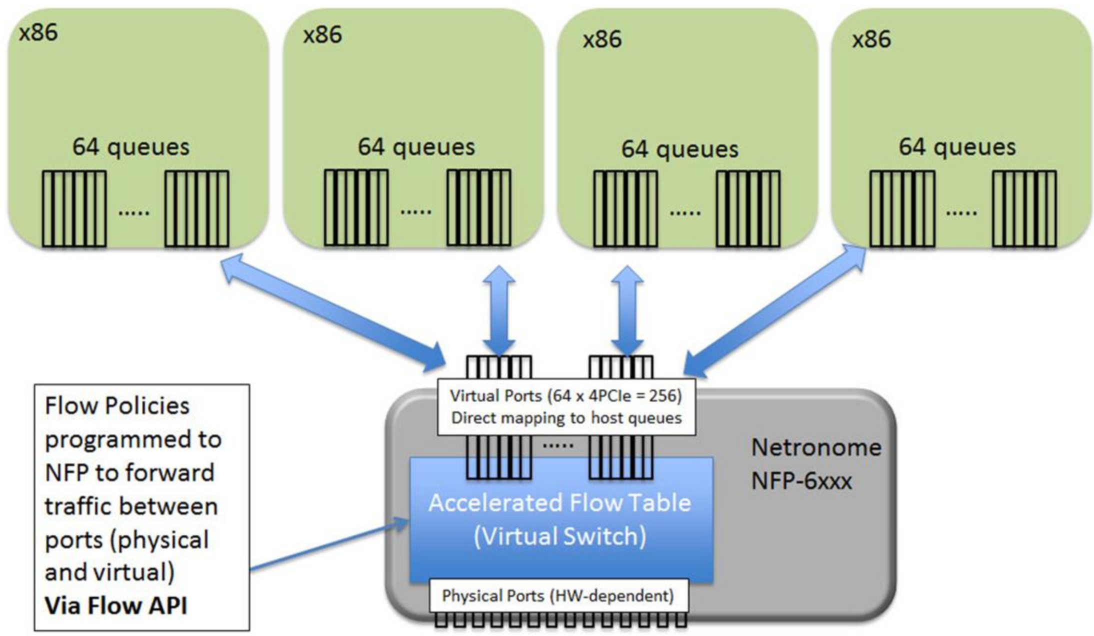

Source: Advantech

The Advantech FWA-6512C network application platform with FlowNIC cards is a highly-integrated appliance that allows SDN and NFV applications to manage virtual switching, tunneling and intelligent load balancing at performance levels that previously required specialist hardware.

Conclusions

Networks are shifting to virtual infrastructure so that service providers can deliver flexible services and meet dynamic customer needs. Network bandwidths continue to grow, driven by mobile data on smartphones and tablets, as well as by cloud services. Network connections that have already shifted from GE to 10GE will be moving on to 40GE and 100GE in the foreseeable future.

SDN and NFV provide complementary approaches to delivering on this virtual infrastructure. The challenge facing service providers is sourcing the right platforms and software to support flexible services on a virtual infrastructure with line-rate performance on the underlying hardware.

Network appliances based on a 2U rack-mount chassis can deliver the right performance for the virtual infrastructure when using the latest multicore processors and flexible high-speed network interfaces. Quad processor implementations with up to 48 cores and many PCIe interfaces provide a compelling balance of performance, power consumption and I/O. The option to add high-speed network interfaces up to 100 Gbit/s and network processor-based acceleration for load balancing and virtual switches dramatically increases the performance of the network appliance in SDN and NFV applications.

This white paper covers a flexible solution for SDN and NFV based on standard server technology with multiple I/O and hardware acceleration options. Outstanding performance is achieved with a combination of standard multicore processors, high-speed network interfaces and software that achieves line-rate performance in virtualized environments.