Bending the Cost Curve for Cryptography and Compression Acceleration on OEM Servers

Cost-effective common building blocks facilitate the move to virtual network functions.

Today's service providers want to control their own destiny in providing products and services to their customers. This includes not being tied to proprietary equipment, and more importantly, being able to lower CapEx and OpEx, improve time-to-market and deliver leading-edge performance. Voicing these desires, thirteen of the world's largest service providers wrote a joint white paper[1] encouraging the development of interoperable solutions based on industry-standard servers. They proposed an approach called network functions virtualization (NFV), which is applicable to any data plane packet processing and control plane function in fixed and mobile network infrastructures.

NFV, and a complementary initiative called software-defined networking (SDN), decouple software workloads from a particular hardware platform, allowing them to be controlled centrally and deployed dynamically throughout the network as needed. The resulting network architecture facilitates the consolidation of various network equipment types onto high-volume OEM servers, switches and storage, thereby reducing equipment footprint, complexity, power consumption and cost. This transformation should enable service providers to save money and increase revenue while maintaining a positive user experience.

Along these lines, 6WIND, Dell and Intel are enabling service providers to deploy applications on OEM Servers with high-performance security and compression acceleration at much lower cost than comparable add-in acceleration cards. This solution provides common building blocks that can be used by telecom equipment manufacturers (TEMs), network equipment providers (NEPs), service providers and appliance makers moving to virtual functions. This paper discusses the relevant technologies and how they are supported by various software options.

Transforming the Network with COTS, SDN and NFV

The cloud, mobility, social media and analytics are increasing data traffic and driving tremendous change in the way networks are being used. This isn't necessarily good news for network operators who have seen profitability slide as revenue growth significantly lags traffic growth (Figure 1), a trend that is changing the rules for future infrastructure investments. Cost must be driven down to increase profitability, thus network operators feel the need to achieve costs comparable to the IT industry that benefits from large economies of scale and innovation curves. At the same time, flexible and agile infrastructure is required to address the competitive landscape and quickly changing needs of the market.

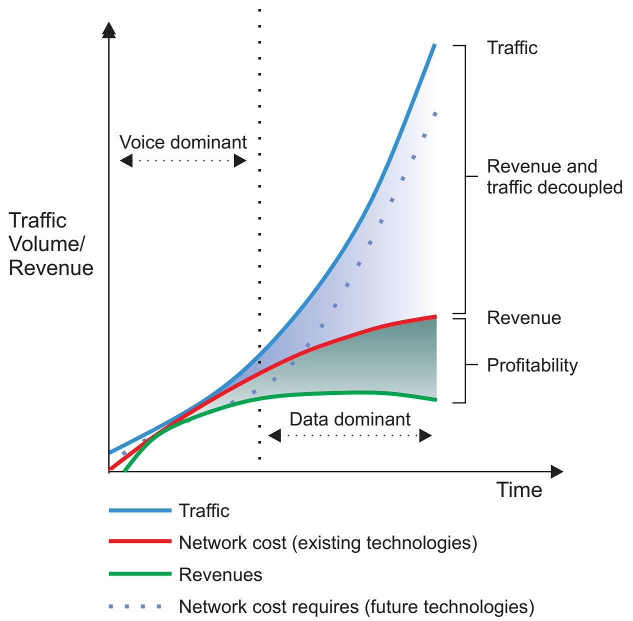

Figure 1: Revenue Growth Trails Traffic Growth

The telecommunications industry is at an inflection point, facing a future based on software-built networks with the potential to reduce costs and increase innovation. Service providers around the world see the potential for SDN and NFV to rein in both capital and operational expenditures, as well as enable new revenue-generating services, such as content storage and distribution for enterprises, and geo-fenced, audience analytics based retail advertising content distribution. Virtualized network equipment based on commercial, off-the-shelf (COTS) servers, instead of numerous purpose-built boxes, can reduce CapEx through lower hardware costs and network element consolidation. OpEx savings are attainable through increased efficiency in managing and reconfiguring network functions, a homogeneous tool set and lower overall energy consumption delivered by power management features on standard servers. Service providers can also bring new capabilities to market faster by enlisting an ecosystem of independent software vendors (ISVs), open-source developers and academia on the cutting edge of virtual appliances. Moreover, service providers will have greater network agility since SDN and NFV facilitate the rapid scaling, up or down, of applications and bandwidth.

COTS Technology Is Ready

It's now possible to transition from purpose-built network elements to virtual network functions running on COTS servers, in large part due to exceptional CPU performance increases predicted by Moore's Law as well as innovative software that solves virtual networking performance challenges. In particular, x86-based platforms have advanced steadily, with the ability to deliver high-speed switching, improved network intelligence via deep packet inspection (DPI), virtualization, data center convergence and high availability.

Solution Overview

6WIND, Dell and Intel combined their manufacturing, hardware and software expertise to develop a high-performance platform that runs virtualized network elements, thus enabling service providers to more easily realize the benefits from COTS, SDN and NFV. The hardware platform is the award-winning[2] Dell™ PowerEdge™ R720 Rack Server built with the Intel® Xeon® processor E5-2600 series and up to 24 dual in-line memory modules (DIMMs). The COTS server also supports the Intel® QuickAssist Server Adapter 8920 Family, which performs crypto and compression acceleration. Packet processing performance is greatly improved by the 6WINDGate™ networking software, which includes the Intel® Data Plane Development Kit (Intel® DPDK) library. In all, the solution can deliver up to 40 Gbps of IPSec throughput and 16 Gbps of compression.[3,4,5]

OEM Server

The Dell PowerEdge R720 Rack Server, pictured in Figure 2, features the latest computing and networking technologies, and is part of a large Dell portfolio available to help equipment manufacturers get to market faster. There are several models: standard (Dell branded), unbranded, NEBS and ETSI-compliant, and front-facing PCI Express* ports to accommodate network appliance configurations. All versions are designed for easy lifecycle manageability with intelligent hardware-driven systems administration, extensive power management and other innovative management tools.

Dell can take a developed solution and provide the hardware, manufacturing, integration, customization, delivery, support and even financial services necessary to deliver solutions to the market. Their technical staff is there to assist in getting everything working.

Figure 2. Dell™ PowerEdge™ R720 Rack Server

Supporting Network Transformation

In support of the network transformation enabled by SDN and NFV, Intel envisions a three stage transition to virtualized network functions as shown in Figure 3:

- **4:1 Workload Consolidation:** Consolidate all network workloads (application, control plane, packet and signal processing) onto a single platform based on Intel® architecture processors.
- **Virtualize the network workloads:** Use optimized software switches for VM to VM and VM to network data paths.
- **Architect the software platform:** Use the appropriate standard, open protocols and APIs from the SDN and NFV "Industry" to enable programmability, automation and cost effective software based network elements. Run these software network elements on standard, high-volume servers.

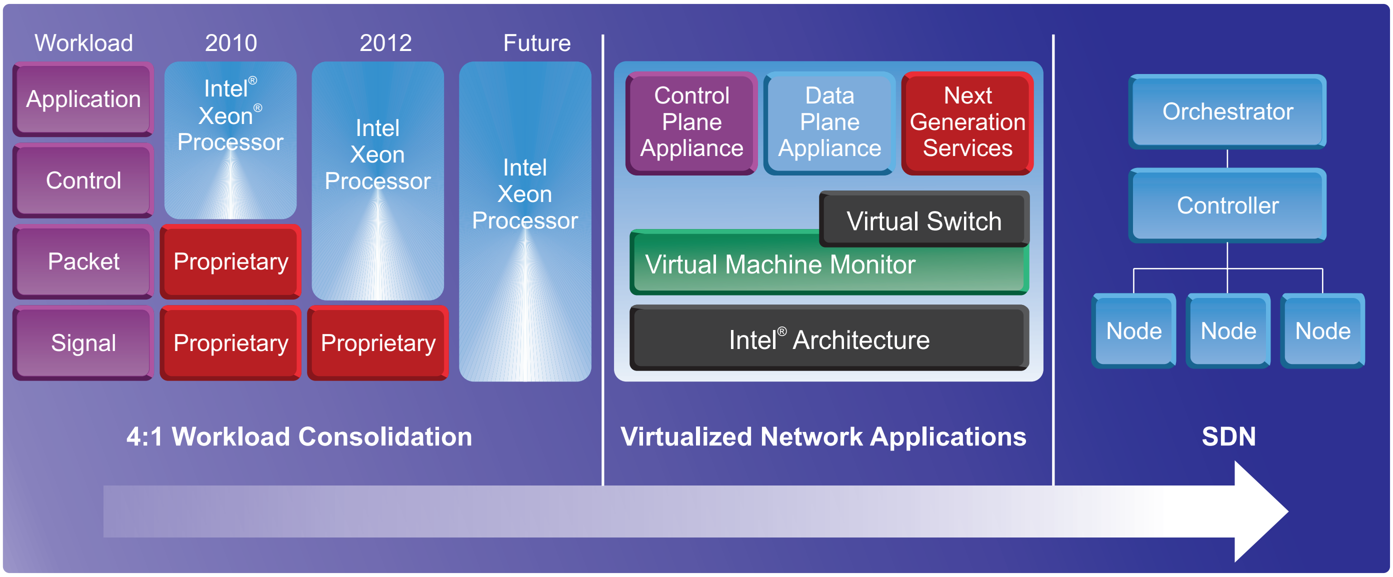

Figure 3: How Intel® Architecture Is Supporting Network Transformation

Virtual Switch

Equipment manufacturers designing virtual switches for deployment in the SDN node layer can take advantage of an Intel reference design, called the Intel® Open Network Platform Server Reference Design (Intel® ONP Server Reference Design). This reference design, diagrammed in Figure 4, runs on nearly any Intel® Xeon® or Intel® Core™ processor-based hardware platform. The KVM hypervisor[6] and Intel® Virtualization Technology (Intel® VT)[7] provide a flexible, high performance and robust virtualization environment. In the future, additional hypervisor options will be available. A reference implementation of a high performance version of Open vSwitch, accelerated by 6WINDGate and the Intel DPDK, is included. In addition, optimizations will be provided to facilitate remote management and integration into the orchestration infrastructure. For some workloads, the use of PCI-SIG Single Root I/O Virtualization (SR-IOV) could be used to provide direct access to virtual appliances.

Fundamental to the high performance of the Intel reference design is the accelerated packet forwarding enabled by the Intel DPDK. It also allows the consolidation of data and control planes on a general-purpose processor. Thanks to the architecture of 6WINDGate, the packet processing performance of the Dell PowerEdge R720 Rack Server scales linearly with the number of processor cores configured to run the 6WINDGate fast path. Pre-integrated into 6WINDGate, the Intel DPDK provides Intel architecture-optimized libraries to accelerate L3 forwarding and other packet processing functions, compared to native Linux.

The Intel DPDK contains a growing number of libraries, whose source code is available for developers to use and/or modify in a production network element. Likewise, there are various use case examples, like L3 forwarding, load balancing and timers, that help reduce development time. The libraries can be used to build applications based on "run-to-completion" and/or "pipeline" models, enabling the equipment provider's application to maintain complete control.

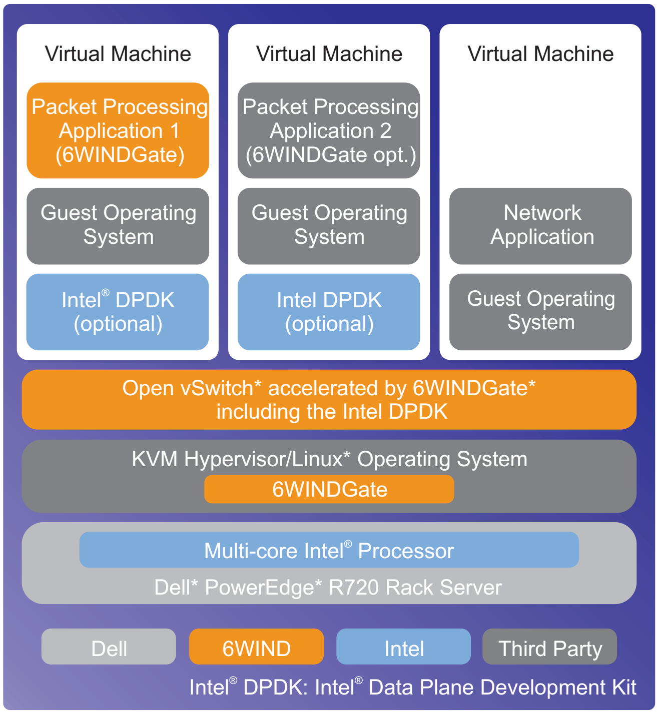

Figure 4: Virtual Switch based on the Intel® Open Network Platform Server Reference Design (Intel® ONP Server Reference Design)

Accelerated Crypto and Compression

For applications requiring crypt or compression acceleration, the Intel QuickAssist Adapter 8920 Family can be added to the Dell PowerEdge R720 Rack Server via the x8 PCI Express Gen 3 interface. The card takes advantage of software and hardware modules incorporated in the Intel processor and chipset that accelerate encryption and data compression workloads. The acceleration features are accessed via a unified set of industry-standard application programming interfaces (APIs), which provides consistent conventions and semantics across multiple accelerator implementations and future software investments. Future communications chipsets will support these APIs.

There are two versions of server acceleration cards supporting 20 and 40 Gbps of IPSec, and 8 and 16 Gbps compression, as shown in Table 1.[3,4,5]

SKU Option	Performance Bulk Crypto	PKE (Public Key Inspection)	Compression Deflate, Stateless, Dynamic	Max Total Throughput (when services used simultaneously)	Power
Intel® QuickAssist Server Adapter 8920-SCC	20 Gbps	28K Ops/s 1K Keys 5.5K Ops/s 2K Keys	8 Gbps	20 Gbps	25W
Intel® QuickAssist Server Adapter 8920-DCCP	40 Gbps	56K Ops/s 1K Keys 11K Ops/s 2K Keys	16 Gbps	40 Gbps	40W

Table 1: Two Versions of the Intel® QuickAssist Adapter 8920 Family

Overheads and Latencies in Linux

Widely used in network appliances, standard operating systems and hypervisors present performance challenges for networking and communications applications. For instance, the Linux standard networking stack uses services provided by the kernel, and consequently may experience significant overheads associated with functions such as preemptions, threads, timers and locking. These processing overheads are imposed on each packet passing through the system, resulting in a major performance penalty on the overall throughput.

Another challenge is getting packet processing applications to scale linearly as the number of cores increase. This is because the overhead from some processor operations, like for memory coherency and interrupts, tends to compound as more cores are added. For example, a processor with eight cores may not process packets significantly faster than one with two cores. In addition, a standard OS stack may not exploit the data plane performance potential of a multi-core platform and thus fail to deliver the networking performance required for cost-effective SDN architectures.

Three Significant Challenges

The transition to SDN and NFV will drive network equipment to incorporate virtualization technology, which could introduce additional I/O and communication bottlenecks. The following describes three potential issues (Figure 5) and solutions developed by 6WIND:

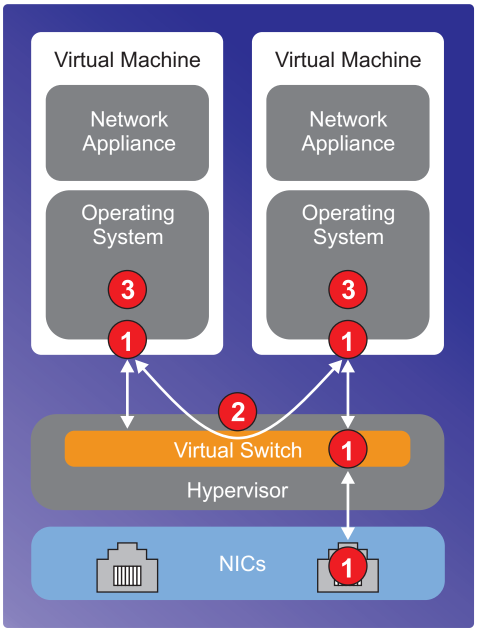

Figure 5: Potential Challenges for a Virtual Network Appliance

***Issue 1*:** The overall throughput between the physical network interface card NIC and the virtual appliance (aka guest) is constrained by performance limitations within the host and guest networking stacks as well as in the virtual switch. The problem is amplified when un-optimized NIC drivers are used.

> ***Solution*:** Bypass the virtual switch within the hypervisor using I/O virtualization, thus removing the I/O performance constraints imposed by the virtual switch. Also, optimize physical and virtual NIC drivers to maximize I/O performance.

***Issue 2*:** VM-to-VM bandwidth and scalability is impacted by bottlenecks in the virtual switch. This problem is especially important for data center applications with high "East-West" (VM-to-VM) traffic as a result of VM sprawl and multi-tenant environments.

> ***Solution*:** Implement a virtual NIC (vNIC) driver that leverages communication between VMs via the virtual switch, enabling the efficient development and provisioning of systems with multiple VMs and significant East-West network traffic. Also, use a VM-to-VM (VM2VM) driver to provide direct VM-to-VM communication, bypassing the virtual switch while remaining fully compatible with industry-standard hypervisors.

***Issue 3*:** A standard OS stack running in a VM (or in a physical appliance) typically limits packet processing performance due to poor scalability over multiple cores, and significant overhead and interrupt latency.

> ***Solution*:** Use high-performance software solution that provides networking protocols optimized for maximum performance on multi-core platforms while also delivering linear performance scalability across multiple cores.

Optimized Architecture for Packet Processing

The 6WINDGate™ packet processing software solves the issues previously discussed, as well as many others. It is based on a fast path-based architecture and incorporates a comprehensive set of high-performance networking protocols fully optimized for Intel Xeon processor-based platforms.

6WINDGate processes the vast majority of packets in a fast path environment, which executes outside the OS kernel in Linux userspace. By avoiding typical OS overheads associated with functions, such as preemptions, threads, timers and locking, this architecture maximizes data plane processing performance. Only those rare packets that require complex processing are forwarded to the OS networking stack, which performs the necessary management, signaling and control functions.

Most of the processor cores can be dedicated to running the fast path in order to maximize the overall throughput of the system, while at minimum, only a single core is required to run the OS, the OS networking stack and the application's control plane. The fast path cores run in a Linux userspace environment, so the system can be reconfigured dynamically as traffic patterns change in order to optimally allocate CPU resources to the control plane and the fast path.

Figure 6 shows that 6WINDGate delivers approximately ten times greater IP forwarding throughput than standard Linux in a physical appliance. Also, 6WINDGate with IOMMU enables cost-effective NFV implementations by delivering comparable performance in a virtual appliance.

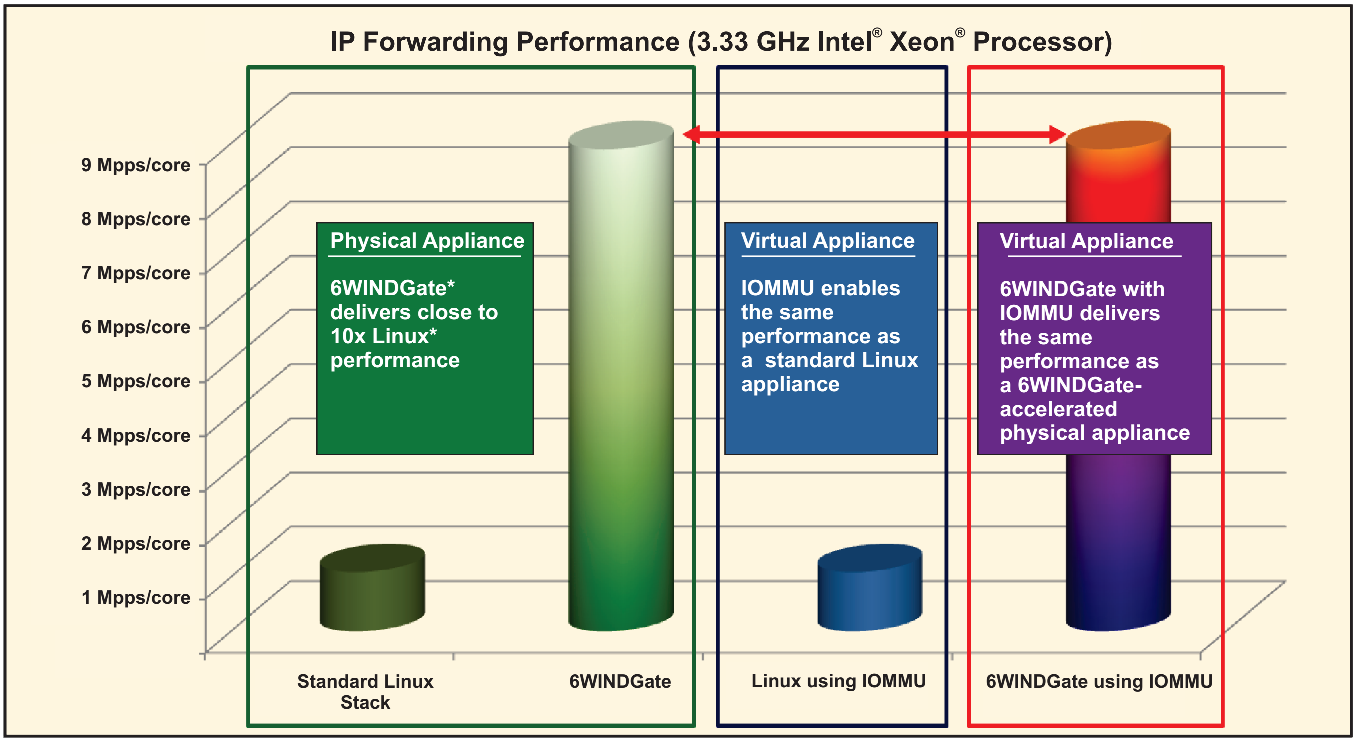

Figure 6: IP Forwarding Performance

Use Case: Firewall

A common application for the Dell PowerEdge R720 Rack Server running 6WINDGate within an SDN ecosystem is in network equipment (gateways, firewalls, routers etc), whether located in mobile infrastructure, in enterprise IT environments or in data centers. Figure 7 illustrates a virtual firewall or intrusion prevention system (IPS). The comprehensive set of networking protocols delivers approximately 10 Mpps per core of IP forwarding performance on a dual Intel Xeon processor E5-2600 series platform running at 2.7 GHz.

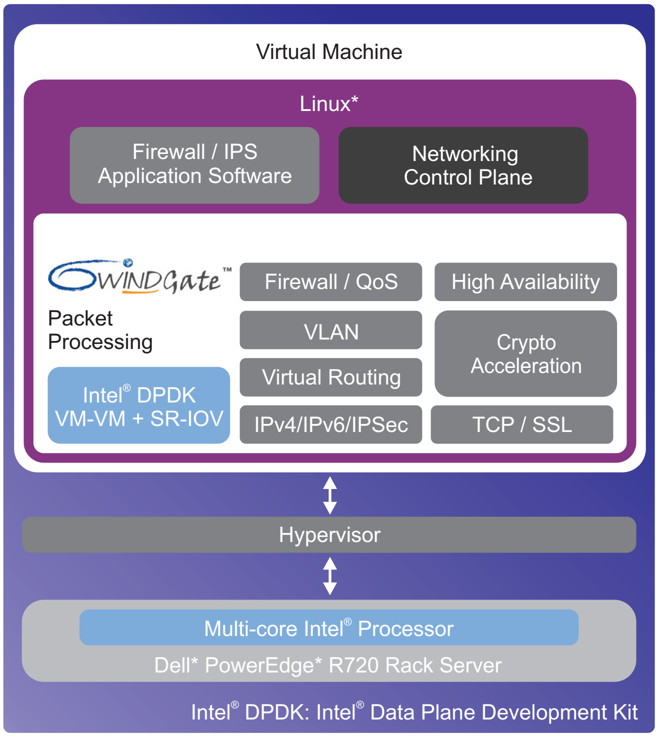

Figure 7: Firewall Example

As mentioned above, full support is provided for industry-standard hypervisors, and the architecture is fully scalable across processors, boards and racks. 6WINDGate is fully-compatible with OpenFlow and OpenStack, enabling its use in emerging enterprise and data center applications.

Use Case: Cloud and Telecom Data Centers

A second SDN-oriented application is in data center networking. 6WINDGate solves network performance problems on the application server blades, such as when East-West traffic requires VM-to-VM communication at a level of performance beyond the capabilities of existing protocols. At the same time, as processor performance increases so does the number of VMs instantiated on each server blade. For multi-tenant environments, this drives a need for an accelerated virtual switch that also provides high-performance tunneling protocols for VM-to-VM communication, as shown in Figure 8.

A virtual switch based on the Dell PowerEdge R720 Rack Server and accelerated by 6WINDGate, retains full compatibility with the standard Open Virtual Switch (OVS) and with OpenFlow, as depicted in Figure 8. It is completely transparent to the applications running on the platform, which do not need to be recompiled or re-verified in order to work with this high-performance solution. The solution provides a data plane solution that delivers 5 to 10 times acceleration (depending on the number of fast path cores) for the baseline Layer 2 switching function. At the same time, it supports the necessary secure tunneling protocols such as IPsec, GRE, NVGRE, VLAN and VxLAN.

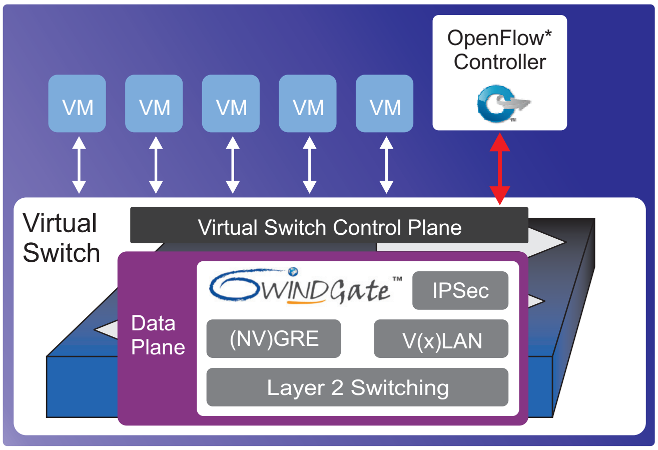

Figure 8: Cloud and Telecom Data Center Example

Summary

COTS, SDN and NFV have the potential to pave the way to significant benefits for enterprises and service providers, including lower CapEx/OpEx, faster time-to-market and easier deployment of new features and services. The ensuing network transformation is expected to enable flexible, centrally-managed networks based on economical, industry-standard OEM servers, switches and storage.

A prime example is the high-performance Dell PowerEdge R720 Rack Server based on the Intel Xeon processor and accelerated by 6WINDGate, in all providing exceptional packet processing performance. Moreover, the Intel QuickAssist Server Adapter 8920 Family can be added to servers for cost-effective crypto and compression acceleration. The solution solves critical networking performance challenges, such as overhead in OS stacks and virtual switches, that must be addressed in order for SDN and NFV architectures to be cost-effective.

[1] "Network Functions Virtualisation – Introductory White Paper," published at the October 22-24, 2012 at the "SDN and OpenFlow World Congress", Darmstadt-Germany, http://portal.etsi.org/portal/server.pt/community/NFV/367.

[2] CRN 2012 Products of the Year, http://www.crn.com/slide-shows/components-peripherals/240144534/the-2012-products-of-the-year.htm?pgno=3.

[3] The maximum total throughput when crypto and compression services are used simultaneously is 40 Gbps.

[4] Performance estimates are based on internal Intel analysis and are provided for informational purposes only.

[5] Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit http://www.intel.com/performance/resources/limits.htm.

[6] See www.linux-kvm.org for more information.

[7] Intel® Virtualization Technology (Intel® VT) requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM), and for some uses, certain platform software enabled for it. Functionality, performance, or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

Printed in USA 0913/MS/TM/PDF ♻ Please Recycle 329129-002US