WHITE PAPER

Two Network Functions Virtualization Infrastructure Platforms Demonstrate Predictable Performance and Continuous Improvement

Intel and HPE* collaborate to simplify the evaluation of Network Functions Virtualization Infrastructure (NFVI) for service providers

Hewlett Packard Enterprise

Table of Contents

Overview

Founded in September 2014, the Open Platform for Network Functions Virtualization* (OPNFV*) is an open source community project whose objective is to develop a carrier-grade, integrated, open source reference platform intended to accelerate the introduction of new products and services.[1] The OPNFV community currently has over 60 members, including many of the world's leading network equipment providers, service providers, server OEMs, and OS vendors.

In support of this effort, Hewlett Packard Enterprise* (HPE*) and Intel are designing open platforms for Network Functions Virtualization Infrastructure (NFVI), and measuring and correlating the performance of these platforms. This white paper provides the results of these activities, along with key performance indicators (KPIs) that can be used to demonstrate predictable performance and continuous improvement of packet processing performance, and guidance on how to maximize the performance of industry-standard, high-volume servers for NFV.

This information is intended to help network equipment developers and service providers better understand the repeatable and achievable data plane performance on similar NFVI, using the principles laid out by the OPNFV community.

Driving Network Transformation

Network virtualization and cloud-enabled technologies are steadily reshaping network architectures of telecommunication providers, creating strong incentives to demonstrate the viability of these technologies for commercial use. Even as the first commercial deployments occur, additional technical work still needs to be done to help communications service providers (CoSPs) deploy new solutions faster. Through an open ecosystem and open standards approach, stakeholders are collaborating to improve interoperability at each layer of NFVI. One of the primary objectives in this collaborative work is to measure performance in a consistent and predictable way to help CoSPs evaluate platform alternatives. The current and completed work of the Benchmarking Methodology Working Group in the IETF is working to build a test framework, called vsperf, in the OPNFV community and aims to have a consistent measurement framework for NFVI.[2]

Another objective is to minimize the complexity in planning and implementing the infrastructure deployed in support of Software-Defined Networking (SDN) and NFV implementations. To these ends, Intel and HPE are collaborating to simplify test integration and performance measurements on an HPE commercial platform based on the Intel® Open Network Platform (Intel® ONP) reference architecture.

A SERVER REFERENCE ARCHITECTURE OPTIMIZED FOR SDN/NFV

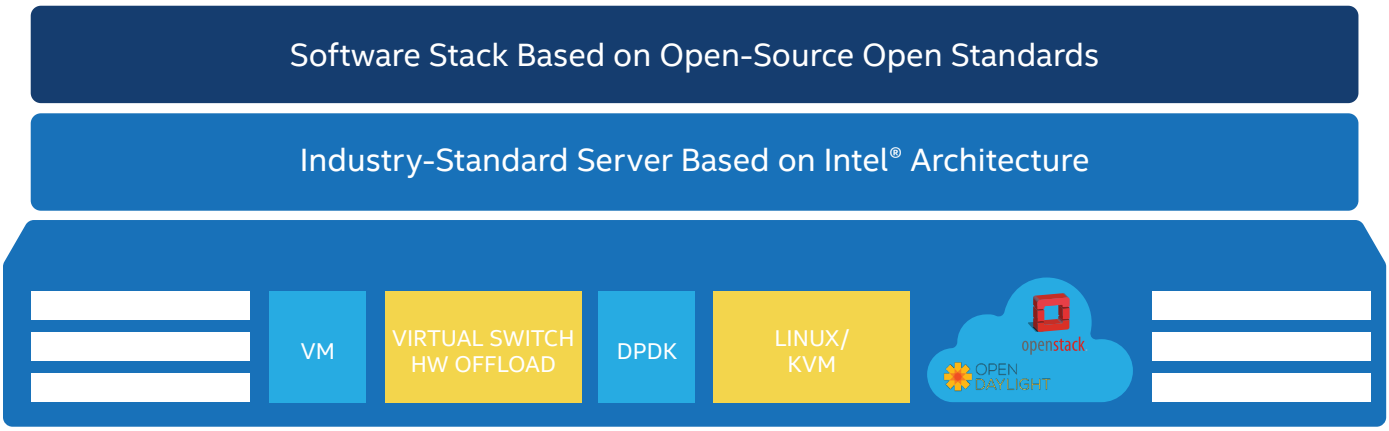

Figure 1. The Intel® Open Network Platform reference architecture enables efficient development of software-defined networking (SDN) and network functions virtualization (NFV) solutions.

First released in September 2014, the Intel Open Network Platform defines a common infrastructure and test framework to allow comparison of different NFVI platforms. The latest Intel ONP 2.1 release (March, 2016) is enhanced to incorporate the most recent server hardware based on the Intel® Xeon® processor E5-2600 v4 product family.

What is the Intel® Open Network Platform Reference Architecture?

To streamline the evaluation, design, and deployment of open SDN and NFV solutions, the Intel ONP reference architecture provides a blueprint for accelerating the development of commercial hardware and software platforms (Figure 1), and redefining network architectures around the principles of virtualization.

This reference architecture decouples network functions from the underlying hardware components, establishing an infrastructure for NFV, and providing a model that enables centralized provisioning and management of network resources. With Intel ONP, OPNFV solution developers can more easily build solutions using an open-source software stack running on industry-standard, high-volume servers. The reference architecture gives solution providers a way to plan, evaluate, and benchmark components in advance of full NFVI deployments.

HPE has advanced the architecture further with the HPE OpenNFV Reference architecture, consisting of high performance, carrier-grade NFV, and end-to-end service orchestration products and solutions, as shown in Figure 2. They are supported by a rich ecosystem of VNF vendors, including the traditional NEPs and innovative ISVs that build

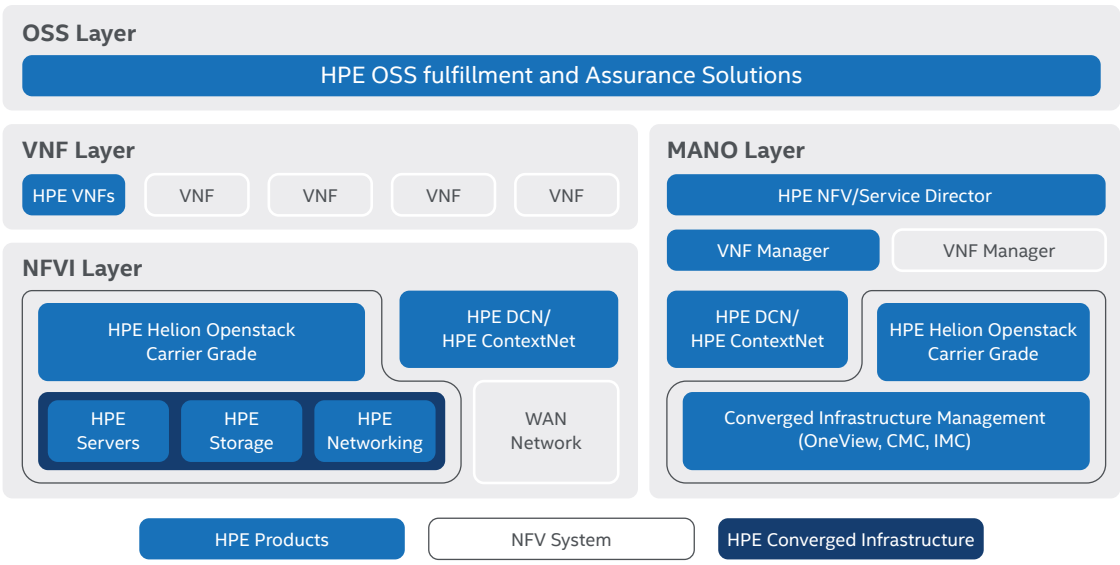

Figure 2. HPE* OpenNFV* reference architecture

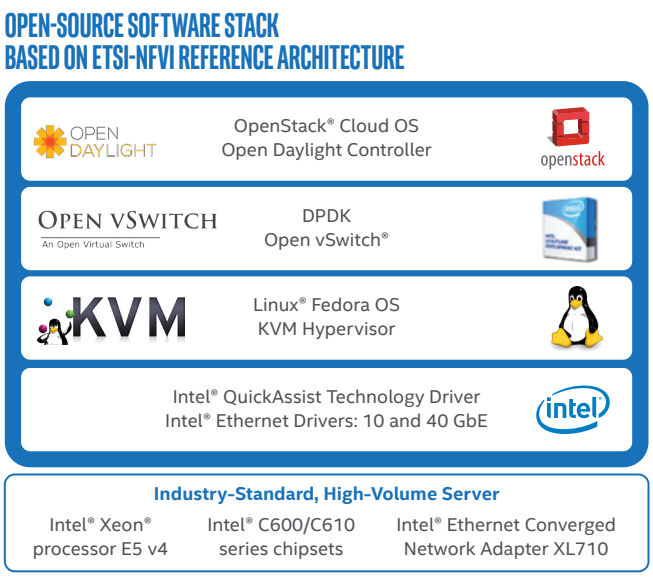

Figure 3. Software stack and hardware platform for the Intel® Open Network Platform reference architecture, release 2.1.

multi-vendor solutions, and a network of global labs to test and validate these solutions for a faster deployment.

Upstream contributions provided by HPE, Intel, and others into NFV and SDN open-source projects drive the technology advances in this sector forward. The open-source software stack runs on top of the hardware platform specified by Intel ONP, as shown in Figure 3, based on standard, high-volume servers (SHVSs) powered by the Intel Xeon processor E5-2600 v4 product family.

The value of Intel ONP lies in the demonstrated, wide-scale interoperability of proven components in support of commercial implementations. Through tested reference architecture ingredients and industry standardization work, NFV developers have a solid infrastructure and the compatible building blocks they need to confidently work on solutions that further advance the development of SDN.

Both Intel and HPE have made substantial contributions along these lines. HPE is the largest single contributor to the OpenStack* Mitaka release, and both companies are helping to drive standardization and SDN initiatives through participation and leadership in OpenDaylight*, Open vSwitch*, and other open community projects.

By integrating the latest SDN/NFV technology into a standardized architectural framework, HPE has engineered a flexible, proven NFVI to support NFV developers and their customers. Intel ONP serves as the baseline platform for developing HPE OpenNFV, providing a collection of NFV components that can be rapidly assembled to construct complete NFVI solutions. HPE's comprehensive approach to the SDN/NFV transformation into a standardized architectural framework provides support to NFV developers and tangible business outcomes to their CoSP customers.

Test Cases Considered

As proposed by the IETF and adopted by the vsperf project in OPNFV, there are a number of test cases used to measure virtual switching and virtual machine performance. The vSwitch performance (PHY-OVS-PHY) tests measure the Layer 3 forwarding throughput across four ports at various packet sizes, ranging from 64 to 1518 bytes, as shown in Table 1. In this white paper, this first test case is measured for Key Performance Indicators (KPIs) #1 and #2. For KPIs #3 and #4, virtual machine (VM) performance (PHY-VM-PHY) is measured via Layer 3 forwarding performance testing across four ports with packet sizes ranging from 64 to 1518 bytes. Table 1 provides other test case details.

Table 2 lists the major hardware and software components of the HPE and Intel OPNFV platforms. There are some hardware and software variances, but they did not result in significantly different performance. The key platform differences are the amount of memory (64 GB vs 256 GB) and the operating systems; the Intel platform runs Fedora* and the HPE platform is built on a commercial off-the-shelf Linux* release Red Hat Enterprise Linux (RHEL) 7.2. The processor is the same in both cases.

Test Description	Metrics	Packet Size (Bytes)	Test Duration	Flow per Port
vSwitch Performance (PHY-OVS-PHY)				
L3 Forwarding 4 Ports	Throughput Latency (min, max, avg)	64, 72, 128, 256, 512, 768, 1024, 1280, 1518	60 sec	One flow/port in both directions; One thousand flows/port in both directions
One VM Throughput (PHY-VM-PHY)				
Single VM (vhost-user) L3 Forwarding 4 Ports	Throughput Latency (avg)	64, 256	60 sec	One flow/port in both directions; One thousand flows/port in both directions

Table 1. Summary of Test Cases

Item	NFVI Developed and Tested by Intel	NFVI Developed and Tested by HPE
Server Platform	Supermicro* X10DRH-I http://www.supermicro.com/products/motherboard/Xeon/C600/X10DRH-i.cfm Dual Integrated 1GbE ports via Intel® i350-AM2 Gigabit Ethernet	HPE* ProLiant* DL380 Gen9 Server http://www8.hp.com/us/en/products/proliant-servers/product-detail.html?oid=7271241#!tab=features
Chipset	Intel® C612 chipset	Intel® C610 Series Chipset
Processor	Dual Intel® Xeon® processor E5-2695 v4 (formerly Broadwell-EP) 2.10 GHz; 120 W; 45 MB cache per processor 18 cores, 36 hyper-threaded cores per processor	Dual Intel® Xeon® processor E5-2695 v4 (formerly Broadwell-EP) 2.10 GHz; 120 W; 45 MB cache per processor 18 cores, 36 hyper-threaded cores per processor
Memory	64 GB Total; Samsung* 8 GB 2Rx8 PC4-2400MHz, 8 GB per channel, 8 Channels	256 GB memory (128 GB per processor) HPE 32 GB 2Rx4 PC4-2400T-L Kit
Host Operating System	Fedora* 23 x86_64 (Server version) Kernel version: 4.2.3-300.fc23.x86_64	RHEL* 7.2 (Commercially supported Linux* OS) Kernel Version: 3.10.0-327.el7.x86_64
VM Operating System	Fedora 23 x86_64 (Server version) Kernel version: 4.2.3-300.fc23.x86_64	RHEL 7.2 Kernel Version: 3.10.0-327.el7.x86_64
QEMU-KVM	QEMU-KVM version 2.5.0 libvirt version: 1.2.18.2-2.fc23.x86_64	QEMU-KVM version 2.5.1 1.5.3-105.el7.x86_64
Open vSwitch*	Open vSwitch 2.4.9 Commit ID:53902038abe62c45ff46d7de9dcec30c3d1d861e	Open vSwitch 2.5.0 Commit ID:53902038abe62c45ff46d7de9dcec30c3d1d861e and
DPDK	DPDK version: 2.2.0 http://www.dpdk.org/browse/dpdk/snapshot/dpdk-2.2.0.tar.gz	DPDK version: 2.2.0 http://www.dpdk.org/browse/dpdk/snapshot/dpdk-2.2.0.tar.gz
Local Storage	500 GB HDD Seagate* SATA Barracuda 7200.12 (SN:Z6EM258D)	HP* Smart Array* P440ar/2G FIO Controller 2x HP 600GB 12G SAS 15K 2.5in SC ENT HDD (RAID 1)
PCI Express*	2 x PCI-E 3.0 x8 slot	2x PCIe 3.0 x8 slots
Network Interface Cards	2 x Intel® Ethernet Converged Network Adapter X710-DA4. Total: 8 Ports; 2 ports from each NIC used in tests.	2x Intel Ethernet X710-DA4 FH. Total: 8 Ports; 2 ports from each NIC used in tests.
BIOS	AMIBIOS* Version: 2.0 Release Date: 12/17/2015	HP ProLiant System BIOS P89 v2.00 (12/27/2015)

Table 2. HPE and Intel NFVI Configurations

KPI #1: Predictable NFVI Performance on Industry-Standard, High-Volume Servers

One of the goals of NFV is to have an interoperable infrastructure based on industry-standard, high-volume servers. Also important to advancing NFVI for large-scale, commercial CoSP developments is a reliable, hardened platform based on OPNFV principles. These NFVI platforms must integrate and ensure full interoperability of the latest components from standards bodies and open-source projects, and enable ecosystem vendors to deliver predictable performance.

This KPI for NFVI interoperability is intended to demonstrate predictable packet processing throughput by measuring the maximum packet processing throughput in a virtualized network environment. HPE and Intel have shown how two different platforms could satisfy these requirements by comparing the performance of two different software instantiations based on the same server hardware platform, in this case, based on the Intel Xeon processor E5-2600 v4 product family. As shown in Figures 4, packet processing throughput test results for the two platforms differed by around one percent across the range of packet sizes tested. This testing demonstrates how platforms built by different vendors could produce very similar performance.

These test results confirm a baseline NFVI platform can support line rate speeds approaching 10 Gbps for packet sizes 128 bytes and larger. The data also suggests a Data Plane Development Kit (DPDK)-enabled Open vSwitch platform has the theoretical bandwidth capability to support demanding NFV application environments.

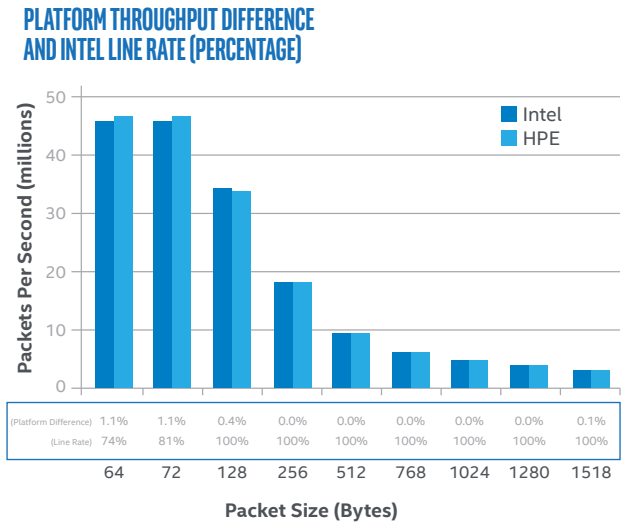

Figure 4. Throughput comparison: HPE* and Intel OPNFV* platforms (see Appendix A)

KPI #2: Continuous Performance Improvement - Generation to Generation

Developers and service providers can expect the performance of NFVI platforms to increase over time through a combination of hardware and software innovation. Hardware-based improvements, delivered at the pace of Moore's Law, will continue to increase the packet throughput of future generations of server platforms. Software enhancements by an active, open-source ecosystem will further increase performance, as seen by the Open vSwitch adoption of the DPDK.

The virtual switching tests attempt to achieve aggregated system throughput of 40 Gbps using four ports to ascertain whether there are any performance gains from successful generations of Intel® Xeon® processors. In 2015, Intel introduced the Intel® Xeon® processor E5-2600 v3 product family, which was based on server processor microarchitecture. Typically used in a data center server configuration, the product family has 14 physical cores per processor socket or 28 virtualized cores with Intel® Hyper-Threading Technology (Intel® HT Technology)[3] enabled.

The next generation of this processor, the Intel Xeon processor E5-2600 v4 product family was introduced in 2016, delivering 18 cores per processor socket, or 36 virtualized cores with Intel HT Technology enabled. The ability to integrate additional cores was made possible via a 14 nanometer process technology for the Intel Xeon processor E5-2600 v4 product family as compared to the 22nm process technology in the previous generation.

The packet processing performance improvement is partly due to the larger last level cache size and the higher main memory controller frequency (2400 MHz versus 2133 MHz) on Intel Xeon processor E5-2600 v4 product family. A three percent higher throughput for small 64 byte packets in a single core configuration is observed, and a scalable gain in throughput performance is seen as the number of cores increases. The throughput performance scales linearly as the number of processor cores dedicated to poll mode driver threads doubles through the use of the Linux FIFO scheduling policy.

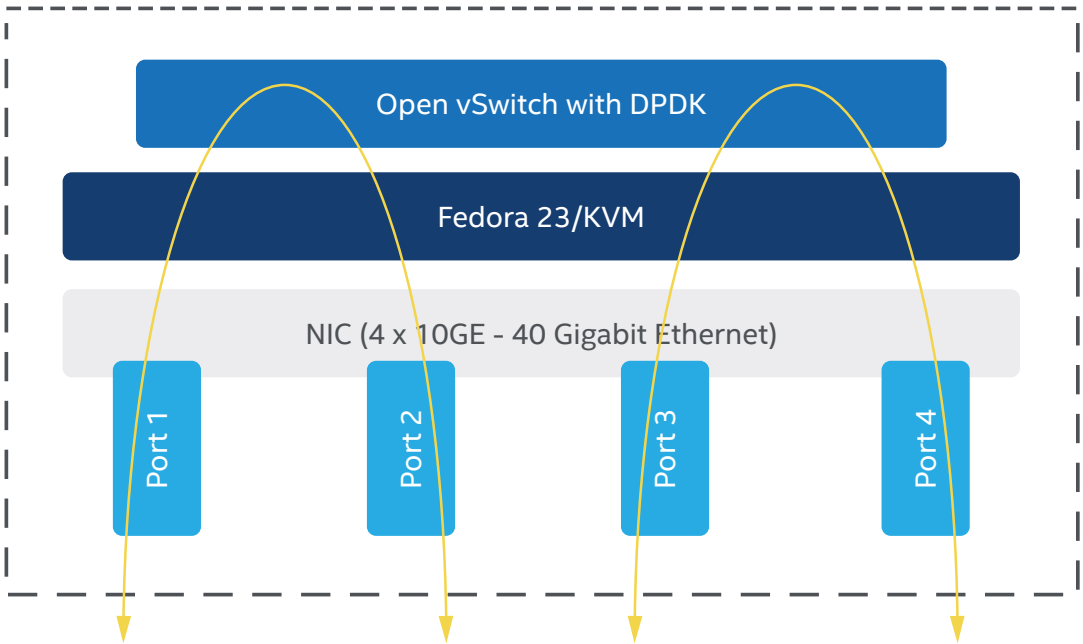

Figure 5. Virtual switching performance test setup (PHY-OVS-PHY)

SYSTEM CONFIGURATION		
Hardware		
Platform	Intel® Server Board S2600WT2	Supermicro* SX10DRH
CPU	Intel® Xeon® processor E5-2697 v3 @ 2.10 GHz	Intel® Xeon® Processor E5-2695 v4 @ 2.10 GHz
Intel® C610 series chipset	Intel® C612 chipset	Intel C610 series chipset
No of CPU	1	1
Cores per CPU	14 (Intel® Hyper-Threading Technology (Intel® HT Technology) enabled. Total: 28)	18 (Intel HT Technology enabled. Total: 36)
LL CACHE	35 MB	45 MB
QPI/DMI	Auto	Auto
PCIe	Port3A and Port3C(x8)	Port3A and Port3C(x8)
MEMORY	Micron* 16GB 1Rx4 PC4-2133 MHz, 16 GB per channel, 4 Channels, 64 GB Total	Samsung* 8GB 2Rx4 PC4-2400 MHz, 8 GB per channel, 4 Channels, 64 GB Total
NIC	2 x Intel® Ethernet X710-DA2 Adapter (Total: 4 ports)	2 x Intel Ethernet X710-DA2 Adapter (Total: 4 ports)
NIC Mbps	10000	10000
IOS	Version: SE5C610.8 6B.01.01.0008.021120151325 & Date: 02/11/2015	AMIBIOS* Version 2.0 Release date 12/17/2015
Software		
OS	Fedora* 23	Fedora 23
Kernel version	4.2.3	4.2.3
Host Machine boot setting	Hugepage size = 1G ; No. of Hugepages = 16 Hugepage szie=2MB; No. of Hugepages = 2048 isolcpus=1-13,15-27,29-41,43-55	 isolcpus=1-9,21-29
Software version	DPDK 2.2.0	DPDK 2.2.0
	OVS 2.5.0	OVS-2.5.0
Compilation Flags	DPDK compiled with "-Ofast -g -march=native" OVS configured and compiled as follows: ./configure --with-dpdk=<DPDK SDK PATH>/x86_64-native-linuxapp CFLAGS="-Ofast -g" make CLFAGS="-Ofast -g -march=native"	
Host settings	firewall, iptables, Selinux, network Manager disabled ip_forward = 0 set uncore frequency to the max ratio kill -9 dhclient rmmod ipmi_si ipmi_devintf ipmi_msghandler lpc_ich bridge	
IXIA TEST	RFC 2544 0% PACKET LOSS, 4 flows total/four ports	RFC 2544 0% PACKET LOSS, 4 flows total/four ports
BIOS settings	P-state Disabled Intel HT Technology enabled C-state Disabled Turbo Boost Disabled	P-state Disabled, Intel HT Technology enabled C-state Disabled Turbo Boost Disabled
Fortville NIC FW Version	FW 4.33 API 1.2 NVM 04.04.02 eetrack 8000191c	FW 5.0 API 1.5 NVM 05.00.02 eetrack 80002282

Table 3. Setup details for servers based on the Intel® Xeon® Processor E5-2600 v3 and v4 product families

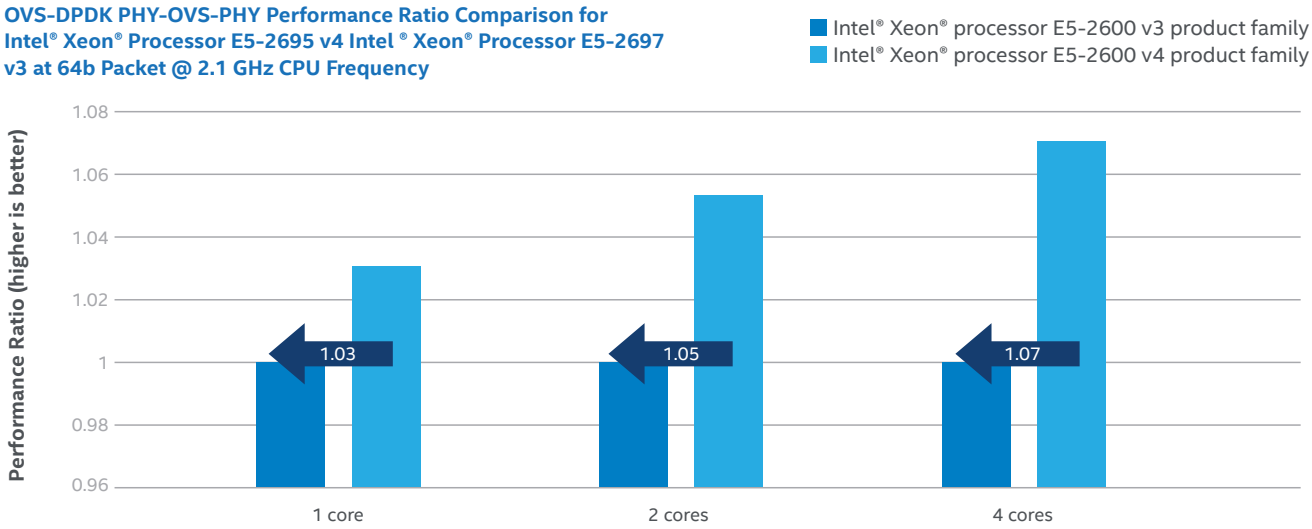

Figure 6. Throughput increase over successive generations of server platforms

KPI #3: Allocation of Processor Cores to Packet Processing

Processor cores are a precious resource in a computing system, and how they are partitioned among different workloads is an important decision made by system developers. When assigning cores to virtualized network functions (VNFs), one of the partitioning tradeoffs is the allocation of cores between packet processing (e.g., Tx/Rx) and application processing. Assigning more cores to packet processing will increase throughput until line rate is achieved; however, this means fewer cores are available for applications.

Figure 7 shows Intel and HPE OPNFV platforms with two and four physical processor cores dedicated to L3 packet forwarding. The four-core configurations achieved line rate for both large and small packet sizes, whereas the two-core configurations achieved line rate for medium and large packets (256 byte and higher). For example, the 4-core HPE OPNFV platform throughput for 64 byte packets was almost 45 Mpps; while with two cores, about 12 Mpps was achieved, or roughly half the throughput of the Inel platform. The results indicate systems that are required to support applications with a large proportion of small packets should allocate four cores to packet processing. For systems handling a mix of packet sizes, it may be possible to assign just two cores and allocate the other two cores to application processing.

Note that when the systems used two cores for packet processing, their performance was boosted by as much as 30 percent by enabling Intel HT Technology. Shown in more detail, the left side of Figure 8 indicates there is an approximately 30 percent performance gain in L3 forwarding throughput when one and two cores process four flows with Intel HT Technology is turned on. When the number of flows is increased to 4K (1K flow per port), the right side of Figure 8 shows the one core throughput increases by approximately 60 percent, and the two core case increases by about 40 percent with Intel HT Technology turned on.

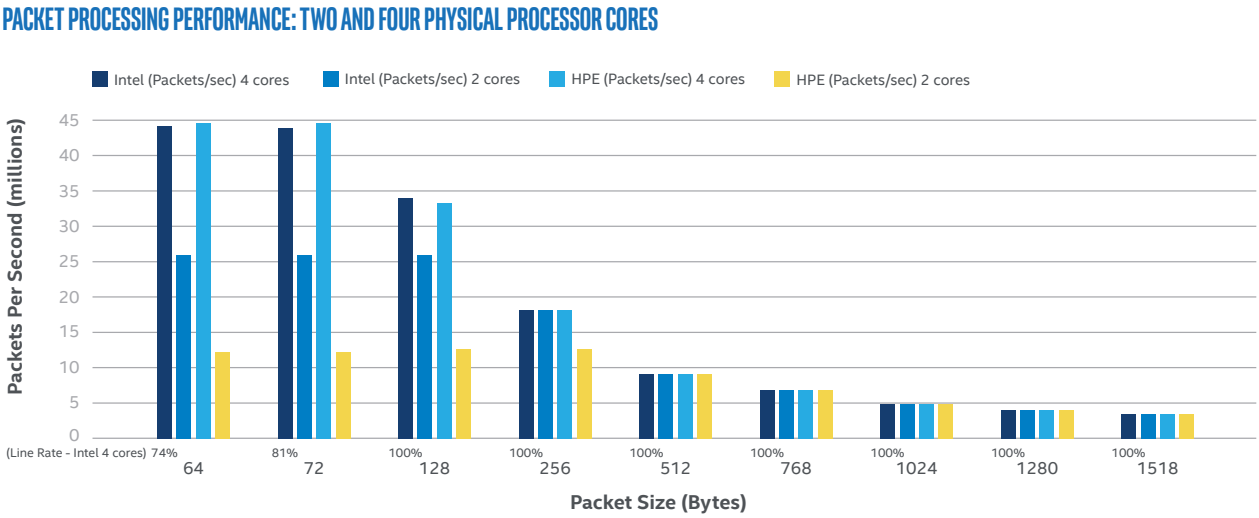

Figure 7. Packet processing throughput for different packet sizes and number of processor cores

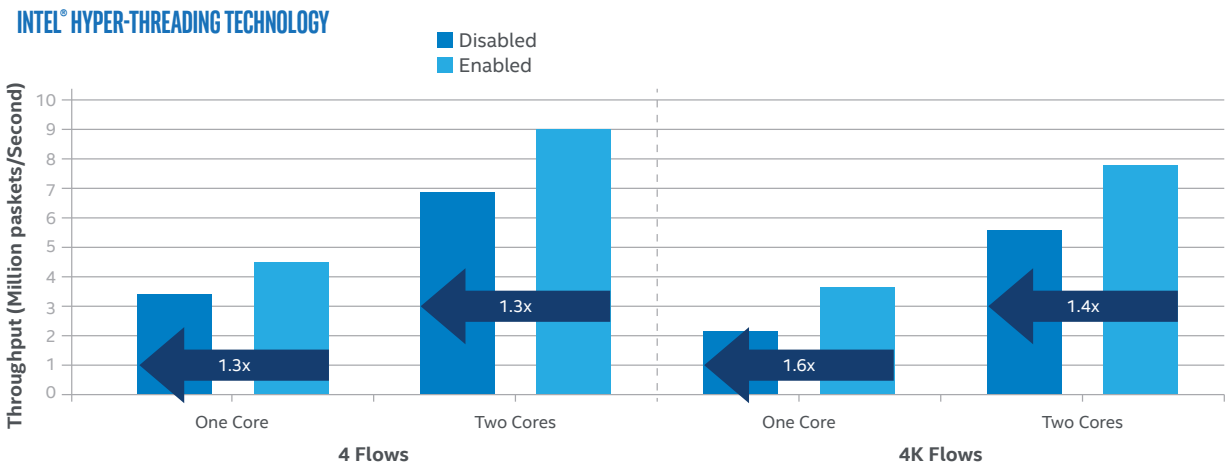

Figure 8. Layer 3 forwarding throughput for one and two processors cores and with Intel® Hyper-Threading Technology disabled/enabled.

An important platform configuration decision is whether to enable Intel HT Technology, which uses processor resources more efficiently by allowing two execution threads to run on each core. Although Intel HT Technology allows the CPU to execute instructions at a higher rate, the drawback is instruction execution is somewhat less deterministic, which can be an issue for time-critical applications.

The previous results indicate that assigning additional cores, whether physical or hyper-threaded, yields meaningful packet throughput performance improvement.

KPI #4: Open vSwitch* Performance with Multiple Flows

The data presented in Figure 9 is the Layer 3 packet forwarding throughput of one flow, but a real-world NFV scenario is likely to have multiple flows. To simulate this scenario, Intel configured a system with a VM and Open vSwitch to forward up to 4K bidirectional flows across four 10 GbE ports, as shown in Figure 10. The theoretical maximum throughput for the system is 40 Gbps, an aggregation of the four flows containing small, 64 byte traffic.

The test results for the PHY-VM-PHY use case with one core and one flow per core are shown in Figure 10 and Table 4, which show better performance for the HPE server than the Intel server, especially for the larger packet sizes. This may be due to more efficient packet scheduling achieved with the commercially available RHEL 7.2, as compared to the developer-oriented Fedora Linux distribution used on the Intel server platform.

The tests were conducted with one, two, and four processor cores performing Layer 3 forwarding, and the results are shown in Table 4 and Figure 10. The left side of the figure shows the throughput with one flow per port and Intel HT Technology disabled. There is almost a linear performance increase when more cores are used to forward packets in and out of the VM. The data also shows that to achieve better line rate in the system, dedicating additional cores to the workload yields better system throughput. The linear relationship holds whether there is only one flow per port or 1K flows per port.

Using the same setup (Figure 9), one, two, and four cores were tested on the Intel server platform, with measurements taken for one flow per port (four flows total) or 1K flows per port (4K flows total). As shown in Figure 11, adding more cores improves the throughput. Once again, adding additional physical or virtualized cores to packet forwarding improved small packet performance. For comparison purposes, Appendix C has the test results for both the Intel and the HPE server configurations.

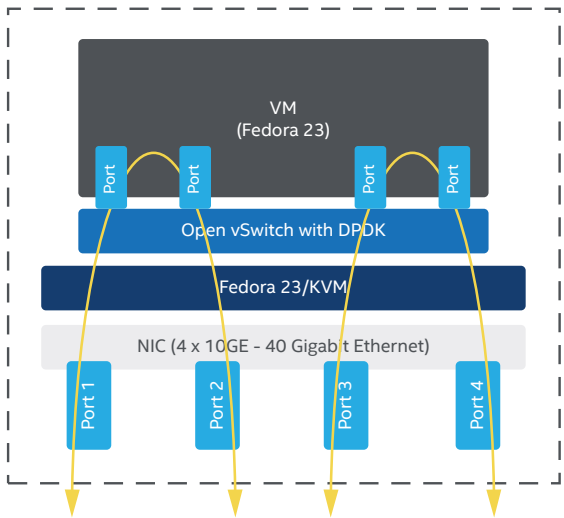

Figure 9. System with Four Packet Processing Flows

PACKET SIZE	PACKETS/SEC	
	Intel	HPE
64	3,472,608	3,591,954
72	3,328,272	3,481,894
128	3,342,744	3,491,620
256	3,211,155	3,229,974
512	2,865,347	2,983,293
768	2,130,776	2,450,980
1024	1,862,949	2,090,301
1280	1,685,732	1,903,730
1518	1,531,733	1,700,680

Paving the Way to Large-Scale NFVI Deployment

HPE and Intel plan to continue work in this area, with ongoing use of Intel ONP as a means for HPE to build commercial NFVI solutions. The companies will also be incorporating the latest versions of major open-source projects – including OpenStack*, Open Daylight, and Open vSwitch. Interested parties can access the knowledge and expertise accumulated from all of the collaborative projects through the Intel® Network Builders program.

Table 4. Intel and HPE* test Results for PHY-VM-PHY, 1 core, 1 flow per port, and Intel® Hyper-Threading Technology disabled

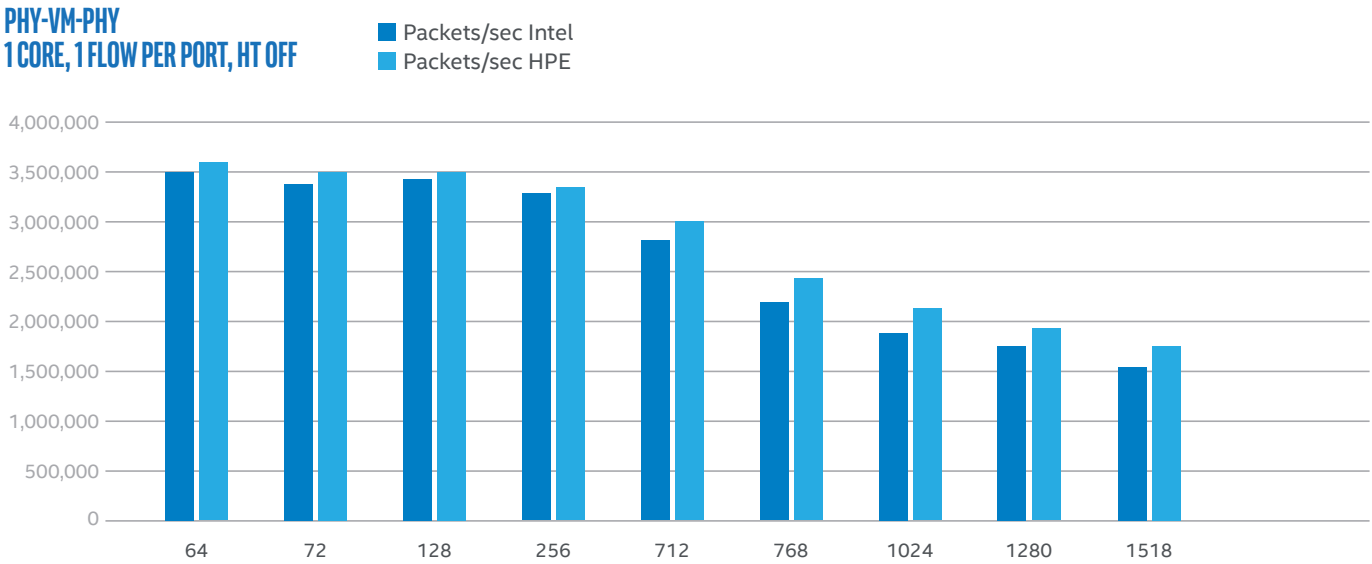

Figure 10. Chart for the PHY-VM-PHY 1 core, 1 flow per port, Intel® Hyper-Threading Technology disabled

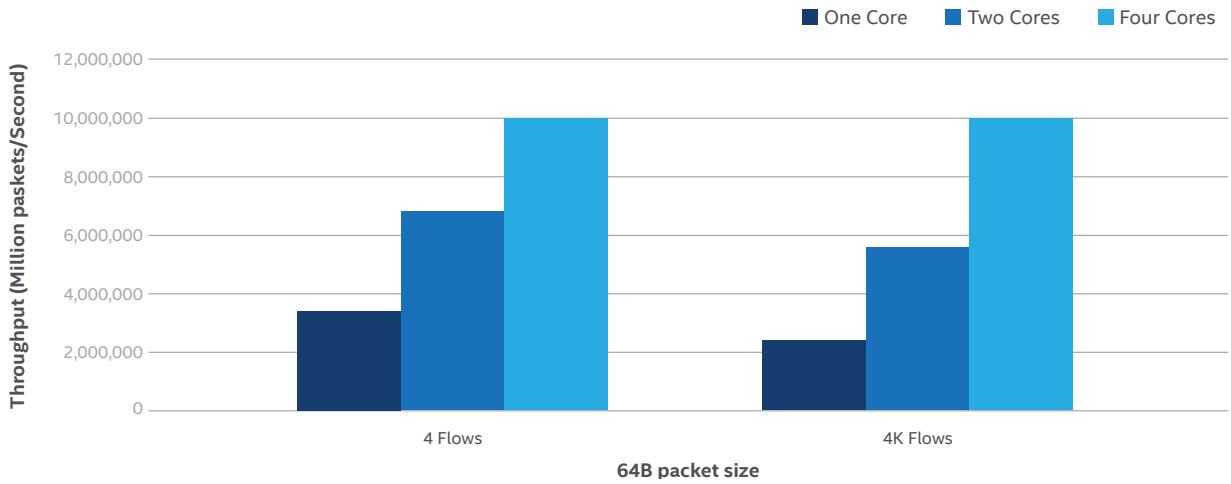

Figure 11. Layer 3 forwarding throughput for one, two, and four processors cores with Intel® Hyper-Threading Technology Enabled, Intel test set up.

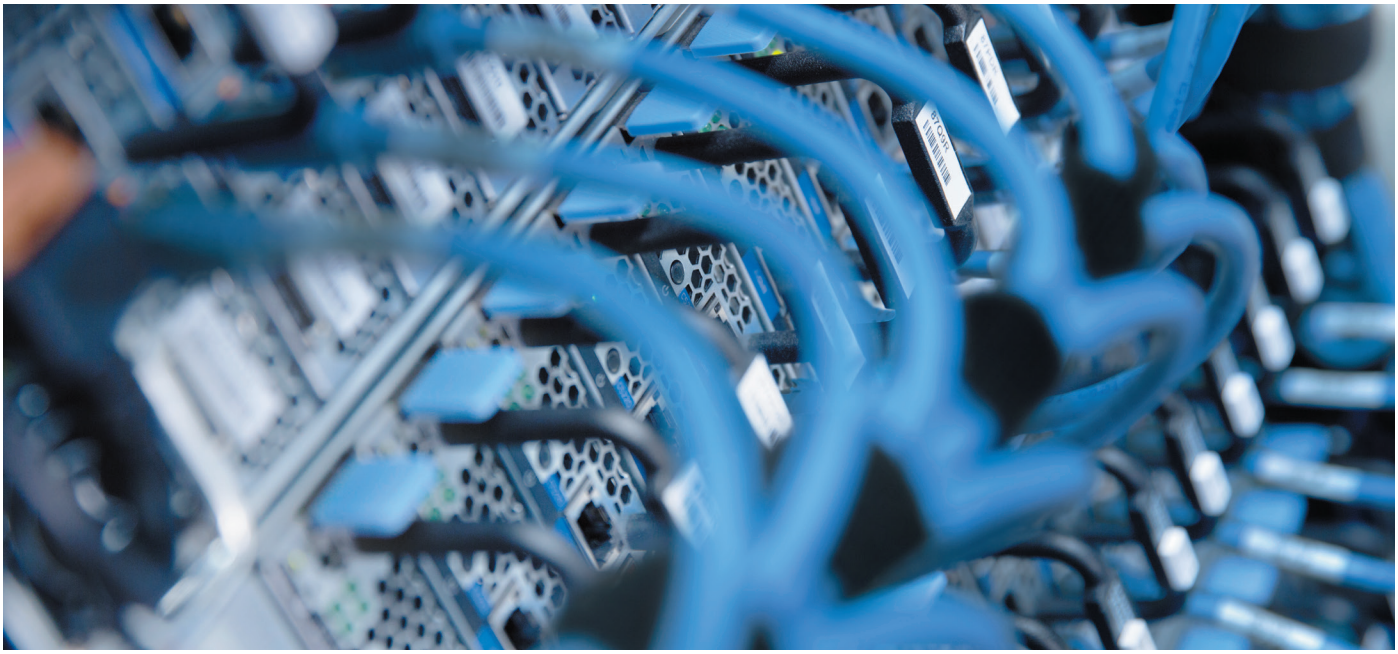

The adoption of NFV technology adds agility and manageability to networking environments and the ability to create new services for customers more efficiently at lower costs. Work by OPNFV members, like HPE and Intel, helps make large-scale NFV deployments possible, and gives telecom service providers and enterprises the confidence needed to move ahead with the virtualization of network services.

Next Steps

Service providers can take advantage of the NFVI test environments that HPE and Intel have developed. The Intel Open Network Platform reference architecture is a good starting point to create an NFVI test environment to compare and contrast different KPIs for VNFs. The testing performed for this white paper provides useful information, but it is still critical to test different NFVI configurations against actual VNF workloads.

The OPNFV community is progressing with multiple projects – Pharos*, Yardstick*, and Dovetail* – to develop open testing and interoperability. Communications service providers contribute the next level of KPIs via test scripts and use cases, and NFVI providers contribute infrastructure and testing methods to the community so VNFs can be dimensioned more quickly.

In either instance – via the OPNFV open community projects or in-house labs – testing and evaluating NFVI is a key foundation for developing a virtualized network based on industry-standard, high-volume servers and networking software from the open community.

More Information

HPE NFV solutions and HPE OpenNFV
www.hpe.com/us/en/networking/nfv.html

Intel® Open Network Platform
www.intel.com/ONP

Intel® Open Network Platform Release 2.1 Performance Test Report

Download the Intel Open Network Platform Server Reference Architecture for NFV and SDN
01.org/packet-processing.

Intel Network Builders
networkbuilders.intel.com

Appendix A

PHY-OSV-PHY test results for OvS with DPDK (four cores, one flow per port, one core disabled
One TxRx queue per core
Four logical core on Four physical core (SMT not used)
4PMD thread based OVS and 0.0 percent Loss resolution
/ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=3c
4P FVL Card

PACKET SIZE	PACKETS/SEC		% LINE RATE	
	Intel	HPE	Intel	HPE
64	44,216,458.844	44,694,077.867	74	75.086
72	44,154,661.894	44,642,857.200	81	82.143
128	33,783,781.576	33,653,513.067	100	99.614
256	18,115,943.388	18,115,942.067	100	100
512	9,398,496.911	9,398,496.267	100	100
768	6,345,178.098	6,345,177.667	100	100
1024	4,789,272.054	4,789,272.067	100	100
1280	3,846,153.858	3,846,153.867	100	100
1518	3,250,975.201	3,246,753.267	100	100

Appendix B

One Tx/Rx queue per core
Two logical cores on two physical core, Intel HT Technology disabled
Four poll mode driver (PMD) threads based OVS and 0.0 percent Loss resolution
./ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=c
4P FVL Card

PACKET SIZE	PACKETS/SEC	
	Intel	HPE
64	21,945,431	23,245,451
72	21,928,701	23,268,124
128	21,992,819	23,358,516
256	18,115,940	18,115,942
512	9,398,497	9,398,496
768	6,345,178	6,345,177
1024	4,789,272	4,789,272
1280	3,846,154	3,846,153
1508	3,250,975	3,246,753

Appendix C

PHY-VM-PHY (1 core)
1 flow per port, Intel HT Technology disabled
One poll mode driver (PMD) threads based OvS and 0.0% Loss resolution
/ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=4
On a VM:
./testpmd -c 0x6 -n 4 -- --burst=64 -i --txd=2048 --rxd=2048 --txqflags=0xf00 --disable-hw-vlan
Two Quad Port Intel® Ethernet Converged Network Adapter X710-DA4; two ports in use per each NIC

PACKET SIZE	PACKETS/SEC	
	Intel	HPE
64	3,472,608	3,591,954
72	3,328,272	3,481,894
128	3,342,744	3,491,620
256	3,211,155	3,229,974
512	2,865,347	2,983,293
768	2,130,776	2,450,980
1024	1,862,949	2,090,301
1280	1,685,732	1,903,730
1508	1,531,733	1,700,680

Phy-VM-Phy (2 cores)
Two poll mode driver (PMD) threads based OvS and 0.0% Loss resolution
./ovs-vsctl set Open_vSwitch. other_config:pmd-cpu-mask=c
On a VM:
./testpmd -c 0x6 -n 4 -- --burst=64 -i --txd=2048 --rxd=2048 --txqflags=0xf00 --disable-hw-vlan
Two Quad Port Intel® Ethernet Converged Network Adapter X710-DA4; two ports in use per each NIC

PACKET SIZE	PACKETS/SEC	
	Intel	HPE
64	6,925,459	7,238,868
72	6,585,960	6,759,637
128	6,641,618	6,868,131
256	6,153,582	6,250,000
512	4,991,573	5,081,300
768	3,940,452	4,139,072
1024	3,529,840	3,768,140
1280	3,124,770	3,453,038
1508	2,773,231	3,104,089

PHY-VM-PHY (four cores)
Four poll mode driver (PMD) threads based OvS and 0.0 precent loss resolution
/ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=3c (four cores)
On a VM:
./testpmd -c 0x6 -n 4 -- --burst=64 -i --txd=2048 --rxd=2048 --txqflags=0xf00 --disable-hw-vlan
Two Quad Port Intel® Ethernet Converged Network Adapter X710-DA4; two ports in use per each NIC

PACKET SIZE	PACKETS/SEC	
	Intel	HPE
64	10,148,118	10,461,110
72	10,106,372	10,274,627
128	10,103,798	10,245,901
256	9,919,190	10,162,601
512	9,398,488	9,398,496
768	6,345,175	6,345,177
1024	4,789,270	4,789,272
1280	3,846,152	3,846,153
1508	3,250,974	3,246,753

1. "Telecom Industry and Vendors Unite to Build Common Open Platform to Accelerate Network Functions Virtualization," September 30, 2014, https://www.opnfv.org/news-faq/press-release/2014/09/telecom-industry-and-vendors-unite-build-common-open-platform.
2. "Benchmarking Virtual Switches in OPNFV," https://datatracker.ietf.org/doc/draft-vsperf-bmwg-vswitch-opnfv/00.
3. Available on select Intel® processors. Requires an Intel® Hyper-Threading Technology (Intel® HT Technology)-enabled system. Consult your PC manufacturer. Performance will vary depending on the specific hardware and software used. For more information including details on which processors support Intel HT Technology, visit http://www.intel.com/info/hyperthreading.

Printed in USA 0916/TT/ICMCSW/PDF ♻ Please Recycle 335016-001US