TECHNICAL BRIEF

QxStack NFV Infrastructure with Red Hat OpenStack Platform

A Carrier-Grade Infrastructure Pre-Integrated and Validated for Network Service Providers

Version 1.0

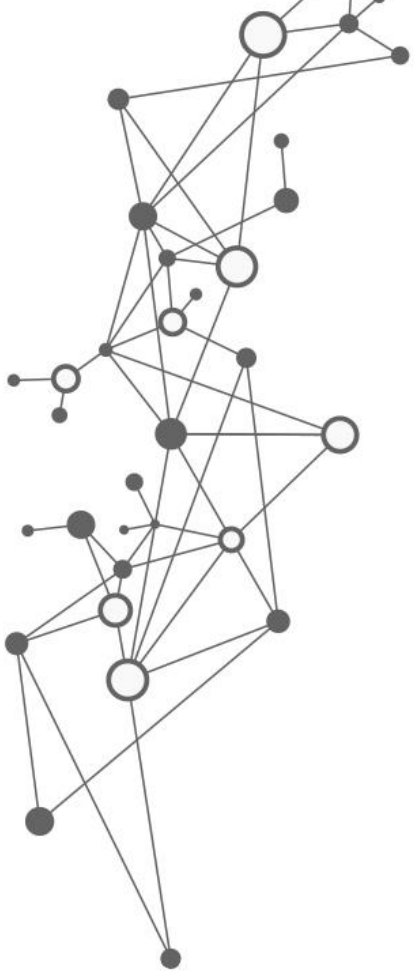

EXECUTIVE SUMMARY

Demand for digital media and mobile applications is growing worldwide, creating an increasing need for high-speed networks. To meet this demand and remain successful, communications services providers (CSPs) must operate high-speed networks and continuously innovate to deliver compelling new services, faster.

However, traditional network infrastructure can hold CSPs back. Legacy appliances and systems often cannot cost-effectively handle today's growing quantities of data and complex workloads. As a result, digital transformation and adoption of new technologies are critical.

Network functions virtualization (NFV) can help CSPs transform their network by decoupling network functions from proprietary appliances and running them as virtualized workloads on a common infrastructure. A programmable network lets CSPs independently design, deploy, and operate complex network functions with greater agility and lower total cost of ownership (TCO). Additionally, integration of network virtualization technology, resource allocation, and performance optimization helps CSPs enhance network functions for increased customer satisfaction.

Quanta Cloud Technology (QCT) and Red Hat offer an NFV infrastructure solution that meets carrier-grade requirements for datacenter and network innovation. With integrated hardware and software, a non-uniform memory access (NUMA) balanced design, and accelerated data plane packet processing, this cloud-based solution delivers exceptional network performance to satisfy digital demands and prepare CSPs for future innovation.

Designed to comply with strict CSP requirements, the QCT QxStack NFV Infrastructure with Red Hat® OpenStack® Platform solution provides a flexible, modular platform based on industry-standard x86 hardware, allowing CSPs to resize compute, storage, and networking resources independently with minimal overhead. High system reliability and availability ensure virtualized network functions (VNFs) always have access to the resources they need to perform optimally.

QCT and Red Hat also simplify and streamline deployment of your NFV infrastructure. The QxStack Auto-Deployment Tool uses scripts to customize OpenStack Heat orchestration templates, reducing installation time from weeks to hours.

Finally, QCT used the industry-standard Open Platform for NFV (OPNFV) Yardstick framework to validate the performance of the solution. This framework validates infrastructure compliance from the VM perspective to provide an assessment of the performance benefits of deploying the QCT and Red Hat solution with Enhanced Platform Awareness (EPA) capabilities.

QxStack NFV Infrastructure with Red Hat OpenStack Platform

CSPs are increasingly adopting open source solutions to reduce vendor lock-in and gain flexibility. As the leading provider of open source software, Red Hat works closely with CSPs to develop the products and functionality needed for modern telecommunications datacenters. As a result, Red Hat products form an ideal foundation for NFV deployments. One example is Red Hat OpenStack Platform.

OpenStack is the de facto cloud platform standard for NFV deployments. Red Hat is a leading contributor to the OpenStack community and provides intensely validated and hardened versions of OpenStack releases for production use. Beginning with Red Hat OpenStack Platform 10, based on the Newton community release, Red Hat offers three years of production support, including unlimited technical support, security and bug patches, and backporting of select features. Two additional years of extended life-cycle support (ELS) can be purchased.[1]

As a global datacenter solution provider, QCT combines the efficiency of hyperscale hardware with infrastructure software from industry leaders to solve modern datacenter challenges. With original design manufacturer (ODM) expertise, QCT has successfully designed, customized, and manufactured servers for large internet companies.

Transforming CSP datacenters with open source cloud technologies can be challenging. Many components must be tested and integrated into a flexible, reliable, and high-performance infrastructure. To help CSPs overcome these challenges, QCT and Red Hat collaborate to build integrated, readily deployable solutions. The QCT QxStack NFV Infrastructure with Red Hat OpenStack Platform solution accelerates NFV adoption and deployment through preintegration, simplified installation tooling, and performance optimization.

The minimum recommended hardware infrastructure includes:

- ***1x Red Hat OpenStack Platform director node.***
- ***3x Red Hat OpenStack Platform controller nodes.***
- ***3x Red Hat OpenStack Platform compute nodes with DPDK or SR-IOV.***
- ***3x Red Hat Ceph Storage nodes.***

Solution and test environment overview

The default SKU consists of QCT D52B 1U servers for controller and compute nodes and QCT D51PH 1U servers storage nodes. All the nodes are equipped with Intel 25G network interface cards (NICs), in which the compute nodes are configured to support either Data Plane Development Kit (DPDK) or single-root I/O virtualization (SR-IOV). QCT QuantaMesh BMS T4048-IX2 and QuantaMesh T3048-LY9 switches are recommended for data and management switching, respectively.

Red Hat OpenStack Platform 10 provides the cloud operating system for the solution. To meet CSP requirements, Red Hat OpenStack Platform supports several EPA features, including memory huge pages, CPU pinning, NUMA awareness, and DPDK and SR-IOV network adapters. The QCT and Red Hat NFV infrastructure solution enables all of these features and tailors Red Hat OpenStack Platform 10 and QCT hardware for CSP workloads.

[1] *Learn more about support for Red Hat OpenStack Platform at https://access.redhat.com/support/policy/updates/openstack/platform*

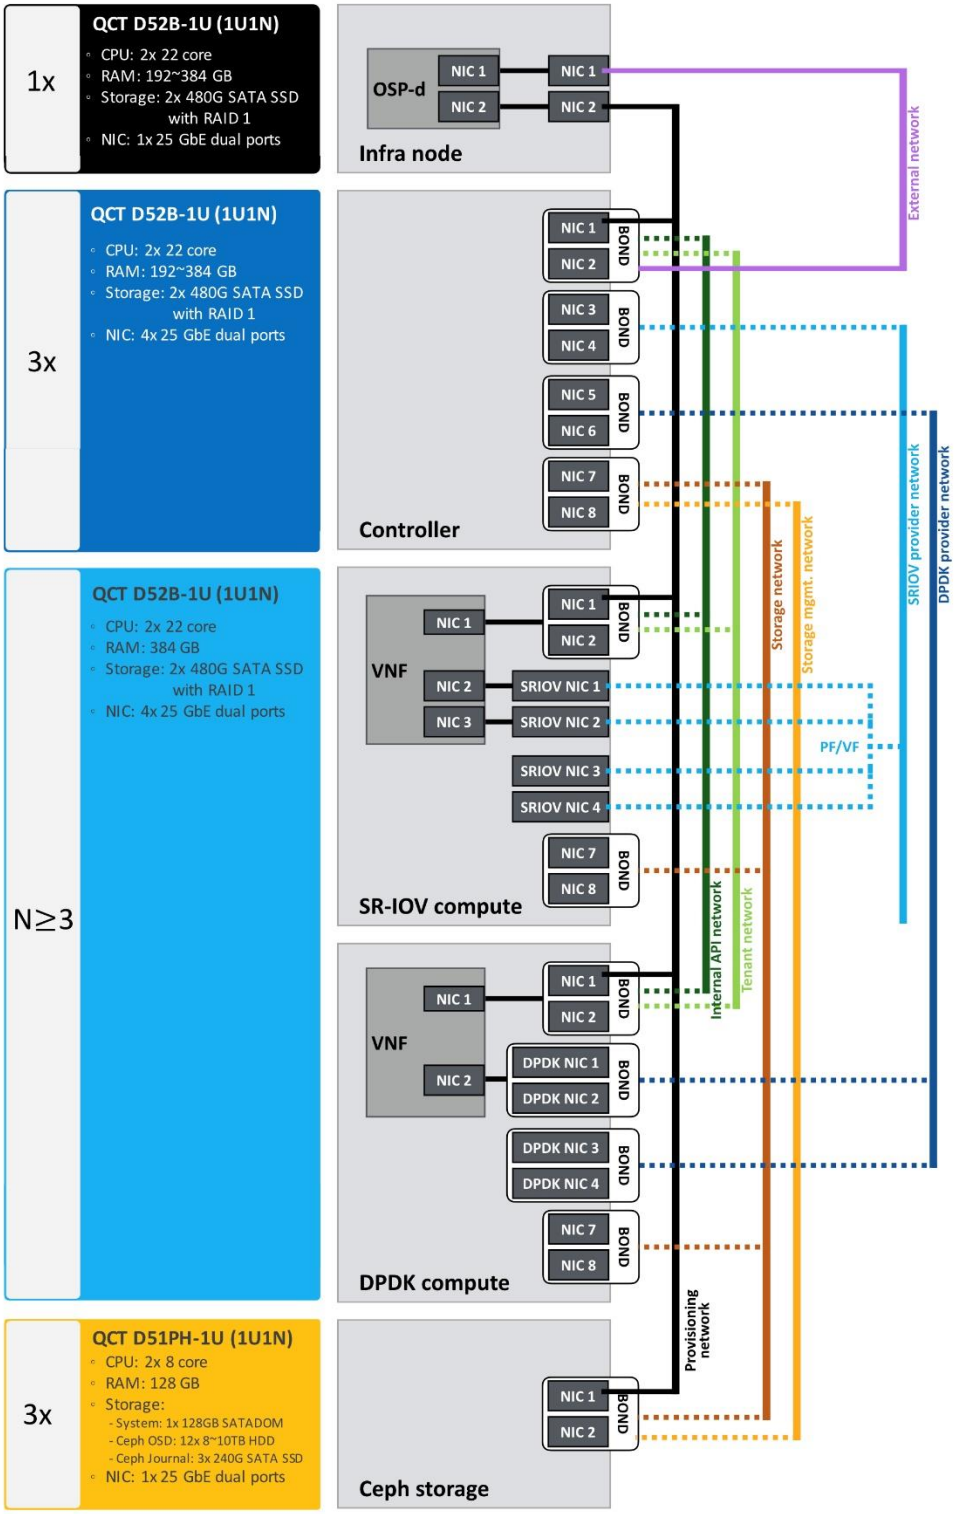

Figure 1. QCT QxStack NFV Infrastructure with Red Hat OpenStack Platform solution architecture

QCT used the OPNFV Yardstick framework to verify and evaluate the performance of this solution. As a testing project sponsored by the Linux Foundation, Yardstick implements system-level validation aligned with the European Telecommunications Standards Institute (ETSI) TST 001 specification.[2] It verifies the underlying infrastructure performance of an NFV deployment from the perspective of VM running on the virtual infrastructure manager (VIM) platform, in this case, Red Hat OpenStack Platform. To accommodate a variety of NFV use cases, Yardstick test cases decompose typical workload performance metrics into several characteristics and performance vectors. Yardstick performance evaluations of the optimized QCT and Red Hat solution are detailed in the following sections of this document.

"To accelerate data plane packet processing, the QCT QxStack NFV Infrastructure with Red Hat OpenStack Platform solution supports either ***DPDK or SR-IOV*** in compute nodes with memory huge pages."

Test results: Data plane packet processing

DPDK consists of a set of data plane libraries and user-space network drivers for accelerated packet processing. It provides a programmable framework that implements a run-to-completion model, eliminates packet interrupt processing overhead, and enables applications to perform packet processing operations directly from and to the NIC. This significantly improves network throughput and latency performance in Red Hat OpenStack Platform. Shown in Figure 2, the QCT and Red Hat NFV infrastructure solution uses a DPDK-accelerated version of Open vSwitch (OVS-DPDK) to enhance network performance. In this case, OVS-DPDK replaces the standard OVS kernel datapath with a DPDK-based datapath, creating a user-space Open vSwitch (OVS) for packet forwarding. OVS-DPDK efficiently allocates virtual host (vhost) memory across NUMA nodes while remaining transparent in the overall architecture and exposing the same interfaces—including OpenFlow, Open vSwitch Database (OVSDB), and command lines—as the standard OVS implementation.

SR-IOV is a specification that allows physical PCI devices to be shared between multiple virtual machines (VMs) for increased network performance. SR-IOV virtualizes PCI hardware devices to create multiple virtual functions (VFs)—lightweight functions that can be assigned to specific VMs—on top of a physical functions (PFs)—full-feature physical hardware ports. A VF driver is required to implement SR-IOV. This driver resides in the VM, presents VFs to the VM as physical NICs, and allows the VM to communicate directly with the physical device. Network traffic from a VM with a direct-attached VF bypasses the software switching layer to achieve near line-rate performance.

Both OVS-DPDK and SR-IOV take advantage of memory huge pages. Physical memory is typically segmented into 4KB pages. Memory huge pages increases the size of these memory blocks to either 2MB or 1GB, reducing the number of pages needed for a given amount of data. This increases the amount of memory that can be mapped by the translation lookaside buffer (TLB), reducing the potential for TLB misses and improving computational performance.

---

[2] *ETSI TST 001 is an informative report on methods for pre-deployment testing of the functional components of an NFV environment. For further information, http://www.etsi.org/deliver/etsi_gs/NFV-TST/001_099/001/01.01.01_60/gs_nfv-tst001v010101p.pdf*

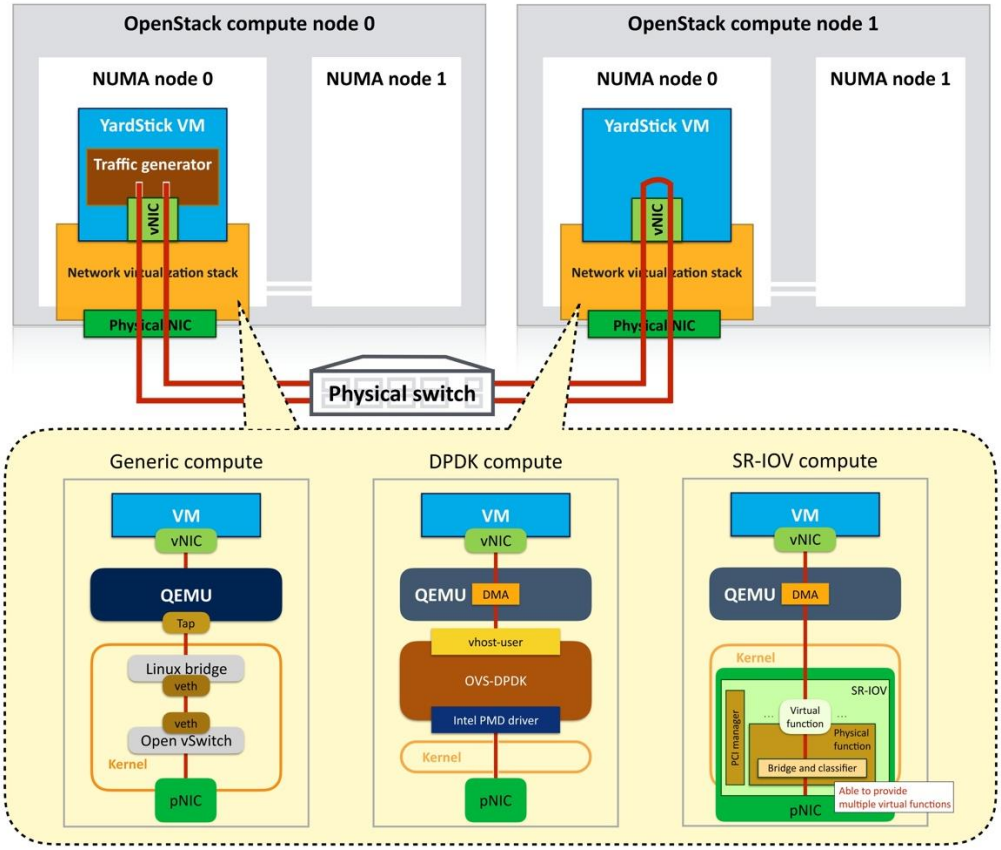

Figure 2. Testing topology and comparison of network virtualization stacks

Generic network virtualization stacks can be complicated, relying on the Linux® and OVS bridges for traffic flow control, shown in Figure 2. Both the OVS-DPDK and SR-IOV improve the network virtualization stack to accelerate data plane processing. QCT used the OPNFV Yardstick framework to evaluate the performance of each of these network virtualization stack designs. In this evaluation, two VMs were launched on different compute hosts and assigned two virtual CPU (vCPU) cores, one dedicated to host processes and the other to packet generation and transmission.

In each test, a traffic generator transmits traffic from one VM's virtual network interface. This traffic then flows through the given network virtualization stack to the other VM. Figure 3 shows the throughput performance for generic, OVS-DPDK, and SR-IOV network virtualization stacks, while Figure 4 shows the corresponding latency performance. These test results demonstrate that both OVS-DPDK and SR-IOV can significantly improve network performance regardless of packet size.

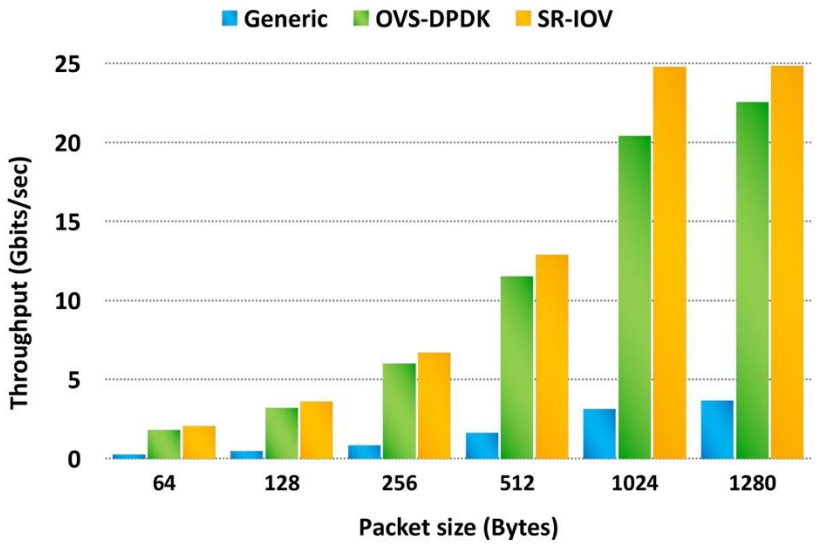

Figure 3. Throughput performance of generic, OVS-DPDK, and SR-IOV network virtualization stacks

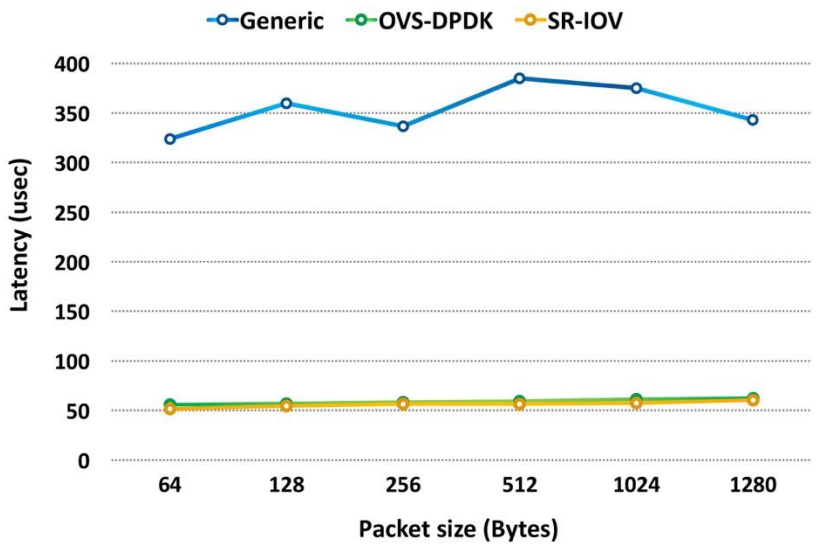

Figure 4. Latency performance of generic, OVS-DPDK, and SR-IOV network virtualization stacks

"To optimize resource allocation, the QCT QxStack NFV Infrastructure with Red Hat OpenStack Platform solution supports ***CPU pinning*** and features a ***NUMA-aware design***."

Test results: Resource allocation optimization in NUMA-balanced designs

In virtualized infrastructures, a pool of physical CPUs (pCPUs) on a host are shared across multiple vCPUs associated with VMs. CPU pinning enables one-to-one mapping between vCPUs and pCPUs to increase VM performance. Because VMs run as user-space tasks within the host operating system, CPU pinning provides similar advantages to task pinning. Shown in Figure 5, CPU pinning dedicates specific compute resources to specific VMs and increases cache efficiency.

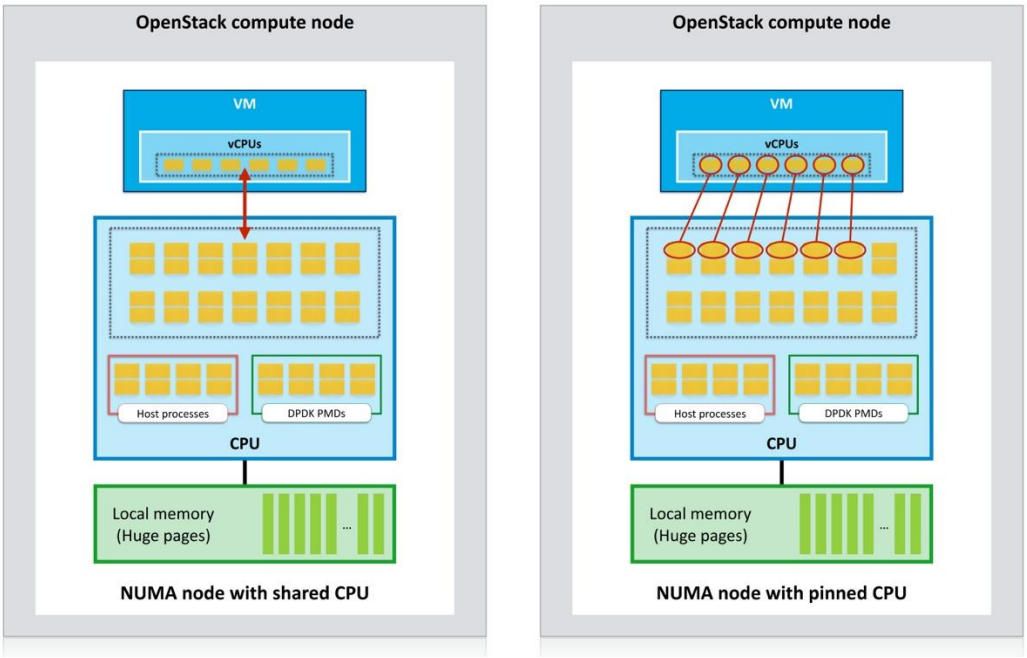

Figure 5. CPU topology and core allocation

Traditional uniform memory access (UMA) architecture models share memory resources evenly across all CPUs and sockets in a multiprocessor system. This often results in long memory access times, regardless of the location of the memory in relation to the CPU or socket. NUMA architecture models geographically distribute system memory in a manner that takes into account its location in relation to each CPU, speeding access to memory that is closer to the CPU. Processes can then access local CPU memory—rather than another CPU's local memory or shared memory—to improve computational performance. In Red Hat OpenStack Platform, OpenStack Compute (Nova) intelligently schedules and places memory when launching instances. Administrators can create instance configurations customized for specific performance levels to target specialized workloads like NFV and high-performance computing (HPC).

The QCT QxStack NFV Infrastructure with Red Hat OpenStack Platform solution uses a NUMA-balanced design that supports local memory access and distributes NICs across CPUs and sockets. Shown in Figure 6a, a NUMA-balanced design uses CPU pinning to place vCPUs, memory, and NICs on the same local sockets, providing consistent, high performance for VMs. In contrast, a non-NUMA-balanced design using the default OpenStack CPU sharing policies is shown in Figure 6b.

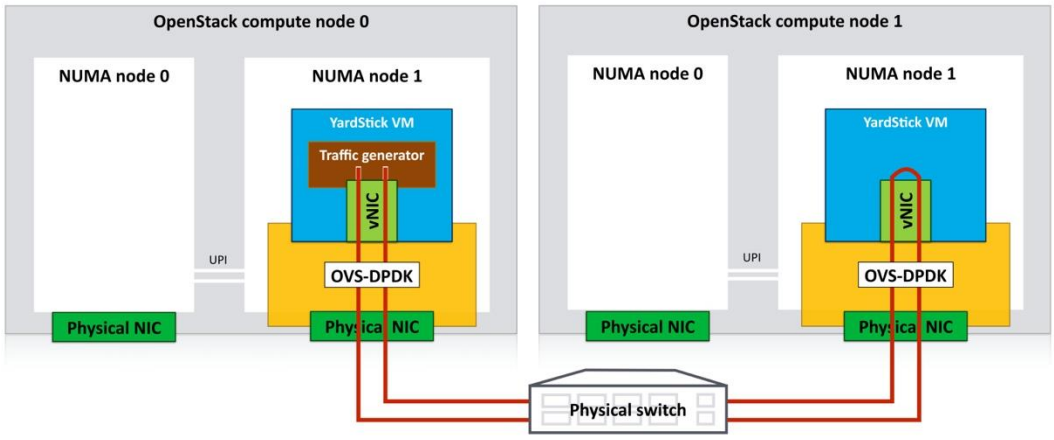

Figure 6a. NUMA-balanced design with OVS-DPDK

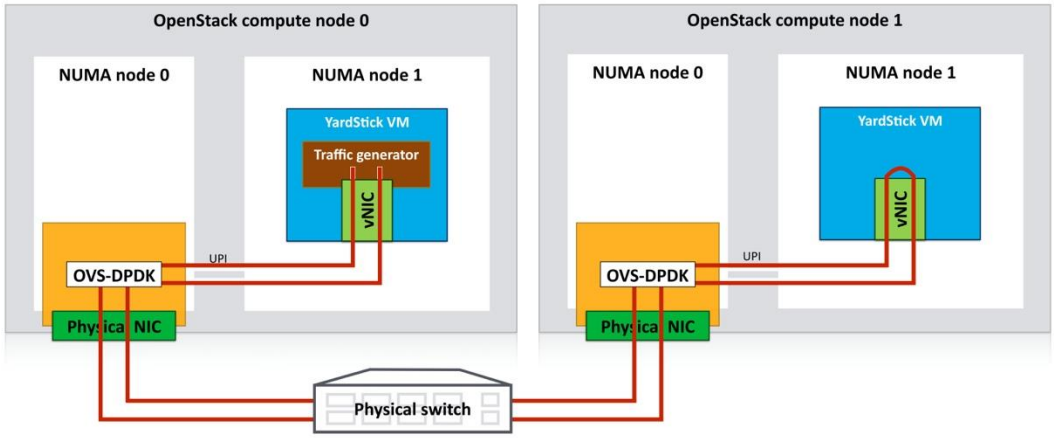

Figure 6b. Non-NUMA-balanced design with OVS-DPDK

QCT evaluated the performance of NUMA-balanced and non-NUMA-balanced designs using OPNFV Yardstick framework. As in the packet processing tests, two VMs were launched on different compute hosts and assigned two cores, one dedicated to host processes and the other to packet generation and transmission. Traffic was generated and transmitted from one VM to the other and network throughput and latency were measured. Shown in Figures 7 and 8, the NUMA-balanced design demonstrated 5% to 15% better performance for both network throughput and latency.

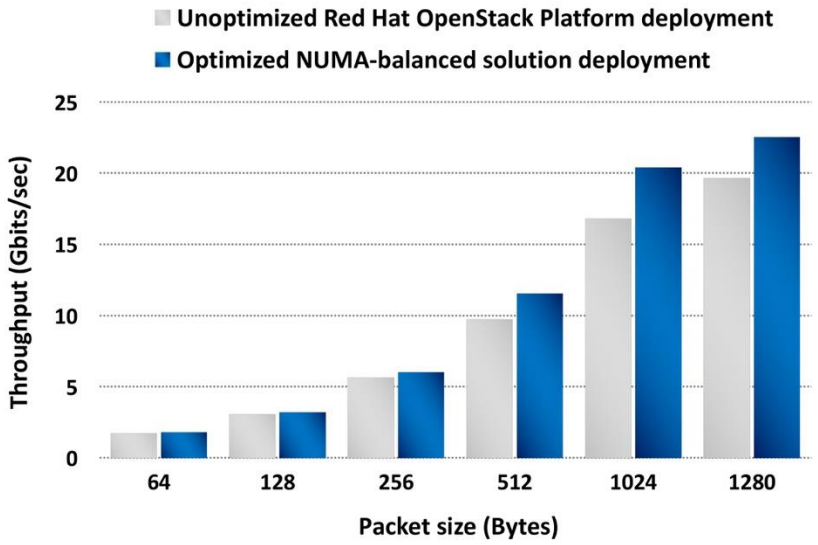

Figure 7. Throughput performance of an unoptimized Red Hat OpenStack Platform deployment compared to the optimized NUMA-balanced solution deployment

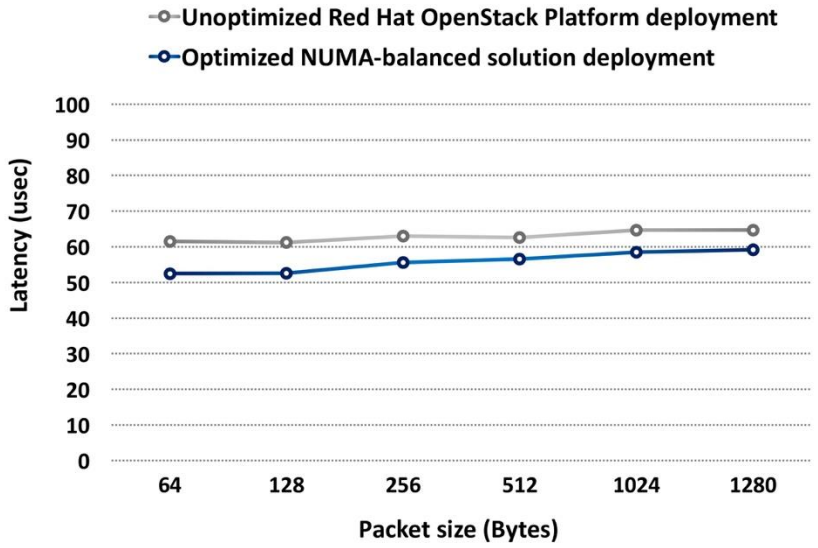

Figure 8. Latency performance of an unoptimized Red Hat OpenStack Platform deployment compared to the optimized NUMA-balanced solution deployment

CONCLUSION

CSPs must digitally transform their network infrastructure to meet growing customer demand for mobile applications and data. The QxStack NFV Infrastructure with Red Hat OpenStack Platform solution provides a flexible, reliable, high-performance foundation for NFV. The solution preintegrates QCT's NUMA-balanced hardware with Red Hat's open source software to reduce the time, effort, and expense of deploying large NFV environments. This allows CSPs to quickly and cost-effectively implement the infrastructure they need today and expand easily over time. Optimizations across the physical and virtual solution stack—including EPA features and data plane enhancements—provide consistent workload and VM performance. Solution evaluation using the OPNFV Yardstick framework demonstrates improvements in both network throughput and latency compared to unoptimized, standard implementations.

To learn more about the QxStack NFV Infrastructure with Red Hat OpenStack Platform solution and how QCT and Red Hat can help you transform your network infrastructure, visit www.qct.io/q/NFVI and redhat.com/openstack.

LEGAL DISCLAIMER

ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.

ABOUT QCT

QCT (Quanta Cloud Technology) is a global datacenter solution provider extending the power of hyperscale datacenter design in standard and open SKUs to all datacenter customers.

Product lines include servers, storage, network switches, integrated rack systems and cloud solutions, all delivering hyperscale efficiency, scalability, reliability, manageability, serviceability and optimized performance for each workload.

QCT offers a full spectrum of datacenter products and services from engineering, integration and optimization to global supply chain support, all under one roof.

The parent of QCT is Quanta Computer Inc., a Fortune Global 500 technology engineering and manufacturing company.

http://www.QCT.io

UNITED STATES

QCT LLC., Silicon Valley office
1010 Rincon Circle, San Jose, CA 95131
TOLL-FREE: 1-855-QCT-MUST
TEL: +1-510-270-6111
FAX: +1-510-270-6161
Support: +1-510-270-6216

QCT LLC., Seattle office
13810 SE Eastgate Way, Suite 190, Building 1,
Bellevue, WA 98005
TEL: +1-425-633-1620
FAX: +1-425-633-1621

CHINA

云达科技, 北京办公室（Quanta Cloud Technology）
北京市朝阳区东大桥路 12 号润诚中心 2 号楼
TEL +86-10-5920-7600
FAX +86-10-5981-7958

云达科技, 杭州办公室（Quanta Cloud Technology）
浙江省杭州市西湖区古墩路浙商财富中心 4 号楼 303 室
TEL +86-571-2819-8650

JAPAN

Quanta Cloud Technology Japan 株式会社
東京都港区芝大門 2-5-8 芝大門牧田ビル 3F, 105-0012
TEL +81-3-5777-0818
FAX +81-3-5777-0819

GERMANY

Quanta Cloud Technology Germany GmbH
Hamborner Str. 55, 40472 Düsseldorf
TEL +492405-4083-1

TAIWAN

雲達科技（Quanta Cloud Technology）
桃園市龜山區文化二路 211 號 1 樓
1F, No. 211 Wenhua 2nd Rd., Guishan Dist., Taoyuan City 33377, Taiwan