

QxStack NFV Infrastructure Reference Architecture

VMware vCloud NFV OpenStack Edition



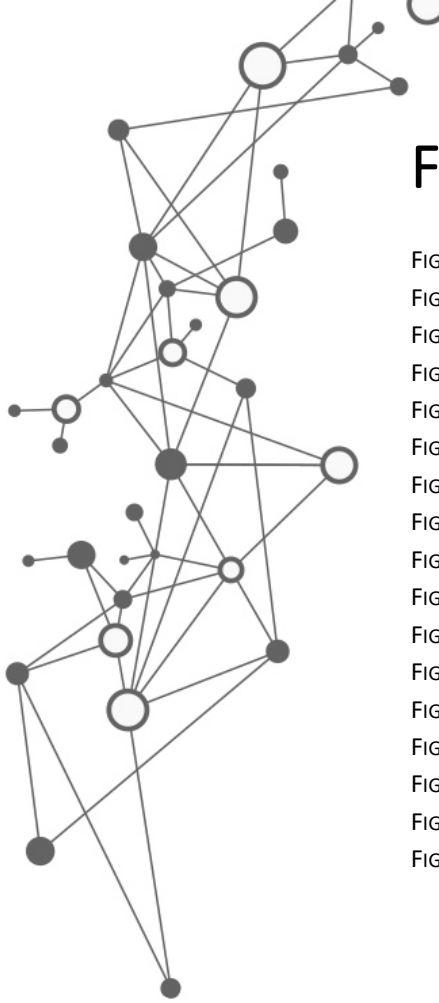
CONTENT

1. EXECUTIVE SUMMARY	1
2. SOLUTION OVERVIEW	2
3. SOLUTION ARCHITECTURE.....	3
3.1. THREE-POD AND TWO-POD DESIGNS.....	3
3.1.1. <i>Management Pod</i>	3
3.1.2. <i>Edge Pod</i>	3
3.1.3. <i>Resource Pod</i>	4
3.2. NETWORKING TECHNOLOGY AND ARCHITECTURE	4
3.2.1. <i>SR-IOV</i>	4
3.2.2. <i>Resource Allocation Strategy</i>	4
3.2.3. <i>Network Design Concept</i>	5
3.2.4. <i>Network Design for Management Pod</i>	6
3.2.5. <i>Network Design for Edge Pod</i>	6
3.2.6. <i>Network Design for Resource Pod</i>	6
3.3. SOLUTION HARDWARE SUGGESTION	7
3.3.1. <i>QCT HCI Server</i>	7
3.3.2. <i>QCT QuantaMesh Switch</i>	10
3.4. SOLUTION SOFTWARE	11
3.4.1. <i>VMware vSphere®</i>	11
3.4.2. <i>ESXi™</i>	12
3.4.3. <i>VMware vCenter Server®</i>	12
3.4.4. <i>vSAN™</i>	12
3.4.5. <i>VIO</i>	12
3.4.6. <i>NSX-T</i>	12
3.4.7. <i>vRealize® Operations Manager™</i>	13
3.4.8. <i>vRealize® Log Insight™</i>	13
3.4.9. <i>vRealize® Network Insight™</i>	13
4. SOLUTION CONFIGURATION	14
4.1. HARDWARE CONFIGURATION	14
4.2. NETWORK PLANNING.....	14
4.2.1. <i>Logical Network Topology</i>	14
4.2.2. <i>Logical Network Layout in Compute Nodes</i>	15
5. SOLUTION PERFORMANCE	16
5.1. DATA PLANE PERFORMANCE	16
5.1.1. <i>Test Scenario Brief</i>	16
5.1.2. <i>Test Method and Tools</i>	16
5.2. SOLUTION PERFORMANCE EVALUATION.....	16
5.2.1. <i>Network I/O Performance</i>	16
5.2.2. <i>Multi-Host Data Plane Throughput Performance</i>	18
5.2.3. <i>Scalable Data Plane Performance</i>	19
6. CONCLUSION	22
REFERENCE	23



TABLES

TABLE 1. SOLUTION HARDWARE – QUANTA vSAN READY SERIES.....	8
TABLE 2. SOLUTION HARDWARE – QUANTAMESH SWITCH SERIES	11
TABLE 3. HARDWARE CONFIGURATION FOR TWO-POD DESIGN.....	14
TABLE 4. TEST CONFIGURATION OF NETWORK I/O PERFORMANCE.....	17
TABLE 5. TEST CONFIGURATION OF MULTI-HOST DATA PLANE PERFORMANCE.	19
TABLE 6. TEST CONFIGURATION OF SCALABLE DATA PLAN PERFORMANCE.....	20



FIGURES

FIGURE 1. OVERVIEW OF QXSTACK NFV INFRASTRUCTURE ARCHITECTURE.....	2
FIGURE 2. VIRTUAL BUILDING BLOCKS FOR EACH POD.	3
FIGURE 3. COMPARISON OF CPU CORE ALLOCATION BETWEEN SHARED CPU AND PINNED CPU.	4
FIGURE 4. NUMA-BALANCED DESIGN FOR QXSTACK NFV INFRASTRUCTURE.	5
FIGURE 5. VIRTUAL NETWORK TOPOLOGY ACROSS EACH POD.....	6
FIGURE 6. OVERVIEW OF QUANTAPLEX T42S-2U EXPANSION SLOTS.....	8
FIGURE 7. OVERVIEW OF QUANTAGRID D52B-1U EXPANSION SLOTS.	10
FIGURE 8. OVERVIEW OF QUANTAGRID D52BQ-2U EXPANSION SLOTS.....	10
FIGURE 9. NETWORK LAYOUT FOR EACH POD.	15
FIGURE 10. NETWORK LAYOUT OF COMPUTE NODE.....	15
FIGURE 11. TEST TOPOLOGY OF NETWORK I/O PERFORMANCE.	17
FIGURE 12. TEST RESULT OF NETWORK I/O PERFORMANCE.	18
FIGURE 13. TEST TOPOLOGY OF MULTI-HOST DATA PLANE PERFORMANCE.	18
FIGURE 14. TEST RESULT OF MULTI-HOST DATA PLANE PERFORMANCE.	19
FIGURE 15. TEST TOPOLOGY OF MULTI-CORE N-VDS (E).	20
FIGURE 16. TEST TOPOLOGY OF N-VDS (S).....	20
FIGURE 17. TEST RESULT OF SCALABLE DATA PLAN PERFORMANCE.....	21

1. Executive Summary

In the era of technology development, the high requirement of network speed on media, mobile, and high-tech applications accelerates the transformation of network architecture. Communication Service Provider (CSPs) are facing challenges on managing complicated and legacy architecture. To overcome the insufficient network scalability, efficiency, and manageability, CSPs are now looking towards Network Function Virtualization (NFV) and recognize it as a key enabler to modernize the network functions running on standard platforms.

Network Function Virtualization (NFV) is a novel technology adopted to decouple the network functions from proprietary hardware and to run on standard platforms such as x86 servers. NFV provides abstract network functions such as virtual switch, virtual firewall, virtual router, and load balancer. VMware vCloud NFV OpenStack Edition is a fully-integrated modular which combines a carrier-grade NFV infrastructure with VMware® Integrated OpenStack. Building on top of VMware platform, this edition integrates OpenStack API and is compliant with Europe Telecommunication Standard Institution (ETSI) standard. The suite features the automation of service deployment, multi-tenancy for security, operational intelligence for monitoring the performance events, and carrier-grade networks to meet customers' expectation.

Quanta Cloud technology (QCT) provides a carrier-grade solution for CSPs to immediately deliver virtual network functions (VNFs) and launch services on NFV infrastructure. QCT conducted a series of tests to verify the solution performance and ensure the network functions. With QCT's solid knowledge and rich experience, **QxStack NFV Infrastructure for VMware vCloud NFV OpenStack Edition** is a pre-validated and performance-optimized solution for CSPs to pursue unlimited new business opportunities. Moreover, the preloaded and preinstalled services can highly shorten customers deploy time and accelerate time to value.

2. Solution Overview

QxStack NFV Infrastructure solution provides a carrier-grade NFV infrastructure which is compliant with ETSI NFV architectural framework. This solution adopts QCT Hyper-Converged Infrastructure (HCI) servers certified by VMware to deliver infrastructure hardware resources. QCT HCI servers feature Non-Uniform Memory Access (NUMA)-balanced hardware design, and provide high availability and scalability to build reliable infrastructure.

The VMware vCloud NFV OpenStack Edition platform is composed of the vSphere® virtualization platform, vSAN™, and NSX® for virtualization resources. In this platform, the two pods are designed based on the three pods, including management pod, edge pod, and resource pod. VMware® Integrated OpenStack (VIO) in the pods is utilized as the NFV Virtualized Infrastructure Manager (VIM). OpenStack API is embedded in VMware vCloud NFV OpenStack Edition platform to manage and operate the NFV infrastructure.

In addition, a series of fully-integrated software in VMware vCloud NFV OpenStack Edition platform include vRealize® Operations™, vRealize® Log Insight™, and vRealize® Network Insight™ for operational intelligence and monitoring. Enhanced Platform Awareness (EPA) is another feature in this platform to improve network performance and satisfy modern NFV requirements. The industry-leading tool with RFC2544 standard, Spirent TestCenter, is utilized to evaluate the NFV infrastructure networking performance, as shown in *Figure 1*.

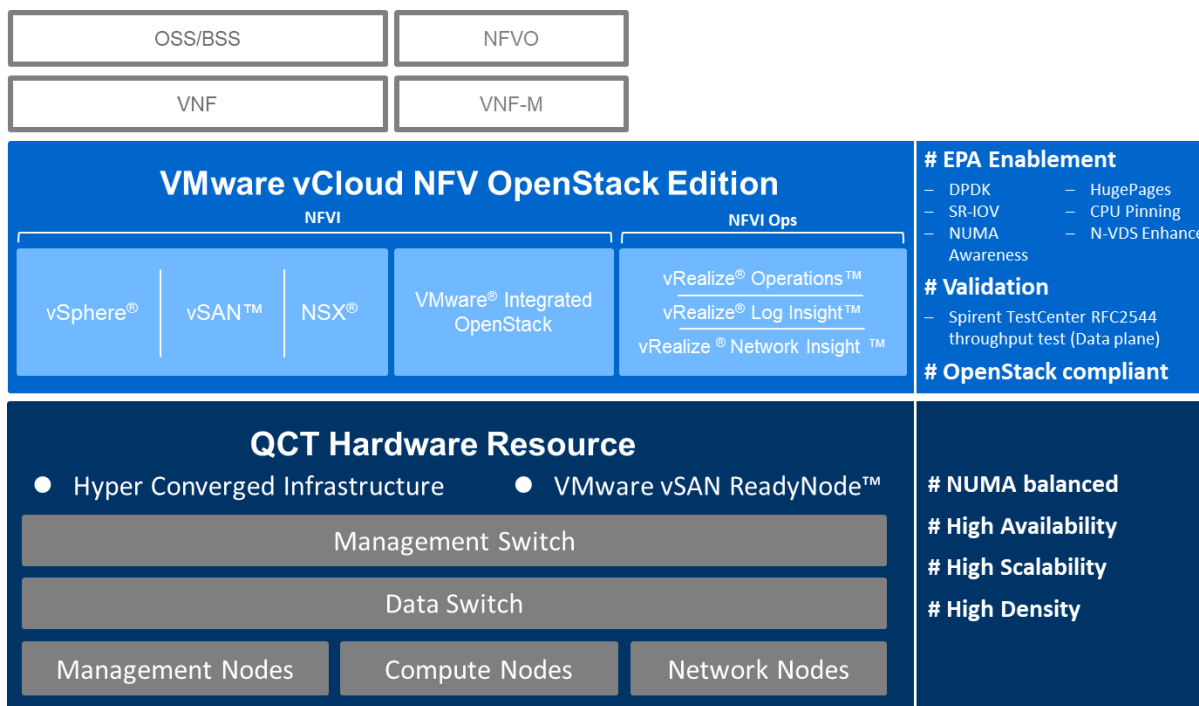


Figure 1. Overview of QxStack NFV Infrastructure Architecture

3. Solution Architecture

3.1. Three-Pod and Two-Pod Designs

In QxStack NFV Infrastructure solution, QCT hardware resources are divided into three pods, namely, management pod, edge pod, and resource pod. Each pod corresponds to one vSphere cluster and can be scaled up by adding a node. The design can separate functional blocks for better operational flexibility, as shown in *Figure 2*. The functions in the three pods can be separated for CSPs to flexibly allocate the resources according to the specific functionality of each pod. Based on the three-pod design, the two-pod design combines edge pod and resource pod into a single pod which requires fewer hosts, racks, and switches. CSP can benefit from the two-pod design for gaining flexible operational experience.

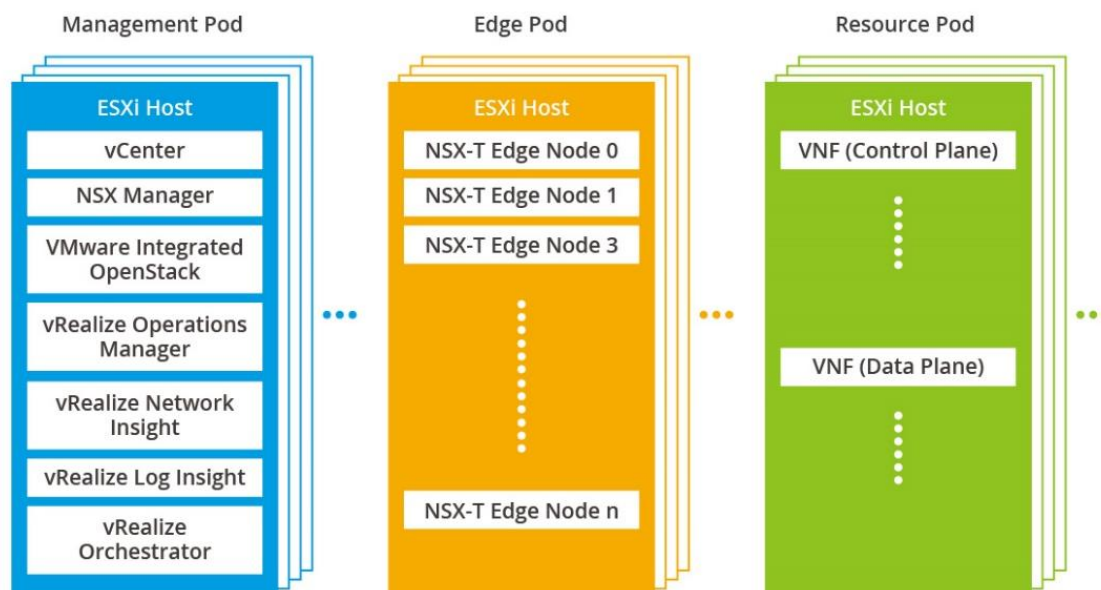


Figure 2. Virtual Building Blocks for Each Pod.

3.1.1. Management Pod

Management pod hosts the software components to centrally manage the NFV environment and orchestrate the resources. The software such as vCenter Server® Appliance™, NSX® Manager™, and VMware® Integrated OpenStack can efficiently orchestrate the resources for VNFs. The management pod also covers monitoring and analytical software, including vRealize® Operations Manager™, vRealize® Network Insight™, vRealize® Log Insight™, and vRealize® Orchestrator™.

3.1.2. Edge Pod

Edge pod hosts NSX-T Edge nodes for providing north-south traffic connectivity between VNF and external network. The NSX-T Edge can be deployed in a Virtual Machine (VM) or bare metal. The Edge node facilitates the network traffic transition between the physical domain and the virtual domain. It also provides the centralized edge services, including load balancer, Network Address Translation (NAT), and edge firewall.

3.1.3. Resource Pod

Resource pod provides virtualized resources, including compute, network, and storage for VNFs. The EPA features PCI passthrough technology and resource allocation strategy implemented in resource pod to achieve the network acceleration on VNF. Resource pod also provides virtual network services and virtualized runtime environment for network functions. In the following sections, resource pod will specifically be elaborated.

3.2. Networking Technology and Architecture

3.2.1. SR-IOV

Signal Root I/O Virtualization (SR-IOV) is a technology that allows physical PCI devices to be allocated for multiple VMs to increase network performance. SR-IOV virtualizes PCI hardware devices to create multiple virtual functions (VFs)—lightweight functions that can be assigned to specific VMs. A VF driver is required to implement SR-IOV. This driver resides in the VM, introduces 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. Although the traffic would not go through VMkernel stack, the ability to insert VLAN ID tags to each SR-IOV VF depends on the vSphere® Distributed Switch™ (vDS) port group ports (VLAN IDs) on which the VMs reside.

3.2.2. Resource Allocation Strategy

To optimize resource allocation, CPU pinning and a NUMA-aware design are adopted in the QxStack NFV Infrastructure. In the virtualized infrastructure, a pool of physical CPUs (pCPUs) in an ESXi host is shared across multiple virtual CPUs (vCPUs) associated with VMs. CPU pinning enables one-to-one mapping between vCPUs and pCPUs to increase VM performance. It could be used to eliminate the extra latency that is imposed by the virtualization. As shown in *Figure 3*, CPU pinning dedicates specific compute resources to specific VMs and increases cache efficiency.

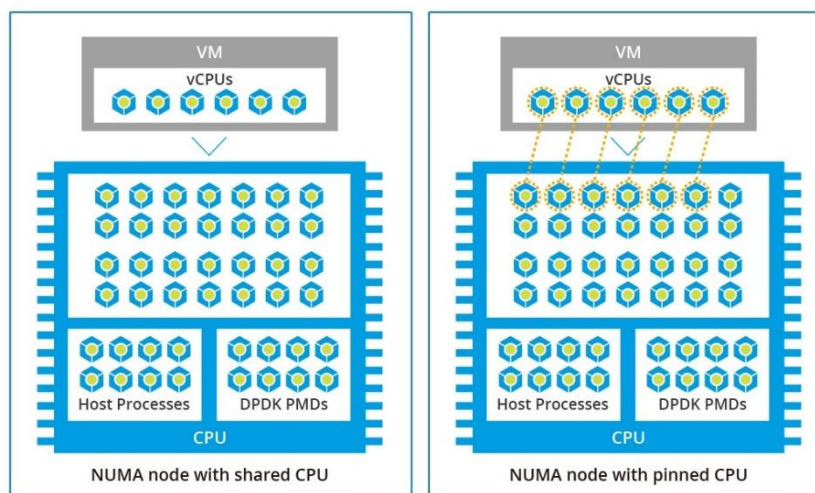


Figure 3. Comparison of CPU core allocation between shared CPU and pinned CPU.

Traditional Uniform Memory Access (UMA) architectures share memory resources evenly across all CPUs and sockets in a multiprocessor system. This often results in long memory access time, regardless of the location of the memory in relation to the CPU or socket. NUMA architecture is a system with geographically-distributed memory by considering its location in relation to each CPU, speeding the access to memory closer to the CPU.

Processes can then access local CPU's memory—rather than another CPU's local memory or shared memory—to improve computational performance. The QxStack NFV Infrastructure uses a NUMA-balanced design to support local memory access and distributes NICs across CPUs and sockets, as shown in *Figure 4*. To ensure the optimal networking performance, the completely-vertical NUMA alignment is a prerequisite. That is, CPU, memory, and physical NIC should be aligned in the same NUMA-node boundary.

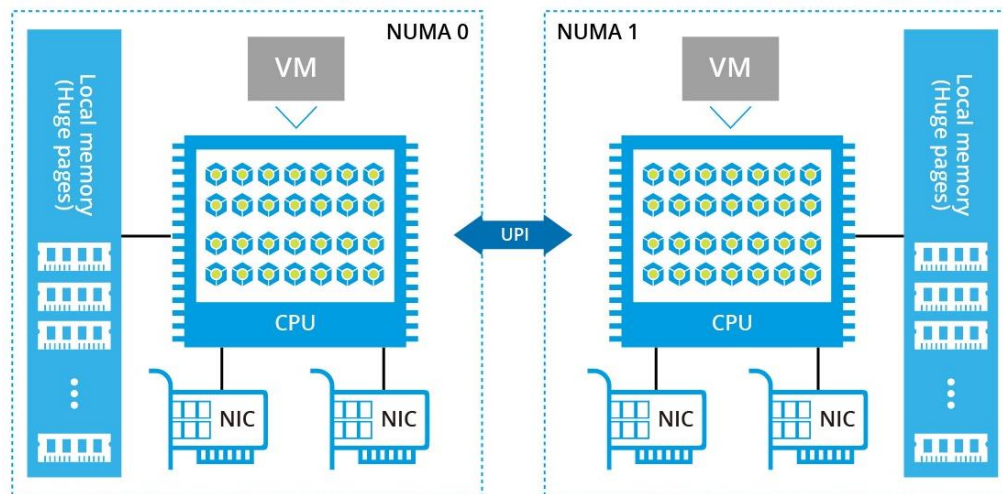


Figure 4. NUMA-balanced design for QxStack NFV Infrastructure.

3.2.3. Network Design Concept

QxStack NFV Infrastructure provides consistent network configuration of infrastructure across multiple hosts in each pod. The traffic for infrastructure network services such as ESXi management, vMotion network, VM management, and vSAN network are isolated using different VLAN. All these infrastructure network traffics run on the vDS to provide consistent network configuration across pods. vDS implements the VLAN tagging for different port groups to logically separate the traffic for infrastructure network services.

The configuration of infrastructure network services running on vDS do not change across each pod while the configuration of VM networks varies based on the requirement of each pod. NSX-T consists of two types of vDS, N-VDS (Standard) and N-VDS (Enhanced Data Path), for different traffic types in VM network. N-VDS (S) provides VLAN and Geneve-based overlay tunnels network for the connectivity between VMs while the N-VDS (E) provides enhanced networking service for data plane intensive workload in the resource pod. *Figure 5* shows the virtual network design across each pod.

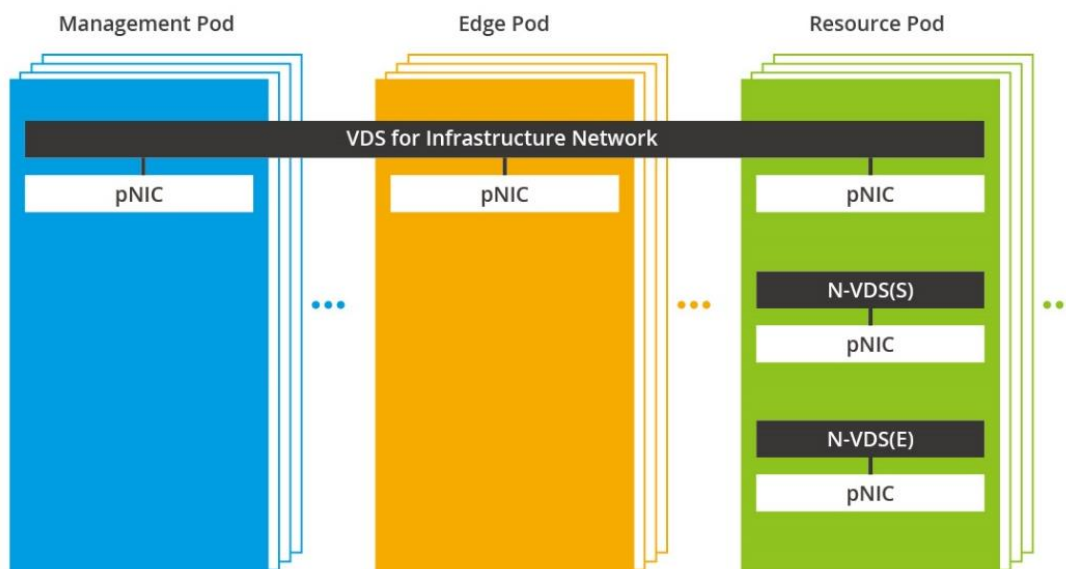


Figure 5. Virtual network topology across each pod.

A well-designed network topology ensures that the traffic for infrastructure network services and VMs can run efficiently, correctly, and available between hosts. In order to ensure network availability for the uplink connectivity on vDS or N-VDS, NIC teaming is a way to bind two NICs together. The two uplink adapters, vmnic1 and vmnic2, are teamed on a vDS or N-VDS to either two active uplinks or an active/standby uplink for failover and redundancy purposes.

3.2.4. Network Design for Management Pod

Management pod contains a vDS to provide the network connectivity for infrastructure network service and VM network. The distributed port group for each infrastructure network service consists of ESXi management, vMotion network, VM management, and vSAN network in management pod. Network traffic for different services is separated by using VLAN tagging on the vDS port group. To increase the availability of networking, teaming and failover policy can be applied to a distributed port group.

3.2.5. Network Design for Edge Pod

The NSX-T Edge Nodes hosted in the edge pod provide the north-south network connectivity for the VNF workloads. Before installing NSX Edge Nodes, three port groups named Management Network, Overlay Network, and External Network are created on vDS to connect the Edge VM. The port groups associated with Overlay Network and External Network are VLAN trunking type. An NSX-T Edge VM has four internal interfaces, namely, eth0, fp-eth0, fp-eth1, and fp-eth2. The eth0 interface is connected to vDS Management Network for managing host. The fp-eth0 interface is connected to Overlay Network, providing overlay network connectivity between VMs. The fp-eth1 and fp-eth2 interfaces are connected to External Network port group, providing an external uplink to the physical network. Moreover, the infrastructure network service in edge pod runs on the same vDS in the management pod to provide a consistent infrastructure network configuration across hosts.

3.2.6. Network Design for Resource Pod

NFV workloads in resource pod are separated into control plane and data plane. Control plane refers to all the functions and processes that determine which network route to use. Data plane refers to all the functions and

processes that forward packets from one interface to another. The network design in resource pod is highly dependent on the VNF network topology, particularly for the data plane networking acceleration.

Data Plane Development Kit (DPDK) consists of a set of data plane libraries and user-space network drivers for accelerating packet processing. Based on the programmable framework, administrators can implement a run-to-completion model, eliminate packet interrupt processing overhead, and enable applications to perform packet processing directly from and to the NIC. In order to accelerate package processing, N-VDS (E) in the resource pod is leveraged to implement DPDK technology and efficiently support data plane intensive VNFs. N-VDS (E) provides key DPDK features such as Poll Mode Driver (PMD), flow cache, and optimized packet copy for various packet sizes. The data plane network traffic also leverages SR-IOV for VM to communicate directly with the physical device, bypassing the software switching layer to achieve near line rate performance.

N-VDS (S) can provide the generic network connectivity such as overlay and VLAN network between control and management plane workloads. The remaining VMkernel traffic and VM management traffic for the infrastructure network service in resource pod share the same vDS in management pod.

3.3. Solution Hardware Suggestion

3.3.1. QCT HCI Server




HCI is a novel technology which can highly integrate compute, storage, and network in a single hardware box. Every single node is capable of delivering compute and storage resources at the same time. When a single node fails, VMware vSAN™ on HCI servers can be utilized to handover the workloads to other nodes. VMware vSAN™ is a mature software-defined storage technology for hyper-converged solutions. Uniquely embedded in the hypervisor, vSAN™ delivers high-performance, flash-optimized, easily-scaled hyper-converged storage for any virtualized applications. By adopting VMware vSAN™, the high availability design for both compute and storage resources can be easily implemented on the hosts.

To ensure the performance and stability, QCT strictly selects hardware components, firmware and driver, and software stack to guarantee the compatibility between servers and software. Three QCT servers that have passed VMware vSAN ReadyNode™ certification and listed on VMware Compatibility Guide¹ are recommended for VMware vCloud NFV OpenStack Edition, as shown in *Table 1*. With the certified infrastructure, administrators can save considerable time on verifying hardware performance and compatibility.

¹ Visit VMware Compatibility Guide for more QCT vSAN Ready servers.

https://www.vmware.com/resources/compatibility/search.php?deviceCategory=vsan&details=1&vsan_type=vsanreadynode&vsan_partner=114&page=1&display_interval=10&sortColumn=Partner&sortOrder=Asc

Table 1. Solution hardware – Quanta vSAN Ready series.

	Resource Pod and Edge Pod		Management Pod
Server Model	QuantaGrid D52B-1U	QuantaGrid D52BQ-2U	QuantaPlex T42S-2U
			
Form Factor	1U	2U	2U
CPU per Node	2 x Intel® Xeon® Scalable Processors Platinum/Gold/Silver Series		
Memory	Up to 24 slots		Up to 16 slots per node
Driver Bays	12 x 2.5" hot-plug (NVMe support)	24 x 2.5" hot-plug 12 x 3.5" hot-plug (NVMe support)	(24) 2.5" hot-plug
vSAN Profile ²	AF – 4,6,8 HY – 2,4,6,8	AF – 4,6,8 HY – 4,6,8	AF – 4,6 HY – 4,6

Hardware for Management Pod

To select suitable servers for building management pod, the redundant design to tolerate the failures of compute, storage, and networking is inevitable. QCT multi-node server, QuantaPlex T42S-2U, is a 2U server with four nodes in one chassis. Compared to regular 1U rackmount servers, this 2U server provides double compute density which makes it an ideal choice to build the HCI solution. From the VMware® vSAN™ Design and Sizing Guide³, four or more nodes in a vSAN cluster can provide great flexibility to accommodate storage failure. QuantaPlex T42S-2U with four-node design is easy to implement a vSAN cluster in just one server chassis which makes it an ideal choice to build the management pod. The server is also equipped with 2 PCIe expansion slots to adopt physical NICs which can be used for NIC teaming to achieve network redundancy, as shown in Figure 6.

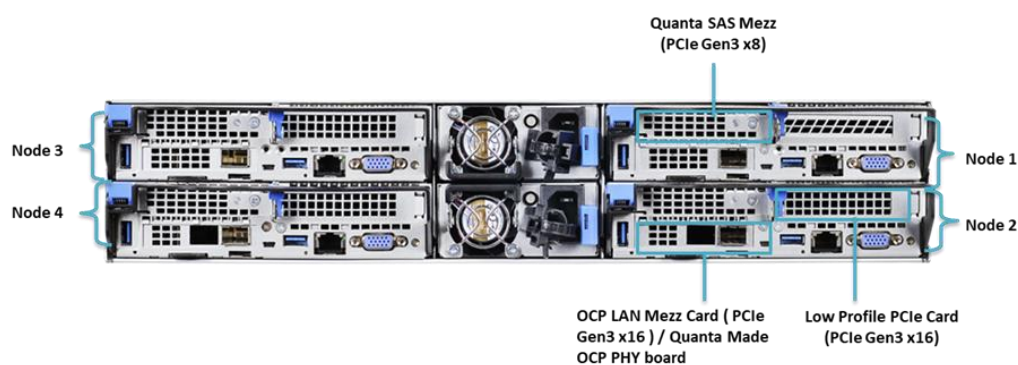


Figure 6. Overview of QuantaPlex T42S-2U expansion slots.

² Visit the link for vSAN Hardware Quick Reference Guide

https://www.vmware.com/resources/compatibility/vsan_profile.html?locale=en

³ Visit the link for VMware® vSAN™ Design and Sizing Guide

<https://storagehub.vmware.com/t/vmware-vsantm-design-and-sizing-guide-2/>

Hardware for Edge Pod

The NSX-T Edge node can be deployed in either a VM or bare metal form factor in the edge pod. Compared to the VM form factor, NSX-T Edge node running on bare metal form factor in a production environment can provide faster failover time and greater throughput. When NSX-T Edge node is deployed in bare metal, the use of NIC and CPU should fulfill the requirements of NSX-T according to NSX® Edge™ and Bare-Metal NSX® Edge™ Physical Hardware Requirements⁴. At least three physical NICs are suggested to be installed when the bare metal Edge node is deployed. One dedicated NIC is retained for management while the other two NICs are respectively allocated for overlay tunneling traffic and VLAN uplink connectivity. The server with multiple PCI Express (PCIe) expansion slots for physical NIC can ensure flexibility and redundancy for uplink assignment. QuantaGrid D52B-1U is an ultra-dense server with up to 4 PCIe expansion slots per chassis which can fulfill the Edge node uplink assignment, as shown in *Figure 7*. In order to provide better flexibility and redundancy on the uplink, D52BQ-2U provides three additional PCIe expansion slots, as shown in *Figure 8*.

Hardware for Resource Pod

To achieve networking acceleration, the principle for building the resource pod infrastructure is aimed at hosting intensive workloads in data plane. The three features including NUMA-node design, NIC with DPDK support, and the number of CPU core and cycle are taken into consideration.

NUMA Node Design

Modern high-performance servers typically have more than one NUMA nodes in a single chassis. In order to improve the networking performance, the resources including physical CPUs, memory, physical network interfaces, and virtual machines should be vertically allocated to the same NUMA node.

Choice of NIC

The capability (e.g., offloading CPU and overlay traffic) and performance (e.g., PCIe speed and throughput) of physical NIC can influence a data plane intensive workload. Most servers today support multiple PCIe slots to accommodate various ethernet NICs. The more the number of PCIe expansion slots is, the more flexible for uplink on N-VDS (S) and N-VDS (E) is. With multiple PCIe expansion slots, the passthrough technologies such as DirectPath I/O and SR-IOV acceleration can provide relatively more passthrough functions. In addition, the NIC used for N-VDS (E) must use the driver to support N-VDS Enhanced Data Path. To identify the driver with N-VDS Enhanced Data Path support for uplink NIC on N-VDS (E), please refer to VMware Compatibility Guide for more detail.

CPU Speed and Density

Choosing CPU type in resource pod is a critical factor to influence the overall networking performance. Packet processing and network traffic rely heavily on CPU clock cycles and CPU core count in the ESXi host. These CPU resources are used for the data plane intensive workload to perform its task and for the hypervisor layer to transport network traffic from the physical NIC to the VNF.

⁴ NSX-T according to NSX® Edge™ and Bare-Metal NSX® Edge™ Physical Hardware Requirements

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.3/com.vmware.nsx.install.doc/GUID-14183A62-8E8D-43CC-92E0-E8D72E198D5A.html>

QuantaGrid D52B-1U is an ultra-dense server with up to 4 PCIe expansion slots per chassis which makes it an ideal choice to build the resource pod. The server features flexible I/O options, including a variety of SAS Mezzanine and OCP NIC/ PHY Mezzanine cards in three different SKUs. With NUMA-balanced design, QuantaGrid D52B-1U supports data plane acceleration technologies like DPDK and SR-IOV to accommodate the most demanding Telco workloads. This 1U server with Intel® Xeon® Processor Scalable Family provides fast socket interconnection and strengthens memory bandwidth to fulfill the requirement of Telco demand.

In order to accommodate different speeds of the NIC for data plane workload, QuantaGrid D52BQ-2U features 3 additional PCIe expansion slots. The 2U chassis provides up to 7 PCIe expansion slots with different PCIe lane widths which can accommodate more PCIe NICs for passthrough technology and uplink redundant design. Moreover, QuantaGrid D52BQ-2U also provides 3 Intel® Ultra Path Interconnect (UPI) to fast interconnect the socket and fulfill the NUMA balance design.

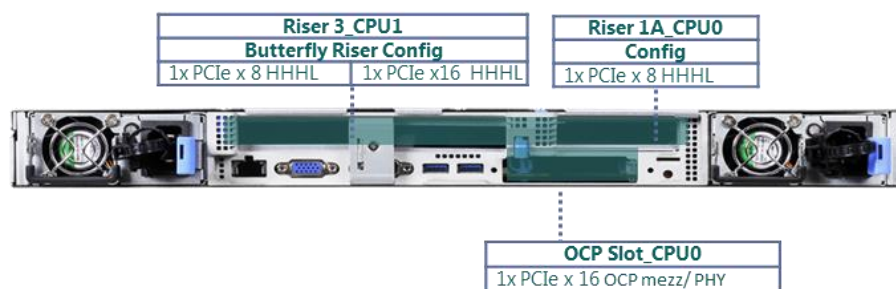


Figure 7. Overview of QuantaGrid D52B-1U expansion slots.

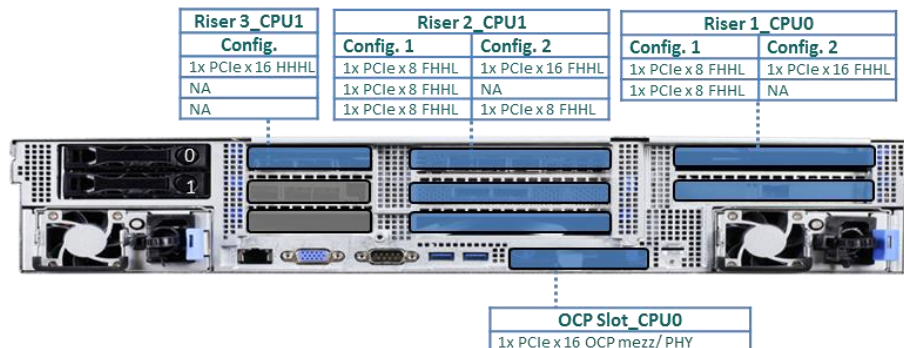


Figure 8. Overview of QuantaGrid D52BQ-2U expansion slots.

3.3.2. QCT QuantaMesh Switch

To design an impeccable NFV infrastructure, the selections of management switch and data switch are also critical. QuantaMesh T1048-LY4R and T4048-IX8D are respectively recommended for management switch and data switch. The hardware configuration is shown in *Table 2* below.



Management Switch

The QuantaMesh T1048-LY4R is a new generation of 1GBASE-T solution for data center rack management, which provides 48 triple speed (10/100/1000Base-T) ports and 4 SFP+ (1/10Gbps) ports in a 1U form factor. Equipped with Baseboard Management Controller (BMC), T1048-LY4R provides IPMI protocol via out-of-band management port (10/100/1000BASE-T). T1048-LY4R also supports out-of-band management port via 1000BASE-X SFP port for long distance support. Two-PSU SKU provides redundancy feature while one-PSU SKU provides best TCO.

Data Switch

QuantaMesh T4048-IX8D is a layer 3 managed switch with high-performance, high port density, and low latency for Data Center Leaf Switch or Top of Rack (ToR) application. T4048-IX8D supports 48 SFP28 and 8 QSFP28 ports and is equipped with BMC in a 1U switch that supports 2.0 Terabits per second (Tbps) bandwidth. The 48-port downlink ports on the T4048-IX8D can be configured to work as 1/10/25-Gbps ports, offering deployment flexibility and investment protection. The uplink can support up to eight 100-Gbps or eight 40-Gbps ports, or a combination of 10-, 25-, 40-, and 100-Gbps connectivity, offering flexible migration options. By leveraging merchant silicon chip, T4048-IX8D is fully compliant with IEEE 802.3by, and is a high performance and high density Ethernet switch with advanced features such as large table size and large packet buffers.

Table 2. Solution hardware – QuantaMesh switch series.

	Data Switch	Management Switch
Switch Model	QuantaMesh T4048-IX8D 	QuantaMesh T1048-LY4R 
Feature	The Next Wave 25G Ethernet Switch for Data Center and Cloud Computing <ul style="list-style-type: none"> • 25G & 100G Ethernet Switch • VXLAN • Multi-Chassis Link Aggregation (MLAG) • OSPF, BGP4 with ECMP • Network Automation 	The Next Wave Data Center Rack Management Switch <ul style="list-style-type: none"> • ONIE Pre-load • x86 CPU Design • BMC Built-in • IPMI Management Support • Optical OOB Port Support
Chip set	Broadcom Trident3	Broadcom Hurricane2
Port configuration	48 25G SFP28 and 8 100G QSFP28	48 10/100/1000BASE-T and 4 1/10GbE SFP+ ports
Management Port	Out-of-band management port: RJ-45, 10/100/1000Base-T	Out-of-band management port: 1 x RJ-45 (10/100/1000BASE-T) 1 x SFP (1000BASE-X)
Switching capacity	4Tbps	176Gbps
Maximum forwarding rate	2Bpps	131Mpps

3.4. Solution Software

3.4.1. VMware vSphere®

VMware vSphere® is an industry-leading virtualization platform which virtualizes and aggregates the physical hardware to provide a virtual resource pool to the VNF component. By leveraging the virtualization technology, vSphere® can provide a highly available, efficient, and centralized infrastructure for IT administrators to deliver flexible and reliable services. VMware vSphere® provides some benefits:

- Using proactive high-availability technology to prevent machine downtime.
- Using predictive load balancing technology to fully exert the datacenter resources.
- Simplifying user experience to deliver a large-scale automation and management.
- Leveraging virtual machine level encryption technology to reduce risk.
- Using REST API to promote IT automation and business flexibility.

VMware vSphere® is composed of ESXi™ and vCenter® which will be explicated in detail below.

3.4.2. ESXi™

VMware ESXi™ is an industry-leading hypervisor, installed on a bare-metal physical server. ESXi™ has its own kernel, called VMkernel based on Linux kernel. ESXi™ enables the virtualization technology to break the traditional hardware architecture. By sharing the resources of a single hardware across multiple environments, a physical system (x86 server) is capable of executing multiple VNFs. VMware proposed that “virtualization is the process of creating a software-based (or virtual) representation of something rather than a physical one.” Virtualization can be applied to applications, servers, storage, and networks. It is an effective way to reduce IT expenses and boost efficiency and agility for all-sized businesses.

3.4.3. VMware vCenter Server®

vCenter server® is a centralized management and operation platform for QxStack NFV Infrastructure environment. By aggregating all virtual resources, vCenter server® can execute resource provisioning and monitoring. It empowers the capability to provision compute, storage, and other resources to the VNF component and enables High Availability (HA), Distributed Resource Scheduler (DRS), vMotion®, etc.

3.4.4. vSAN™

vSAN™ is a software-defined storage built in the vSphere kernel. It is tightly integrated with the hypervisor to minimize the CPU and memory overhead and optimize the data I/O path to deliver an outstanding performance. vSAN™ is a hyper-converged solution, particularly designed for virtual machines. It minimizes the effort to configure the storage and simplifies the virtual machine deployment in the NFVi virtualization layer. By using the virtualization technology, the underlayer physical disk is a concept of abstract resource pool, aggregated into a virtual resource for providing a shared storage between hosts in the vSphere® cluster.

3.4.5. VIO

VMware® Integrated OpenStack (VIO) is a distribution of OpenStack running on vSphere® environment to simplify installation and operation. From a single and programmable interface, VIO is implemented as compute clusters to handle tenant workloads and as management clusters to handle OpenStack components (e.g., Horizon, Keystone, Nova, and Neutron) and other services (e.g., load balancing and database, DHCP).

In this reference architecture, VIO 5.1 based on the OpenStack Queens release is adopted. By adopting vCenter Server® Appliance™ and NSX-T, it is capable of orchestrating compute, storage, and networking. vCenter Server® Appliance™ orchestrates compute and storage resources while NSX-T is responsible for tenant networking. VIO provides administrators with easy-management OpenStack cloud on VMware infrastructure which highly boosts developer productivity and flexibility.

3.4.6. NSX-T

NSX-T is a technology used to communicate between VNF components and to dynamically control network service environments. NSX-T contains NSX-T Manager and NSX-T Controller. NSX-T Manager is a management plane for NSX-T to create, configure, and monitor NSX-T components (e.g., logical switch and NSX edge nodes). By monitoring and troubleshooting network, NSX-T can configure and orchestrate logical network components such as logical switching and routing, and security services. NSX-T Controller is a management system used to

control virtual network and overlay transport tunnels. It serves as a cluster to programmatically deploy virtual networks across the overall NSX-T architecture. NSX-T can facilitate CSPs to programmatically create, delete, and manage software-based virtual network.

3.4.7. vRealize® Operations Manager™

vRealize® Operations Manager™ is an intelligent appliance resides in the NFV infrastructure which can used to monitor platform and collect data including health, capacity, availability, and performance. According to the collected data, the resources can be automatically optimized and planned. vRealize® Operations Manager™ can proactively identify and remedy the issues, and continuously enforce the standard of compliance. From a single console, the visibility is comprehensive across VMs between virtualized and physical infrastructures. An API can retrieve performance and health data for both NFV infrastructure and VNF virtual resources. With vRealize® Operations Manager™, CSPs can adequately plan and monitor the solution environment based on the existing infrastructure.

3.4.8. vRealize® Log Insight™

vRealize® Log Insight™ is an appliance used to automatically identify and collect massive amounts of log data and to build an index for performance analysis. The API of the appliance provides administrators to programmatically access to the functionality and its datastore. vRealize® Log Insight™ is integrated to MANO for monitoring the events and logs. Different VM sizes are provided for CSPs to choose according to the requirements of the configuration; hence, the resources of compute and storage can be effectively managed. By adopting vRealize® Log Insight™, administrators can easily collect logs, organize unstructured logs, troubleshoot vSphere®, audit security, etc.

3.4.9. vRealize® Network Insight™

vRealize® Network Insight™ provides intelligent operations for software-defined network and security which can be installed in either two-pod or three-pod designs. It can constantly monitor and audit the compliance status. With micro segmentation planning, the micro segmentation deployment can be accelerated based on the firewall rules recommendation. To simply the NSX® operation, the appliance can be adopted to scale NSX® across multiple NSX® Manager™ and avoid configuration issues. Through intuitive UI and language search, administrators can rapidly identify and troubleshoot issues. With the powerful network visibility, administrators can also troubleshoot the connectivity between VMs across virtual and physical network.

4. Solution Configuration

4.1. Hardware Configuration

In this reference architecture, the two-pod design is adopted for QxStack NFV Infrastructure since the resource requirement is relatively small, compared to the three-pod design. The hardware configuration consists of four management nodes in the management pod, and three compute nodes in the edge and resource pod. In the management pod, Intel® Xeon® Silver series processor is selected to provide sufficient processing power for the resource orchestration while in the resource and edge pod, Intel® Xeon® Gold series processor is chosen to provide outstanding performance on packet processing and packet forwarding.

Three QuantaGrid D52BQ-2U servers with DPDK and SR-IOV enabled are configured as a vSAN cluster, providing network traffic acceleration for data plane intensive workload. In addition, each of the vSAN cluster is equipped with 2 disk groups, including one cache tire drive and three capacity tier drives. The selected servers are detailed in *Table 3*.

Table 3. Hardware configuration for two-pod design.

	Edge and Resource Pod	Management Pod
Server Model	QuantaGrid D52BQ-2U	QuantaPlex T42S-2U
Specification per Node	<ul style="list-style-type: none"> - CPU: Intel® Xeon® Scalable Processors Gold series - RAM: 384 GB - Cache: Intel® S4610 960GB x 2 - Capacity: Seagate ST1800MM0129 1.8TB x 6 - NICs: Intel X710 dual port DA2 x 2 	<ul style="list-style-type: none"> - CPU: Intel® Xeon® Scalable Processors Silver series - RAM: 384 GB - Cache: Intel® S4610 960GB x 1 - Capacity: Seagate ST1800MM0129 1.8TB x 5 - NIC: Quanta OCP Mezz 82599 dual port x1
Server Quantity	4	1
Node Quantity	4	4

4.2. Network Planning

A well-designed network topology ensures the efficiency, correctness, and availability of communication among running services. The infrastructure network services are configured to be distributed port groups including ESXi management, vMotion network, VM management, and vSAN network in each pod. To separate the network traffic for different services, VLAN tagging is configured in the distributed port group.

QxStack NFV Infrastructure leverages NSX-T to manage the software-based networks, static and floating IP addresses, and the DHCP service. To accelerate the networking for data plane intensive workload, passthrough and DPDK acceleration technologies are implemented in the edge and resource pod to optimize the overall data plane path.

4.2.1. Logical Network Topology

Each pod implements different networking functions according to the functionality design. Different network services are assigned for each pod to easily manage and optimize system resource utilization. In QxStack NFV Infrastructure, the logical network topology for each pod is illustrated in *Figure 9*.

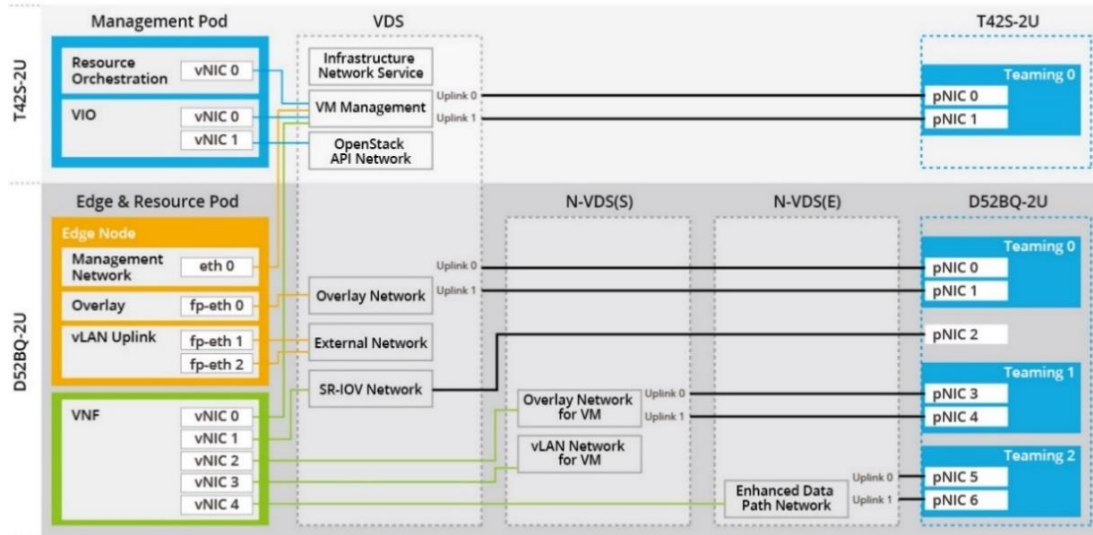


Figure 9. Network layout for each pod.

4.2.2. Logical Network Layout in Compute Nodes

To separate the network resources and traffic, control plane and data plane resources are associated with different NUMA nodes. Control plane intensive VNFs are allocated to NUMA node 0 while data plane intensive VNFs are allocated to NUMA node 1, as shown in *Figure 10*. NUMA affinity should be configured to ensure that the VNF resources remain within the NUMA boundary to achieve optimal performance.

N-VDS (S) in resource pod is used to facilitate control plane and overlay network traffic. In NUMA node 0, a teaming policy is configured on two dual-port 10G NICs as redundant uplinks to an N-VDS (S). In NUMA node 1, two networking acceleration technologies, SR-IOV and DPDK, are adopted.

From the perspective of the server, two NICs which support N-VDS Enhanced Data Path drivers are configured as redundant uplinks to an N-VDS (E). By connecting the VNF to N-VDS (E), vertical alignment is established among VNF, logical CPU core assigned to the NIC, and N-VDS (E) uplink on the same NUMA node. The rest one dual-port 10G NIC is configured as SR-IOV to provide a virtual function for data plane intensive VNFs

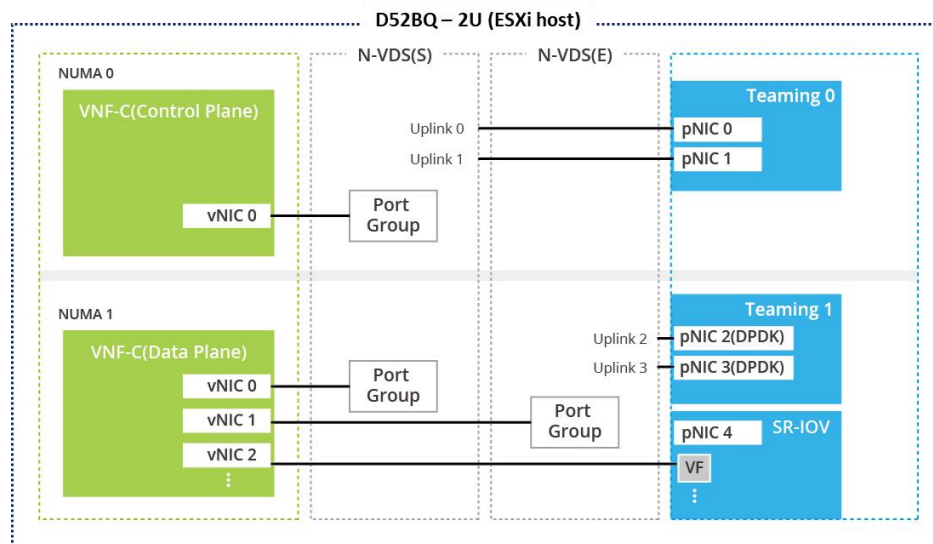


Figure 10. Network Layout of Compute Node.

5. Solution Performance

5.1. Data Plane Performance

With the release of NSX-T 2.2, N-VDS (E) is introduced. The new networking stack of N-VDS leverages DPDK technology, bringing the network performance of data plane into a new milestone and fully emancipating the CPU power to boost the I/O throughput in terms of network. Further, the release of NSX-T 2.3 incrementally upgrades the manageability and performance. In this reference architecture, QCT QxStack NFV Infrastructure adopts NSX-T 2.3. In the following sections, the purpose and the result of the three test cases will be introduced.

5.1.1. Test Scenario Brief

There are three test cases conducted in this reference architecture. First, the I/O performances of three network interfaces, SR-IOV, N-VDS (S), and N-VDS (E), are analyzed on QxStack NFV Infrastructure. Second, the performance of N-VDS (E) is evaluated in the multiple VNFs within multiple hosts to simulate the data plane performance in the real-world use cases such as service function chaining. Third, the scalability of N-VDS (E) is evaluated since flexible resource arrangement is critical to optimize the performance of dynamic workloads.

5.1.2. Test Method and Tools

In QxStack NFV Infrastructure, the data plane performance is evaluated according to the industry-standard RFC 2544. The RFC 2544 benchmarking methodology is designed for the measurement of network performance. Widely adopted by NFV community, the RFC 2544 benchmarking methodology is applied to NFV test tools and test framework (e.g., OPNFV's NSB, NFVbench, VSperv, and Yardstick) and is used to measure fundamental network performances, including network packet throughput, packet loss ratio, packet delay (latency), and delay variation (jitter). Spirent Communication, an industry-leading provider of network test solution, has been developing test solutions for networks such as SDN, NFV, IOT, and 5G. In this reference architecture, Spirent TestCenter is selected to be the testbed to execute the RFC 2544 benchmarking test.

5.2. Solution Performance Evaluation

5.2.1. Network I/O Performance

In QxStack NFV Infrastructure, the three major network interfaces, N-VDS (S), N-VDS (E), and SR-IOV are utilized to build management plane network, control plane network, and data plane network. Among the three networks, the N-VDS (E) and SR-IOV can be accelerated by using DPDK and can be configured with dedicated resources to ensure expected performance. Hence, these two network types are perfectly utilized in data plane intensive workload which requires low network latency in Control and User Plane Separation (CUPS) scenarios such as vEPC and vBNG. Compared with N-VDS (E) and SR-IOV, the N-VDS (S) shows rather low network throughput and latency; nevertheless, it can still provide benefit to advanced network services, network security, and multi-cloud networking integration.

To measure and compare network throughput among the three network interfaces, the test topology is configured, as detailed in *Figure 11*. In this test topology, the Spirent TestCenter is used to perform the network

throughput test in physical-to-virtual-to-physical (PVP) packet path and to verify the three network performances in north-south data plane. The loop VM (DPDK L3FWD) and Spirent TestCenter configurations are respectively shown in *Tables 4 and 5*.

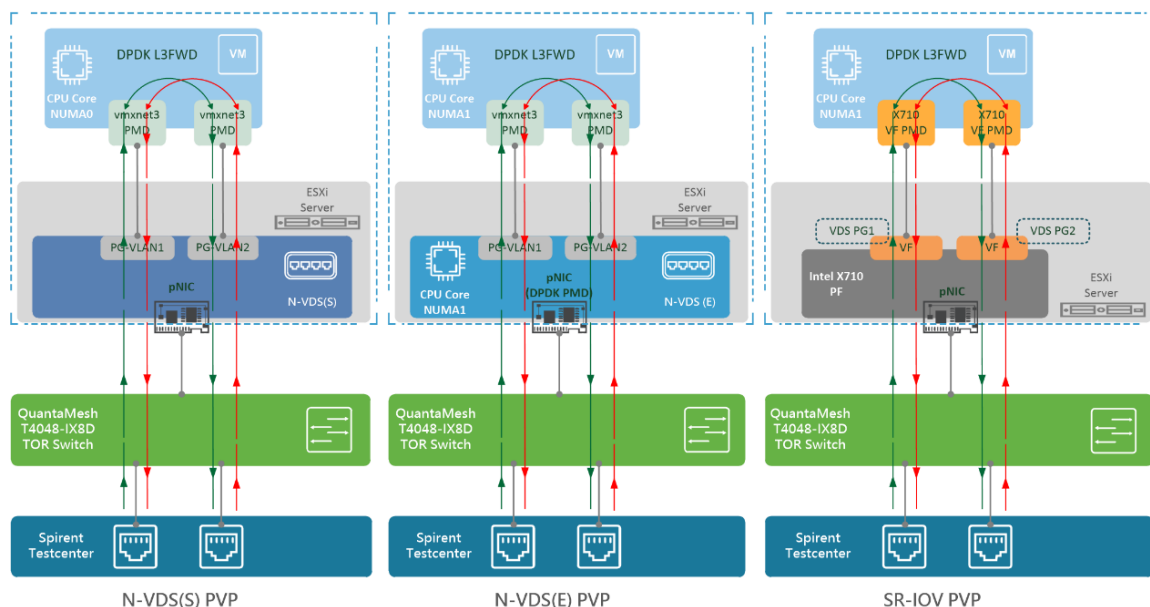


Figure 11. Test topology of network I/O performance.

Table 4. Test configuration of network I/O performance.

Item	Specification
VNF (DPDK L3FWD)	
DPDK version	17.11.4 (LTS)
NIC (N-VDS (S))/(E))	VMXNET3 Paravirtualized NIC
Driver (N-VDS (S))/(E))	VMXNET3 PMD driver
NIC (SR-IOV)	Intel i40evf
Driver (SR-IOV)	i40evf PMD
Configuration	Total 4 logical cores, 4 queues for 2 ports
Spirent TestCenter	
TestCenter version	4.96
Interfaces	Two 10Gbps ports
Rate upper limit (%)	100
Frame size (bytes)	64, 128, 256, 512, 1024, 1280, 1518
Acceptable frame loss (%)	0

In PVP test topology, the test result shown in *Figure 12* reveals that SR-IOV achieves over 97% offered load in most of the cases except for the one in 64-byte packet size. The throughput of SR-IOV can reach the line rate in almost all packet sizes while N-VDS (E) reaches the line rate only when the packet size is larger than 512 bytes. Even though the N-VDS (E) is not as ideal as SR-IOV, it still improves the performance 5.15, 4.14, and 2.32 times respectively in 64, 128, and 256 bytes packet sizes, compared to N-VDS (S).

This result illustrates that SR-IOV reaches theoretical maximum throughput for data plane intensive workload. The performance of N-VDS (E) is similar to that of SR-IOV. Moreover, N-VDS (E) inherits the benefits from NSX-T such as security, traffic monitoring, and resilience manageability. Therefore, N-VDS (E) is considered to be the best choice for workload acceleration in virtualized network environment.

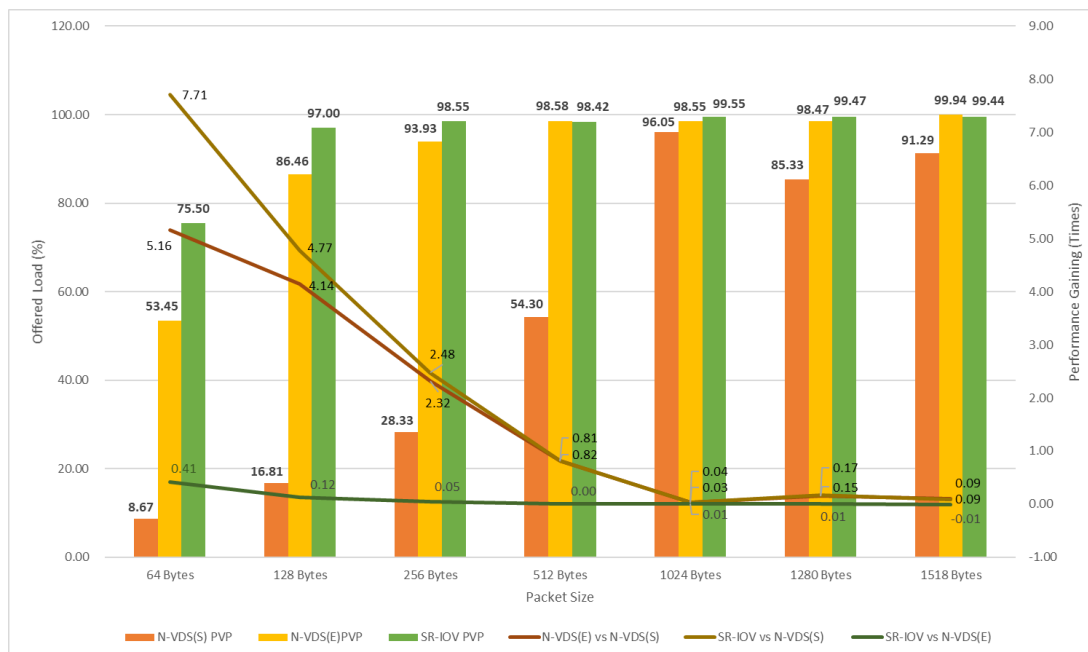


Figure 12. Test result of Network I/O performance.

5.2.2. Multi-Host Data Plane Performance

In this section, one multi-host test topology is built to test the packet path across multiple hosts and evaluate the throughput and latency in data plane across multi-host environment. This test simulates the packet path that passes through multiple VNFs within multiple hosts such as service function chaining in the real SDN/NFV environment, as shown in *Figure 13*. To ensure sufficient resource and bandwidth in the entire packet path, additional N-VDS (E) instances and physical NIC are allocated for the interconnect network in ESXi™ host. The traffic passes through the N-VDS (E) to DPDK L3FWD virtual machine. It is assumed that the throughput in this topology might be similar to the throughput of N-VDS (E) in PVP packet path.

The configurations of DPDK L3FWD virtual machine and Spirent TestCenter are respectively shown in *Table 6* and *7* below. Since the packet path passes back and forth through both physical and virtual networks, the zero packet loss is inevitable; thus, the tolerance packet loss is changed from 0% to 0.0001% in Spirent TestCenter.

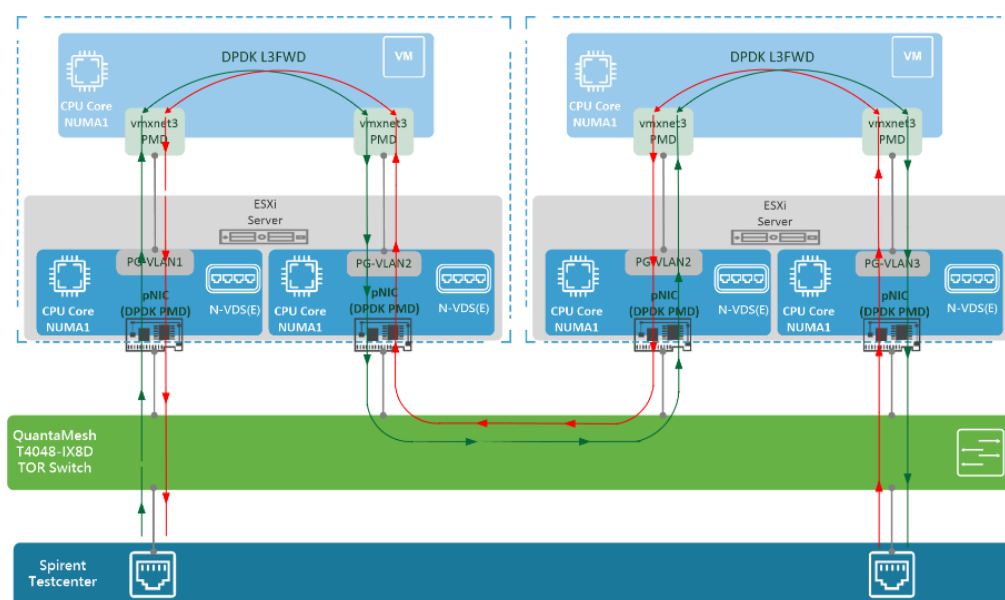


Figure 13. Test topology of multi-host data plane performance.

Table 5. Test configuration of multi-host data plane performance.

Item	Specification
VNF (DPDK L3FWD)	
DPDK version	17.11.4 (LTS)
NIC (N-VDS(E))	VMXNET3 Paravirtualized NIC
Driver (N-VDS(E))	VMXNET3 PMD driver
Configuration	Total 4 logical cores, 4 queues for 2 ports
Spirent TestCenter	
TestCenter version	4.96
Interfaces	Two 10Gbps ports (bidirectional traffic)
Rate upper limit (%)	100
Frame size (bytes)	64, 128, 256, 512, 1024, 1280, 1518
Acceptable frame loss (%)	0.0001
Latency measurement method	LIFO (Forwarding Delay)

In multi-host packet path test, N-VDS (E) shows outstanding performance in 1518 bytes packet size at 50% offered load where the latency is merely 0.157ms, as shown in *Figure 14*. Nevertheless, the latency of N-VDS (E) in multiple hosts in rest of the test cases are considerably similar to that of N-VDS (E) in single host. On the other hand, even though the throughput of N-VDS (E) in multiple hosts is slightly lower than the throughput of N-VDS (E) in single host in 128 and 256 bytes packet sizes, the test result shows the consistency of performance in the two scenarios.

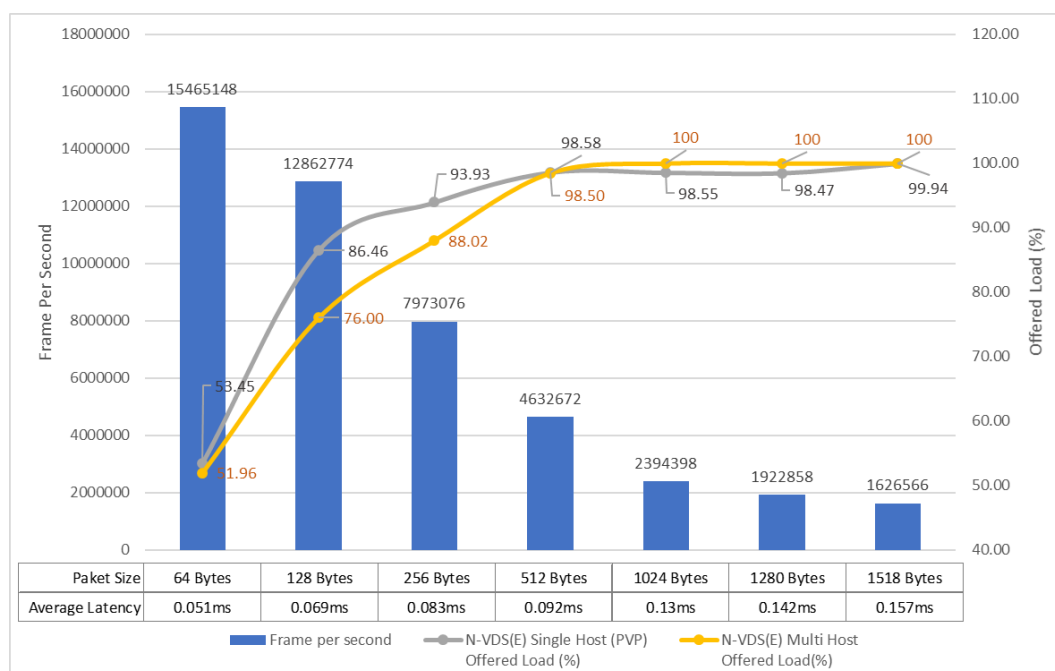


Figure 14. Test result of multi-host data plane performance.

5.2.3. Scalable Data Plane Performance

QxStack NFV Infrastructure leverages the NSX-T Enhanced Data Path technology. By allocating different number of CPU logical cores to the N-VDS (E), the performance of network throughput in data plane is scalable for different workloads. This feature shows the resource usage elasticity and system adaptability.

To demonstrate the scalable performance, the test topologies for N-VDS (E) and N-VDS (S) in east-west packet path are respectively configured for comparison, as shown in *Figure 15 and 16*. The RFC 2544 testing on the

Spirent TestCenter Virtual is implemented to evaluate the throughput performance as the logical core configuration changes. *Table 8* below illustrates the configuration of Spirent TestCenter Virtual for RFC 2544 throughput test.

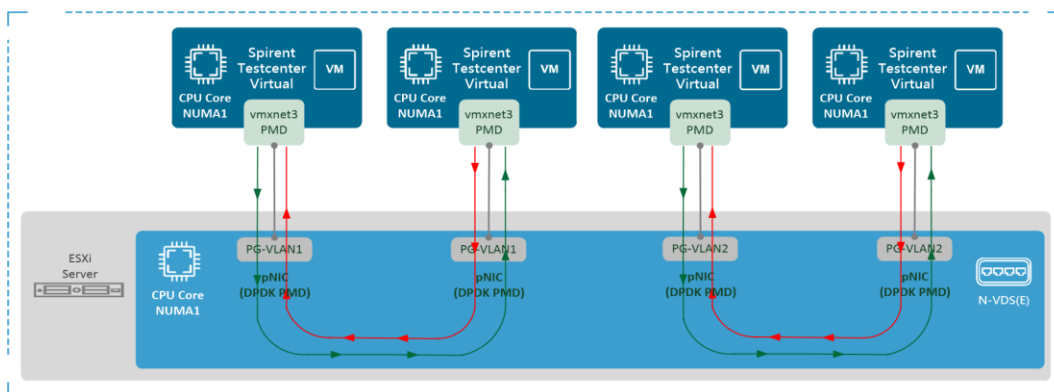


Figure 15. Test topology of multi-core N-VDS (E).

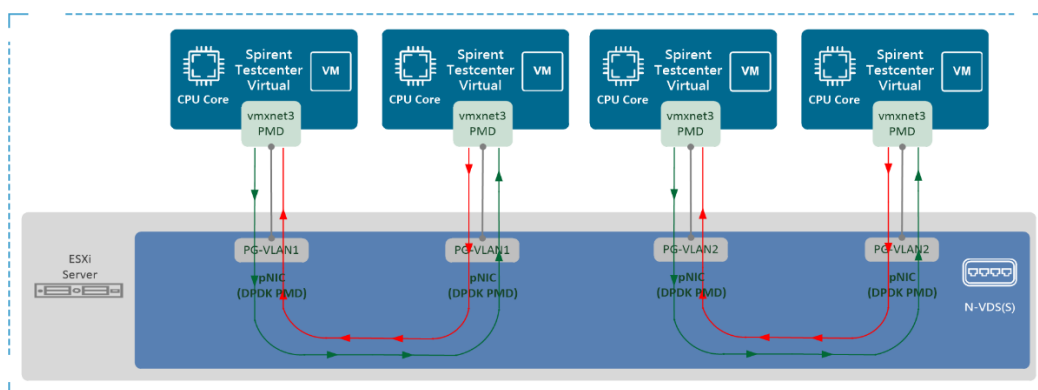


Figure 16. Test topology of N-VDS (S).

Table 6. Test configuration of scalable data plan performance.

Item	Specification
TestCenter Virtual version	4.96
NIC	VMXNET3 Paravirtualized NIC
STCv Driver mode	DPDK
Logical cores	3
Interfaces	Two 10Gbps ports (Fixed speed on 10Gbps)
Rate upper limit (%)	100
Frame size (bytes)	64, 128, 256, 512, 1024, 1280, 1518
Acceptable frame loss (%)	0

Compared with N-VDS (S), N-VDS (E) is configured respectively with 1, 2, and 4 logical cores to observe the data plane performance in different N-VDS networks and logical core configurations, as shown in *Figure 17*. N-VDS (E) configured with 4 logical cores reaches the line rate in all packet sizes which is suitable for ultimate performance required workloads. Since the N-VDS (E) configured with 2 logical cores has similar performance to the one configured with 4 logical cores except for the case in 64 bytes packet size. As a result, it can be considered to be a cost-effective selection. The result reveals that the more logical CPUs are assigned, the higher throughput performance is expected, indicating the scalability of N-VDS (E).

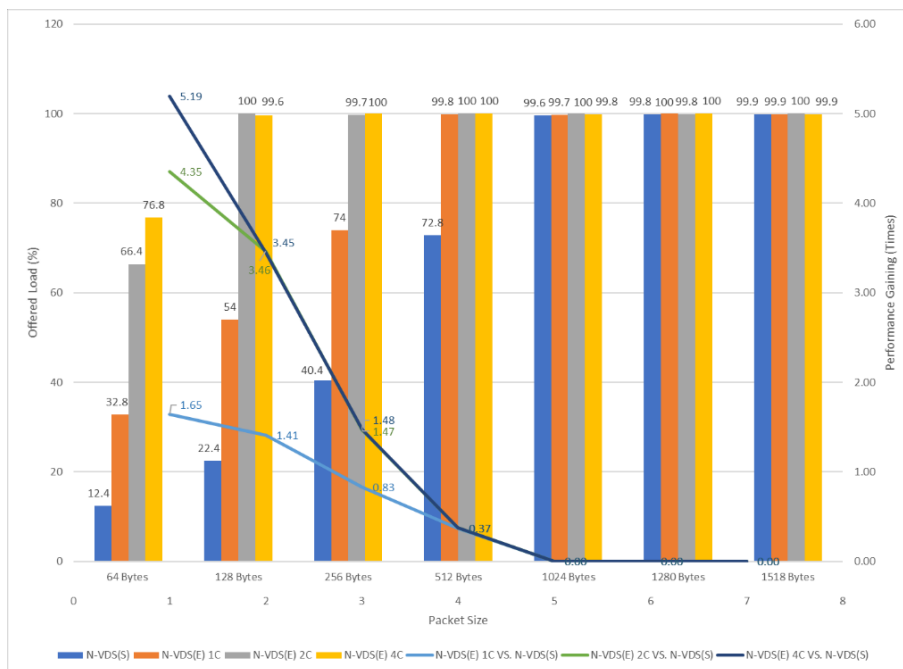


Figure 17. Test result of scalable data plan performance.

6. Conclusion

Modernizing the network architecture for rapidly processing massive data drives CSPs toward novel cloud transformation strategies. QxStack NFV Infrastructure for VMware vCloud NFV OpenStack Edition is a powerful solution, providing scalable, cost-efficient, and easy-managed NFV infrastructure to fulfill CSPs demands.

The solution integrates QCT NFV performance-optimized hardware with VMware vCloud NFV OpenStack Edition. The characteristics of ETSI compliance and OpenStack compatibility on VMware platform further brings the advantages of reliability and flexibility to the network system.

Meanwhile, the Spirent TestCenter with industry-leading RFC 2544 standard is adopted to verify the performance in various scenarios. QCT conducted a series of tests and verified the capability of this solution from three different perspectives, outstanding network I/O performance, multi-host throughput performance, and high data plane scalability.

With comprehensive validation and performance optimization, CSPs can leverage **QxStack NFV Infrastructure for VMware vCloud NFV OpenStack Edition** as a turnkey solution or a reference architecture to minimize the effort and quickly deploy NFV infrastructure and deliver services to win more businesses.

Reference

- [1] VMware official website
<https://www.vmware.com/>
- [2] vCloud NFV OpenStack Edition Reference Architecture – VMware vCloud NFV OpenStack Edition 3.1
<https://docs.vmware.com/en/VMware-vCloud-NFV-OpenStack-Edition/3.1/vmware-vcloud-nfv-open-stack-edition-ra31.pdf>
- [3] Tuning vCloud NFV for Data Plane Intensive Workloads – VMware vCloud NFV OpenStack Edition 3.0
<https://docs.vmware.com/en/VMware-vCloud-NFV-OpenStack-Edition/3.0/vmwa-vcloud-nfv30-performance-tuning.pdf>
- [4] Tuning VMware vCloud NFV for Data Plane Intensive Workloads – Technical White Paper
<https://docs.vmware.com/en/VMware-vCloud-NFV/2.0/vmware-tuning-vcloud-nfv-for-data-plane-intensive-workloads.pdf>
- [5] DPDK Sample Application User Guide, L3FWD-POWER.
http://dpdk.org/doc/guides/sample_app Ug/l3_forward_power_man.html
- [6] DPDK Sample Application User Guide, Skeleton.
http://dpdk.org/doc/guides/sample_app Ug/skeleton.html
- [7] Network Functions Virtualisation (NFV); Pre-deployment Testing; Report on Validation of NFV Environments and Services
https://www.etsi.org/deliver/etsi_gs/NFV-TST/001_099/001/01.01.01_60/gs_NFV-TST001v010101p.pdf
- [8] Network Functions Virtualisation (NFV) Release 3; Testing; Specification of Networking Benchmarks and Measurement Methods for NFVI
https://www.etsi.org/deliver/etsi_gs/NFV-TST/001_099/009/03.02.01_60/gs_NFV-TST009v030201p.pdf
- [9] How to Efficiently Evaluate NFVi Performance by Leveraging OPNFV Testing Projects_kubi_trevor
https://wiki.opnfv.org/download/attachments/5734608/How%20to%20Efficiently%20Evaluate%20NFVi%20Performance%20by%20Leveraging%20OPNFV%20Testing%20Projects_kubi_trevor.pdf?version=1&modificationDate=1498464116000&api=v2
- [10] Ensuring Open vSwitch performance for a predictable NFV infrastructure - A Linux Foundation Collaborative Project
https://www.opnfv.org/wp-content/uploads/sites/12/2016/12/opnfv_vswitch_final.pdf
- [11] Spirent TestCenter RFC 2544 benchmarking test package
https://www.spirent.com/-/media/datasheets/broadband/pab/spirenttestcenter/stc_rfc-2544_benchmarking_test_package_datasheet.pdf?la=en
- [12] RFC 2544 – IETF
<https://www.ietf.org/rfc/rfc2544.txt>

LEGAL DISCLAIMER

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH QUANTA CLOUD TECHNOLOGY (QCT) PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN QCT'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, QCT ASSUMES NO LIABILITY WHATSOEVER AND QCT DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF QCT PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY QCT, THE QCT PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE QCT PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Quanta Cloud Technology (QCT) may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." QCT reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice. Contact your local QCT sales office or your distributor to obtain the latest specifications and before placing your product order.



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

All specifications and figures are subject to change without prior notice. Actual products may look different from the photos.

QCT, the QCT logo, Rackgo, Quanta, and the Quanta logo are trademarks or registered trademarks of Quanta Computer Inc.

All trademarks and logos are the properties of their representative holders.

Copyright © 2020-2021 Quanta Computer Inc. All rights reserved.