SOLUTION BRIEF

QxStack VMware Edition - NSX-T

The Validated Carrier-Grade Solution to Virtualize Your Network

Main Challenges

- High complexity of heterogeneous infrastructure that impacts the visibility of network topology
- Difficult manageability across different workloads
- Discrete security coverage for entire workloads and platforms

QCT Value

- Rich experience in integrating software and hardware
- Validated solutions that benefits customers to accelerate time to value
- Optimized SKU to build modern cloud
- The latest technology to eliminate customers' difficulties

INTRODUCTION

Software-defined data center (SDDC) platform has been widely adopted to manage and deploy data centers to boost resource utilization and manageability. The workloads extending to different platforms such as ESXi, KVM, and bare-metal machines have become the trend so that enterprises tend to adopt multicloud for better DevOps and services. With lightweight and standalone structure, container nowadays is a popular software technology utilized to run workloads on various platforms. However, the adoption of various platforms involves complicated integration between workloads and platforms, indicating the difficulties in unifying networking, centralizing security policies, and automating the deployment on infrastructures. Therefore, a cross-platform software-defined network (SDN) is considered to be an appropriate technology to solve the integration of network and security issues between workloads across multi-infrastructures.

VMware NSX-T™ Data Center is an SDN software that enables IT administrators to construct singularized networks across workloads and platforms. The key benefits that NSX-T Data Center brings include extending network across heterogeneous platforms, supporting networking and security on container-based workloads, providing multi-tenancy routing model for administrators, delivering high performance edge node on bare-metal server, standardizing REST API for automation, etc. Moreover, it is independent from compute managers; thus, it can be simultaneously performed with multiple vCenters.

QxStack VMware Edition solutions are integrated with NSX-T Data Center on discreetly-selected QCT servers. With rich experience in integrating hardware and software, QCT validated the functions of NSX-T Data Center on ESXi and KVM platforms to ensure infrastructure and services are fully compatible and available.

NSX-T ARCHITECTURE

NSX-T Data Center is a software-defined network with the NSX managed virtual distributed switch (N-VDS) enabled on hosts to implement the hardware-agnostic overlay networks on clouds in which the overlay network services from L2 to L7 can be produced on demand in a single user interface (UI). NSX-T Data Center introduces logical network functions such as routers, switches, and firewalls. These logical functions are distributed to the transport nodes' system so that the networks in data center can be automatically singularized and deployed via NSX-T manager or API. It also brings significant security controls to exceed the traditional ways. The distributed firewall features micro-segmentation and optimizes traffics on individual object, even containers or virtual machines (VMs) are in the same network segment. The guest introspection integrates third-party anti-virus solutions to offload and centralize the protection so as to simplify the anti-virus tasks. Moreover, NSX-T Data Center combines the security benefits across infrastructures, bringing high operational simplicity to administrators.

NSX-T Data Center is formed by management plane, control plane, and data plane. In the three-plane structure, the components, including manger nodes, transport nodes, and edge nodes, are distributed to construct the fundamental infrastructure of the NSX-T Data Center, as show in Fig. 1.

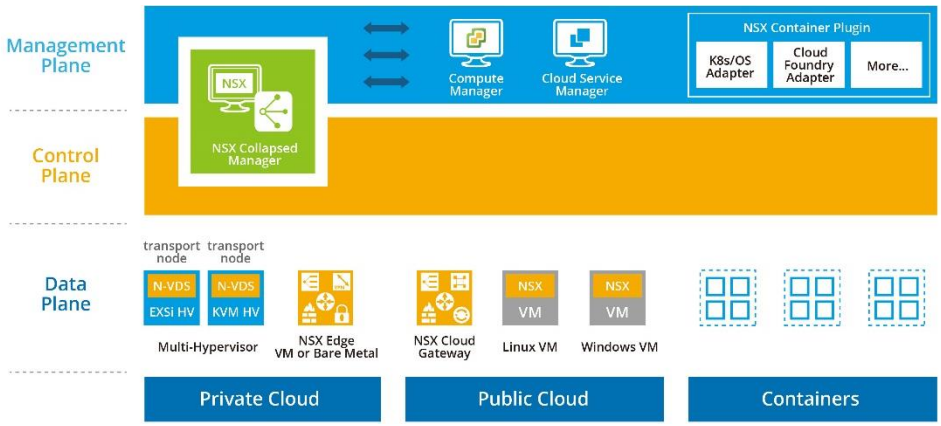

Figure 1. NSX-T three plane architecture.

NSX-T Benefits

- Decouple the network functions from dedicated network equipment
- Enhance security via policy-based control and logical network functions
- Support multi-cloud environments
- Accelerate workload with N-VDS(E)

Management Plane

The management plane provides an API entry point for the whole NSX networks and implements the management control for NSX operational tasks. Covering both management plane and control plane, the NSX Manager appliance is the central control node of NSX-T utilized to configure logical networks, transport nodes, edge nodes, security policies, etc. The appliance is also responsible for communicating with compute managers and container plugin. To integrate the public cloud with NSX-T Data Center, it can be deployed as Cloud Service Managers.

Control Plane

The control plane receives the information delivered from the management plane, distributes the network topology sent from the data plane, and transmits the settings to the forwarding engines. The NSX-T collapsed manager appliance is a distributed system used for providing information to logical networks and edge nodes. For redundancy, three controllers are typically deployed.

Data Plane

The data plane is in charge of transforming and forwarding packets, establishing virtual tunnel, and handling physical uplinks. The transport and edge nodes in this plane perform packet forwarding under N-VDS according to the topology that control plane reports.

In SDN, the three planes implemented in software are decoupled from each other to brings flexibility on network operations. The NSX-T Data Center further extends this feature to multi-hypervisors and Linux bare-metal machines; thus, the overall network topology across platforms and workloads can be unified. Furthermore, NSX-T Data Center is independent from VMware vSphere®, indicating that it can operate with other infrastructure managers as a standalone product. Based on these benefits, NSX-T Data Center is the most suitable SDN to be integrated with QxStack VMware Edition solutions for the future cloud environment.

THREE CLUSTER DESIGN

To streamline the operations in complicated multi-cloud environment, the three-cluster architecture, is designed to group the nodes with similar roles. The three clusters, namely management, compute, and edge clusters, consist of management-related appliances, production-related appliances, and edge-related appliances (physical or virtual) respectively and they can be easily scaled out by adding more nodes. As shown in Figure 2, the three-cluster design spans across QCT Hyper-Converged Infrastructure (HCI) and heterogeneous platform on NSX-T Data Center.

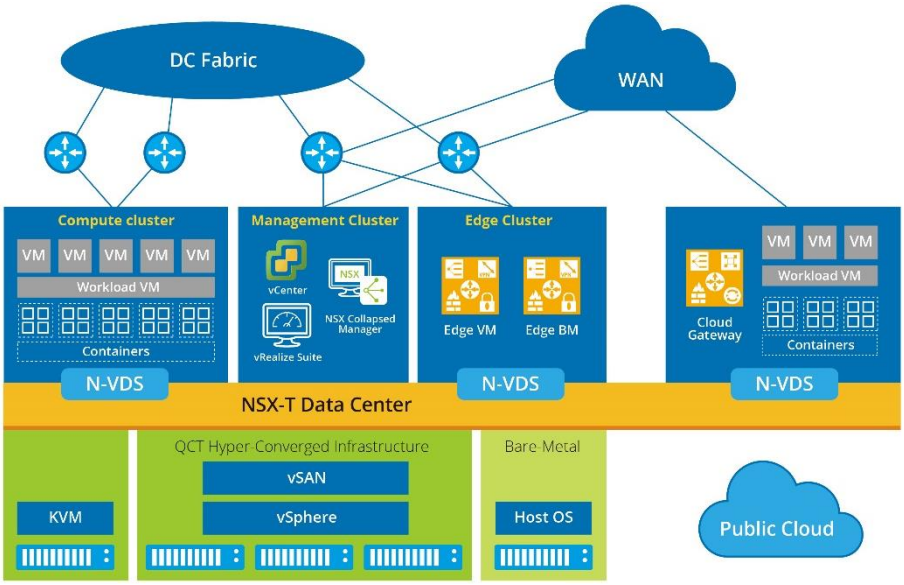

Figure 2. Three cluster design

Management Cluster

The management cluster is in charge of cloud administration and NSX-T's central control. This cluster typically includes management appliances such as vCenter server, virtual infrastructure manager (SDDC manager, VIO, or PKS), vRealize appliances, and NSX manager collapsed nodes, which provides overall management, monitoring, and operations on the infrastructure and workloads to orchestrate resources.

Compute Cluster

The compute cluster formed by NSX-T transport nodes is in charge of allocating compute resources. This cluster consists of hypervisors or bare-metal hosts configured with NSX managed virtual distributed switch which can provide on-demand resource for services. The compute cluster serves the compute, storage, and network resources in multi-cloud environment.

Edge Cluster

The edge cluster is in charge of the north-south traffic for the QxStack + NSX-T infrastructure and it consists of virtual-appliance or bare-metal based edge nodes. Virtualized network functions (VNFs) such as DHCP, NAT, N-S firewall, multi-tenancy routing model, VPN, and load balancer are combined to provide critical routing and services for data centers that brings the flexibility of network design for cloud-native workloads.

Under such architecture, the NSX-T components based on different functions are grouped into three clusters to optimize the overall performance and operation in data center. QCT provides diverse validated and optimized servers for deploying the three clusters. The hardware selection for deploying management and compute clusters highly depends on the workload type. For building the two clusters, please refer to QCT website or QxStack NFV Infrastructure for VMware vCloud NFV OpenStack Edition for hardware suggestions. Since edge cluster is responsible for north-south ingress and egress traffic, the hardware selection mainly lays stress on the network workload. The edge cluster can be either deployed on VM or bare-metal server. The approaches of deploying edge node along with the benefits will be explicated in the following section.

EDGE CLUSTER

The scale of north-south traffic that processes in edge node varies based on the size of data center. Two options either "Collapsed Edge" or "Bare-Metal Edge" on VM or bare metal form factors can be selected to deploy edge clusters. The comparison between the two deploy methods is detailed in Table 1.

In the "Collapsed Edge" design, a virtual appliance is executed on ESXi servers to provide network services. It can be automatically deployed and can be used to share virtualized resources with other workloads. The existing CPU, disk, NIC, and vSphere settings on ESXi servers can be leveraged to

optimize resource utilization, VM mobility, and CapEx that brings economic and scalable advantages. Therefore, this design is suitable for relatively low traffic environment such as lab, POC, and environment with less than 64 hypervisors.

In the "Bare-Metal Edge" design, the NSX Edge is implemented directly on QCT servers to provide services. The performance of packet processing in the Bare Metal Edge is better than that of in the Collapsed Edge. Thus, this design is suitable for large-scale use cases since it works on a dedicated bare metal server to deliver native performance. Besides, the implementation on bare metal server indicates the independence of ESXi hypervisor that brings enterprises the advantage to delimit the edge clusters.

Table 1. Comparison between Collapsed Edge and Bare-Metal Edge.

Form Factor	Scenario	Benefit
Collapsed Edge (Virtual Machine)	- POC or lab environment - Environment less than 64 hypervisors - ESXi server shared with other workloads - Automatic deployment	- Cost-effective infrastructure - High redundancy under vSphere - High resource utilization - Flexible deployment options
Bare-Metal Edge (Bare Metal)	- High throughput environment - High routing redundancy - Dedicated server and components	- Dedicated server and hardware - Native performance to process packets - High bandwidth together with stateful services - Fast failure convergence for routing path

HARDWARE SUGGESTION

Equipped with different form factors, drive counts, and PCI Express (PCIe) expansion slots, two QCT HCI servers, QuantaGrid D52B-1U and D52BQ-2U, are recommended to fulfill diverse network performance requirements, as shown in Table 2.

When NSX Edge is deployed either on VM or bare-metal server, the selections of CPU and NIC should be in accordance with NSX Edge system requirements[1]. The selected CPU must feature AESNI capacity and 1 GB Huge Page support. When the edge VM is implemented, at least two physical NICs are suggested to be installed for redundancy. When the edge node is implemented in a bare-metal sever, at least three physical NICs are suggested to be installed. One NIC is dedicated for management while the other two are respectively allocated for overlay tunneling traffic and VLAN uplink connectivity.

QuantaGrid D52B-1U is an ultra-dense server with fast socket interconnect, 1.5x memory bandwidth, 2x FLOPs peak performance with Intel® Xeon® Processors, and up to 5 PCIe expansion slots per chassis which can fulfill the uplink requirement of edge node. D52BQ-2U is a 2-unit sever with multiple storage configurations and more PCIe expansion slots to idealize a high-performance edge node. For more details, please visit QCT website[2].

Table 2. Recommended servers for edge cluster.

	QuantaGrid D52B-1U	QuantaGrid D52BQ-2U
Server Model		
Form Factor	1U	2U
CPU	2 x Intel® Xeon® Scalable Processors Platinum/Gold/Silver Series	
Memory	Up to 24 slots	
Driver Bays	12 x 2.5" hot-plug (NVMe support)	24 x 2.5" hot-plug or12 x 3.5" hot-plug (NVMe support)
PCIe Slots	Up to 5 slots	Up to 10 slots

[1] VMware NSX® Edge™ System Requirements in NSX-T Data Center.
https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.4/installation/GUID-22F87CA8-01A9-4F2E-B7DB-9350CA60EA4E.html
https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.4/installation/GUID-14C3F618-AB8D-427E-AC88-F05D1A04DE40.html
[2] Please visit QCT website for more detail. https://www.qct.io/product/index/Server/rackmount-server

Summary

QCT verified the NSX-T Data Center on QxStack VMware Edition to ensure customers a singularized and hardware-agnostic network in cloud environment. In the three-cluster design, NSX-T Data Center components are leveraged in the scenarios such as NFV and PKS to optimize cloud operations. Furthermore, QCT foresees the demands of NSX-T Data Center edge and offers deployment approaches to completely solve cloud traffic sizing problem. By adopting QxStack VMware Edition – NSX-T, customers can greatly accelerate time to value, save a considerable amount of time and efforts, and leverage high-performance computing in the multi-cloud environment.

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

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 株式會社
日本国東京都港区芝大門二丁目五番八号牧田ビル 3 階
TEL:+81-3-5777-0818
FAX: +81-3-5777-0819

Germany
Quanta Cloud Technology Germany GmbH
Hamborner Str. 55, 40472 Düsseldorf
TEL: +49-2405-4083-1300

Korea
QCT Korea, Inc. (주식회사 큐씨티코리아)
서울특별시 영등포구 의사당대로 97 교보증권빌딩 10 층, 07327
TEL:+82-10-5397-1412
FAX: +82-2-6336-6710

Other regions
Quanta Cloud Technology
No. 211 Wenhua 2nd Rd., Guishan Dist., Taoyuan City 33377, Taiwan
TEL:+886-3-327-2345
FAX: +886-3-397-4770

ABOUT QCT

Quanta Cloud Technology (QCT) is a global data center solution provider. We combine the efficiency of hyperscale hardware with infrastructure software from a diversity of industry leaders to solve next-generation data center design and operation challenges. QCT serves cloud service providers, telecoms, and enterprises running public, hybrid and private clouds.

Product lines include hyperconverged and software-defined data center solutions as well as servers, storage, switches and integrated racks with a diverse ecosystem of hardware components and software partners. QCT designs, manufactures, integrates and services cutting-edge offerings via its own global network. The parent of QCT is Quanta Computer, Inc., a Fortune Global 500 corporation.

http://www.QCT.io

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 © 2019 Quanta Computer Inc. All rights reserved.