



**QxStack VMware® Edition**

**-vSAN ReadyNode™**

Performance Optimized SKU

**Reference Architecture**



## Table of Contents

<b>Legal Disclaimer</b> .....	<b>ii</b>
<b>1 Executive Summary</b> .....	<b>1</b>
<b>2 Reference Architecture Overview</b> .....	<b>2</b>
2.1 Purpose .....	2
2.2 Scope .....	2
2.3 Audience .....	2
<b>3 Solution Overview</b> .....	<b>3</b>
3.1 Intel® Select Solution .....	3
3.2 QxStack vSAN ReadyNode™ – Performance Optimized SKU .....	4
<b>4 Solution Architecture</b> .....	<b>8</b>
4.1 Hardware Architecture .....	8
4.2 Software Configuration .....	9
4.3 QxStack VMware® Edition – Auto-Deployment Manager .....	12
<b>5 Solution Validation</b> .....	<b>16</b>
5.1 Test Overview .....	16
5.2 Test Configuration .....	18
5.3 Test Result .....	24
<b>6 Conclusion</b> .....	<b>26</b>
<b>7 References</b> .....	<b>27</b>
<b>About QCT</b> .....	<b>28</b>



## 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 INTEL<sup>®</sup>LECTUAL 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 INTEL<sup>®</sup>LECTUAL 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.

Copyright<sup>®</sup> 2017-2018 Quanta Cloud Technology Inc. All rights reserved.

Other names and brands may be claimed as the property of others.



## 1 Executive Summary

Enterprises around the world are facing drastic information explosion. To address the challenges, it is critical for enterprises to rearrange their data centers to support new ways of doing business. Software-defined data center can be an answer to the market demand and help organizations stay ahead by delivering simplified management, agility, and lower Total Cost of Ownership (TCO). However, hundreds of solutions in the market complicate solution selection, not to mention the burden of cost and efforts in installation and performance tuning.

In order to assist customers and partners to accelerate data center transformation, Quanta Cloud Technology (QCT), a global data center solution provider, participates in the Intel® Select Solution Program and collaborates with Intel® to provide a ready-to-use software-defined data center solution **QxStack vSAN ReadyNode™ – Performance Optimized SKU** with the following benefits:

- Simplify modernization: reduce time to evaluate, select, and purchase necessary hardware components.
- Strengthen reliability: provide a confident choice by passing rigorous VMware® vSAN ReadyNode™ certification.
- Accelerate time to value: minimize time to deploy new infrastructure.
- Ensure the performance in business-critical scenarios: pre-optimize the setting parameters for business common use cases such as E-commerce or scalable web scenario to reach outstanding performance, exceeding Intel® Select program's threshold.

With the aforementioned benefits and outstanding performance, **QxStack vSAN ReadyNode™ – Performance Optimized SKU** is believed to be a valid choice for partners and customers to construct software-defined data center and stay in a leading position.



## 2 Reference Architecture Overview

### 2.1 Purpose

The purpose of this reference architecture is to introduce **QxStack vSAN ReadyNode™ – Performance Optimized SKU** and provide ultra-performance for customers. This reference architecture will illustrate the validated configuration and parameters tuning which can highly minimize the deployment time and simplify the deployment process.

### 2.2 Scope

This reference architecture:

- Introduces Intel® Select Solution program and the benefits of adopting **QxStack vSAN ReadyNode™ – Performance Optimized SKU**.
- Illustrates the hardware configuration and software components discreetly selected by QCT and Intel® in the solution.
- Simulates the commonly-used workloads in the data center, and demonstrates the ultra-performance and scalability of **QxStack vSAN ReadyNode™ – Performance Optimized SKU**.

### 2.3 Audience

The intended audience of this document are IT professionals, technical architects, and sales engineers. The document is to assist them in selecting, implementing, and tuning SDDC products.



### 3 Solution Overview

This section introduces Intel® Select Solution Program and **QxStack vSAN ReadyNode™ – Performance Optimized SKU**, and highlights the benefits of this solution for customers and partners.

#### 3.1 Intel® Select Solution

In order to address the challenges end customers faced in a data center transformation journey, Intel® collaborates with several partners including QCT to launch a “Intel® Select Solution”. The Select Program provides rich system solutions for the future-forward, agile data center. The solutions in Intel® Select Program fulfill diverse scenarios in the data center use cases including virtualized infrastructure, database, NVFI, and hybrid cloud with the key benefits below:

##### **Simplify Modernization:**

The pre-defined and workload-oriented solutions in Intel® Select Program can simplify and speed up the data-center modernization process from evaluation and procurement to integration and deployment.

##### **Workload-Optimized performance:**

Developed by Intel® and solution partners, Intel® Select Solution is based on the latest Intel® technologies to provide an accelerated, simplified path to unlock the overall performance.

##### **System-level Benchmarking and Verification**

With pre-defined settings and rigorous system-wide tuning, Intel® Select Program respectively sets up performance thresholds for different workloads to ensure reliable, workload-optimized performance for complex data center applications.



### 3.2 QxStack vSAN ReadyNode™ – Performance Optimized SKU

In Intel® Select Solution, QCT collaborates with Intel® to develop **QxStack vSAN ReadyNode™ – Performance Optimized SKU** - a total solution for software-defined data center transformation. As shown in Fig. 1, **QxStack vSAN ReadyNode™ – Performance Optimized SKU** is composed of QCT’s well-designed 2<sup>nd</sup> Generation Purley server platform and marketing leading virtualization software developed by VMware, delivering a reliable and confident choice for customers.

This solution addresses the common business challenges that CIOs face today in the data center such as insufficient resource utility, management inefficiency, reliability issues, and daunting process of deployment and tuning. In the following sections, a solution overview with some benefits will be provided for customers to solve the problems.

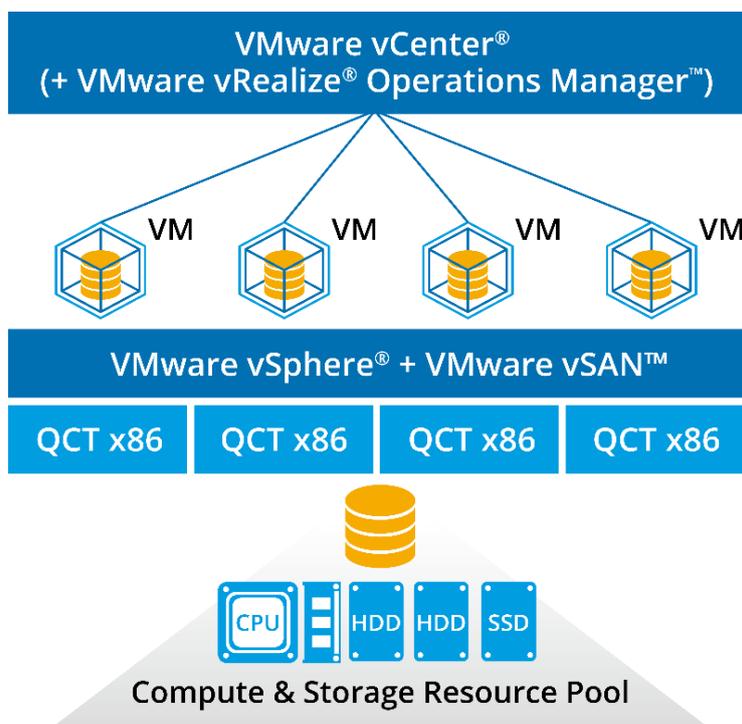


Figure 1. Solution Architecture of QxStack vSAN ReadyNode™ – Performance Optimized SKU.

#### 3.2.1 Manageability, Scalability, and Efficiency- Hyper-Converged Infrastructure

**QxStack vSAN ReadyNode™ – Performance Optimized SKU** is a hyper-converged infrastructure solution. Traditionally, IT technicians face the challenges of resource management and scalability since compute and storage resources are separated. Hyper-Converged Infrastructure (HCI) is a novel technology which can integrate compute, storage, and virtualization resources in a single hardware box. Every single



node is capable of delivering compute and storage resources at the same time. Several benefits are listed in detail below.

### Simplified Management

In legacy-converged architecture, since the compute and storage resources are provided by different servers and storage devices, the configuration settings and operation management are independent which means IT administrators must manage two devices through different management tools. To provide a storage device to a server, users need to configure settings from LUN and Volume, and then mount the storage device to the server host for VM to access. In hyper-converged architecture, compute and storage can be regarded as a system. Users can manage both compute and storage resources with a single management portal. By achieving full “policy-driven management”, IT administrators only need to define their own compute and storage resources. The allocation process can be automatically completed by a single management portal, which significantly reduces the management effort.

### Scalability and Efficiency

The hyper-converged infrastructure integrates compute and storage resources into a basic unit, called building block. By implementing the clustered architecture, users can add more building blocks to the cluster to expand the overall performance and capacity. This also makes the expansion of the hyper-converged architecture simple and predictable, as shown in Fig. 2.

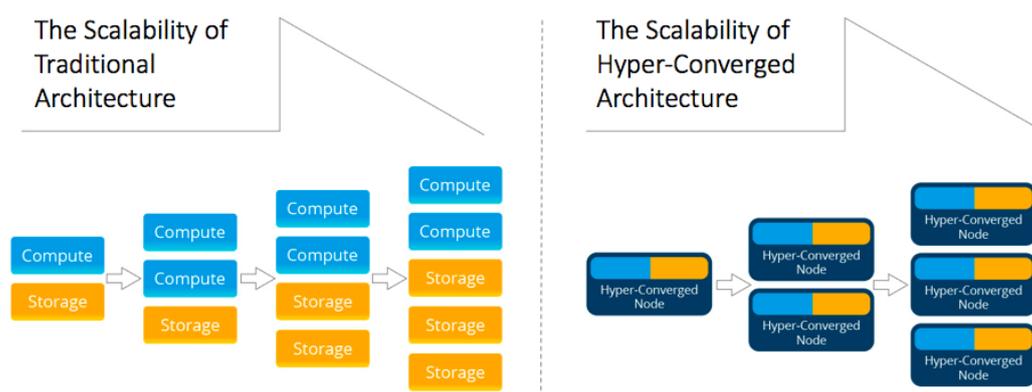


Figure 2. Comparison between Traditional Architecture and HCI.

### 3.2.2 Reliability - vSAN ReadyNode™ Certification

vSAN ReadyNode™ is a program created by VMware® to verify the compatibility between server platform and VMware®-developed software and to guarantee the performance and stability of a solution. To pass vSAN ReadyNode™ certification, all



the details of a solution including hardware components, firmware and driver, and software stack should be strictly examined to meet the rigorous requirements. QCT made lots of efforts in certification validation of **QxStack vSAN ReadyNode™ – Performance Optimized SKU**. The certified solution configuration is listed in [VMware® Compatibility Guide](#), as shown in Table 1. In the past, it is common for administrators to spend weeks researching and struggling against the compatibility issues to deploy a new system. Now, with the solution validated by QCT and VMware®, customers can rest assured of the solution reliability and focus on strategic and productive tasks.

Table 1. QxStack vSAN ReadyNode™ – Performance Optimized SKU.

vSAN ReadyNode™ Details		
Model: AF4-QCT-QuantaGrid D52B-1U Profile: AF-4 Series Type: All Flash Partner Name: Quanta Computer Inc Generation: Gen3 - Xeon Scalable		
Components	Details	Quantity
SKU	QuantaGrid D52B-1U_AF4 -with TPM1.2	
System	Model: QuantaGrid D52B-1U System Type: Rackmount	1
CPU	Intel® Xeon® Gold 6152 CPU @ 2.10 GHz / 22C/44T	2
Memory	32GB 2666MHz 288-pin DDR4 RDIMM	12
Caching Tier	Model : <a href="#">Intel® SSD DC P4800X Series SSDPED1K375GA (375 GB, AIC)</a> Partner Name: Intel Device Type: NVMe Capacity: 375 GB Performance Class: Class F: 100,000+ writes per second TBW Endurance Class: Endurance Class D >=7300 TBW	2
Capacity Tier	Model : <a href="#">Intel® SSD DC P4500 Series SSDPE2KX020T7 (2 TB, 2.5-inch)</a> Partner Name: Intel Device Type: NVMe Capacity: 2000 GB Performance Class: Class E: 30,000-100,000 writes per second TBW Endurance Class: Endurance Class A >=365 TBW	4
NIC	Model : <a href="#">Intel(R) Ethernet Converged Network Adapter X710-2</a>	1
Boot Device	Model: SATADOM 32GB	1



### 3.2.3 Ultra-performance - Selected Hardware and Brilliant Testing Result

**QxStack vSAN ReadyNode™ – Performance Optimized SKU** adopts QCT's 2<sup>nd</sup> Generation Purley server platform which surpasses the physical design limitation and offers high performance and energy efficiency. Besides, the components in this all-NVMe configuration are carefully selected by QCT which boosts the performance to a new level.

Moreover, Intel® Select Solution also sets specified performance benchmark thresholds for solution validation. The benchmark tool used in vSAN ReadyNode™ solution is VMmark, which can stimulate the commonly-used cases such as E-commerce and scalable web scenarios. With QCT's knowledge and tuning experience, the performance of **QxStack vSAN ReadyNode™ – Performance Optimized SKU** exceeds the benchmark thresholds defined by Intel Select Program, as shown in Fig. 3. The test results will be explicitly elaborated in chapter 5.3.

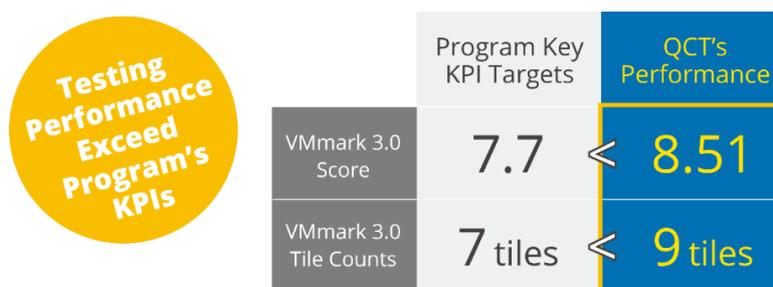


Figure 3. Testing Performance of QxStack vSAN ReadyNode™ – Performance Optimized SKU.

### 3.2.4 Time to Value - Easy Deployment

Data center deployment can be a very complex and daunting process from resource and schedule arrangements to hardware and software installations, not to mention networking topology and performance tuning. QCT understands all the difficulties and takes the following steps to solve customers' problems.

QCT integrates VMware® vSAN™ software stack into QCT's server platform to save your efforts and accelerate the installation process. In addition, this reference architecture in chapter 5.2 describes the parameter setting of network connectivity, firmware and software optimized configuration, and suggested policy setting that can significantly minimize the guesswork in the deployment process and reduce deployment time and expense.

QCT developed an auto-deployment tool, named QCT auto-deployment manager, to simplify and automate the deployment process in merely a few steps. The detailed deployment process is described in chapter 4.3.



## 4 Solution Architecture

In this session, we focus on illustrating the architecture of **QxStack vSAN ReadyNode™ – Performance Optimized SKU**, including hardware architecture, software architecture, and QxStack VMware® Edition – Auto-Deployment Tool.

### 4.1 Hardware Architecture

QuantaGrid D52B-1U server features both extreme storage density and compute power for the hyper-converged solution. The hardware configuration in this reference architecture is elaborately designed to fulfill Intel® Select Program’s requirements. In this design, four QuantaGrid D52B-1U servers contribute its local storage to build a vSANDatastore. Each server uses two Intel® Xeon® Gold 6152 CPUs with 22 cores for compute and storage services. Each node uses 384GB memory capacity. Each host contains 2 NVMe Intel® Optane™ SSD DC P4800X (375GB) for cache tier and 4 NVMe Intel® SSD P4500 for the capacity tier under vSAN™ architecture, as shown in Fig. 4. Each QuantaGrid D52B-1U host contributes its local disks to a vSANDatastore and the disks are organized into disk groups. Each host consists of two disk groups. A disk group can be seen as a “fault domain.” Generally, if the cache device fails, all HDDs and SSDs in the same disk group will be impacted. However, the design of multiple disk groups can highly reduce the impact when cache device fails on overall disk group.

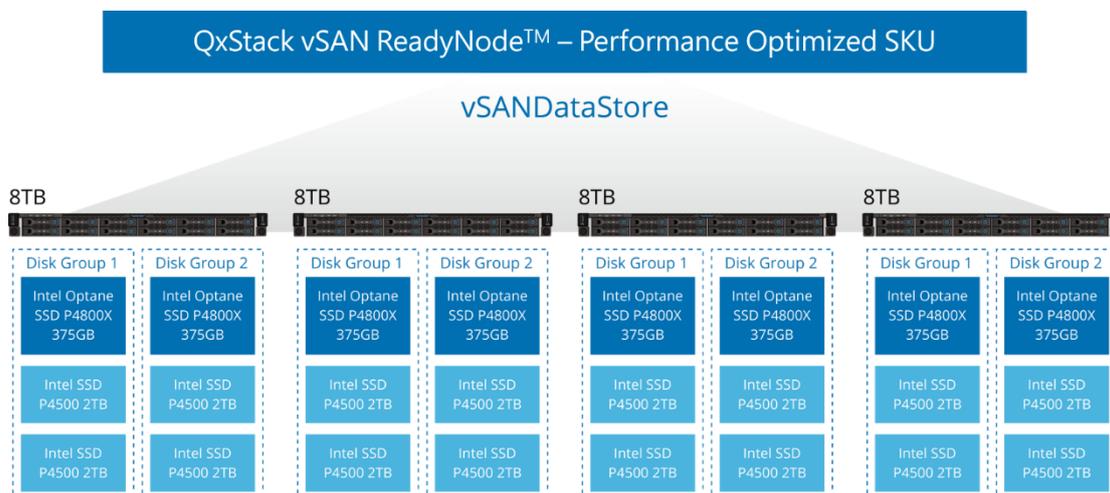


Figure 4. vSANDataStore Configuration of QxStack vSAN ReadyNode™ – Performance Optimized SKU.

#### 4.1.1 Compute: Intel® Xeon® Scalable Processor

Tailored for a software-defined storage solution, QuantaGrid D52B-1U features ultimate compute and storage density in a 1U platform which redefines the physical



limitation. QuantaGrid D52B-1U is a rackmount server with Intel® Xeon® Processor Scalable Family, delivering fast socket interconnection, 1.5x memory bandwidth, and 2x Floating-Point Operations per second (FLOPs) peak performance. Intel® Xeon® Scalable processor can deliver up to 112 vCPUs per server and 3.9x higher virtualized throughput compared to the previous platforms based on the Intel® Xeon® processor E5. The compute capability can empower VMware® vSAN™ by supporting a wide range of critical workloads with low latency in the modern data center.

#### **4.1.2 Storage: Intel® Optane™ NVMe SSD**

Caching is prerequisite not only in today's software-defined storage solution but also virtualization applications to boost performance. Enterprises that value performance can benefit from the cache tier using the Intel® SSD Data Center (DC) Family with Non-Volatile Memory Express (NVMe). Adopting Optane™ SSD DC P4800X as the cache tier in vSAN™ can deliver an extremely high-performance and ultra-low latency while running the write-intensive workload.

1. Intel® Optane™ SSD DC P4800X combines the attributes of both memory and storage which significantly reduces the latency of sensitive workload and accelerates the overall deployment.
2. Intel® Optane™ SSD DC P4800X is able to deliver 5-8X fast performance to low queue depth workload which exhibits exceedingly high throughput.

#### **4.1.3 Network: Intel® X710-DA2**

According to the best practice of VMware® vSAN™ configuration, at least 10 gigabit Ethernet (GbE) is recommended on all flash vSAN™ configuration. Discreetly selected to fulfill the bandwidth demand between the server for vSAN™ traffic, Intel® Ethernet Converged Network Adapter X710 DA-2 not only optimizes the performance but also improves the reliability for data center networking.

## **4.2 Software Configuration**

### **4.2.1 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 data center. By leveraging the virtualization technology, vSphere® can provide a highly-available, efficient, centralized infrastructure for IT administrators to deliver flexible, reliable services. VMware® vSphere® provides the following benefits:



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

#### **4.2.2 VMware® vCenter Server®**

vCenter server® is a platform which provides centralized management and operation for VMware® virtualization environment. By aggregating all the virtual resources, vCenter server® can provide resource provisioning and monitoring. It is capable of provisioning compute, storage, and other resources to the virtual machine and enabling High Availability (HA), Distributed Resource Scheduler (DRS), vMotion, etc. Currently, in vCenter® 6.5, vCenter server® can support up to 2,000 hosts.

#### **4.2.3 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 which is different from the traditional hardware architecture including compute, storage, and networking. By sharing the resources of a single hardware across multiple environments, a physical system (x86 server) is capable of executing multiple virtual machines with different operation systems (OS). 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 a single effective way to reduce IT expenses and boost efficiency and agility for all-sized businesses.

#### **4.2.4 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 optimizes 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.

vSAN™ is built on the industry-standard x86 server and leverages the local storage on the server. By using the virtualization technology, the under-layer physical disk is a concept of abstract resource pool, aggregated into a virtual resource for providing a shared storage to the server in the cluster. It also makes non-disruptive expansion



capacity possible by adding hosts to a cluster or adding disks to a host. SSDs are used as a cache tier to accelerate the I/O performance and HDDs as capacity tier to store data. By leveraging the flash device on server side, vSAN™ can accelerate the read/write I/O processing speed and minimize the storage latency. Administrators can use vSAN™ to define the requirements of virtual machine storage such as performance and availability. The policy requirements are delivered to vSAN™ layer through the Storage Policy-Based Management (SPBM). SPBM, hence, plays an important role to bridge the upper layer application and the under-layer storage devices. Through pre-defined storage policy, SPBM will drive the vSAN™ layer to adopt these policies when users provision the virtual machine.



## 4.3 QxStack VMware® Edition — Auto-Deployment Manager

### 4.3.1 Solution Deployment

QxStack VMware® Edition — Auto-Deployment Manager simplifies the deployment process and provides an easy way to deploy **QxStack vSAN ReadyNode™ – Performance Optimized SKU**.

### 4.3.2 Deployment Process

In order to facilitate the deployment process of **QxStack vSAN ReadyNode™ – Performance Optimized SKU**, QCT developed a QxStack VMware® Edition — Auto-Deployment Manager to dramatically reduce time and minimize guest efforts for building hyper-converge solutions. Auto-Deployment Manager plays a coordinate role such as initializing ESXi™ and deploying vCenter server®.

Administrators can initiate the auto-deployment tool on a single node and type the default IP. Once the tool is powered on, the deployment tool is connected and the view of GUI can show administrators the deployment process. Subsequently, DHCP service needs to be enabled and applied to each server that is going to be deployed, as shown in Fig. 5. After clicking “Save and Next”, three different options including “System Initial Only or Scale Out”, “Install vCenter® and Prebuild Cluster”, and “Install vCenter® and VDI environment” can be selected. To initiate and deploy the environment, “Install vCenter® and Prebuild Cluster” can be selected to execute the deployment process.



Please input Out Of the Band(OOB) management network DHCP range:

Blocksystem initialized: False  
DHCP Status: Running

You can change the DHCP settings only before the first host is initialized.

OOB DHCP Scope: 172.24.0.0  
OOB DHCP Range Start:   
OOB DHCP Range End:   
OOB DHCP Netmask:   
OOB DHCP service:

- Enable DHCP service
- Disable DHCP service

---

Select All

IP:172.24.0.21    MAC:54:ab:3a:3d:2a:71    Model:QuantaPlex T41S-2U    Storage:True    Status:New

If you choose host "In-Use", it might be already used by existing vCenter.

You might need to disassociate host from vCenter controller before you can modify ESXi IP configuration

---

Figure 5. DHCP Service Enablement of QxStack VMware® Edition — Auto-Deployment Tool.

Administrators need to enter a general system information including inband management IP information, ESXi™ hosts information, and vCenter® information, as shown in Fig. 6. Once administrators enter all the information and click the “Install vCenter® and Prebuild Cluster”, the system will initiate each ESXi™ host and automatically deploy the vCenter® and vSAN™ service. The entire deployment process will take around 20 to 30 minutes, as shown in Fig. 7.

Please enter the QCT Block System information

Inband management IP:

Inband management netmask:

Inband management gateway:

---

Please enter the ESXi hosts information

ESXi hosts management IP(s):

Selected Node(s): 1

Please use "," to separate or use "-" to set the IP range.  
 e.g. 192.168.100.21,192.168.100.25-192.168.100.29  
 Maximum number of IPs should not be higher than number of Select Node(s).

ESXi hosts management VLAN:

ESXi hosts management netmask:

ESXi hosts management gateway:

Enable VSAN:       Enable Jumbo Frame(9000):

---

Please enter new vCenter information

vCenter Appliance Name:

vCenter IP:

vCenter Netmask:

vCenter Gateway:

vCenter DNS:

vCenter Password:

SSO Password:

The SSO password must be 8-20 characters long.  
 It must include:  
 - 1 special character(s)  
 - 1 upper case character(s)  
 - 1 lower case character(s)  
 - 2 alphabetic character(s)  
 - 1 numeric character(s)  
 - No more than 3 adjacent characters can be identical.

---

Please enter your first cluster information

Datacenter Name:

Cluster Name:

Install vCenter and Prebuild Cluster

Figure 6. General System Information of QxStack VMware® Edition — Auto-Deployment Tool.



---

System will be deployed automatically, please wait for 20~30 mins.  
Do not close the window during deployment.

**vCenter environment already initialized.**  
**vCenter login:** <https://192.168.33.50:9443>

Figure 7. Deployment Process of QxStack VMware® Edition — Auto-Deployment Tool.

## 5 Solution Validation

In this session, the test methodologies and test results will be demonstrated to guarantee the performance and reliability of this solution.

### 5.1 Test Overview

In today's data center, more and more enterprises adopt virtualization technology. With virtualization technology, enterprise can benefit from the flexibility and scalability. VMmark is a benchmark tool provided by VMware to evaluate the performance and scalability of virtualization platforms. It consists two main workloads including applications and infrastructure operation behaviors to evaluate the overall system performance. VMware choses three subworkloads to simulate the applications including standby system, scalable web simulation, and E-commerce simulation, as shown in Fig. 8.

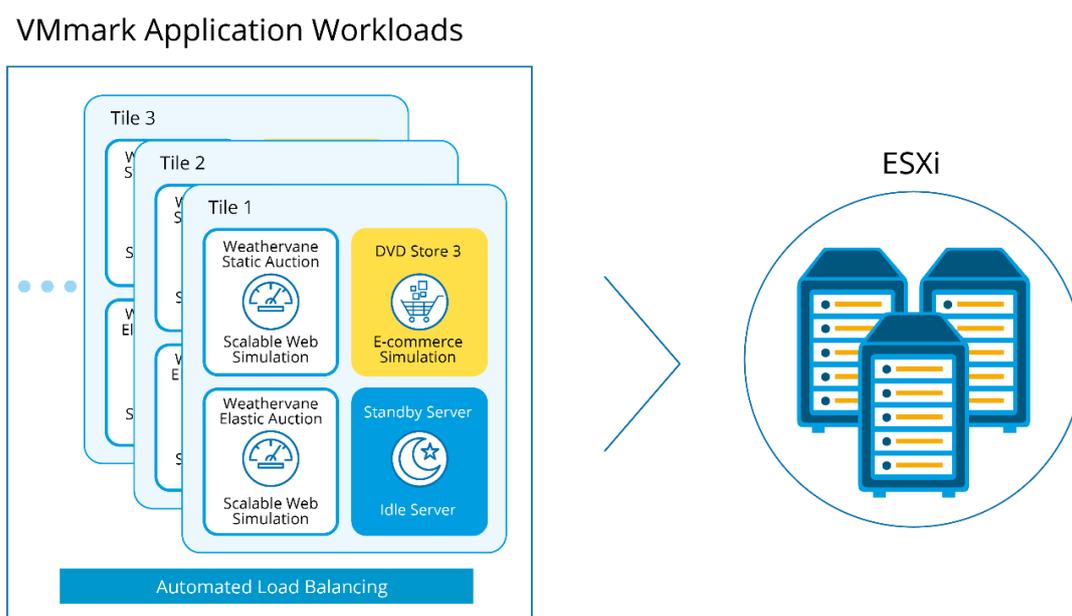


Figure 8. VMmark Application Workloads.

Scalable Web applications are used to simulate the workloads such as on-line auction websites or social networking behavior. Typically, the web simulation behaviors including the services such as load balancers, message servers, web servers, and database consume different resources in the infrastructure. VMmark leverages Weathervane, an application-level benchmark tool, to test and evaluate virtual infrastructure performance. The static and elastic applications are used. For the static application instance, all of the services mentioned above are run on their own virtual

machine to generate consistent load to the system. As for the elastic application instance, all services are shared on a single virtual machine to generate elastic load to the system.

Database is a typical resource intensive application commonly used in a virtualized data center. Most of the transactional workloads are running on the database which is a reasonable reason to evaluate such workloads on the virtualization platform. VMmark leverages the open source application, DVD Store, to simulate the on-line transaction behavior. It provides an E-commerce application with a back-end database component, a web application layer, and driver programs. The web server communicates with the back-end database to query user account and inventory data, and simulates the behaviors such as users' web login and product browse.

In addition to the application workloads, VMmark triggers some common virtualization procedures running on the virtual machines, as shown in Fig. 9. The procedures include:

- Virtual machine cloning and deployment.
- Dynamic virtual machine relocation between servers.
- Dynamic virtual machine relocation across storages.
- Simultaneous server and storage virtual machine relocation.
- Automated load balancing.

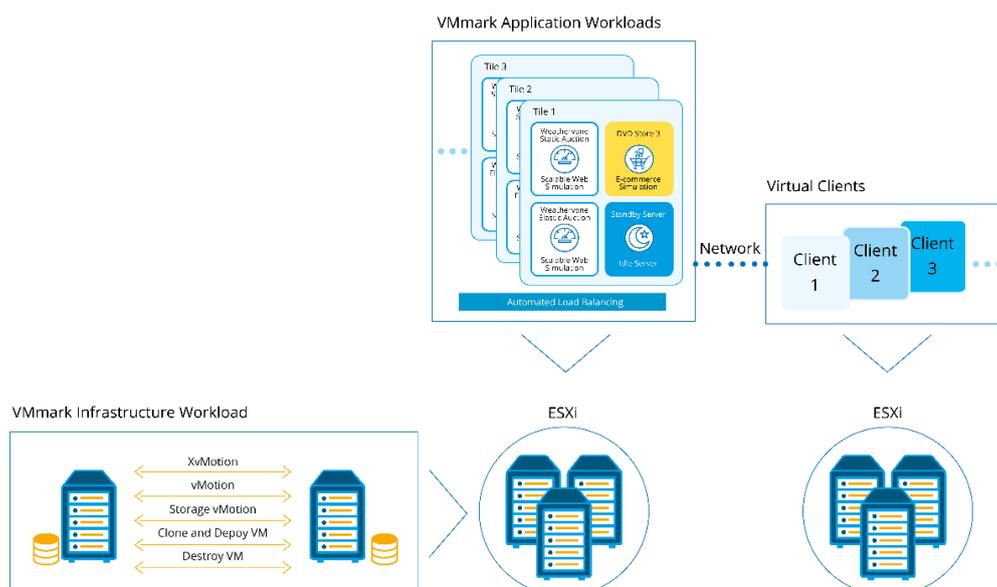


Figure 9. VMmark Workload Summary.

## 5.2 Test Configuration

### 5.2.1 System in Detail

Virtualization Software	
Hypervisor Vendor, Product, Version, and Build/Availability Date (MM-DD-YYYY)	VMware® ESXi 6.5 U1, Build 5969303 / 07-27-2017
Datacenter Management Software Vendor, Product, Version, and Build /Availability Date (MM-DD-YYYY)	VMware® vCenter Server 6.5 U1, Build 5973321/ 07-27-2017
Supplemental Software	None

Server	
Number of Servers in System Under Test (all subsequent fields in this section are per server)	4
Server Manufacturer and Model	QuantaGrid D52B-1U
Processor Vendor and Model	Intel® Xeon® Gold 6152
Processor Speed (GHz)	2.1 GHz
Total Sockets/Total Cores/Total Threads	2 Sockets/ 44 Cores/ 88 Threads
Primary CPU Cache	32 KB I + 32 KB D on chip per core
Secondary CPU Cache	1 MB I+D on chip per core
Other CPU Cache	30.25 MB I+D on chip per chip
BIOS Version	3A08
Memory Size (in GB, Number of DIMMs)	384GB, 12
Memory Type and Speed	32GB DDR4 2666MHz RDIMM
Disk Subsystem Type	VMware® vSAN™
Number of Disk Controllers	N/A
Disk Controller Vendors and Models	N/A
Total Number of Physical Disks for Hypervisor	1
Disk Vendors, Models, Capacities, and Speeds	Innodisk, 32G MLC SATADOM, 32GB, 6.0Gb/s
Number of Host Bus Adapters	N/A
Host Bus Adapter Vendors and Models	N/A
Number of Network Controllers	1
Network Controller Vendors and Models	Intel® Ethernet Converged Network Adapter X710-DA2 10 GbE
Other Hardware	None
Other Software	None
Hardware Availability Date (MM-DD-YYYY)	07-12-2017
BIOS Availability Date (MM-DD-YYYY)	07-12-2017
Software Availability Date (MM-DD-YYYY)	07-11-2017

Network	
Network Switch Vendors and Models	QuantaMesh BMS T3048-LY8
Network Speed	10Gb

Storage	
Array Vendors, Models, and Firmware Versions	4 x QuantaGrid D52B-1U, with VMware® vSAN™ 6.6.1 1 x QuantaPlex T41S-2U, with Windows Server 2012



Fibre Channel Switch Vendors and Models	No Fibre Channel Switch
Disk Space Used	32TB (4 x QuantaGrid D52B-1U) 2.6TB (1 x QuantaPlex T41S-2U)
Array Cache Size	3TB (4 x QuantaGrid D52B-1U) No cache used (1 x QuantaPlex T41S-2U)
Total Number of Physical Disks Used	24, (4 x QuantaGrid D52B-1U) 3, (1 x QuantaPlex T41S-2U)
Total Number of Enclosures/Pods/Shelves Used	Internal: 8 NVMe SSD, Shelf: 16 NVMe SSD (4 x QuantaGrid D52B-1U) Shelf: 3 SATA SSD (1 x QuantaPlex T41S-2U)
Total Number of Storage Groups Used	0
Number of LUNs Used	Refer section: Storage Notes
LUN Size and Number of Disks Per LUN	Refer section: Storage Notes
RAID Type	Refer section: Storage Notes
Disk Vendors, Models, and Speeds	8 x Intel® SSD DC P4800X Series SSDPED1K375GA (375 GB, AIC) NVMe SSD, 16 x Intel® SSD DC P4500 Series SSDPE2KX020T7 (2 TB, 2.5-inch) NVMe SSD, 1 x Intel® SSD DC S3710 Series (200 GB, 2.5-inch) SATA SSD, 2 x Intel® SSD DC S3710 Series (1.2TB, 2.5-inch) SATA SSD

Datacenter Management Server	
System Model	QuantaPlex T41S-2U
Processor Vendor and Model	Intel® Xeon® E5-2620 v4
Processor Speed (GHz)	2.1GHz
Total Sockets/Total Cores/Total Threads	2 Sockets / 16 Cores / 32 Threads
Memory Size (in GB, Number of DIMMs)	256GB,16
Network Controller(s) Vendors and Models	QN 10GbE 82599ES
Operating System, Version, Bitness, and Service Pack	VMware® ESXi 6.5 U1, Build 5969303
Virtual Center VM Number of vCPUs	4
Virtual Center VM Virtual Memory (in GB)	16
Virtual Center VM Operating System, Version, Bitness, and Service Pack	VMware® vCenter Server 6.5 U1, Build 5973321
Other Hardware	N/A
Other Software	N/A

Clients	
Total Number of Virtual Clients / Virtual Client Hosts	10/3
System Model(s)	QuantaGrid D51BP-1U
Processor Vendor(s) and Model(s)	Intel® Xeon® E5-2699 v4
Processor Speed(s) (GHz)	2.2 GHz
Total Sockets/Total Cores/Total Threads	2 Sockets/ 44 Cores/ 88Threads
Memory per Virtual Client Host	256GB
Network Controller(s) Vendors and Models	QN 10GbE 82599ES
Virtual Client Networking Notes	1 vmnic for management and 1vmnic for workload traffic
Virtual Client Storage Notes	2TB Intel® SSD DC P3500
Other Hardware	N/A



## Notes for Workload

Virtualization Software Notes
Logging was disabled for all VMs except the PrimeClient.
CDROM & floppy devices were removed on all VMs except the PrimeClient.
Logical CPU layout changed for all multi-cpu VMs to 1 socket with multiple cores.
All DS3DB, Clients, and PrimeClient VMs had CPU share set to high.
All DS3DB and Clients VMs had memory share set to high.
All memory reserved for DS3DB, Clients, and PrimeClient VMs.
All VMs are virtual hardware version 13 (default 11).
Cluster DRS Automation Level was set to "Fully Automated" and DRS Migration Threshold was set to level 2.
Config.HostAgent.log.level = warning (default info).
Power management Policy = High Performance (default: Balanced).

Server Notes - Server BIOS setting
Intel® Hyper Threading Technology Enabled (default: Enabled).
Intel® Turbo Boost Technology Enabled (Frequency boost up to 3.70 GHz) (default: Enabled).
Intel® Speed Shift Technology: HWP Native (default: HWP Native).
C-States: Disabled (default: Disabled).
Power Management Settings: Performance (default: Balance).
Intel® TXT Enabled (default: Disabled).

Networking Notes - vSwitch Configuration
vSwitch0 for the management, vMotion, iscsi, and vSAN™ traffic on vmnic0 at 10Gb/s.
vSwitch1 for all VMs on vmnic1 at 10Gb/s.

Storage Notes	
OS was installed on SATADOM for each host.	
Storage box#1 (System Under Test): vsanDatastore (32TB), 3TB cache used.	<b>Hardware Configuration: QuantaGrid D52B-1U.</b>
	2 x Intel® Xeon® Gold 6152.
	384 GB Memory (12 x 32GB DDR4 2666MHz RDIMM).
	2 x Intel® SSD DC P4800X, 375GB, Gen 3x4 Gbps (used in internal).
	4 x Intel® SSD DC P4500, 2TB, Gen 3x4 Gbps (used in shelf).
	<b>Software Configuration</b>
	All Flash vSAN™ with two disk groups. Each disk group includes 1 Intel® SSD DC P4800X 375GB (AIC) for array cache and 2 Intel® SSD DC P4500 Series 2TB (NVMe) for capacity. vSAN™ Default Storage Policy used for all disk objects.
<b>LUN/VM layout</b>	
LUN1(32TB, 3TB cache): DS3DB*, DS3WebB*, DS3WebC*, AuctionLB*, AuctionNoSQL*, AuctionAppA*, AuctionAppB*, AuctionDB*, AuctionMSQ*, AuctionWebA*, AuctionWebB*, ElasticLB*, ElasticAppA*, ElasticAppB*, ElasticDB*, ElasticWebA*, ElasticWebB* for tile 0-9, template, SVmotion source LUN, XVMotion source LUN, Deploy source LUN.	
Storage box#2 (Shared storage): Windows Server 2012 with two	<b>Hardware Configuration: QuantaPlex T41S-2U (only use one node)</b>
	2 x Intel® Xeon® E5-2620 v4 processors.
	256GB Memory (16 x 16 GB DIMMs 2400MHz DDR4).
	1 x Quanta SAS 3008.



volumes shared as iSCSI target to System Under Test (2.6TB). No cache is used.	1 x Intel® SSD DC S3710 200GB, 2.5", SATA 6Gb/s (used in shelf).
	2 x Intel® SSD DC S3710 1.2 TB, 2.5", SATA 6Gb/s (used in shelf).
	<b>Software Configuration</b>
	Storage Pools were created using Windows Server 2012.
	2 x Intel® SSD DC S3710 1.2 TB were built into one storage pool (RAID0). Storage Pool 1, Volume 1(RAID 0), LUN1.
	1 x Intel® SSD DC S3710 200GB was built into one volume. Volume 1, LUN2.
	<b>LUN/VM layout</b>
	LUN1 (2TB, No cache): DS3WebA*, Standby* for tile 0-9, and Deployed LUN. LUN2 (200GB, no cache): SVMotion target LUN, XVMotion target LUN, and Deployed LUN.

<b>Datacenter Management Server Notes</b>
QuantaPlex T41S-2U running VMware® ESXi 6.5U1, Build 5969303 with VMware® vCenter Server 6.5 U1, Build 5973321 for SUT and client hosts.
4 vCPUs.
16GB virtual memory.
vCenter Server 6.5.0-5973321.

<b>Operating System Notes</b>
VMware® ESXi 6.5U1, Build 5969303.

<b>Client Notes</b>
QuantaGrid D51BP-1U was used for client hosts.
Client hosts have updated from ESXi 6.5, Build 4564106 to ESXi 6.5U1, Build 5969303.
Client host1: Client0, Client1, Client2.
Client host2: Client3, Client4, Client5.
Client host3: Client6, Client7, Client8, PrimeClient.

## 5.2.2 Network Connectivity

A well-designed network topology ensures that vSAN™ traffic can run efficiently, correctly, and available between hosts. VMware® suggests 10 Gigabit Ethernet (GbE) network for vSAN™ to avoid the bottleneck of network bandwidth. In our network design, a VMkernel port is created on the vSwitch, dedicated for vSAN™ traffic. If administrators need to run multiple traffic on the same NIC, Quality of Service (QoS) can be implemented using Network I/O Control (NIOC). NIOC can allocate the dedicated network bandwidth to the vSAN™ traffic. To achieve the network redundancy on the uplink network adapter, NIC teaming is a way to “team” two NICs together. The two uplink adapters “vmnic1” and “vmnic2” are teamed on a vSwitch to either “two active uplinks” or “an active/standby uplink” for failover and redundancy purposes. In Fig. 10, two physical switches with link aggregation are prepared on the upper layer. The real network topology is different depending on users’ requirements.



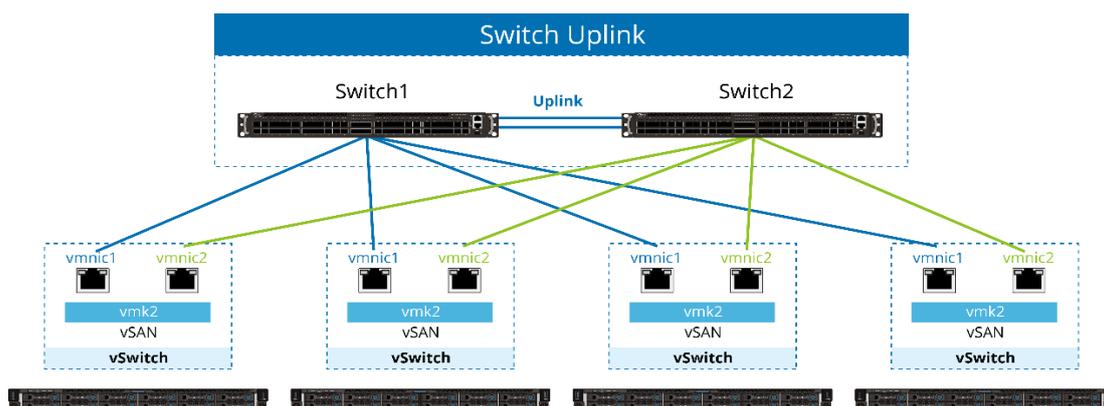


Figure 10. Network Planning of QxStack vSAN ReadyNode™ – Performance Optimized SKU.

### 5.2.3 Software Configuration for VMmark

QCT implemented the firmware and software optimization according to the guidance of Intel® Select before running the VMmark test. The server BIOS setting is elaborated below.

- Intel® Hyper Threading Technology Enabled (Default: Enabled).
- Intel® Turbo Boost Technology Enabled (Frequency boost up to 3.70 GHz) (Default: Enabled).
- Intel® Speed Shift Technology: HWP Native (Default: HWP Native).
- C-States: Disabled (Default: Disabled).
- Power Management Settings: Performance (Default: Balance).
- Intel® TXT Enabled (Default: Disabled).

### 5.2.4 Storage Configuration for VMmark

In the storage configuration, three shared storages listed below are used for different purposes.

#### Shared Storage 1: vSANDatastore

Four QuantaGrid D52B-1U servers contribute their local storage to aggregate into a vSANdatastore. Each server contains two disk groups. Each disk group includes 1 Intel® SSD DC P4800X 375GB (AIC) for cache and 2 Intel® SSD DC P4500 Series 2TB (NVMe) for capacity. The vSANDatastore provides a shared storage for virtual machines when VMmark runs. It also acts as SVMotion source LUN, XVMotion source LUN, and the deployed source LUN.

#### Shared Storage 2: iSCSI volume provided by Windows Server 2012

In Windows Server 2012, 2 Intel® SSD DC S3710 1.2TB are configured to one

storage pool for providing the volume to system under test (SUT) through iSCSI. The volume provides a shared storage for DS3WebA and Standby virtual machines. It also acts as the deployed LUN when VMmark runs the test.

### Shared Storage 3: iSCSI volume provided by Windows Server 2012

In Windows Server 2012, 1 Intel® SSD DC S3710 200GB is configured to one volume for providing the volume to system under test (SUT) through iSCSI. The volume provides a datastore for SVMotion target LUN, XVMotion target LUN, and the deployed target LUN.

### 5.2.5 Configuration Diagram

A logical layout for the test environment is shown in Fig. 11. The layout describes the relationship between the virtualization layer and hardware.

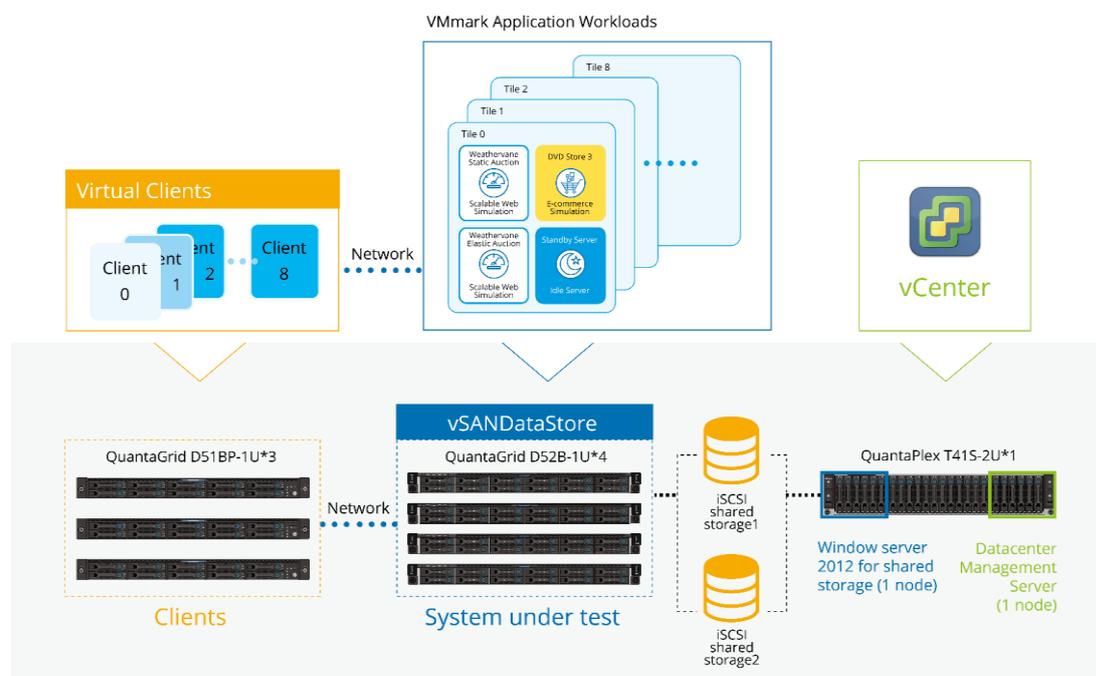


Figure 11. Logical Layout of Test Environment.

### System Under Test

To fulfill Intel® Select Program’s requirement, four QuantaGrid D52B-1U servers are configured to system under test (SUT). The SUT must also be configured to meet or exceed the requirements listed in VMmark User Guide.

### Clients

Each tile requires a client system to provide loads for the application workload



mentioned above. A prime client is needed to run the whole VMmark test process by activating the VMmark Harness software to measure the workload. Each VMmark tile is paired with a virtual client system that drives the tile's virtual machines to perform a variety of tasks. Three QuantaGrid D51BP-1U servers are used to build a cluster for hosting the virtual clients. Each server hosts three virtual clients, that is, nine virtual clients in total.

### Data Center Management Server

One ESXi server is used to install vCenter appliance to manage the entire test system including clients, SUT, and shared storages.

## 5.3 Test Result

### Tile

Tile is a unit in VMmark which contains nineteen virtual machines to simulate a collection of diverse workloads. Each tile includes a scalable web, E-commerce, and standby system running on different virtual machines. The tile consumes either compute or storage resources when running the benchmark test. Each application or infrastructure workload generates its relevant performance metric after the benchmark tool finishes the tests. The performance metric in each workload should meet the minimum quality-of-service to make the result compliant when another tile is added. The total number of tiles can help administrators to evaluate the overall performance and the number of virtual machines on a virtualization platform after VMmark completes the workload tests.

### Score

VMmark collects performance metrics with diverse workloads and normalizes the score based on a reference system provided in [SPEC](#). When the test is finished, VMmark generates a single overall score to be compared in different virtualization platforms. The score is calculated based on a weighted average of the application and infrastructure workload. It gives 80% weights to the application workload and 20% to the infrastructure workload. When VMmark runs multi-tiles, an aggregation score is reported by accumulating individual score in each tile. For more information, please refer to [VMmark User's Guide](#).

To reach an outstanding performance, QCT cooperates with Intel to optimize the setting parameters by adopting Intel® Select solution's base configuration. The test result reaches up to 9 tiles running on the four systems. Each tile contains standby

system, scalable web simulation, and E-commerce simulations. As shown in Table 2, the aggregation score reaches to 8.51, which dramatically exceeds the Intel® select program’s minimum performance threshold. Please refer to [QCT VMmark test result](#) published on VMware VMmark website.

Table 2. VMmark Test Result.

Workload	Score
VMmark3_Applications	10.29
VMmark3_Infrastructure	1.37
VMmark3_Total	8.51

## 6 Conclusion

Nowadays data center transformation is considered to be a main trend to address the dynamic business environment. By joining Intel® Select Program, QCT collaborated with Intel® to provide innovative and flexible solutions to keep your organizations stay in a leading position.

**QxStack vSAN ReadyNode™ – Performance Optimized SKU** is a high-performance data center solution that is rapidly deployed, easy to manage, highly certified, and fully integrated into the industry-leading software-defined storage, vSAN™. This reference architecture has been proven that the outstanding performance of the solution surpasses the specific threshold in Intel® Select Program. This solution can minimize customers' time and expense in evaluation, selection, deployment and tuning of a solution.

With the knowledge of QCT and Intel®, customers can leverage the collective results and have a simplified path to future-defined data center by adopting this solution.

QCT always stay innovative. QCT appreciates any feedback from you. For further inquiry, please visit <http://go.qct.io/solutions/enterprise-private-cloud/qxstor-vmware-edition-vsan-readynode/>



## 7 References

1. **VMware® vSAN™ Design and Sizing Guide 6.5**  
[https://www.google.com.tw/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&uact=8&ved=0ahUKEwiryPOjudjVAhUKvbwKHTGOBpkQFggpMAA&url=https%3A%2F%2Fstoragehub.vmware.com%2Fexport\\_to\\_pdf%2Fvmware-r-virtual-san-tm-design-and-sizing-guide&usg=AFQjCNHSBusC8pMJnrcP0XVX\\_dXBceQ4bQ](https://www.google.com.tw/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&uact=8&ved=0ahUKEwiryPOjudjVAhUKvbwKHTGOBpkQFggpMAA&url=https%3A%2F%2Fstoragehub.vmware.com%2Fexport_to_pdf%2Fvmware-r-virtual-san-tm-design-and-sizing-guide&usg=AFQjCNHSBusC8pMJnrcP0XVX_dXBceQ4bQ)
2. **VMware® VMmark Virtualization benchmark**  
<https://www.vmware.com/products/vmark.html>
3. **Product Brief: Intel Optane™ SSD DC P4800X Series**  
<https://www.intel.com/content/www/us/en/solid-state-drives/optane-ssd-dc-p4800x-brief.html>





## About QCT

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

Product lines include hyper-converged and software-defined datacenter solutions as well as servers, storage, switches, integrated racks with a diverse ecosystem of hardware component 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>

### 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 号润城中心 2 号楼  
Tower No.2, Run Cheng Center, No.12, East Bridge Rd.,  
Chaoyang District, Beijing, China  
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  
Room 303, Building No.4, ZheShang Wealth Center  
No. 83 GuDun Road, Xihu District, Hangzhou, Zhejiang, China  
TEL: +86-571-2819-8650

### Japan

Quanta Cloud Technology Japan 株式会社  
日本国東京都港区芝大門二丁目五番八号  
牧田ビル 3 階  
Makita Building 3F, 2-5-8, Shibadaimon,  
Minato-ku, Tokyo 105-0012, Japan  
TEL: +81-3-5777-0818  
FAX: +81-3-5777-0819

### 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  
TEL: +886-3-286-0707  
FAX: +886-3-327-0001

### Germany

Quanta Cloud Technology Germany GmbH  
Hamborner Str. 55, 40472 Düsseldorf, Germany  
TEL: + 492405-4083-1300

### 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

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 © 2017-2018 Quanta Computer Inc. All rights reserved.