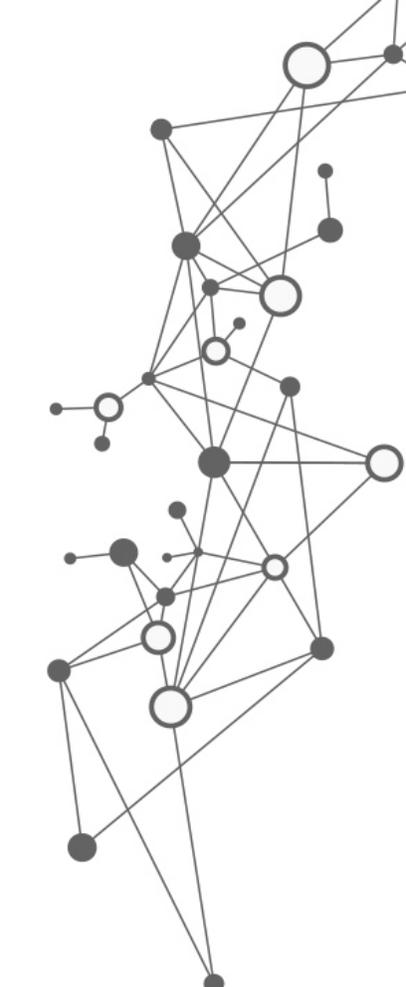


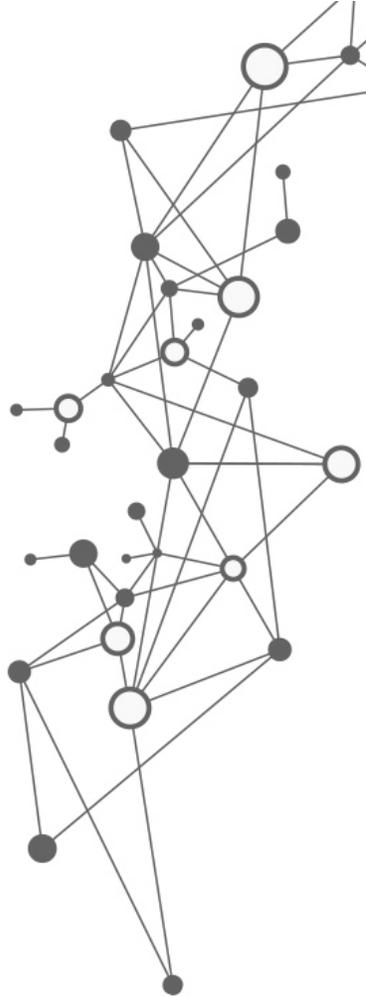
QCT HCI Security Function Reference Architecture





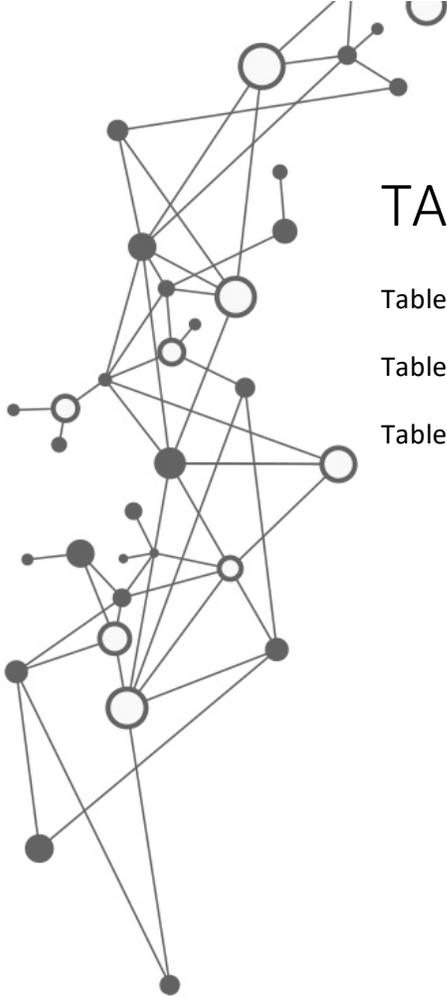
CONTENT

1. EXECUTIVE SUMMARY	5
2. INTRODUCTION	6
2.1. Purpose.....	6
2.2. Scope.....	6
2.3. Audience.....	6
3. SOLUTION OVERVIEW	7
4. SOLUTION ARCHITECTURE AND VALIDATION.....	9
4.1. Hardware Architecture	9
4.2. Software Architecture	9
4.3. Solution Validation.....	10
5. SOLUTION SCENARIO	11
5.1. Host Protection.....	11
5.1.1. <i>Secure Boot in UEFI</i>	11
5.1.2. <i>Trusted Platform Module</i>	15
5.2. Data Protection.....	17
5.2.1. <i>Protection on Virtual Machine</i>	17
5.2.2. <i>Encryption on shared storage</i>	21
5.2.3. <i>Migration with a secured way</i>	23
5.3. Traffic Protection	25
5.3.1. <i>North-South firewall</i>	25
5.3.2. <i>East-West Firewall</i>	29
6. CONCLUSION	35
7. REFERENCE	36
LEGAL DISCLAIMER	36



FIGURES

Figure 1. Overall security protection of QCT HCI Solutions.....	7
Figure 2. QuantaPlex T42S-2U Server.....	9
Figure 3. Process of host boot.....	12
Figure 4. Enablement of Secure Boot on a host.....	13
Figure 5. PSOD of ESXi triggered by Secure Boot.....	14
Figure 6. Halt of unsigned VIB installation and change of host's acceptance level.....	14
Figure 7. Host's boot process with TPM.....	15
Figure 8. Discovery of TPM 2.0 device in BIOS.....	16
Figure 9. ESXi attestation results shown in vCenter.....	16
Figure 10. Detailed encryption process of a VM.....	18
Figure 11. vTPM device and replaceable certificates.....	20
Figure 12. Enablement of Secure Boot for a VM.....	21
Figure 13. Detailed encryption process of data in vSAN storage.....	22
Figure 14. Enablement of vSAN encryption.....	23
Figure 15. Detailed process of encrypted vMotion.....	24
Figure 16. Enablement options of encrypted vMotion.....	25
Figure 17. Edge firewall designed to filter ingress and egress traffic.....	26
Figure 18. Firewall rules established in edge appliances management tablet.....	27
Figure 19. Edge firewall filtering on the connections with specified rules.....	27
Figure 20. Embedded NAT function in the edge appliance.....	28
Figure 21. VPN functions in the management tablet of edge appliance.....	29
Figure 22. Architecture of distributed firewall rules for VMs.....	30
Figure 23. Firewall rules assigned to different objects.....	31
Figure 24. Firewall rules configured on Layer 2 network.....	31
Figure 25. Firewall rules configured on Layer 3 network.....	32
Figure 26. Rules of context-aware firewall configured on applications.....	32
Figure 27. Selection of members for security group.....	33
Figure 28. Created security group with specified AD users.....	34
Figure 29. Identity firewall rule based on AD user groups.....	34



TABLES

Table 1. Testbed configuration.	9
Table 2. Software configuration.....	9
Table 3. QCT HCI Solutions with vSAN ReadyNode certification.....	10

1. Executive Summary

With the development of information technology, our society has become more and more advanced and prosperous. New threats emerge day by day and information security is now a vital management issue for business.

According to “The Global Risks Report 2020”, The World Economic Forum has placed cyberattacks as the seventh most likely and the eighth most impactful global risk. For conducting business globally over the next 10 years, cyberattacks even ranks as the second most concerning risk. The Forum indicates that on the critical infrastructure, cyberattacks have become the new normal across different industries such as healthcare, energy and transportation. For example, the anonymized health data of individuals are facing the threat of cyberattacks with the risks of being identified from the data. Furthermore, large-scale cyberattacks will cause the serious economic damage, geopolitical tensions or even the widespread loss of trust in the internet. From the business perspective, insufficient security control could not only cause tangible damage such as property loss but also intangible loss such as brand image, customer trust, and core competitiveness of a company.

In order to conquer the difficulty of the security and manage the risk of information control, Quanta Cloud Technology (QCT), a global data center provider, discreetly selects security functions and integrates systems to provide solutions with overall protection in diverse scenarios, providing customers end-to-end protection from day 0 to day 2.

In this document, QCT introduces the security structure of QCT HCI Solutions, illustrates the solution integration and validation, and demonstrates the methodology of security functions at the levels of host protection, data protection, and traffic protection. With QCT HCI Solutions, customers can easily take actions to address the challenges in the security control and strengthen overall protection of data centers.

2. Introduction

2.1. Purpose

To introduce several security protection functions integrated in QCT's HCI Solutions, including QxStack vSAN ReadyNode Series, QxStack powered by VMware Cloud Foundation, and QxVDI Series, this RA will demonstrate how the integration and validation works, and the benefits that customers gain from QCT's HCI Solutions.

2.2. Scope

- Introduces overall security structure and benefits in QCT HCI Solutions.
- Demonstrates the architecture of QCT HCI Solutions and validation work.
- Illustrates the detailed security functions applied at different levels or scenarios.

2.3. Audience

The intended audiences of this document are IT professionals, technical architects, and sales engineers who would like to further understand the security functions of QCT HCI Solutions.

3. Solution Overview

Technology brings many benefits to operational practices, but it also poses potential risk to businesses. For instance, banks could bear a great capital loss in a cyberattack; social media companies could destroy users' confidence when personal data leakage occurs; the competitiveness of technology firms could be diminished once the confidential strategy or intellectual property is stolen by internal or former employees. Unexpected data loss or tampering in a hospital could even impact patients' health. With these possible risks, enterprises nowadays are eager to look for a reliable security solution for their system and data center.

QCT, as a global data center solutions provider, foresees the importance of an overall protection in terms of data centers. To address the complicated security threats, QCT discreetly chooses multiple security functions and integrates them into QCT's HCI Solutions. These functions are carefully validated in QCT's HCI solution product lines, including QxStack vSAN ReadyNode Series, QxStack powered by VMware Cloud Foundation, and QxVDI Series to provide powerful and reliable protection for data centers.

The security management process in different scenarios, including host level, data level, and traffic level, is taken into consideration, as shown in Figure 1.

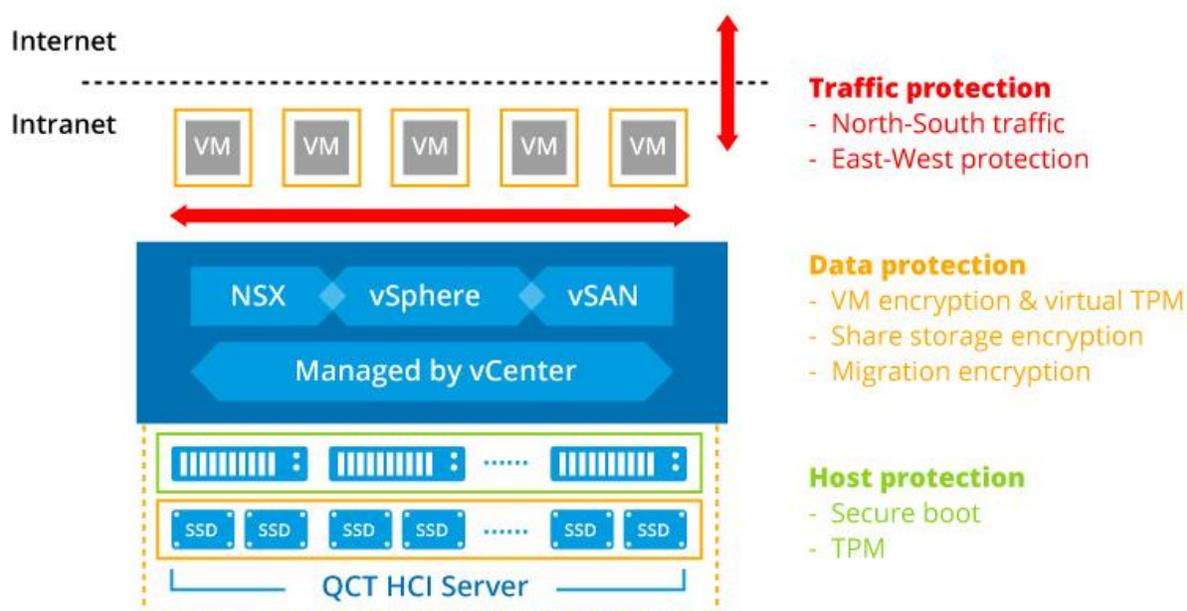


Figure 1. Overall security protection of QCT HCI Solutions.

Host level

A data center may be attacked by some malware when administrators boot a system, which can expose the data center to great threats like confidential data leakage and even malicious control. In order to prevent unauthorized modifications of the boot kernel, the servers that QCT HCI Solutions adopt are all equipped with Secure Boot in the server firmware.

On top of that, the solutions also integrate the latest technology - Trusted Platform Module (TPM)¹ to validate the integrity of booting process, providing a second layer of confirmation. Both functions keep data centers away from malicious attack at the host level.

Data level

Data is a valuable asset for enterprises in all industries. Loopholes in the security design of a system potentially exposes data to loss, leakage, or even tampering, which can result in great loss for enterprises.

QCT HCI Solutions integrate industry-leading virtualized software - VMware vSphere® and VMware vSAN™ to address the risks. With the integration work, valuable data can be well protected in diverse scenarios such as disk storage running apps, and data migration.

Traffic level

Traffic between Internet and data centers is one of the main targets for hackers. To strengthen the protection at different levels, QCT integrates VMware NSX® in QCT HCI Solutions to create virtual firewalls and precisely design security policies on diverse objects such as identity, IP sets, logical switch, resource pool, clusters, and even context.

QCT HCI Solutions provide overall and reliable protection in different scenarios. With QCT's integration and validation working, customers can rest assured that valuable data will be protected. The system validations and security functions are respectively detailed in sections 4 and 5.

¹ Trusted Platform Module (TPM) is an optional component in QCT HCI Solutions. For more information, please click <http://go.qct.io/contact/contact-qct-solutions/> for contact

4. Solution Architecture and Validation

The solution utilizes QCT HCI server T42S-2U as the test platform, which loads hardware components and software to demonstrate the integration of security functions and virtualized infrastructure.

4.1. Hardware Architecture

The four-node QuantaPlex T42S-2U server is utilized as the testbed to prove the availability of security features for QCT HCI architecture, as shown in Figure 2. Some hardware components are chosen to build the hyper-converged data center with the security functions operated, as shown in Table 1.



Figure 2. QuantaPlex T42S-2U Server.

Table 1. Testbed configuration.

Item	Model	Quantity/Per node	Quantity/Per server
CPU	Intel Xeon Gold 5220 2.2GHz	2	8
Memory	DDR4-2666 32GB	8	32
SSD	2.5" SATA 960GB	1	4
HDD	2.5" SAS 1.8TB	5	20
Boot device	M.2 SSD 240GB	1	4
TPM	Intel TPM module	1	4
SAS Mezz.	SAS 3008A mezzanine	1	4

4.2. Software Architecture

The software suite adopted in the solution includes the VMware vSphere, vSAN, and NSX that respectively cover the virtualization of data center, storage, and network. Security functions are enabled via the software. The security functions at the host level target on ESXi and server node; the security functions at data level target on virtual machine (VM); the security functions at network level target on data center traffic. The adopted software suite is shown in Table 2.

Table 2. Software configuration.

Software	Version	License Edition
VMware vSphere®	6.7	Enterprise Plus
Per node ESXi™ hypervisor	6.7	Enterprise Plus
VMware vSAN™	6.7	Enterprise
VMware NSX®	6.4	Enterprise Plus
vCenter Server® Appliance	6.7	Standard

4.3. Solution Validation

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. All solution details, including hardware components, firmware and driver, and software stack should be strictly examined to meet the rigorous requirements. To ensure QCT HCI Solutions pass vSAN ReadyNode™ certification, QCT spent a lot of effort assessing the feature of different server models, processing many phases of standardized validations, selecting components, configuring settings in different scenarios, and optimizing performance. By passing this certification, the HCI solutions are proven to be reliable. QCT HCI Solutions, including QxStack vSAN Ready-Node Series, QxStack powered by VMware Cloud Foundation, and QxVDI Series, pass vSAN ReadyNode's strict certification process. Customers can find QCT HCI Solutions on the [direct link](#) of VMware Compatibility Guide.

Table 3. QCT HCI Solutions with vSAN ReadyNode certification.

vSAN ReadyNode™ Details		
Model: HY4 -QCT-QuantaPlex T42S-2U Profile: HY-4 Series Type: Hybrid Partner Name: Quanta Computer Inc Generation: Gen3 - Xeon Scalable		
Components	Details	Quantity
SKU	QuantaPlex T42S-2U_HY4	
System	Model: QuantaPlex T42S-2U System Type: Rackmount	4
CPU	Intel® Xeon® Silver 4210 Processor (8core, 11M Cache, 2.10 GHz, 85 W)	8
Memory	32GB 2666MHz DDR4 RDIMM	32
Caching Tier	Model: Intel® SSD D3-S4610 Series SSDSC2KG960G8 (960GB, 2.5") (OR) Samsung SATA SSD SM883 Series MZ7KH960HAJR00005 (960GB, 2.5-inch) Partner Name: Intel Device Type: SATA Capacity: 960 GB Performance Class: Class E: 30,000-100,000 writes per second TBW Endurance Class: Endurance Class C >=3650 TBW	4
Capacity Tier	Model: ST1800MM0129 Partner Name: Seagate Device Type: SAS Capacity: 1800 GB	20
Controller	Model: Quanta 3008A Queue Depth: 2936	4
NIC	Model: ON 10GbE 82599ES	4
Boot Device	Model: M.2	4

5. Solution Scenario

This section describes the data center security in different scenarios, including host level, data level, and network level.

5.1. Host Protection

The weaknesses on system software, such as hijacking by a malicious root kit and the unauthorized modification of the boot kernel can compromise the hosts and put the cloud at risk. To prevent the boot process from tampering and deter the operation of untrusted code, the Secure Boot protocol and the TPM microchip are implemented in QCT HCI Solutions to assure that the hypervisor is well protected.

5.1.1. Secure Boot in UEFI

The Unified Extensible Firmware Interface (UEFI) is a specification created by UEFI consortium to standardize an interface between operating system and firmware, and it is also a replacement for the BIOS. The Secure Boot is a protocol of the UEFI designed to ensure that the system's boot loader will not be tampered by comparing its digital signature against a digital certificate stored in the UEFI firmware. The digital signature is embedded with the boot loader and is chained to the certificate in the UEFI firmware. To form the first defense of host's system boot, any unauthorized alteration of the boot loader will change the digital signature and the Secure Boot will discover the mismatch between signature and certificate, and further forbid booting.

ESXi is composed of the elements such as boot loader, vmkboot, vm kernel, Secure Boot verifier, and vSphere installation bundles (VIB). The following sequence will describe how the ESXi boot process is protected by Secure Boot, and the illustration is shown in Figure 3.

1. When the host powers on and boots the ESXi, the hardware will first load the UEFI.
2. The UEFI verifies the boot loader's digital signature against the digital certificate in which the digital certificate is stored in the firmware provided by vendors.
3. The boot loader loads the vmkboot. The VMware public key is stored in vmkboot.
4. The vm kernel is signed using the VMware private key and the boot loader verifies the kernel with the public key.
5. The kernel runs Secure Boot verifier. The VMware public key is stored in the Secure Boot verifier.
6. The Secure Boot verifier validates all the VIBs' digital signature, which is chained to the key in the Secure Boot verifier.
7. The ESXi management applications such as daemon and DCUI begin to run after all the validation processes are finished.

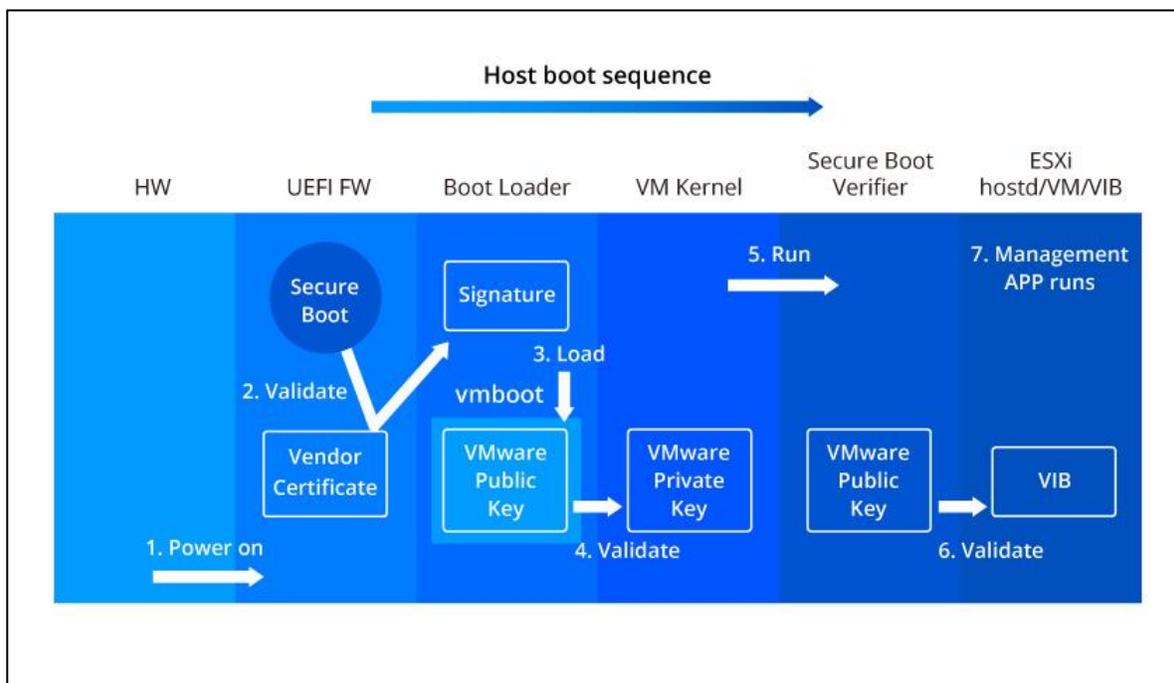


Figure 3. Process of host boot.

To protect ESXi booting using Secure Boot, the following prerequisites are needed:

- Secure Boot is enabled in UEFI.
- The VIBs are signed with at least “partnersupported” VMware acceptance level.
- The hardware supports UEFI Secure Boot.

Administrators can enable the Secure Boot in the BIOS, as shown in Figure 4.

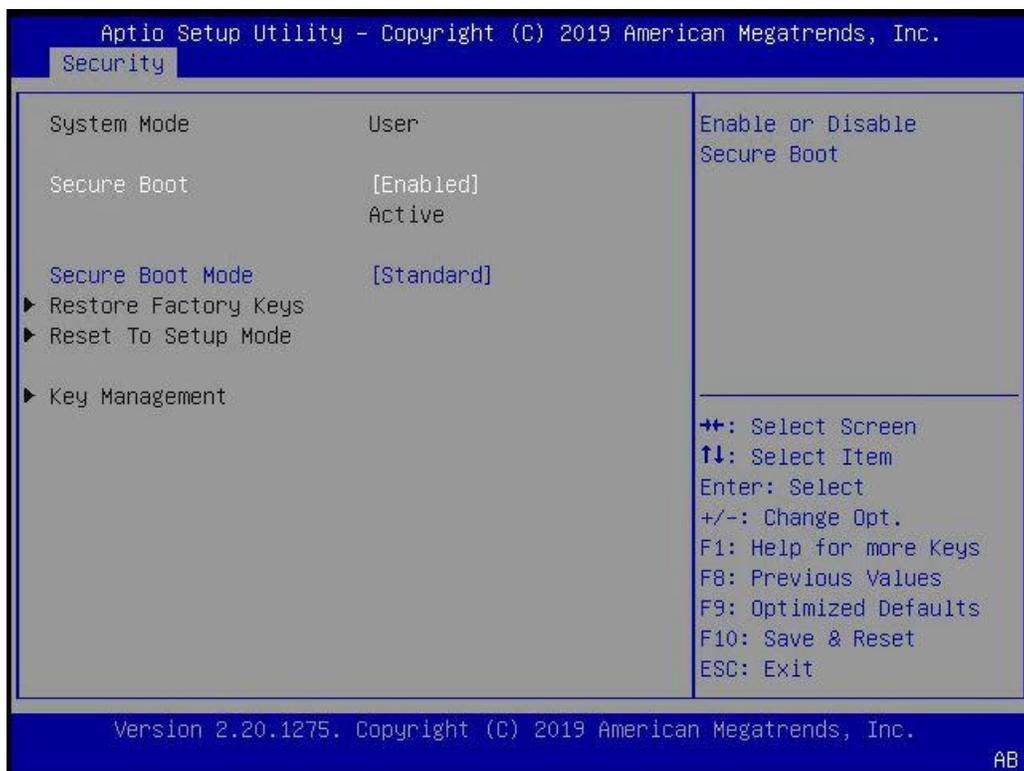


Figure 4. Enablement of Secure Boot on a host.

The VIB is an important ESXi software package that ESXi file system is built upon. VMware defines four “acceptance levels” for VIBs and ESXi hosts for verification standard. The acceptance level of VIBs cannot be changed but the hosts’ acceptance level can be altered. VIBs’ acceptance level has to reach at least the same level as hosts’ acceptance level so that VIBs can be installed. The four certified levels defined by VMware are listed below:

- VMwareCertified: VIBs need to be thoroughly tested and certified by VMware. This is the most rigorous certification standard;
- VMwareAccepted: VIBs are tested by partners and the results need to be verified by VMware;
- PartnerSupported: VIBs are tested by the partners that VMware trusts in. The partners’ test results do not need to be verified by VMware;
- CommunitySupported: VIBs are created by individuals or organizations that are not VMware’s partner and the VIBs do not need to be validated by VMware.

Typically, the protection of Secure Boot is to prevent the use of unsigned VIB. If the unsigned VIBs are inserted into the ESXi installed on the server, the Secure Boot can halt the boot of this ESXi to stop the operation of unsigned VIBs. The purple screen of death (PSOD) triggered by Secure Boot is shown in Figure 5.

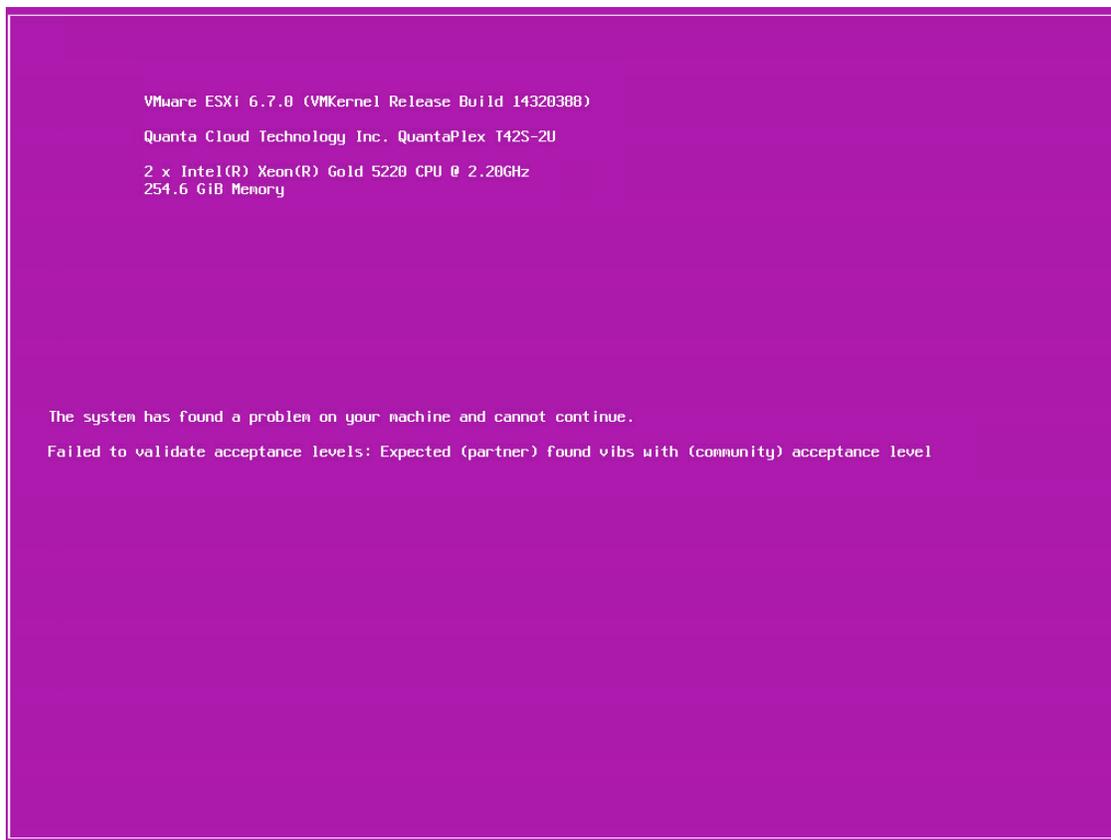


Figure 5. PSOD of ESXi triggered by Secure Boot.

The Secure Boot can also halt the installation of unsigned VIB and the change of VIB acceptance level on the installed ESXi. An example of “community-supported” VIB installation with Secure Boot enabled is shown in Figure 6. The installation of unsigned VIB and the change of acceptance level on the host halted by Secure Boot can protect the host from unknown codes and changes.

```

[root@s5s04node:~] esxcli software vib install -f -v /vmfs/volumes/datastore1\ \ (3)\ /
[redacted].vib
[AcceptanceConfigError]
VIB [redacted]'s acceptance level is community
, which is not compliant with the ImageProfile acceptance level partner
To change the host acceptance level, use the 'esxcli software acceptance set' comman
d.
Please refer to the log file for more details.
[root@s5s04node:~] esxcli software acceptance set --level=CommunitySupported
[AcceptanceConfigError]
Secure Boot enabled: Cannot change acceptance level to community.
Please refer to the log file for more details.

```

Figure 6. Halt of unsigned VIB installation and change of host's acceptance level.

In conclusion, the Secure Boot function provides ESXi a firmware-level boot verification and forms a basic check on the system. The advantages of implementing Secure Boot include:

- Assuring that the hypervisor only boots with a signed boot loader validated by the firmware.
- Verifying if VIBs match the VMware digital certificate.
- Preventing the unauthorized VIBs installation after the boot.

5.1.2. Trusted Platform Module

Trusted Platform Module (TPM)² is a microchip that can securely store values such as measurements, certificates, or encryption keys to authenticate the platform. It can also be utilized to digitally sign content and store platform information to ensure that the system is trustworthy without any unauthorized modification. To form a second defense for the ESXi, TPM can verify if the ESXi boot process is validated by Secure Boot. By combining Secure Boot and TPM, customers can get multiple layers of protection on ESXi boot process.

Before adopting TPM, customers need to prepare an ESXi version with TPM support. ESXi 6.7 with TPM version 2.0 support is adopted in this reference architecture. The measurements stored in TPM will be compared with what ESXi submits by vCenter.

During the boot of ESXi host with Secure Boot and TPM enabled, ESXi will be verified by Secure Boot (as mentioned in section 5.1.1). The component, vmkboot, within the boot loader will input the hash values of modules and settings into TPM chip. After the host completes the boot, vCenter compares the hash values in TPM against the hash values and metadata in ESXi logs, and eventually determines if the attestation passes or fails, as shown in Figure 7.

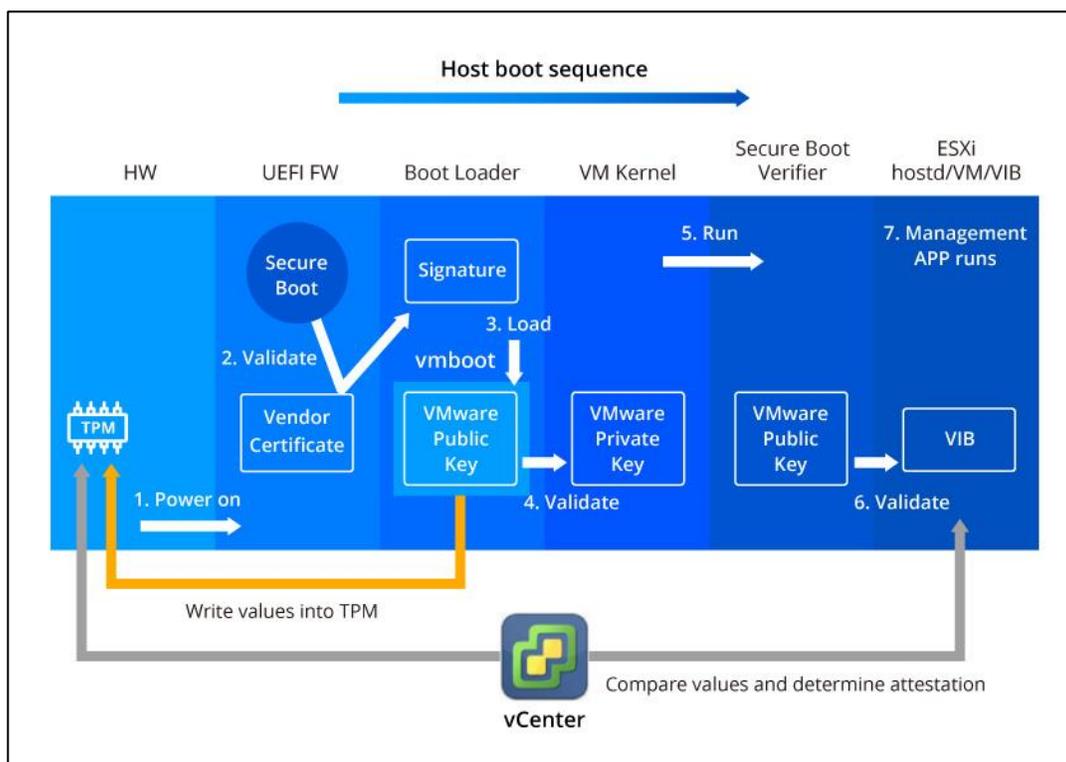


Figure 7. Host's boot process with TPM.

To utilize the TPM chip, the chip is enabled in the BIOS so that the TPM can be detected and operate normally, as shown in Figure 8.

²Trusted Platform Module is an optional component in QCT's solutions, QxStack vSAN ReadyNode Series, QxStack powered by VMware Cloud Foundation, and QxVDI Series.



Figure 8. Discovery of TPM 2.0 device in BIOS.

After the TPM chip on the server node is activated and ESXi boot is finished, the vCenter Server will display if the attestation of the ESXi host is passed. If the host's attestation is failed, the alarm will appear to inform administrators. In this example, the Secure Boot is disabled on a host which makes the attestation failed, as shown in Figure 9.

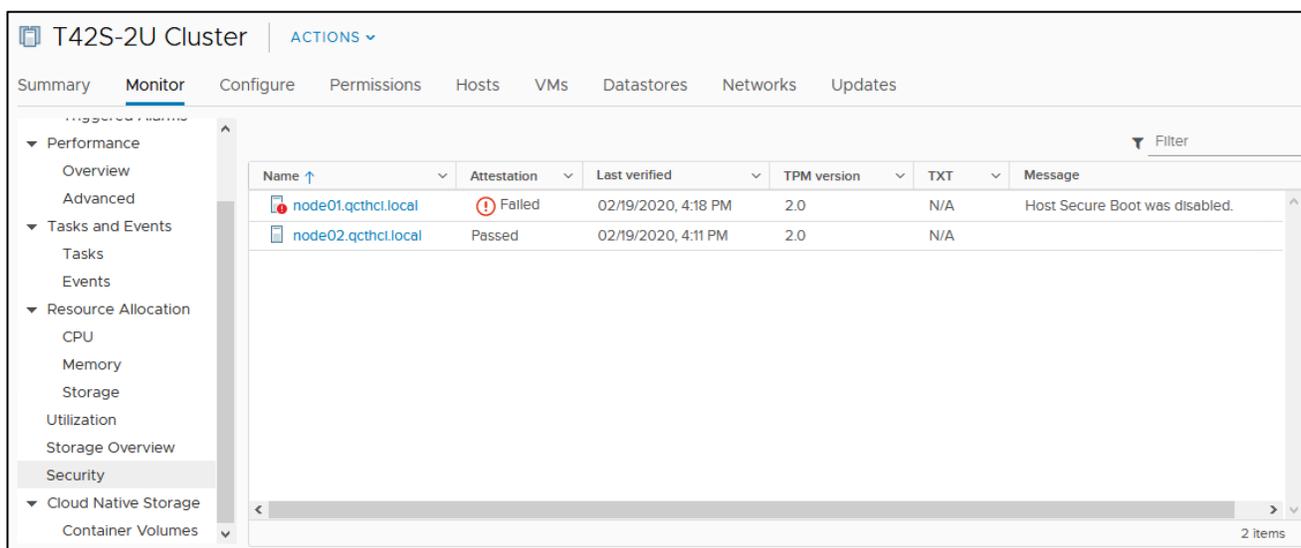


Figure 9. ESXi attestation results shown in vCenter.

In conclusion, TPM offers a mechanism to assure that ESXi is booted with Secure Boot turned on. With the combination of firmware Secure Boot function and TPM microchip, a root of trust can be started on every node in a

data center which not only achieves foundational security on hosts but also establishes effective security policies in the cloud.

5.2. Data Protection

To prevent cloud data from hijack during migration and support Windows in-guest protection features, the VM encryption, vSAN encryption, virtual TPM, VM Secure Boot, and encrypted vMotion are good methods to avoid data from leaking. By adopting these methods, QCT HCI Solutions can provide strong encryption for the valuable data in the cloud.

5.2.1. Protection on Virtual Machine

VM Encryption

Traditionally, the encryption solutions commonly used by enterprises are in-guest and infra-based solutions, such as guest encryption, Self-Encrypting Drives (SED), Host Bus Adapter (HBA) encryption, and switch encryption. The challenges of these solutions are:

- Inconsistency of cross-platform encryption policy.
- Data exposure before entering the encryption medium.
- Complexity of configuring physical host.
- Difficulty in migrating the data between datastores and hosts.

The VM encryption is a data protection method which enforces the encryption on VM files and the mechanism involves three components, Key Management Server (KMS), Key Encryption Key (KEK), and Data Encryption Key (DEK). KMS manages and distributes encryption keys according to hosts' demands; KEK is an AES-256-bit key provided by KMS and utilized by hosts to encrypt the DEKs; DEK is also an AES-256-bit key created by hosts and used to encrypt VMs.

While performing the VM encryption, an ESXi host demands the encryption key from vCenter. The vCenter then demands a KEK from the KMS and passes to the ESXi host. The ESXi host generates DEK to encrypt the VM and their disks, and then encrypt the DEK with KEK from vCenter. To reach solid workload protection, the VM encryption utilizes AES-256-bit key which has high strength to protect VMs from hacking. The VM encryption process in the QCT HCI solution is detailed in Figure 10.

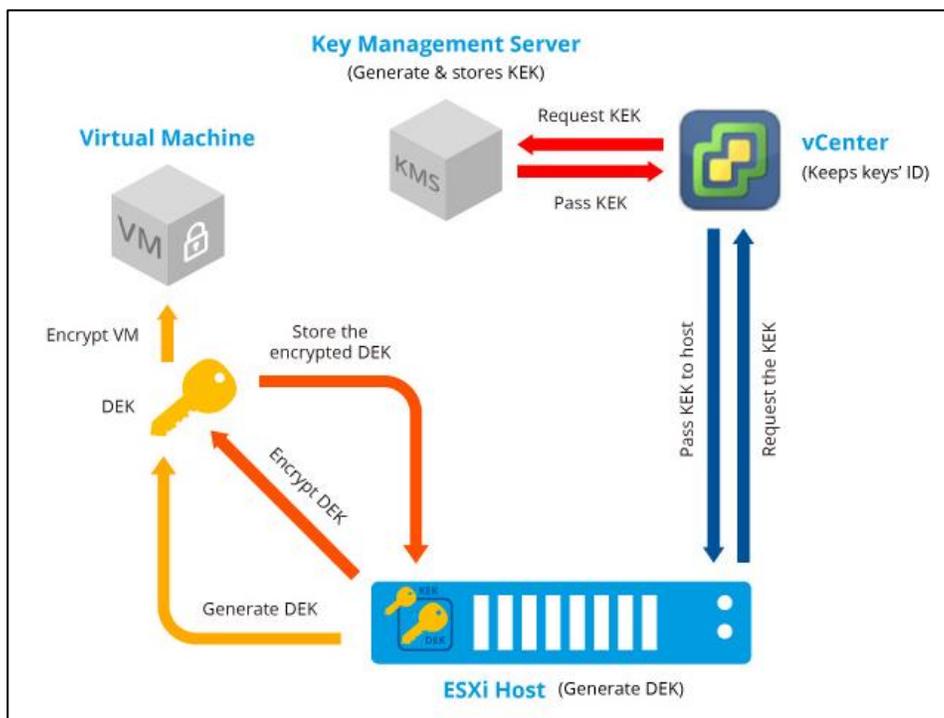


Figure 10. Detailed encryption process of a VM.

Traditional encryptions such as HBA, SED, and in-guest encryption usually focus on specified scenarios or devices, which is lack of consistency and effective authentication. VM encryption is executed on ESXi in the form of software code. The two principles, assurance and attestation, are applied on the encryption in QCT HCI solutions. The assurance promises only the validated code is executed during the encryption tasks. In this case, ESXi hypervisor is verified by the Secure Boot to assure that ESXi is legal to run. The attestation is a method utilized to evaluate the assurance. TPM devices and vCenter Server operate together to validate the assurance and to prove that the ESXi and TPM are not compromised. With the two foundations, the VM encryption can secure services in the cloud effectively and consistently.

The five advantages of VM encryption are listed below.

- VM encryption is OS-agnostic at the VM level which can solve cross-platform issue.
- The encryption happening directly on VMs in the hypervisor can prevent VM data from exposing to risks.
- VM encryption is a software function. Administrators do not need to configure hosts' settings.
- The encrypted VMs can be easily migrated among the hosts in a data center.
- The AES-256-bit encryption key is utilized to protect VMs which is unpractical to hack.

With VM encryption enabled in QCT HCI Solutions, customers' VMs are sure to be solidly protected.

Virtual TPM

Virtual TPM (vTPM) is another virtual device implemented in QCT HCI Solutions which can enhance the guest operation systems such as Window 10 or Windows server 2016 at security level. vTPM performs equivalent functions as a physical TPM device but the difference is that vTPM performs cryptographic processing capabilities in software and allows Windows OS to write the measurements such as boot values or credentials into vTPM for specified use cases.

vTPM exists as a VM hardware device and can be added or removed on demand. When a vTPM device is added, an Endorsement Key by default provides the vTPM a unique identity. Unlike traditional hardware TPM storing the values securely in the “non-volatile secure storage” component, vTPM stores data in a Non-Volatile RAM (NVRAM) file instead. NVRAM file is encrypted by the industry standard AES-256-bit encryption which can create a tamper-resistant storage for data. The benefits of adopting vTPM are listed as follows.

1. With the utilization of vTPM, Windows 10 and Windows Server 2016 guest operating systems in the virtual machines can be enabled to perform authentications on advanced security features such as the smart card, BitLocker encryption, and boot measurement.
2. vTPM can solve the complexity in migrating VMs between hosts. If VMs use physical TPM to store data, the TPM data migration between hosts will need additional tools and operations to move data to another TPM securely, which consumes more time and efforts for administrators. However, a vTPM existing in “file” formation can be migrated within VMs to other hosts via a few clicks which provides great portability.
3. vTPM solves the storage space issue of physical TPM because the non-volatile secure storage, the place where a physical TPM stores data is measured in only kilobytes, is not suitable for a large number of VMs to store measurements. However, each VM can store data in its own vTPM to solve the space issue.
4. Physical TPM is a slow hardware serial device. VMs will need to adopt a tool to perform API calls and execute crypto operations when a physical TPM is adopted. However, vTPM is in a software form which can be added to each individual VM and therefore solves performance and automation problems.
5. The certificates in the vTPM can be changed simply by replacing files issued by other certificate authority, which brings management flexibility and convenience for administrators, as shown in Figure 11.

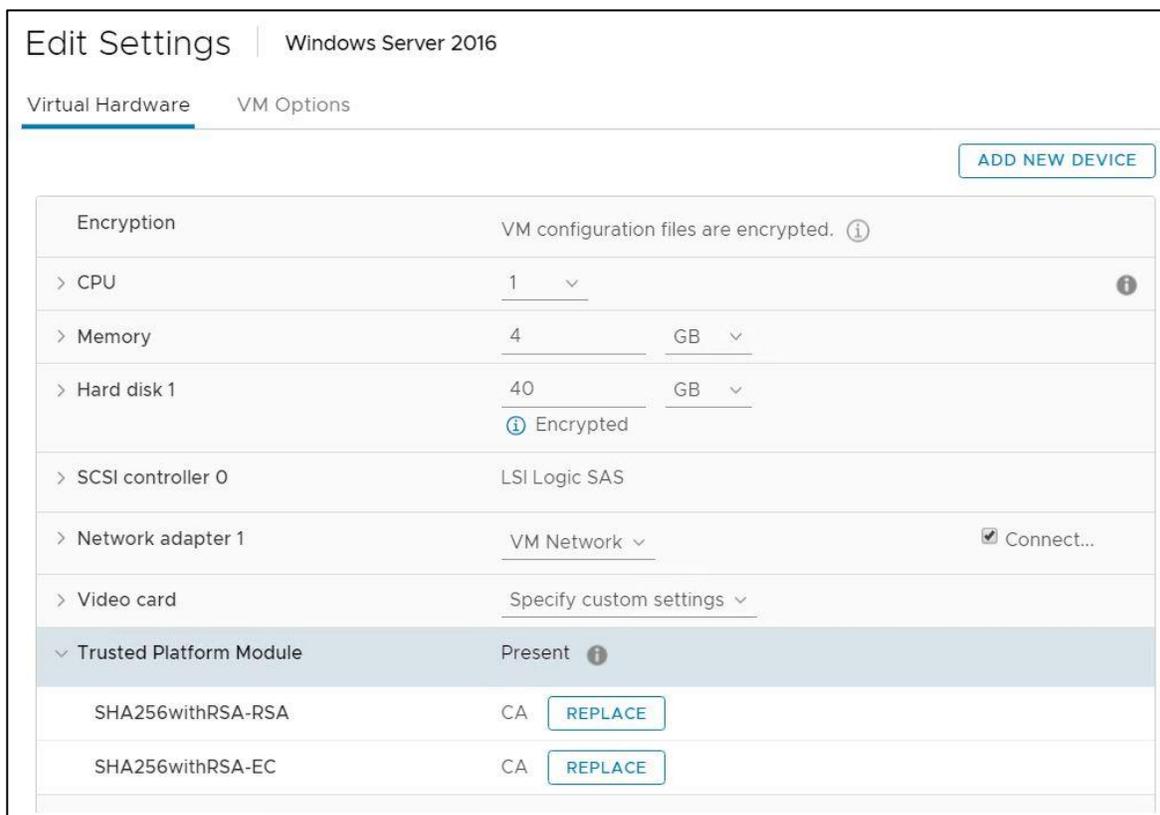


Figure 11. vTPM device and replaceable certificates.

6. The physical TPM can be removed by anyone who can access to the server. However, a vTPM exists in software formation. Administrators can restrict management permissions and prevents the removal of vTPM.
7. The use of vTPM requires the enablement of VM encryption. Therefore, the data in NVRAM file of the vTPM can be secured.

In conclusion, with the adoption of QCT HCI solutions, vTPM enables the advanced security features for Windows 10 and Windows Server 2016 VMs which allows administrators to manage vTPM-related operations easily. This is what traditional datacenter or other solutions fail to achieve.

Virtual Machine Secure Boot

As mentioned in section 5.1.1 above, the Secure Boot implemented in QCT HCI Solutions is able to provide a validation for boot process in a virtual machine which is equivalent to the Secure Boot function on physical machines. In the OS that supports Secure Boot, the boot software, including kernel, bootloader, and drivers are signed. Some default certificates are included in the virtual machines such as a Microsoft (MS) certificate for booting Windows OS, a MS certificate for third-party software, a VMware certificate for nested ESXi, and a MS key exchange certificate for authenticating demands to alter the Secure Boot configuration.

The enablement of the Secure Boot is simplified in the boot options in each individual virtual machine's setting page and the EFI firmware should be chosen for operating Secure Boot, as shown in Figure 12.

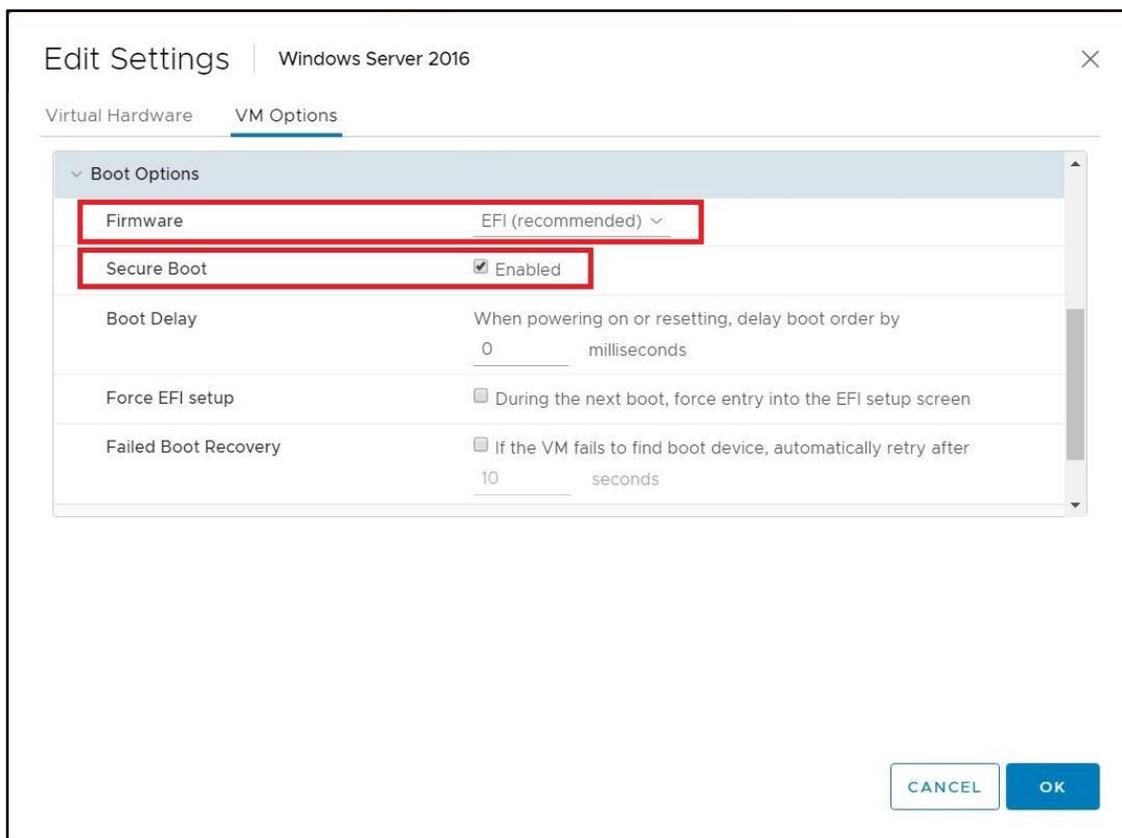


Figure 12. Enablement of Secure Boot for a VM.

The Secure Boot adopted in virtual machines allows the guest OS' boot process to be validated by UEFI Secure Boot for each VM; thus, the possible tampering from malicious software can be prevented and this is what other platforms' VM could never achieve. With the integration of UEFI Secure Boot and VM, the OS that typically supports Secure Boot such as Windows, Windows Server, Ubuntu, Centos, RHEL, VMware Photon, and ESXi can now utilize Secure Boot in a per-VM basis and therefore form a fundamental protection for the services in the QCT HCI Solutions.

5.2.2. Encryption on shared storage

Both individual VM files and data in hyper-converged storage need to be protected in QCT HCI Solutions. To achieve the anti-steal demand, QCT HCI Solutions adopt VMware vSAN encryption. When the encryption is enabled, all the data in the vSAN datastore will be protected via AES-256-bit algorithm.

Enabling the vSAN encryption involves the use of KMS. While performing the vSAN encryption, the vCenter Server asks AES-256-bit KEK from KMS and passes KEKs to ESXi hosts. The ESXi hosts generates DEK for each disk to encrypt data. The KEKs are used to encrypt the DEKs and then ESXi hosts keep the encrypted DEKs on disks, as shown in Figure 13.

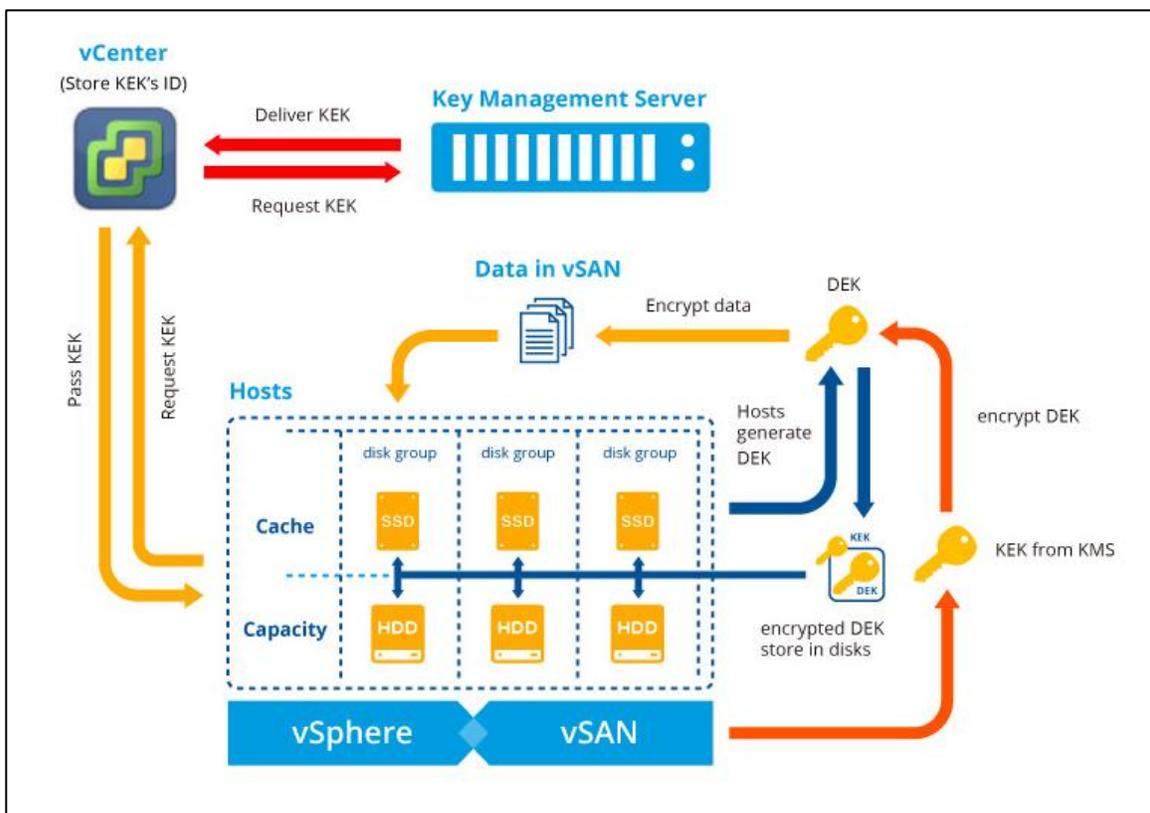


Figure 13. Detailed encryption process of data in vSAN storage.

1. vSAN encryption is a powerful encryption for the data in each disk. The encryption which utilizes the AES-256-bit algorithm is practically difficult to be hacked.
2. vSAN encryption can be executed on any drives certified by VMware without considering hardware brands. No other devices are required due to the pure software characteristic of vSAN encryption.
3. Because vSAN encryption are embedded into ESXi, there is no compatibility issue between hosts and vSAN encryption. Therefore, cloud features of vSphere such as vMotion, High Availability (HA), and Distributed Resource Scheduler (DRS) will operate normally when vSAN encryption is enabled.
4. A KMS is required for the encryption and vSAN encryption supports third-party KMS which utilizes KMIP-1.1 to form a complete Cryptosystem.
5. Traditionally, the security solution such as Host Bus Adapter requires configuration on each host. vSAN encryption is relatively easy to be enabled via merely a few clicks. This greatly reduce time and efforts for administrators, as shown in Figure 14.

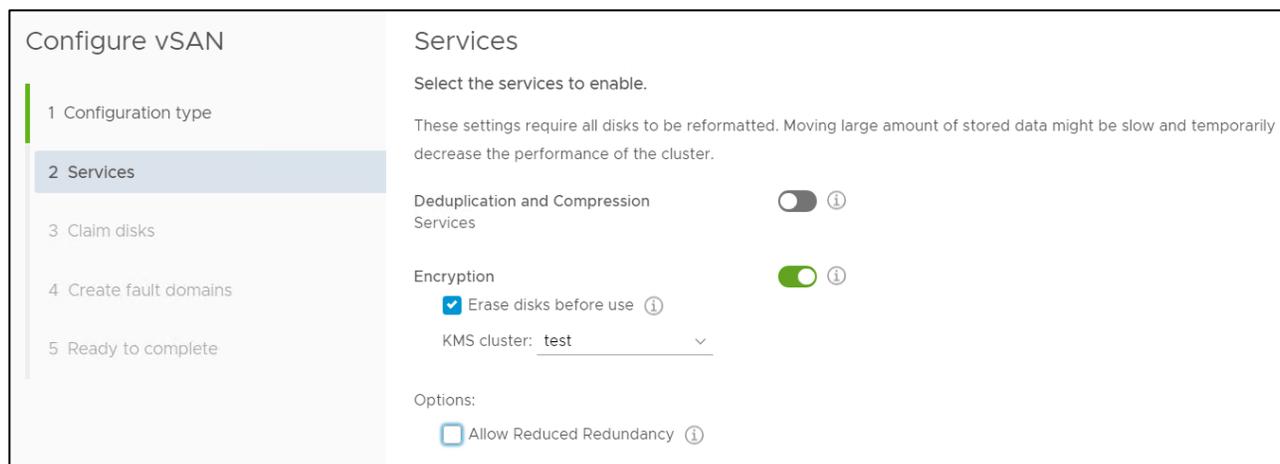


Figure 14. Enablement of vSAN encryption.

In conclusion, with the help of vSAN encryption, QCT HCI Solutions establish a native and secure shared storage in cloud with overall data-at-rest protection. The data leakage on storage will not be a concern.

5.2.3. Migration with a secured way

The migration of VMs between hosts could be a vulnerability of data protection. If malicious snooping behavior occurs in the network part, sensitive data of VM could be altered, and the existing cloud protection will be in vain.

To secure the migration of VMs, the hypervisor adopted in QCT HCI Solutions encrypts all the vMotion traffic. Both the vMotion metadata and the TCP payload of the packet are protected by the Advanced Encryption Standard (AES) algorithm. The method of securing vMotion traffic does not utilize traditional Internet Protocol Security (IPsec) or Secure Sockets Layer (SSL) protocols since SSL implements the encryption and decryption across hosts' kernel and user spaces which cause intensive and unnecessary impact on performance. Also, the hosts support IPsec only for IPv6 traffic and the usability of vMotion is limited. However, vMotion Encryption uses a custom-encrypted protocol on top of the TCP layer formed by vmkernel crypto library and AES-GCM algorithm in ESXi.

Once the encrypted vMotion function is enabled, vCenter generates a one-time use 256-bit key and a 64-bit nonce, and delivers both data to source and destination hosts for encryption and decryption. The hosts use the key to connect to each other over the transmission network. Once the vMotion transmission terminates, the one-time key will be abandoned, as shown in Figure 15.

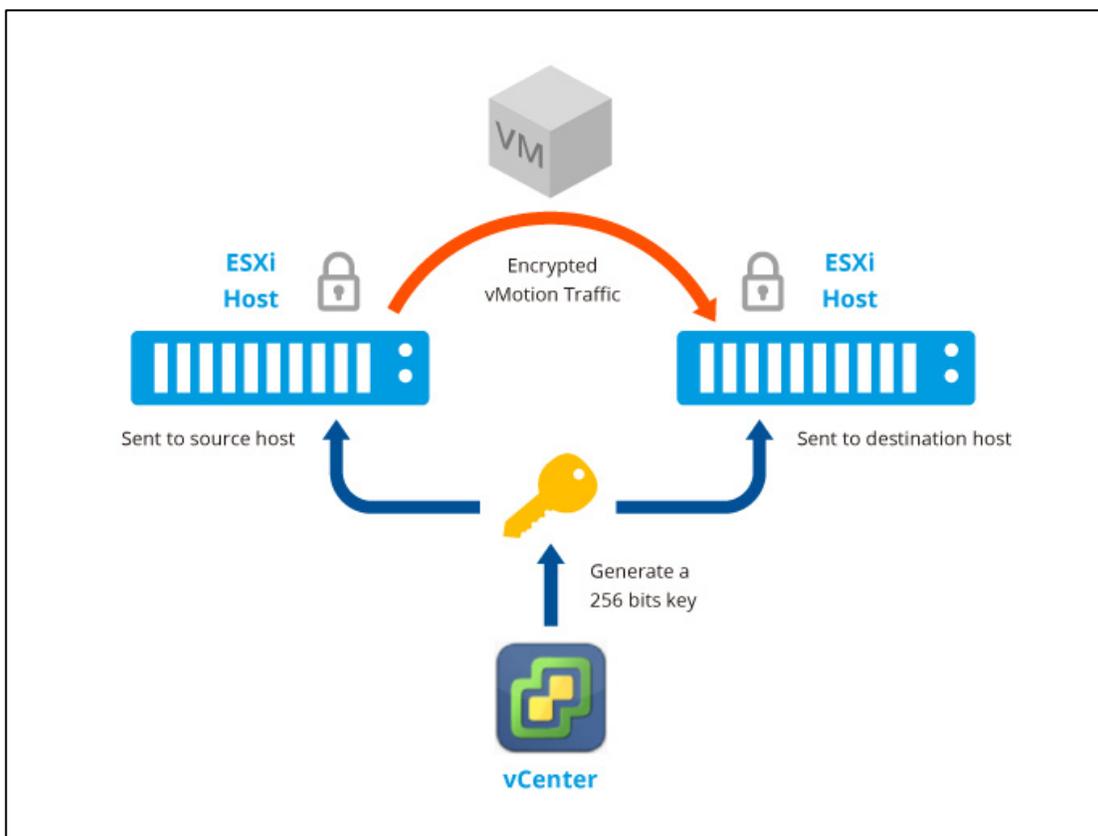


Figure 15. Detailed process of encrypted vMotion.

To enable the encryption on VM's migration, configurations can be set on virtual machines with options such as "required" or "opportunistic." In the "required" mode, both source and destination hosts must use encrypted vMotion. If one of the hosts does not support the encrypted vMotion, the transmission will not be allowed. In the "opportunistic" mode, when either source host or destination host supports the encrypted vMotion, both hosts will use the regular vMotion without encryption. The encrypted vMotion can be configured and enabled on a per-VM basis, as shown in Figure 16.

By adopting vMotion encryption function, the migration of the virtual machines can be protected to deter data leakage during data transmission and therefore establish a secured migration environment in a data center.

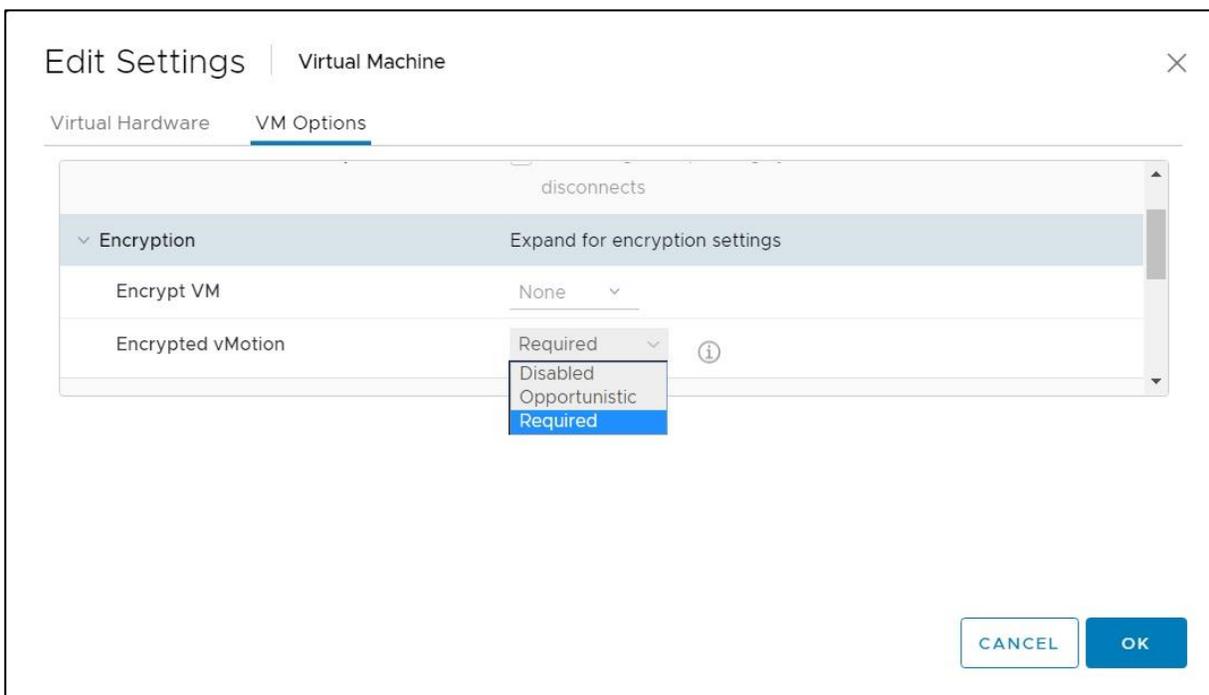


Figure 16. Enablement options of encrypted vMotion.

In conclusion, the overall data protection methods in this section cover VM encryption, vSAN encryption, virtual TPM, VM Secure Boot, and encrypted vMotion. All these functions are virtualized, kernel-embedded, and solid to provide native safety and prevent cloud data from data-theft or malicious alters. On top of that, all the security functions are centralized and can be managed in a single UI. These benefits make the security control of QCT HCI Solutions trustworthy and convincing.

5.3. Traffic Protection

Securing network traffic is always a main topic of data center security. Traditionally, enterprises implement hardware firewalls or virtual firewalls to perform the traffic monitoring which are difficult to bring a complete security for the cloud. To conquer the defects of these traditional solutions and fulfill zero trust protection on each individual virtual service, the software-defined network functions from VMware NSX are utilized in QCT HCI Solutions to provide security on data center traffic.

5.3.1. North-South firewall

The cloud data center always needs overall security to manage any inbound and outbound traffic; however, classic solutions involve multiple hardware devices such as firewalls and routers. These sprawling devices bring challenges not only on management but integration.

The edge firewall is implemented in QCT HCI Solutions to protect the cloud. The edge firewall acts as a gateway to communicate between a data center and an outside network. It exists in the form factor of virtual appliance. All the inbound and outbound traffic in the data center must pass it, as shown in Figure 17. The edge firewall provides embedded network features such as stateful firewall, Network Address Translation (NAT), and Virtual Private Network (VPN) functionality.

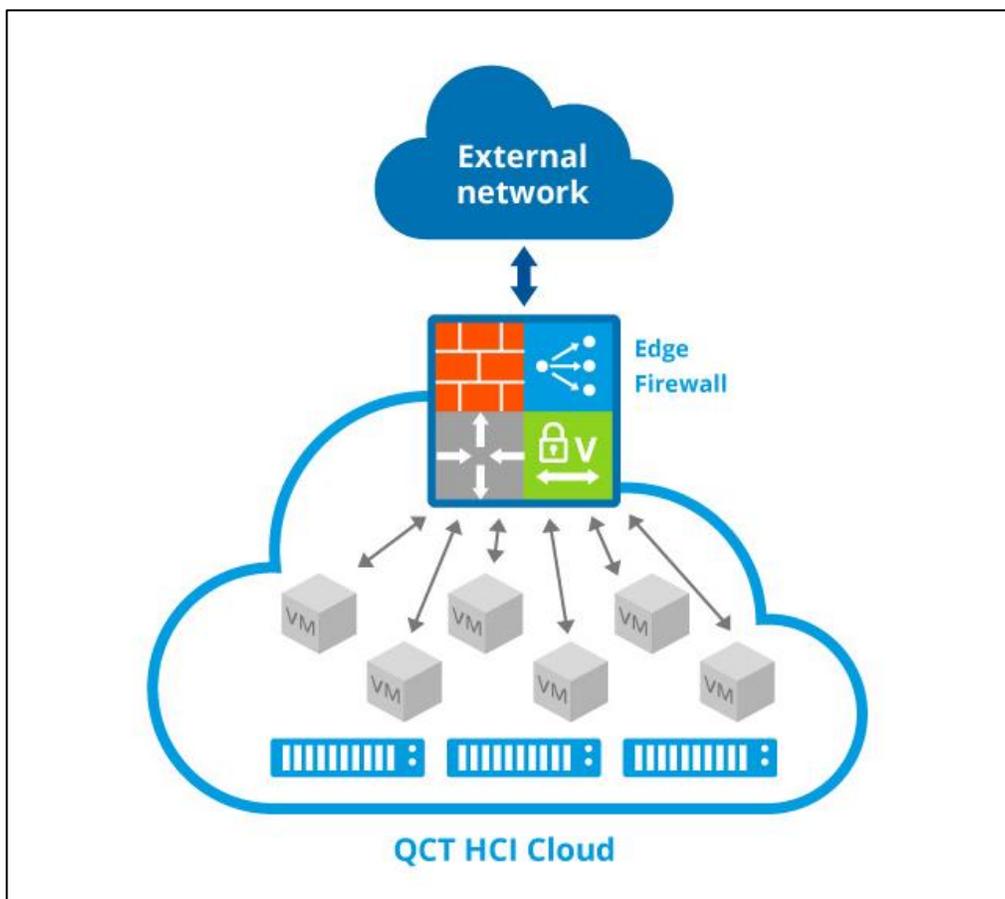


Figure 17. Edge firewall designed to filter ingress and egress traffic.

Stateful Firewall

Built within the virtual appliance, the edge firewall is a stateful firewall which performs stateful network packet inspection and monitors the state of connections. The operation process applies firewall rules which can either allow or deny the north and south network traffic to be connected in a data center. The rules can be established based on source and destination objects like clusters, IP sets, vNIC groups, and so on. The rules are set on services such as HTTP, DHCP, and Active Directory. The specific protocols and network ports are chosen when the rules are set up to provide precise filtering.

When the access requests from the external network reach the edge firewall, the firewall compares the traffic with the rules to determine if the connection is legal, and thus allows or disallows the requests. With the utilization of firewall, customers can centralize the management and design north-south traffic filtering to secure data center, as shown in Figures 18 and 19.

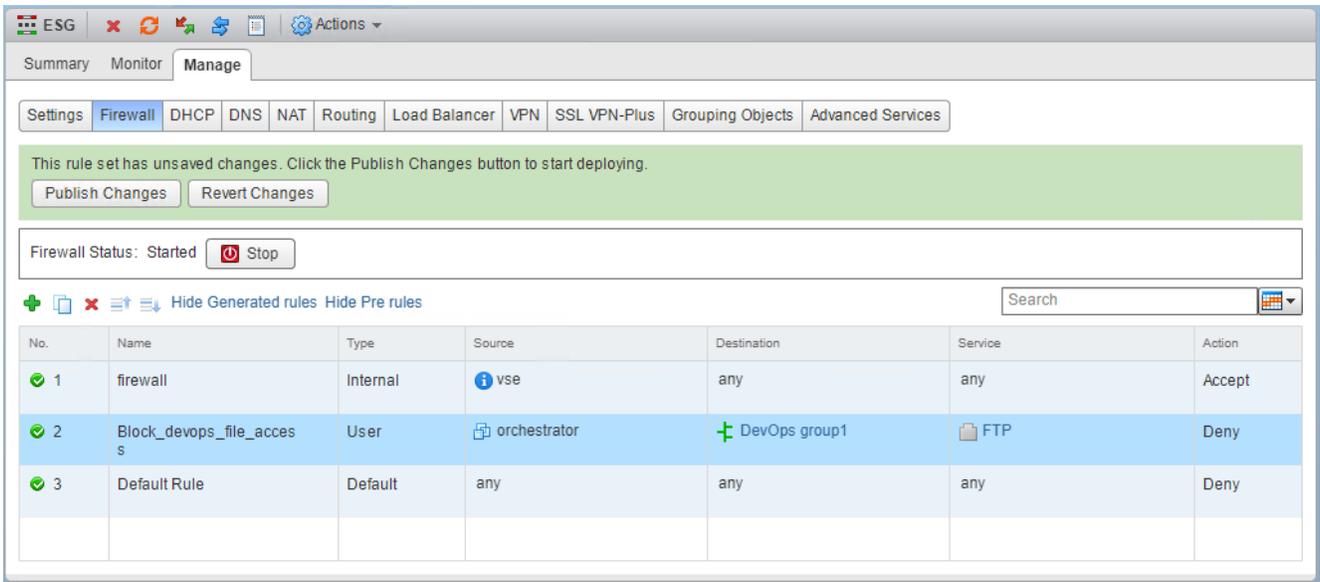


Figure 18. Firewall rules established in edge appliances management tablet.

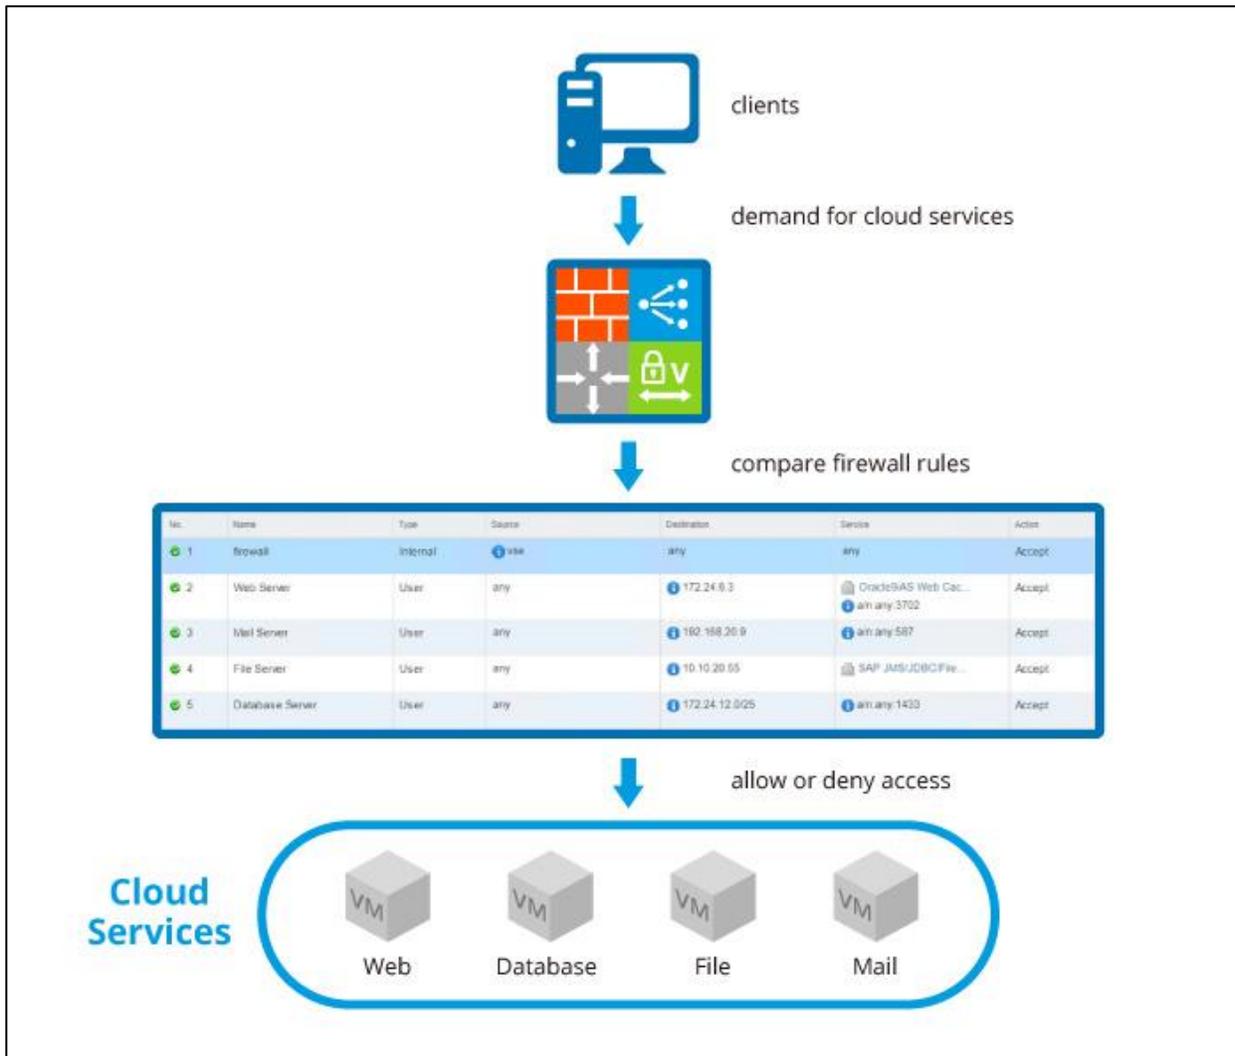


Figure 19. Edge firewall filtering on the connections with specified rules.

Network Address Translation (NAT)

The edge firewall integrates the Network Address Translation (NAT) technique, which remaps IP address space between private IP and public IP by altering network address information in the IP header of packets. The NAT rules can be applied to either source or destination IP addresses. Protocols can also be set specifically with rules to protect intranet information from possible attack from external network, as shown in Figure 20.

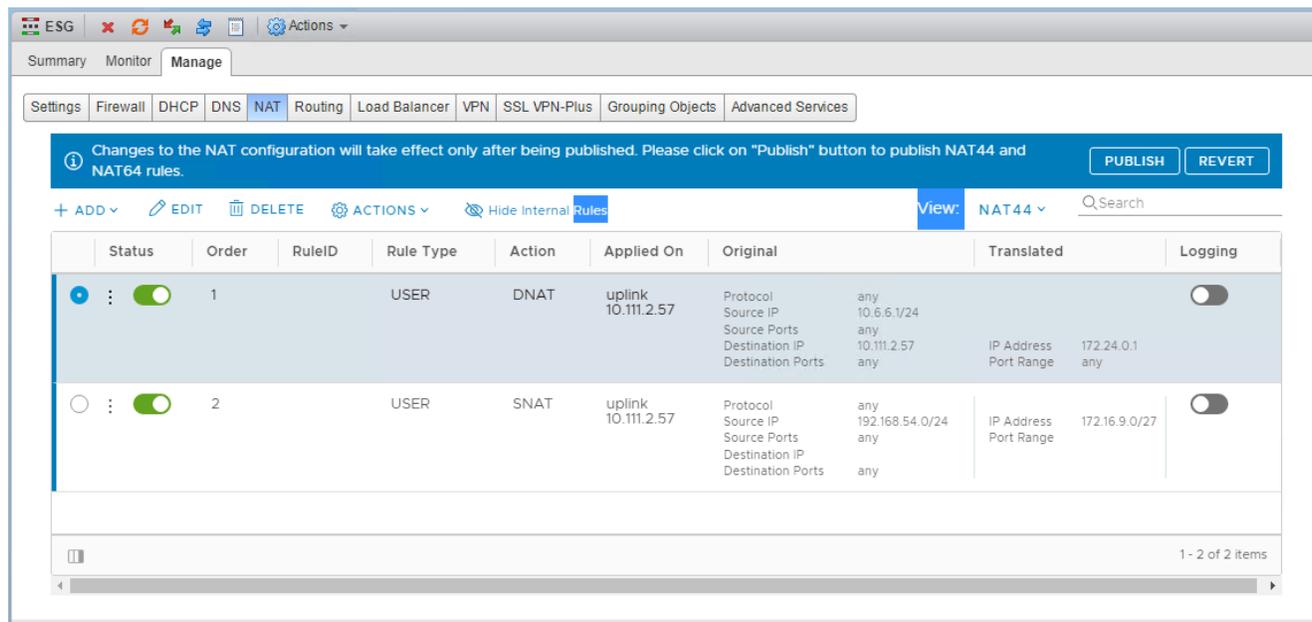


Figure 20. Embedded NAT function in the edge appliance.

Virtual Private Network (VPN)

Virtual Private Network (VPN) is the extension of a private network across a public network. It is integrated into the edge as a built-in function. The encryption method of the VPN traffic matches industrial standards such as SSL, AES, DES, and SHA. Three different types of VPN are adaptable to diverse use cases: (1) IPsec VPN is for site-to-site connection, (2) L2 VPN is for stretching multiple layer-2 networks regardless of geo-locations, and (3) SSL VPN-plus is for the access to private network for remote users. All the three types of VPN are available in an edge appliance, providing solid traffic protection and built-in VPN function for private networks. The VPN service can be configured on the dashboard of edge appliance, as shown in Figure 21.

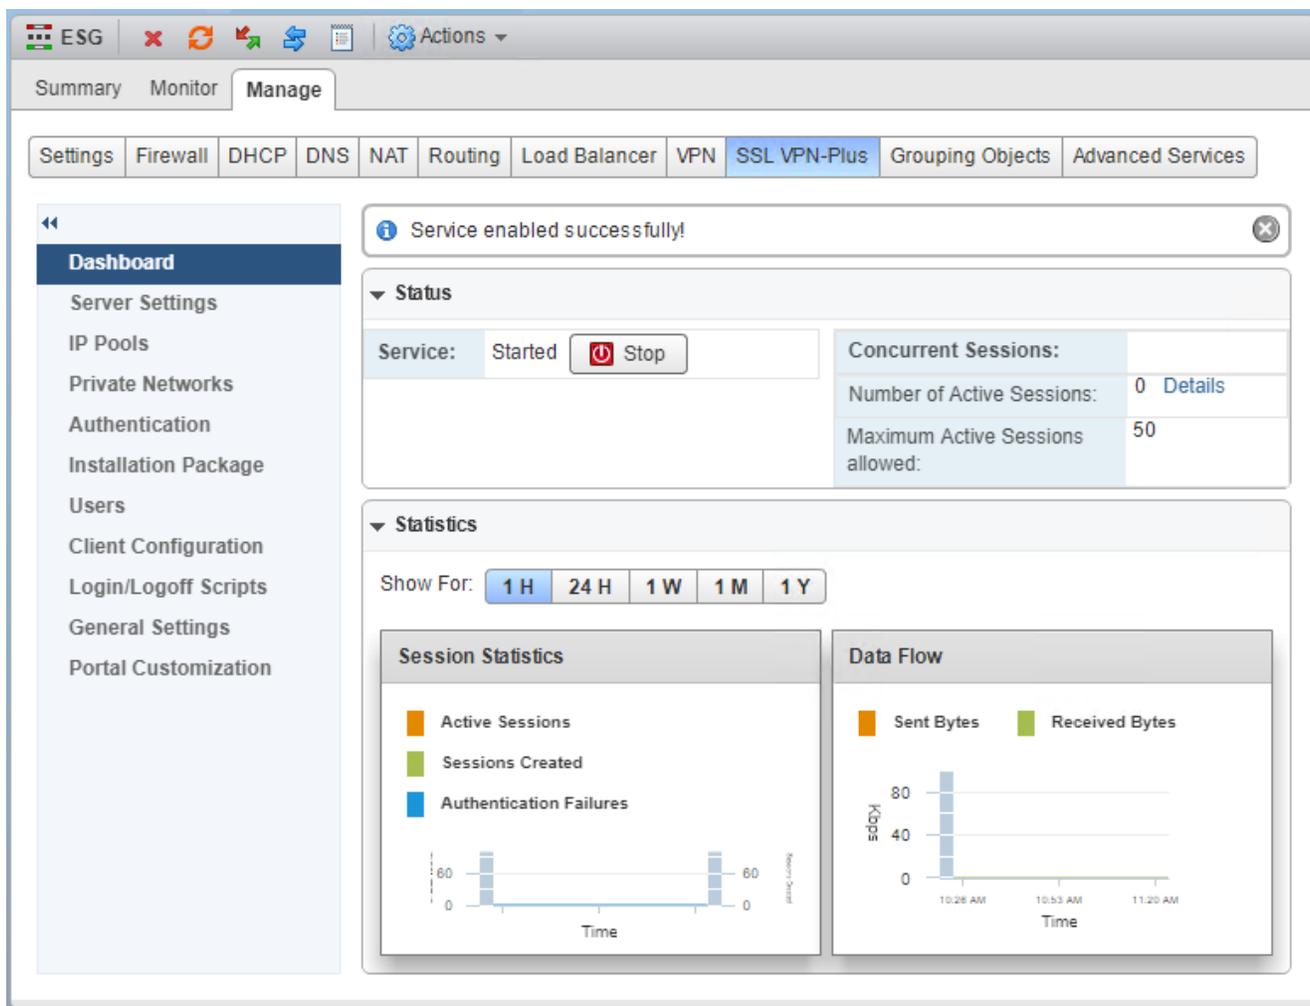


Figure 21. VPN functions in the management tablet of edge appliance.

In conclusion, the north-south firewall functions in NSX allows the traditional N-S traffic to offload to edge virtual appliance, providing the integration of vital features and optimizing the security of network topology of gateway.

5.3.2. East-West Firewall

Typical firewall hardware brings challenges to the services in a cloud such as fixed and sprawling rules, traffic trombone, and lack of judgement of encapsulated traffic. These disadvantages prevent cloud services' traffic control from being optimized because these typical designs are out-of-date and are not suitable for modern cloud environment.

To thoroughly and flexibly protect all the services in the data center, the distributed firewall is implemented in QCT HCI Solutions. The distributed firewall is a kernel-embedded firewall which has full compatibility with the hypervisor and can be applied to virtual machines' virtual network interface card to perform traffic filtering, as shown in Figure 22. The traffic can be checked inside the hypervisor so that the firewall capacity can be extended by simply increasing the number of ESXi hosts. Since the firewall rules are distributed to each host, the filtering performance can reach the line rate and the traffic routes can be optimized as well.

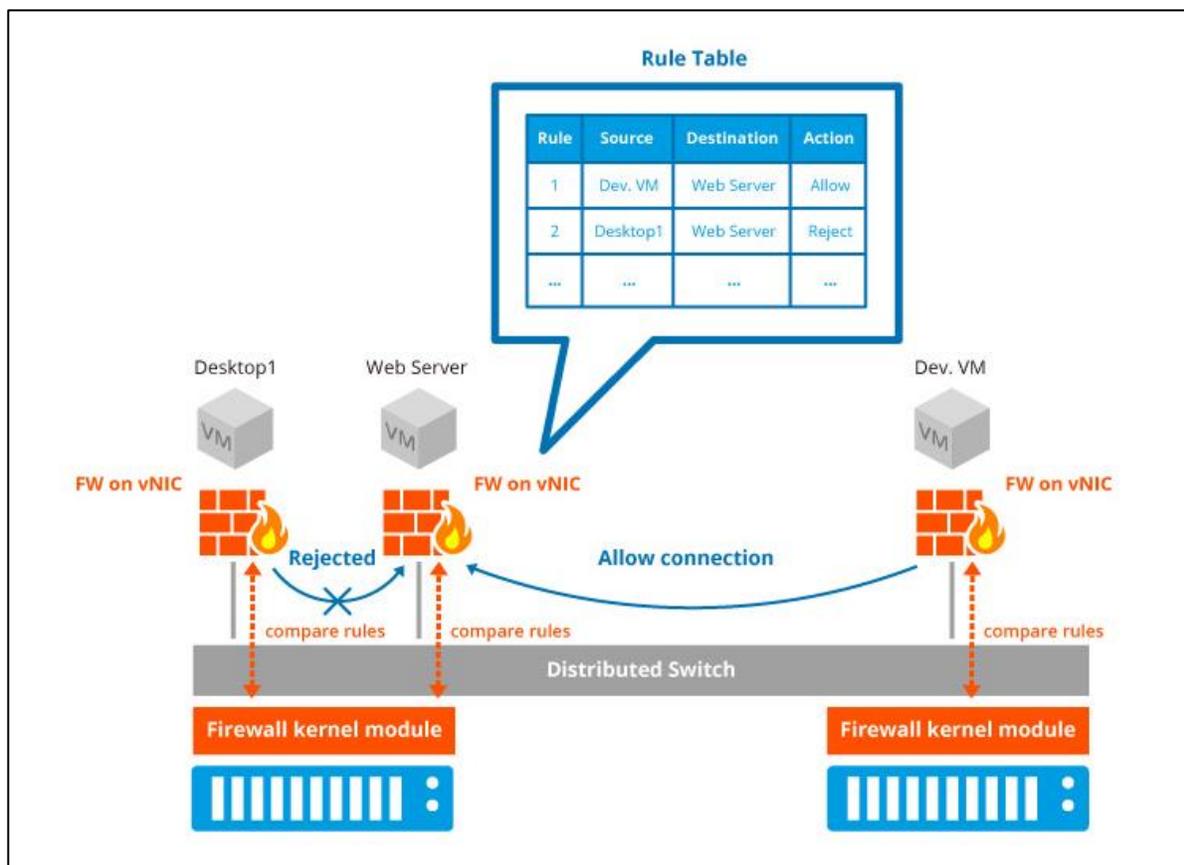


Figure 22. Architecture of distributed firewall rules for VMs.

Stateful firewall

The native distributed firewall is a stateful firewall and the filtering action is based on diverse options like IP, port, and protocols. The firewall rules can be applied to different objects such as IP sets, logical switch, resource pool, and clusters. The rules are stored in the database, the configurations can be set in centralized user interface, and the VM's vNIC is updated with the rules regardless of VMs' location, as shown in Figures 23 to 25. In this way, administrators can save lots of efforts in reconfiguring firewall rules since the rules are distributed among all the hosts to enforce traffic filtering on each vNIC.

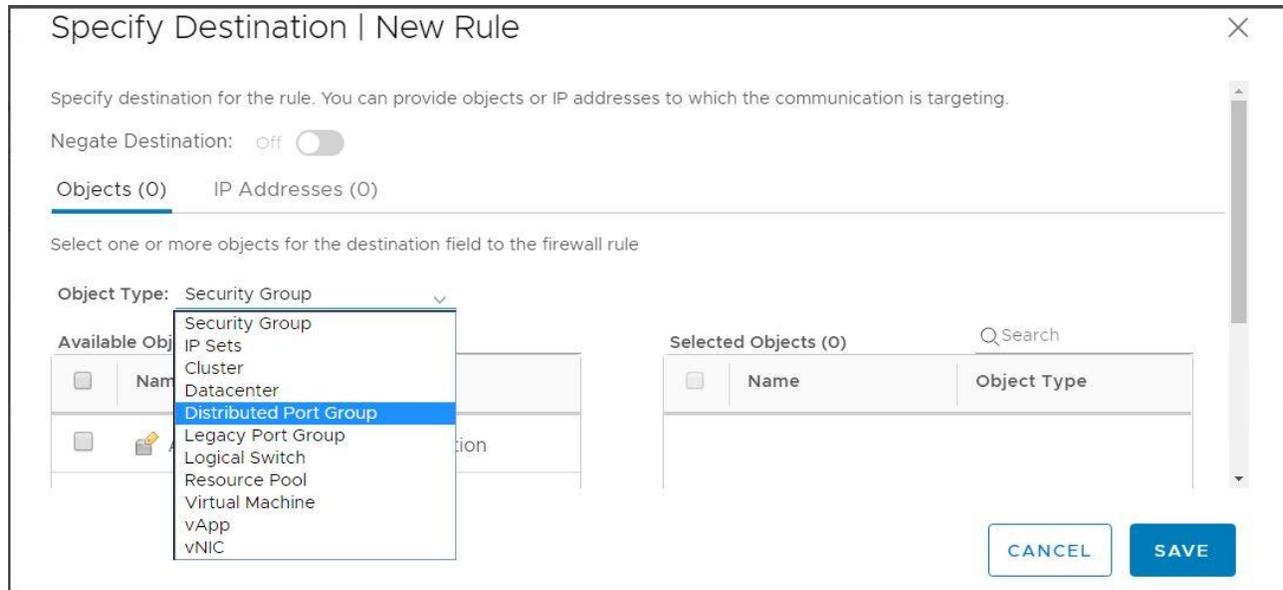


Figure 23. Firewall rules assigned to different objects.

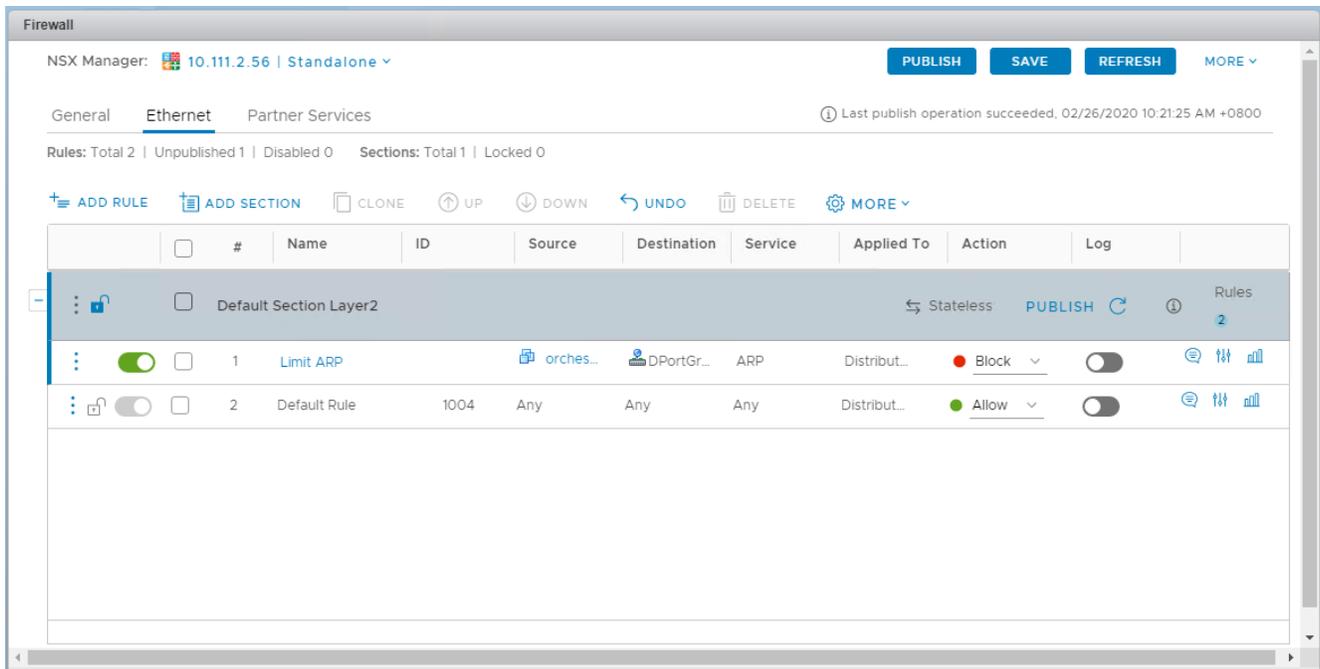


Figure 24. Firewall rules configured on Layer 2 network.

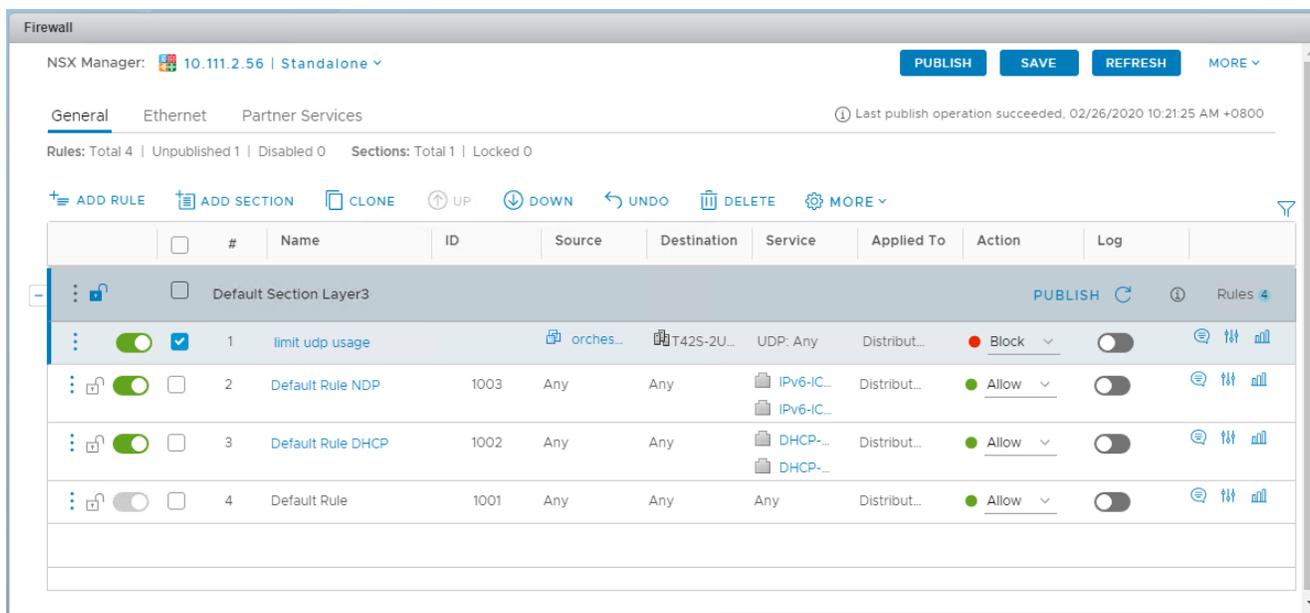


Figure 25. Firewall rules configured on Layer 3 network.

Context-Aware Firewall

The context-aware layer-7 rules can analyze the contents of packets and execute the filtering on application behaviors such as blocking HTTP demands even when the port number is changed. Thus, the accuracy of security control can be greatly improved. The example of context-aware services is shown in Figure 26.

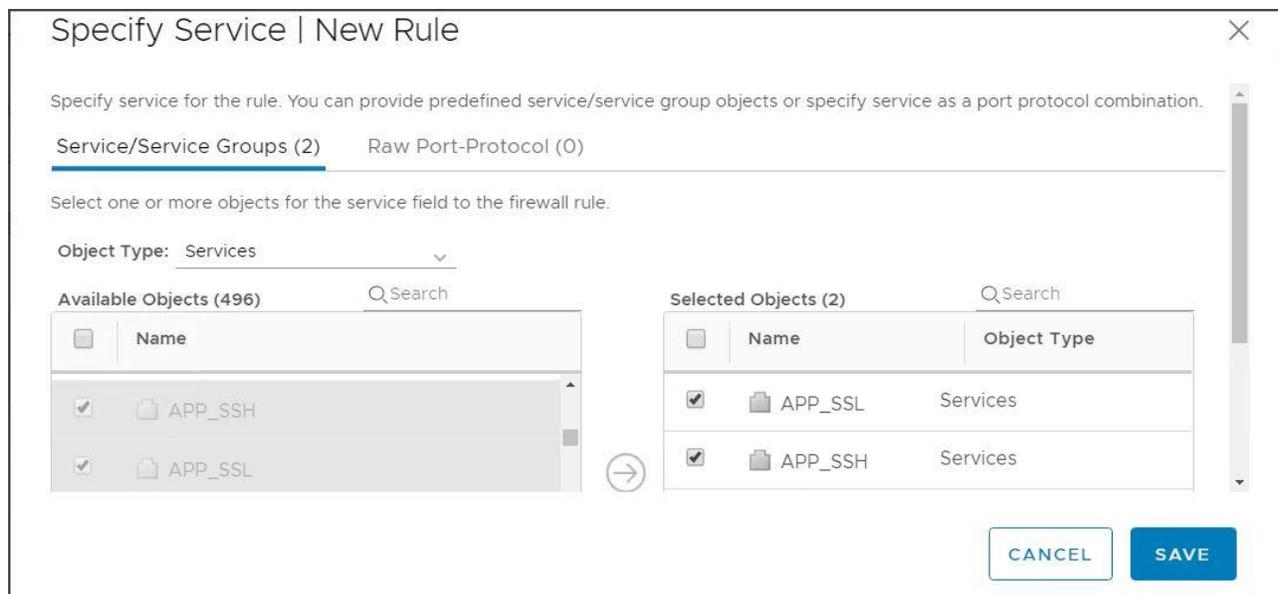


Figure 26. Rules of context-aware firewall configured on applications.

Identity Firewall

This firewall can create user-based distributed firewall rules based on Active Directory (AD). With the integration of AD and the distributed firewall, the distributed firewall rules can be applied to AD users or AD groups, which allow a data center to filter the logon activity based on users' identity. There are two ways to achieve

the detection of logon. One is via Guest Introspection and the other is through the AD Event Log Scraper. The first method utilizes guest agent which is installed in the target VM. When the logon event is triggered by a user, the agent transmits the information to the NSX manager. The second method utilizes existing or newly-created AD service. NSX manager is connected to the AD event log for the logon information. Once the logon event happens, the manager will acquire the events from the log. To precisely filter the identity, administrators can select the users in the AD and create a Security Groups to enforce security policies and firewall rules on the specific security groups.

The identity firewall can be solidly integrated with remote desktop sessions (RDS) to enhance the login security on the session of vital MS Windows remote desktops. The AD users are selected into security groups and the security groups are configured to the distributed firewall to precisely filter the identity on any specified desktop, as shown in Figures 27-29.

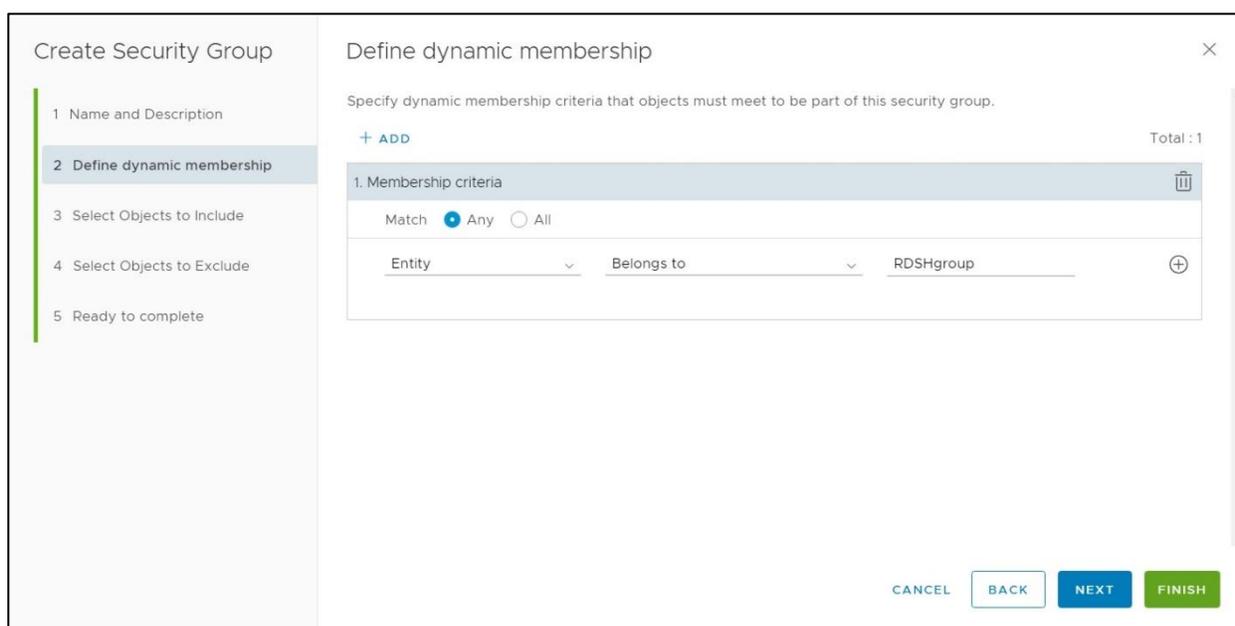


Figure 27. Selection of members for security group.

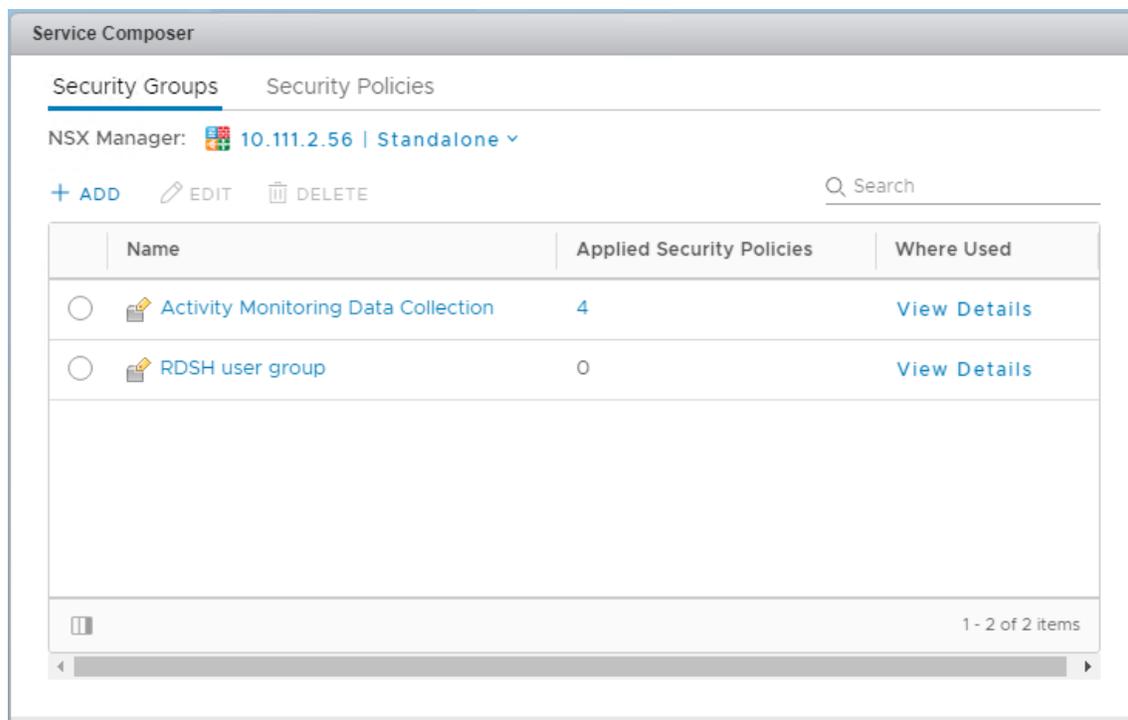


Figure 28. Created security group with specified AD users.

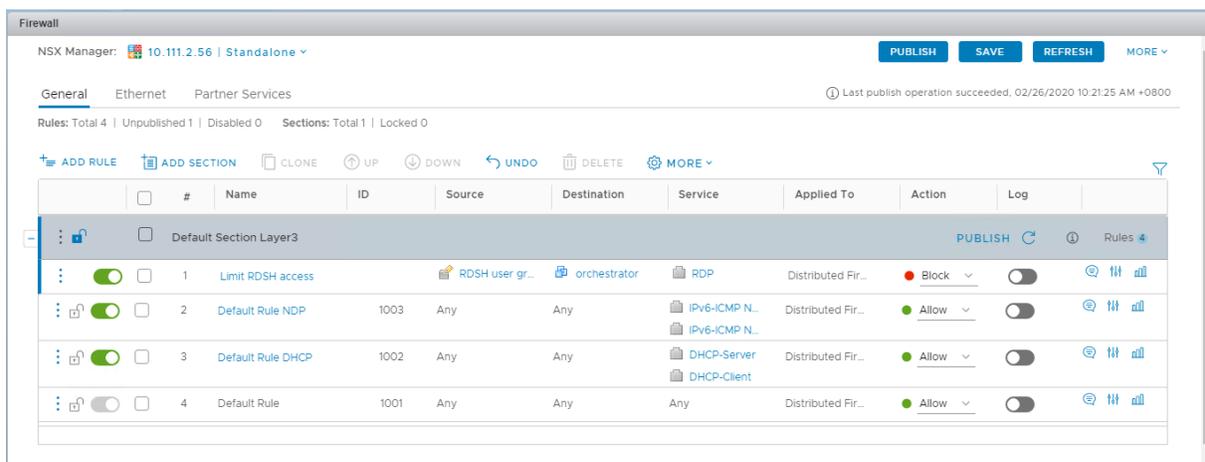


Figure 29. Identity firewall rule based on AD user groups.

With the combination of the stateful, context-aware, and identity firewalls, accurate and inescapable traffic filtering can be executed on specific network layers, virtual machines, and applications according to enterprises' policies.

In conclusion, the network traffic can be offloaded from traditional firewall to logical formation, which brings optimized and precise filtering for network connections. With the utilization of N-S and E-W firewalls, integrated perimeter services, micro-segmentation, and zero-trust protection in a cloud can be implemented. Additionally, with the centralized management and integration, customers can highly minimize the Capital Expenditure (CapEx). These advantages can conquer what traditional firewall could barely achieve and form an ideal software-defined data center in QCT HCI Solutions.

6. Conclusion

Nowadays, more and more devices are connected to the internet. The threats of cyberattacks are relatively increased. Enterprises have no choice but to take actions and transform their data centers to proactively manage these potential risks. To strengthen the protection of data centers in different scenarios, QCT HCI Solutions integrate the industry-leading software, including VMware vSphere®, VMware vSAN™, and VMware NSX® to the high-performance server, and provide customers strictly-validated and native-protected data center.

Conventionally, it is common for administrators to spend weeks or even months researching and struggling against the compatibility issues to deploy a new system. By adopting QCT HCI Solutions validated by QCT and VMware®, customers can highly minimize their time and expense in designing architectures, implementing security functions, and configuring rules so as to reduce the overall Total Cost of Ownership (TCO). On the other hand, customers can rest assured that the solution is reliable and can completely focus on strategic and productive tasks. With the knowledge of QCT, customers can leverage the collective achievement and follow a simplified path to an assured data center.

Just click <http://go.qct.io/contact/contact-qct-solutions/> to contact QCT. Your data center will be well protected right away!

7. Reference

[1] Quanta Cloud Technology (QCT)

<http://www.qct.io/>

[2] The Global Risks Report 2020

http://www3.weforum.org/docs/WEF_Global_Risk_Report_2020.pdf

[3] VMware vSphere®

<https://www.vmware.com/products/vsphere.html>

[4] VMware vSAN™

<https://www.vmware.com/products/vsan.html>

[5] VMware NSX®

<https://www.vmware.com/products/nsx.html>

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 VMware

VMware software powers the world's most complex digital infrastructure. The company's compute, cloud, mobility, networking and security offerings provide a dynamic and efficient digital foundation to over 500,000 customers globally, aided by an ecosystem of 75,000 partners. Headquartered in Palo Alto, California, this year VMware celebrates twenty years of breakthrough innovation benefiting business and society.



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.