

# VMware Setup Guide

## Application Note

November 2024

# ANNOUNCEMENT

## Copyright

© Copyright 2024 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

## Trademarks

- QSAN, the QSAN logo, QSAN.com, XCubeFAS, XCubeSAN, XCubeNXT, XCubeNAS, XCubeDAS, XEVO, SANOS, and QSM are trademarks or registered trademarks of QSAN Technology, Inc.
- Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux is a trademark of Linus Torvalds in the United States and/or other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.
- Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries.
- Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

# TABLE OF CONTENTS

- Announcement..... i**
- Notices.....vii**
- Preface..... viii**
  - Technical Support ..... viii
  - Information, Tip, and Caution ..... viii
- 1. Introduction to VMware ..... 1**
  - 1.1. Recommended Storage for Virtualization ..... 2
- 2. Connect with VMware ESXi 8 ..... 5**
  - 2.1. Introduction to VMware ESXi ..... 5
  - 2.2. Configure Steps..... 5
  - 2.3. Conclusion ..... 12
  - 2.4. Appendix..... 12
- 3. Integration with VMware VAAI .....13**
  - 3.1. Introduction to VMware VAAI ..... 13
  - 3.2. Test Results..... 20
  - 3.3. Conclusion ..... 24
  - 3.4. Appendix..... 25
- 4. Integration with VMware VASA .....26**
  - 4.1. Introduction to VMware VASA ..... 26
  - 4.2. Implementation ..... 28
  - 4.3. Conclusion ..... 37
  - 4.4. Appendix..... 37
- 5. Configure VMware Cluster VMDK .....38**
  - 5.1. Introduction to VMware Cluster VMDK ..... 38
  - 5.2. Installation Tips..... 43
  - 5.3. Test Results..... 47

- 5.4. Conclusion ..... 49
- 5.5. Appendix..... 49
- 6. DR Solution for VMware ..... 50**
  - 6.1. Configure DR Solution..... 50
  - 6.2. Conclusion ..... 60
  - 6.3. Appendix..... 60

# FIGURES

Figure 1-1	Use XCalc. Tool to Obtain Recommended Storages.....	3
Figure 1-2	Select Virtualization Option.....	3
Figure 1-3	Click Proposal Details Button to View More .....	4
Figure 1-4	Click Export Button to Export Result .....	4
Figure 2-1	Demonstration Topology.....	6
Figure 2-2	Create a Block Volume and Add into HostGroup.....	7
Figure 2-3	Add ESXi 8 Host.....	7
Figure 2-4	Discover iSCSI Connections.....	8
Figure 2-5	Add iSCSI Server.....	8
Figure 2-6	Rescan Storage .....	9
Figure 2-7	Add New Datastore 1.....	9
Figure 2-8	Add New Datastore 2.....	10
Figure 2-9	Create VM-1.....	11
Figure 2-10	Create VM-2 .....	11
Figure 3-1	VAAI Full Copy.....	16
Figure 3-2	VAAI Hardware Assisted Locking .....	18
Figure 3-3	Hardware Acceleration Support Status.....	19
Figure 3-4	VAAI Test Diagram .....	20
Figure 3-5	VAAI Storage Configuration.....	21
Figure 3-6	Time Saving of VAAI Full Copy .....	23
Figure 4-1	QSAN VASA Provider System Architecture.....	27
Figure 4-2	Install VASA Provider Step 1 .....	29
Figure 4-3	Install VASA Provider Step 2 .....	29
Figure 4-4	Setting the Port Number .....	30
Figure 4-5	Login to QSAN VASA Provider.....	31
Figure 4-6	Configuring QSAN VASA Provider.....	32

Figure 4-7	Add a VASA Provider.....	33
Figure 4-8	Add a New Storage Provider.....	33
Figure 4-9	Add a Storage Adapter.....	34
Figure 4-10	Add an iSCSI Target.....	35
Figure 4-11	Add a Datastore Step 1.....	36
Figure 4-12	Add a Datastore Step 2.....	36
Figure 5-1	Windows Server Failover Clustering Architecture.....	39
Figure 5-2	Cluster Shared Volume .....	40
Figure 5-3	Raw Device Mapping .....	41
Figure 5-4	Enabling Clustered VMDK.....	42
Figure 5-5	RDM vs. Clustered VMDK .....	42
Figure 5-6	Virtual Machines Clustered Across Hosts.....	45
Figure 5-7	Enable Clustered VMDK.....	46
Figure 6-1	ESXi Server Architecture .....	50
Figure 6-2	Configure a Replication Task .....	51
Figure 6-3	Create a Scheduled Snapshot in the VM .....	52
Figure 6-8	List the Snapshots in the VM .....	56
Figure 6-9	Remote Replication Task .....	57
Figure 6-10	Expose the Snapshot .....	58
Figure 6-11	Snapshot is rolled back.....	58

# TABLES

---

Table 1-1 Storage Options to Enhance VM Performance..... 2

Table 3-1 Hardware Acceleration Status values ..... 19

Table 3-2 Time Taken for Full Copy ..... 22

Table 3-3 Time Taken for Block Zero ..... 24

Table 5-1 RDM vs. Clustered VMDK..... 43

Table 5-2 Performance Results of RDM and Clustered VMDK ..... 48

# NOTICES

---

Information contained in this document has been reviewed for accuracy. But it could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the products. All features, functionality, and product specifications are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.



# PREFACE

---

## Technical Support

Do you have any questions or need help trouble-shooting a problem? Please contact QSAN Support, we will reply to you as soon as possible.

- Via the Web: [https://www.qsan.com/technical\\_support](https://www.qsan.com/technical_support)
- Via Telephone: +886-2-77206355
- (Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support
- (Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summer time: 09:30 - 01:00)
- Via Email: [support@qsan.com](mailto:support@qsan.com)

## Information, Tip, and Caution

This document uses the following symbols to draw attention to important safety and operational information.



### INFORMATION

INFORMATION provides useful knowledge, definition, or terminology for reference.

---



### TIP

TIP provides helpful suggestions for performing tasks more effectively.

---



## CAUTION

CAUTION indicates that failure to take a specified action could result in damage to the system.

---

# 1. INTRODUCTION TO VMWARE

---

VMware, a global leader in cloud infrastructure and digital workspace technology, provides virtualization and cloud computing software and services that enable businesses to optimize their IT environments. In this document, we have consolidated all VMware storage guidance and presented it in separate chapters.

Chapter 2 offers a detailed guide on creating a VM (Virtual Machine) using VMware ESXi 8. It walks through the steps for mounting an iSCSI LUN, configuring it as a datastore, and deploying a virtual machine. By utilizing the iSCSI protocol for block-level storage, it enables efficient resource allocation and management in virtualized environments, making it a preferred choice for high-performance storage solutions.

Chapter 3 introduces the concept of VAAI (vStorage API for Array Integration), which enhances performance by offloading specific storage operations from ESXi hosts to storage arrays. This reduces the workload on the hypervisor and speeds up tasks such as replication, migration and snapshots. Finally, we provide test results to prove it.

In chapter 4, another important feature for storage is VASA (vSphere APIs for Storage Awareness), which allows storage arrays to expose their capabilities to VMware vCenter. VASA gives administrators deeper visibility into the storage infrastructure, helping them make informed decisions and ensuring more efficient management of VMs.

Chapter 5 introduces VMware's Cluster VMDK (Virtual Machine Disk) feature which enables multiple virtual machines to share the same VMDK file across a VMware cluster, providing flexibility for applications that require shared storage access. This is particularly useful in environments running applications like databases that benefit from multiple VMs accessing the same data while maintaining consistency and high performance.

Chapter 6 comes to a VMware's disaster recovery solution. It provides technical guidance for setting up DR (Disaster Recovery) solution in VMware environment and making sure that the replicated data will be consistent with special script implemented in ESXi server, and it leads QSAN storage products being able to achieve real DR with snapshot consistency, it is no longer necessary to install any agent in the environment before achieving this.

In summary, VMware's advanced features such as VAAI, VASA, cluster VMDK, and its comprehensive disaster recovery solutions position it as a critical enabler for businesses seeking to improve performance, optimize storage management, and ensure resilient IT operations.

## 1.1. Recommended Storage for Virtualization

Before starting, first understand which storage is suitable for virtualization. The table below summarizes our findings and provides a clear overview of the maximum number of VMs that each storage type can support, regardless of latency. This comprehensive analysis is designed to assist in selecting the most appropriate storage solution based on specific performance needs and workload requirements, ensuring optimal deployment and scalability of virtual environments.

Table 1-1 Storage Options to Enhance VM Performance

STORAGE TYPE	LATENCY THRESHOLD	ADDITIONAL VMS SUPPORTED UNDER LATENCY	NUMBER OF VMS SUPPORTED
NVMe Storage	< 100 $\mu$ s	50+ VMs	Up to 1,000 VMs (high-end configurations)
SAS SSD Storage	< 500 $\mu$ s	20 ~ 30 VMs	Up to 300 VMs
Hybrid Drive Storage	< 1 ms	10 ~ 20 VMs	Up to 150 VMs
SAS HDD Storage	< 50 ms	3 ~ 4 VMs	Up to 15 VMs

In addition, we provide a tool to select the appropriate storage for virtualization. Here are the steps.

1. Use [XCalc](#) tool on the QSAN website to obtain recommended storages.
2. Enter the **Total Usable Capacity Required** and the desired **RAID Level**.

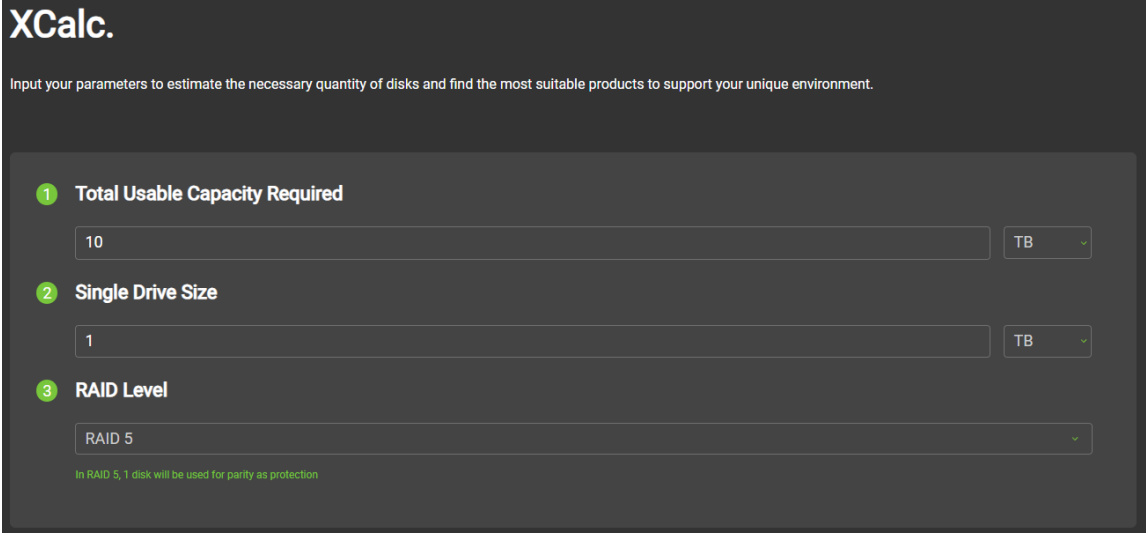


Figure 1-1 Use XCalc. Tool to Obtain Recommended Storages

3. Select the **Virtualization** option.

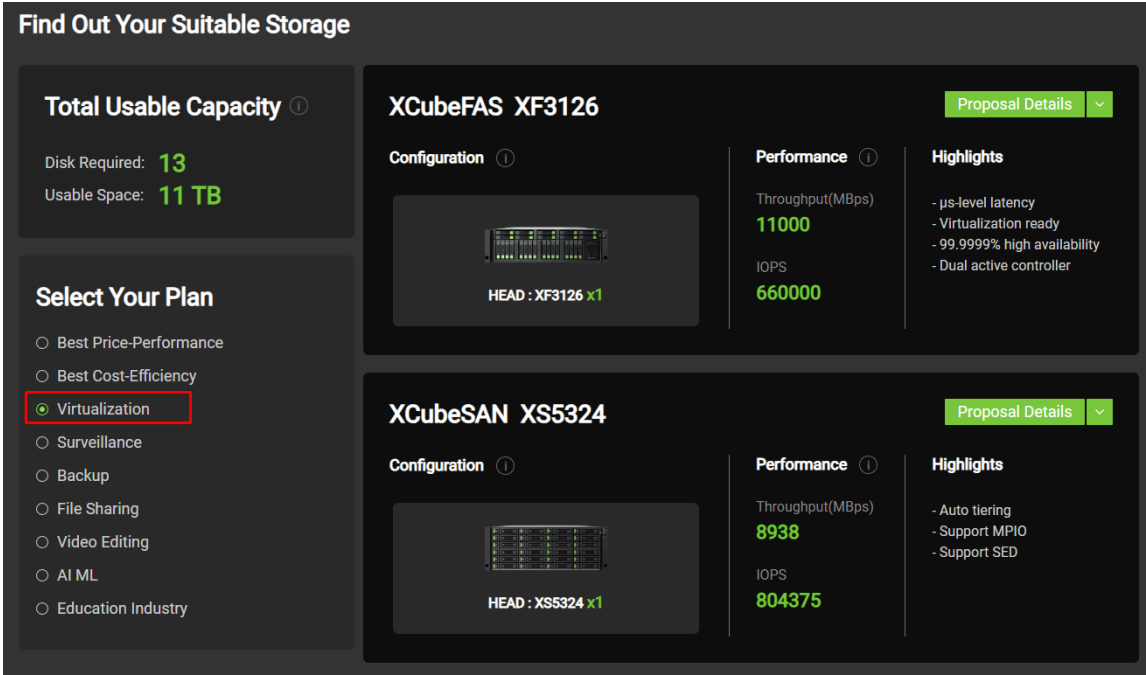


Figure 1-2 Select Virtualization Option

4. Select the model and click the **Proposal Details** button to view more.




Figure 1-3 Click Proposal Details Button to View More

5. If necessary, click the **Export the Result** button to export the report.

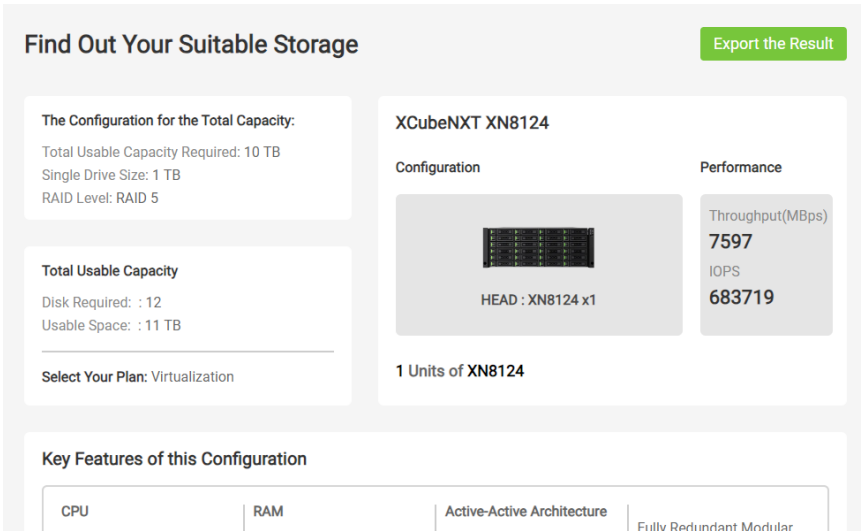


Figure 1-4 Click Export Button to Export Result

## 2. CONNECT WITH VMWARE ESXi 8

---

This document provides comprehensive guidance for setting up virtual machines using VMware ESXi 8.0 and vCenter 8.0. Specifically, it details the process of mounting an iSCSI LUN, configuring it as data storage, and deploying a virtual machine. The iSCSI protocol utilizing block-level storage enables efficient resource allocation and management in virtualized environments, making it ideal for high-performance storage solutions. This guide will serve as a reference for administrators looking to integrate QSAN storage solutions with the latest virtualization technologies from VMware.

### 2.1. Introduction to VMware ESXi

As demand for flexible, scalable storage solutions continues to grow in modern IT infrastructures, the integration of QSAN's iSCSI storage systems with VMware ESXi 8.0 provides a powerful solution for managing virtualized environments. This chapter walks through the steps of connecting an iSCSI LUN from a QSAN storage array to an ESXi 8.0 host, configuring it as a data store, and deploying a virtual machine using vCenter 8.0. This seamless integration enables enterprises to take advantage of the benefits of centralized storage, including easier management, high availability, and enhanced performance for virtualized workloads.

#### **vCenter**

VMware's centralized management utility for managing virtual machines, multiple ESXi hosts, and all related elements from a single, centralized location.

#### **vSphere**

VMware's cloud computing virtualization platform. It includes updated vCenter Configuration Manager, vCenter Application Discovery Manager, and the ability to vMotion multiple virtual machines at a time from one host server to another.

### 2.2. Configure Steps

In this section we will provide an example of setting up in QSM.

## 2.2.1. Environment and Topology

### Demonstration Environment

- Storage
  - Model: XN8116D  
Memory: 16 GB per controller  
Firmware: QSM 4.1.0  
Data Port IP: 192.168.222.91
- Server
  - Model: ASUS Server  
OS: ESXi 8.0  
Server IP: 192.168.202.121

### Demonstration Topology




Figure 2-1 Demonstration Topology

## 2.2.2. Configure Storage

1. In XN8116D, create a pool and a block volume, then create a block HostGroup and add the volume to the HostGroup.



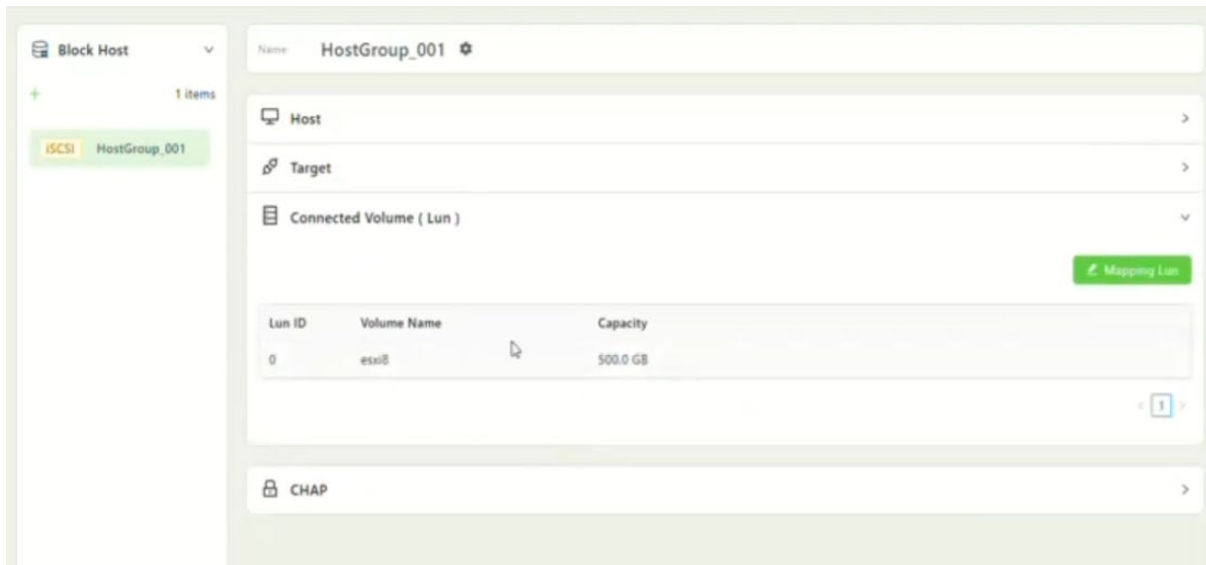


Figure 2-2 Create a Block Volume and Add into HostGroup

## 2.2.3. Configure VMware 8

1. Enter the vCenter 8.0 management page and add an ESXi 8.0 host.

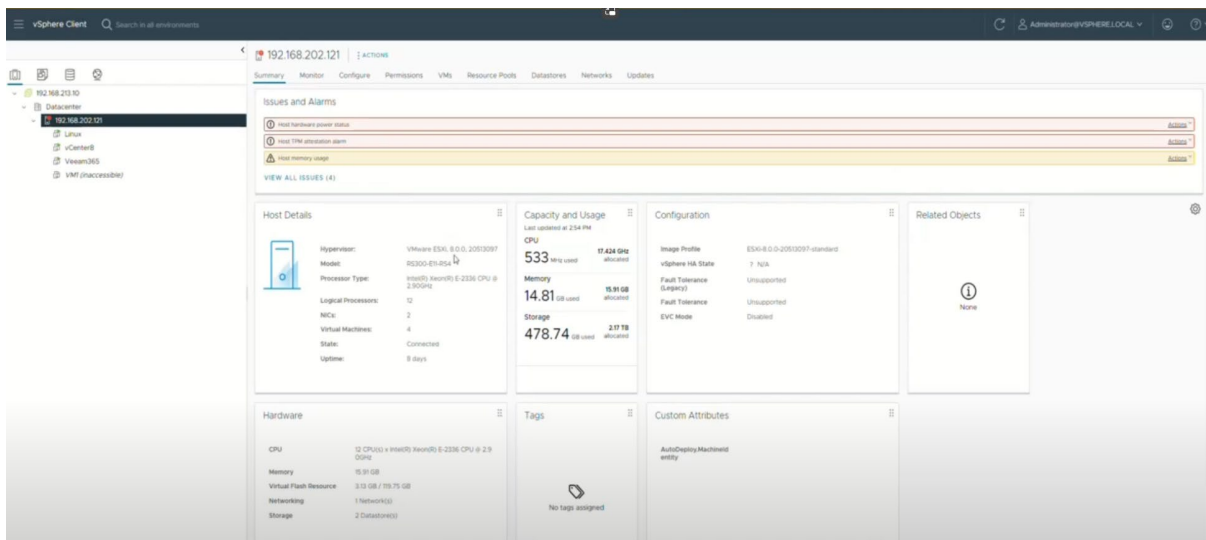


Figure 2-3 Add ESXi 8 Host

2. Go to the **Storage Adapters** menu and check the **vmhba64** item, then select the **Static Discovery** submenu and click the **Add** button to discover iSCSI connections.

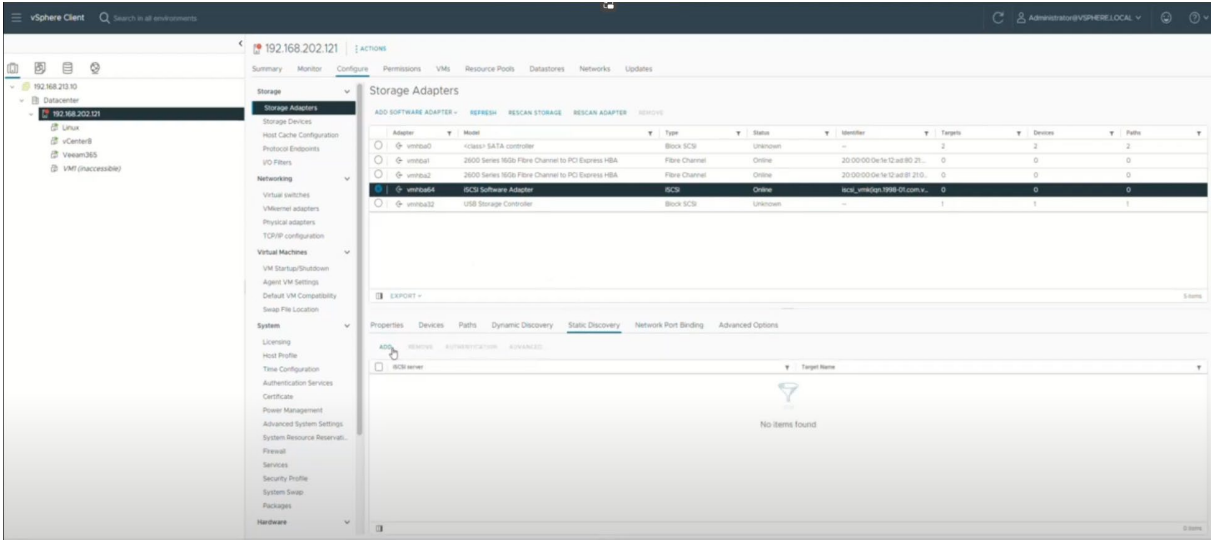


Figure 2-4 Discover iSCSI Connections

- 3. Enter the IP for storage data transfer port, and then copy the host group's IQN from the storage and paste it into the **iSCSI Target Name** field.

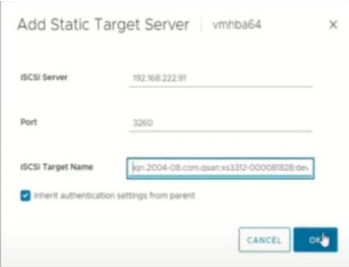


Figure 2-5 Add iSCSI Server

- 4. After rescanning the storage, go to the storage device page and verify whether the iSCSI LUN has been successfully mounted to ESXi.

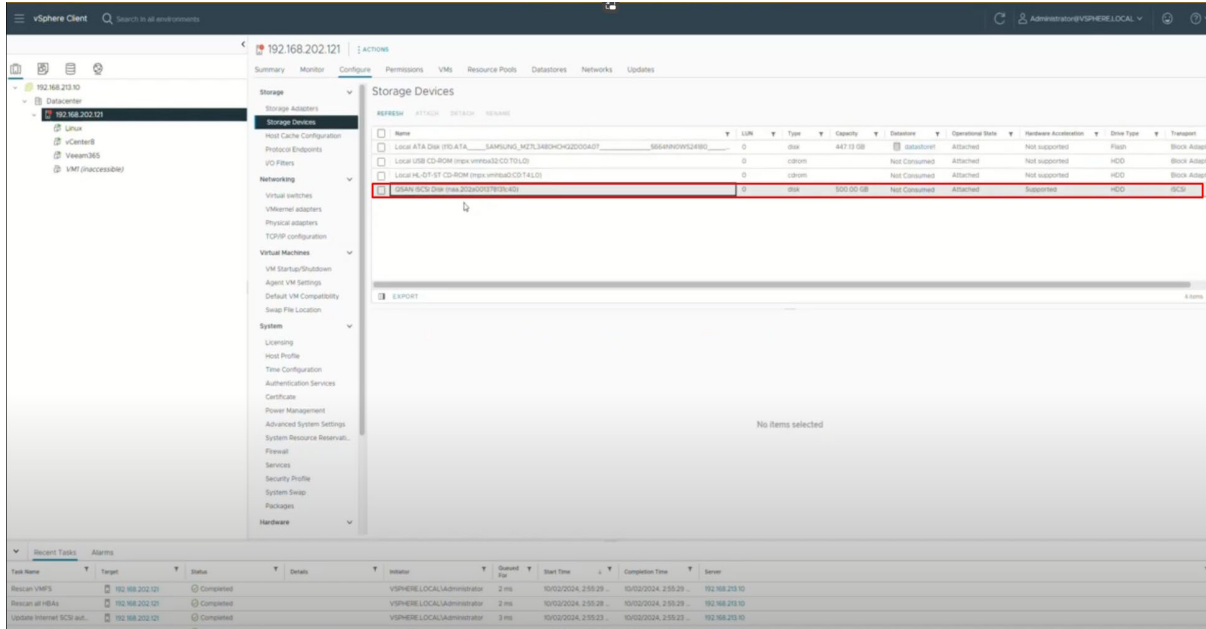


Figure 2-6 Rescan Storage

5. Right-click the ESXi host and add a new Datastore, then select the iSCSI LUN just mounted.

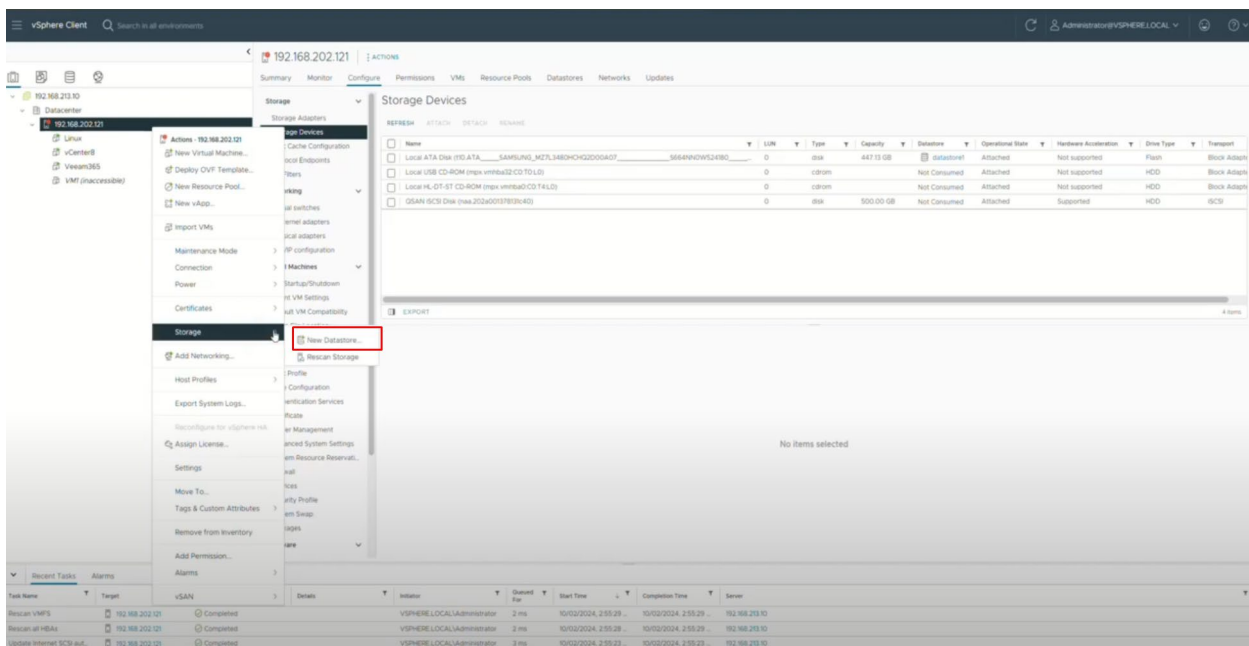


Figure 2-7 Add New Datastore 1

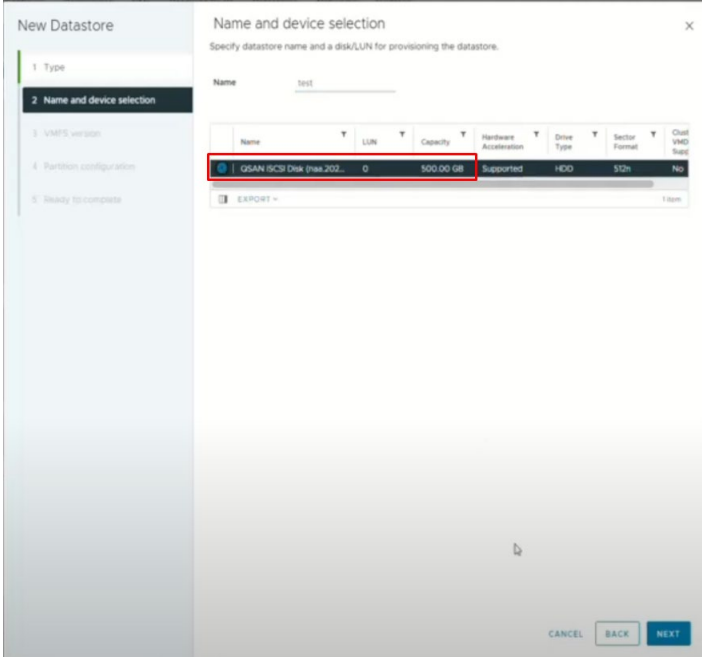


Figure 2-8 Add New Datastore 2

- After creating the datastore, right-click on the datastore and create a new virtual machine.

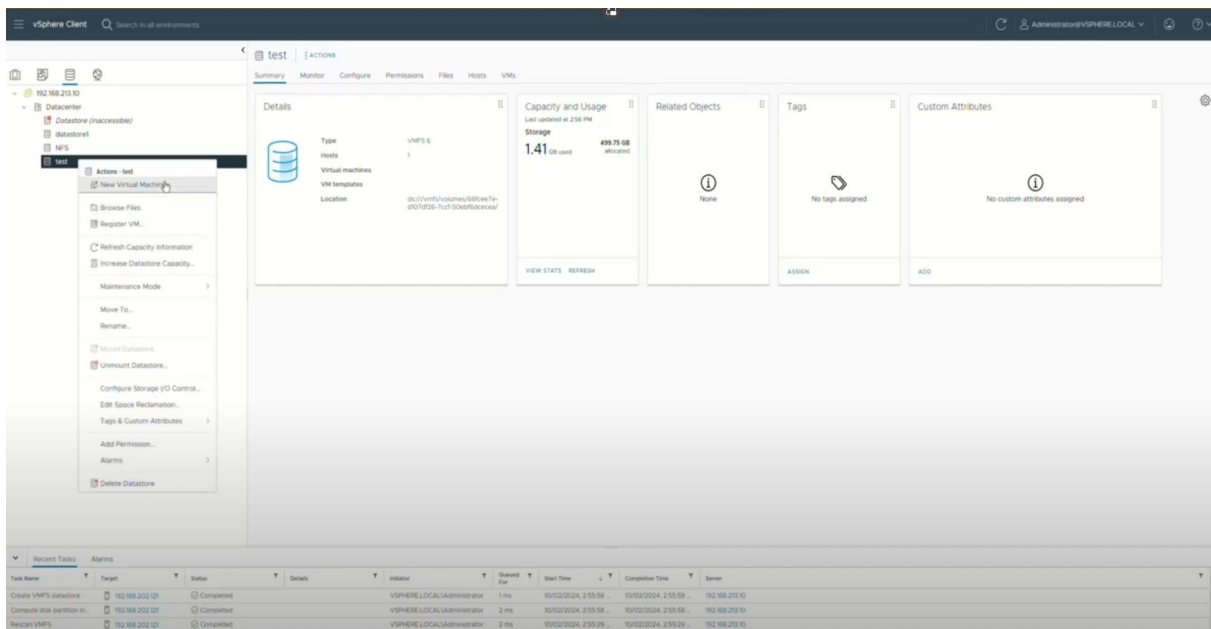


Figure 2-9 Create VM-1

- Configure the virtual machine details settings and click the **Finish** button.

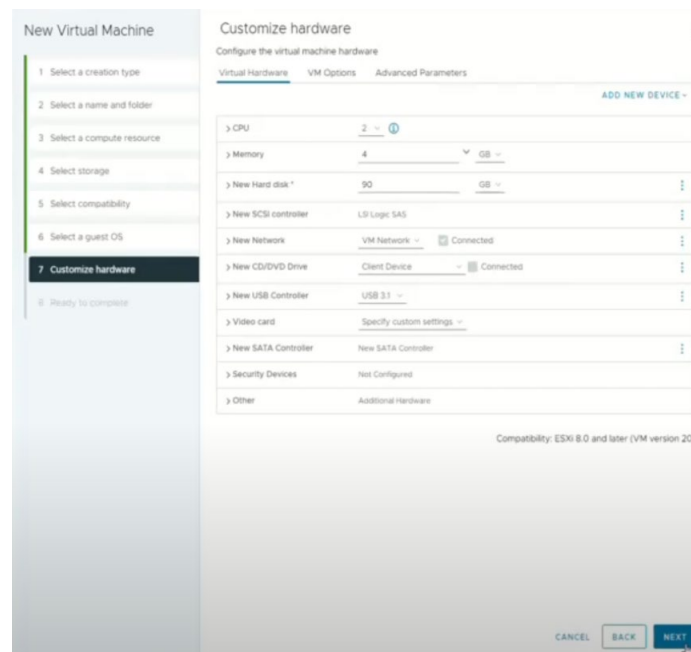


Figure 2-10 Create VM-2

8. Now you can start the virtualized application.

## 2.3. Conclusion

In summary, the integration of QSAN XCubeNXT with VMware ESXi 8.0 simplifies virtualization environment management and provides a scalable, high-performance storage solution. By following the steps outlined in this guide, enterprises can efficiently provision iSCSI LUNs as data storage and deploy virtual machines through vCenter 8.0. This approach enhances flexibility, centralizes storage management, and ensures high availability, making it an ideal solution for organizations looking to optimize their IT infrastructure and improve overall efficiency in managing virtual workloads.

## 2.4. Appendix

### Apply To

- QSM firmware 4.1.0 and later

### Reference

#### Document

- [QSM 4 Software Manual](#)

## 3. INTEGRATION WITH VMWARE VAAI

---

In virtualization and cloud environments, the ever-increasing data production and demand continue to grow, resulting in an increasing demand for high-speed data transmission. Considering the consumption of server and network resources, budget and limited IT resources, it is necessary to find ways to optimize the existing IT resources within the organization.

The VAAI (VMware vSphere Storage APIs for Array Integration) supports direct data transfer in a compatible storage system without data transfer through the host. It can optimize system capacity and performance without increasing cost or complexity. With VAAI, servers can reduce the burden of daily data transmission tasks, thereby reducing the load on servers, SANs (Storage Area Networks), and NASs (Network Attached Storages).

VAAI reduces the burden on the server by using read / write operations to transfer data at the storage array level. Compared with the traditional data transmission method, it also greatly improves the transmission speed.

### 3.1. Introduction to VMware VAAI

VAAI (VMware vSphere Storage APIs for Array Integration) is an API (Application Program Interface) framework that enables many storage tasks, such as Thin Provisioning, Full Copy, Block Zero, and Hardware Assisted Locking. In QSAN storage products with VMware ESXi version 5.x or later, VAAI is supported and fully integrated. When performing storage-related tasks between the VMware ESXi hypervisor and QSAN storage products, this integration can save resources on the VMware ESXi server.

VAAI was introduced in VMware vSphere 4.1 with the following features implemented for achieving offload capabilities:

- Full Copy or Hardware Assisted Move
- Block Zero or Hardware-Assisted Zero
- Hardware Assisted Locking or Atomic Test and Set

Thin Provisioning was introduced in VMware vSphere 5.x. Detailed explanations of these features are presented as following.

### 3.1.1. Thin Provisioning

For scenarios where storage-based Thin Provisioning functions are used, VMware vSphere 5.x implements some VAAI enhancements, and QSAN storage products also support this function. The two main enhancements of VAAI Thin Provisioning are:

- Dead Space Reclamation (also known as UNMAP)
- Out of space conditions

#### Dead Space Reclamation

Traditionally, when a storage volume / LUN was mounted as a datastore, and there were virtual machines stored in the datastore, if any of virtual machines were deleted or migrated, the storage spaces which were occupied by the deleted / migrated virtual machines would still be treated as “in use” from the point of view of storage array. This may lead to a situation where the use of storage space is considered insufficient and the cost of purchasing disks has been wasted.

In the QSAN storage products, this problem can be avoided by providing the function of reclaiming the unused storage space (migrating or deleting virtual machines) when using Thin Provisioning volumes and reflecting it on the management interface of QSAN storage products.

- Advantage

Using this function, the unused storage space can be accurately reported to the QSAN storage system, so that the space can be correctly reclaimed through the space reclamation on the Thin Provisioning volume / LUN in the QSAN storage system.

- Theory

When using Thin Provisioning volumes / LUNs, after deleting the virtual disk or migrating the virtual disk to another datastore, or even after deleting the snapshot, VMware vSphere 5.x will use the SCSI UNMAP command to immediately release the physical space on the volume / LUN.

#### Out of Space Condition

In a Thin Provisioning environment, if the space is insufficient, the datastore space will be over-provisioned by multiple virtual machines, which may lead to a catastrophic situation, which is caused by the lack of space.



In these situations, VMware vSphere 5.x or later will enhance the solution. If Thin Provisioning datastore reaches 100%, virtual machines that only require additional storage space blocks will be suspended, while other virtual machines will remain running.

- **Advantage**

After the Thin Provisioning volume / LUN space is used up, the VMware ESXi server will temporarily suspend the virtual machines that require additional storage space. The administrator can then allocate more storage space by adding other RAID sets to the existing pool.

### 3.1.2. Full Copy

This feature helps storage array to make full copies of data within the array without letting VMware ESXi server physically read or write data to the storage array.

#### Effective Operations

- Clone a virtual machine
- Perform a Storage vMotion
- Deploy virtual machines from a template

#### Advantage

Reduces the CPU loads on VMware ESXi server, and prevent crowded I/Os between VMware ESXi server and storage array.

#### Theory

Without VAAI, when one of the above three operations is performed, the VMware ESXi server will read each block from the storage array and write it to a new location. During this period, a lot of server resources were consumed.

With support for VAAI and based on this feature, the VMware ESXi server sends a single SCSI (Extended Copy) command for a group of consecutive blocks to tell the storage array to copy these blocks from one location to another (new location). The commands on the network (if using iSCSI) are small, and the actual work will be performed within the storage array. This minimizes data transfer and speeds up the copy process. Please refer to the figure below to understand how to perform a Full Copy operation when trying to perform a VM cloning task

from one datastore to another. Of course, these two datastores are based on two volumes/LUNs from the same QSAN storage.

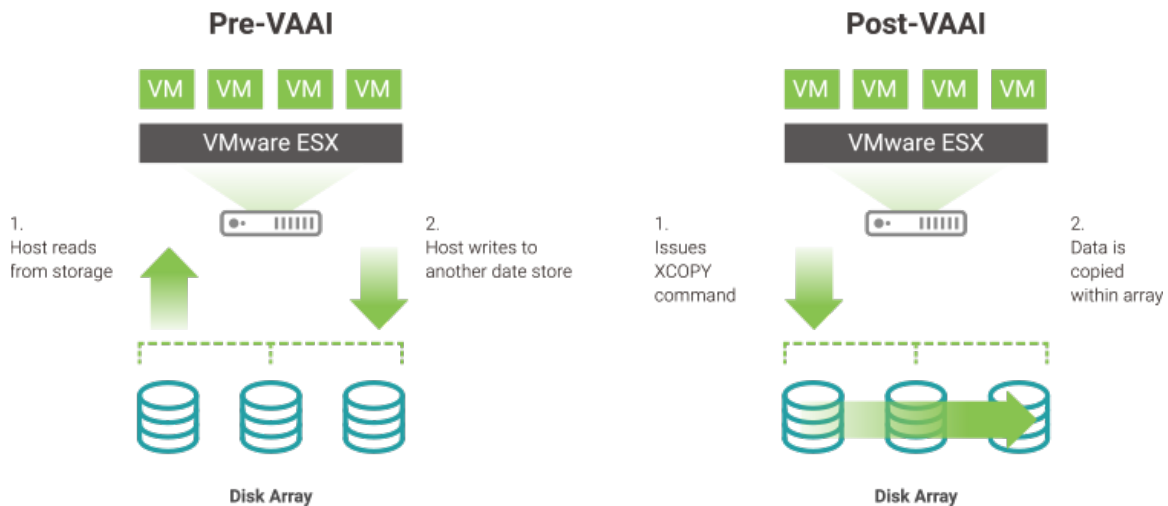


Figure 3-1 VAAI Full Copy

### 3.1.3. Block Zero

This feature helps storage array to zero out a large number of blocks for speeding up virtual machine configuration.

#### Effective Operations

- Create Thin Provisioning Eager Zero virtual disks (Thick Provisioning Eager Zero virtual disks are zeroed out when creating, and are not usable until the process is completed.)
- Write data to an unused area of a Thick Provisioning Lazy Zero virtual disk (Thick Provisioning Lazy Zero virtual disks can be used instantly after they are created.)

#### Advantage

Use this function to offload the process of writing zeros to the storage array. Eliminates repeated repetitive write commands to reduce the load on the VMware ESXi server, thereby greatly improving capacity allocation.

#### Theory

Without VAAI, zeroing disk blocks will send duplicated and repetitive write commands from the VMware ESXi server to each block on the storage array. The VMware ESXi server needs to wait for the completion of the previous write command before sending another command, which will result in huge resource costs and time consumed.

With VAAI enabled, VMware ESXi server uses SCSI Write Same command to tell storage array to write the same data to an amount of blocks. VMware ESXi server then doesn't need to send duplicated write command continuously; instead storage array will return the requesting service as though the process of writing zeros has been completed. QSAN storage finishes the zeroing out internally. Please refer to Figure-2 below, which shows the operation and process how Block Zero is performed between VMware ESXi server and QSAN storage.

After enabling VAAI, the VMware ESXi server uses the SCSI Write Same command to tell the storage array to write the same data to a certain number of blocks. In this way, the VMware ESXi server does not need to continuously send repeated write commands; instead, the storage array will return the requested service as if the process of writing zeros has been completed. QSAN storage finishes the zeroing out internally.

### 3.1.4. Hardware Assisted Locking

Hardware Assisted Locking, also called ATS (Atomic Test and Set), provides another way to protect the metadata of the VMFS cluster file system and improve the scalability of large ESXi servers that share VMFS datastore. ATS helps to lock the blocks in the volume/LUN instead of the entire volume/LUN added as datastore in the VMware ESXi server.

#### Effective Operations

- Create a VMFS datastore
- Expand a VMFS datastore onto additional extents
- Power on a virtual machine
- Acquire a lock on a file
- Create or delete a file
- Create a template
- Deploy a virtual machine from a template
- Create a new virtual machine
- Migrate a virtual machine with vMotion

- Grow a file (e.g., a snapshot file or a thin-provisioned virtual disk)

## Advantage

When multiple VMware ESXi servers share the same datastore, Hardware Assisted Locking (or ATS) provides a more effective way to avoid retries for getting a lock. The locking mechanism is offloaded to the storage array, and the storage array performs locking at a granular level. This is helpful for scalability when sharing a datastore in a VMware cluster environment without compromising the integrity of the metadata in the VMFS shared storage pool.

## Theory

Previously, VMware had a similar mechanism for locking virtual machines to prevent virtual machines from running on them. This mechanism can be modified by multiple VMware ESXi servers at the same time. It is based on the use of SCSI RESERVE and RELEASE commands. This protocol calls the unique access to an entire volume/LUN for the reserving ESXi server until this ESXi server sends a release. Under the protection of the SCSI RESERVE command, the ESXi server can update the metadata records on the storage array to reflect the usage without being disturbed by any other ESXi servers that also call the same part of the same storage array. Please refer to the Figure below, which shows the overall structure of this solution and affects the overall performance of the entire clustered VMware ESXi environment. The performance degradation caused by a large number of RESERVE and RELEASE commands is unacceptable in a VMware cluster environment. The VMware cluster environment accesses shared data storage from different virtual machines exponentially every day.




Figure 3-2 VAAI Hardware Assisted Locking

With VAAI, Hardware Assisted Locking provides a more granular method to protect VMFS metadata than the SCSI RESERVE and RELEASE commands. Hardware Assisted Locking uses the storage array ATS function to enable a fine-grained block-level locking mechanism. First, Hardware Assisted Locking replaces the sequence of RESERVE, READ, WRITE and RELEASE SCSI commands with a single SCSI COMPARE AND WRITE (CAW) request for an atomic read-modify-write operation, based on the presumed availability of the target lock. Then, this new request only requires exclusion of other accesses to the target locked block, not the entire VMFS (which is volume/LUN) which contains the requested lock. When the virtual machine state changes, VMware uses this lock metadata update operation. This may be due to turning ON or OFF the power of the virtual machine, or modifying the configuration of the virtual machine, or even migrating the virtual machine from one ESXi server to another through vMotion.

### 3.1.5. Hardware Acceleration Support Status

After adding any storage volume/LUN through VMware vSphere Client, you can observe the status of hardware acceleration. Please navigate to the Configuration -> Hardware -> Storage, then click Datastores View, and check the Hardware Acceleration column displayed after each added datastore.

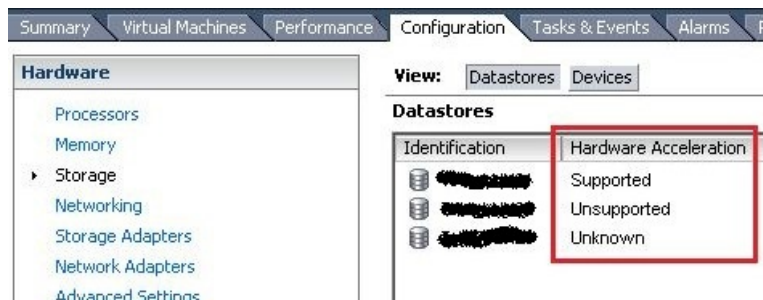


Figure 3-3 Hardware Acceleration Support Status

Table 3-1 Hardware Acceleration Status values

STATUS VALUE	DESCRIPTION
<b>Supported</b>	Storage devices support VAAI
<b>Unsupported</b>	Storage devices do not support VAAI
<b>Unknown</b>	Local datastores

## 3.2. Test Results

The integration of VAAI provides many benefits for improved performance. We have prepared tests and provided some experimental data to prove that VAAI is effective.

### 3.2.1. Environment and Topology

In this test, we use an example to build an environment that connects a VMware ESXi server with a QSAN XS5316D storage array to test the VAAI function.

#### Demonstration Topology

For the connection between XS5316D storage array and VMware ESXi server, please refer to the figure below. In this example, a brief environment will be provided.




Figure 3-4 VAAI Test Diagram

### 3.2.2. Configure Storage

By simulating two FC (Fibre Channel) volumes / LUNs from the XS5316D storage array as two VMFS data stores, this test was performed on a VMware ESXi server to simulate VM cloning and storage vMotion functions. The following Figure provides an idea of how to create pools and volumes.

When verifying the time taken for optimal data protection, the cache mode of these volumes is set to WT (Write-Through, the cache of the storage array on the volume is set to OFF).

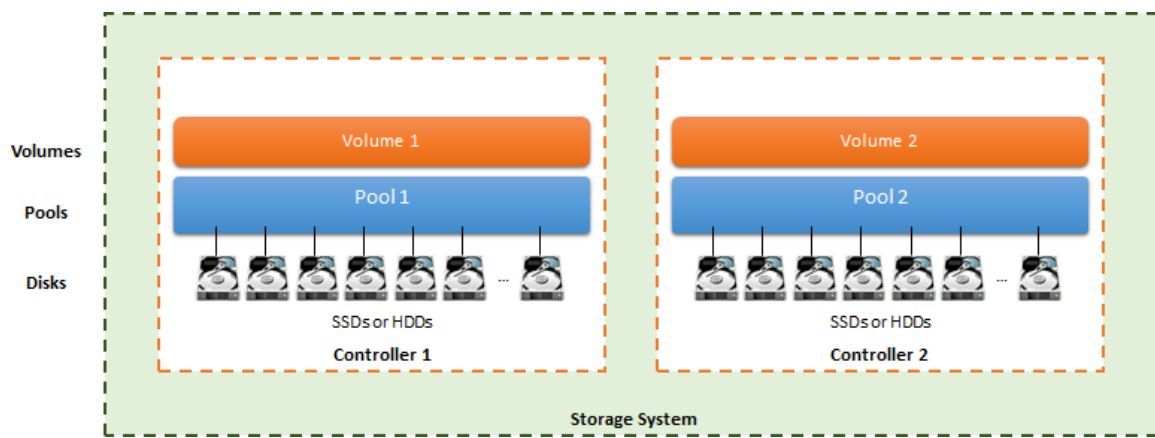


Figure 3-5 VAAI Storage Configuration

### 3.2.3. Thin Provisioning Test

#### Test Scenario

1. Create a Thin Provisioning volume / LUN on the XS5316D storage array, in which 200 GB has been allocated.
2. Create a VMFS datastore on the connected VMware ESXi server.
3. Create a virtual machine based on this VMFS datastore, set the type of Disk Provision to Thin Provision, and set the size to 100 GB.
4. Generate about 50 GB data on the virtual machine. Observe the capacity consumed on VMFS datastore, which uses 50 GB of 200 GB.
5. Performed storage vMotion with VAAI ON to migrate the virtual machine to another VMFS datastore.
6. Observed the consumed capacity on the source VMFS datastore again, you will find that the used capacity is about 0 GB of 200 GB.

7. However, after the virtual machine is migrated, check the Available Capacity (GB) in the WebUI of XS5316D, it may still show 50 GB. This is because the granularity of the QSAN XCubeSAN storage array is 1 GB, and there is only a continuous 1 GB as zero blocks can be reclaimed.
8. Please create a new virtual machine with Thick Provision Eager Zero on this VMFS datastore, and delete it after creation, then execute Space Reclamation in XS5316D storage array, the space shall be able to be reclaimed.

### Summary

The supported granularity in the Thin Provisioning pool of QSAN XCubeSAN series products is 1 GB. Although space reclamation can be enabled when creating a volume, it is sometimes necessary to manually fill zero blocks from the server so that unused blocks can be filled with zeros and reclaimed.

## 3.2.4. Full Copy Test

### Test Scenario

1. Create a virtual machine with a 200 GB Thick Provision Lazy Zero virtual disk on a VMFS datastore. The virtual disk is made of a FC (Fibre Channel) volume / LUN from the XS5316D storage array. The actual storage consumption on the datastore is about 77 GB.
2. Migrate or clone the virtual machine from this datastore to another one which is made by another FC volume / LUN from the same XS5316D storage array.
3. Observe the time it takes to migrate or clone a virtual machine.
4. Repeat the above steps 1 to 3 with disabling VAAI, and compared the time taken.
5. The Table below shows the results of VAAI ON and VAAI OFF.

### Test Result

Table 3-2 Time Taken for Full Copy

FULL COPY USE CASE	VAAI OFF	VAAI ON
Storage vMotion	26 min. 56 sec. = 1,616 sec.	6 min. 5 sec = 365 sec.



<b>Virtual Machine Clone</b>	25 min. 50 sec. = 1,550 sec.	5 min. 59 sec. = 359 sec.
------------------------------	------------------------------	---------------------------

### Summary

Compared with VAAI ON and OFF, the performance is improved by 77.4% when testing storage vMotion. And the performance increased by about 76.8% when testing the virtual machine clone. So, it reduces the CPU loads on VMware ESXi server, and prevent crowded I/Os between VMware ESXi server and storage array.




Figure 3-6 Time Saving of VAAI Full Copy

### 3.2.5. Block Zero Test

#### Test Scenario

1. Measured the time taken to create a 200 GB Thick Provision Eager Zero virtual disk on a virtual machine.
2. Repeated the same step above in comparison with VAAI OFF.
3. The Table below shows the results of VAAI ON and VAAI OFF.

Table 3-3 Time Taken for Block Zero

BLOCK ZERO USE CASE	VAAI OFF	VAAI ON
Thick pool volume	9 min. 6 sec. = 546 sec.	4 min. 0 sec. = 240 sec.

### Summary

When VAAI is enabled and trying to create a 200 GB Thick Provision Eager Zero virtual disk, the performance improves by 56%. The virtual disk is stored in a thick pool in the XS5316D storage array.

## 3.3. Conclusion

The integration of VAAI in all QSAN storage series of products provides many benefits for improved performance and storage array management. The main features are:

- The Dead Space Reclamation capability of the Thin Provisioning feature can reclaim blocks from the thin provisioning volume / LUN on the QSAN storage products. In this way, you can avoid the lack of space by temporarily suspending the virtual machine. When the VMFS datastore space is used up, the virtual machine needs additional space. The administrator can then allocate more capacity by adding more RAID sets to the existing pool.
- The Full Copy feature accelerates the storage vMotion or virtual machine clone operations by transferring operations from the VMware ESXi server to the storage array itself, and greatly reduces resource usage when performing these operations.
- The Block Zero feature speeds up the deployment of Thick Provision Eager Zero virtual disks by offloading a large number of duplicate and repetitive zero blocks to the QSAN storage platform, helps to free the resources of VMware ESXi server for other tasks.
- The Hardware Assisted Locking feature delivers a more effective method to prevent retrying to obtain a lock when multiple ESXi servers share the same datastore. It can offload the locking mechanism to the QSAN storage array, which can be locked at a granular level. This improves the scalability of large ESXi servers that share the same datastore.

## 3.4. Appendix

### Apply To

- XEVO firmware 2.0.0 and later
- QSM firmware 3.3.0 and later

### Reference

#### Document

- [VMware vSphere Storage APIs: Array Integration \(VAAI\)](#)

## 4. INTEGRATION WITH VMWARE VASA

---

In virtualization and cloud environments, the ever-increasing data production and demand continue to grow, resulting in an increasing demand for storage device that offer enough capacity for all application. Considering the consumption of limited IT resources, it is necessary to find ways to help IT manager for easy management and let the IT manager to finish something more value.

### 4.1. Introduction to VMware VASA

VASA (vStorage APIs for Storage Awareness) is a software component first introduced in vSphere 5. It acts as an information conduit between storage systems and vCenter Server, allowing you to monitor related storage system status. VASA Provider collects data from storage systems and communicates information about storage topology, LUN and volume properties, and events and alerts to vCenter Server. Although VASA Provider is transparent to you after installation and configuration, the information VASA Provider sends to vCenter Server can assist you in making critical decisions about placing virtual machines on vCenter Server datastores.

#### **VASA Provider**

VASA (vStorage APIs for Storage Awareness) is a set of APIs that enable vSphere vCenter to recognize the capabilities of storage arrays. This visibility makes it easier for virtualization and storage administrators to make decisions about how to maintain datastores.

#### **vVol**

vVols (Virtual Volumes) enable administrators to apply policies to virtual machines that define various performance and service level agreement requirements, such as RAID levels, replication, or deduplication. Virtual machines are then automatically placed on storage arrays that meet these requirements.

#### **Datastore**

Used to store virtual machine files, templates, and ISO images. They can be formatted using VMFS (Virtual Machine File System, VMware's Cluster File System) or the storage provider's native file system (for NAS / NFS devices).

### 4.1.1. Introduction to QSAN VASA Provider

QSAN VASA Provider is the communication bridge between vCenter and QSAN storage. The following figure shows the QSAN VASA Provider system architecture.

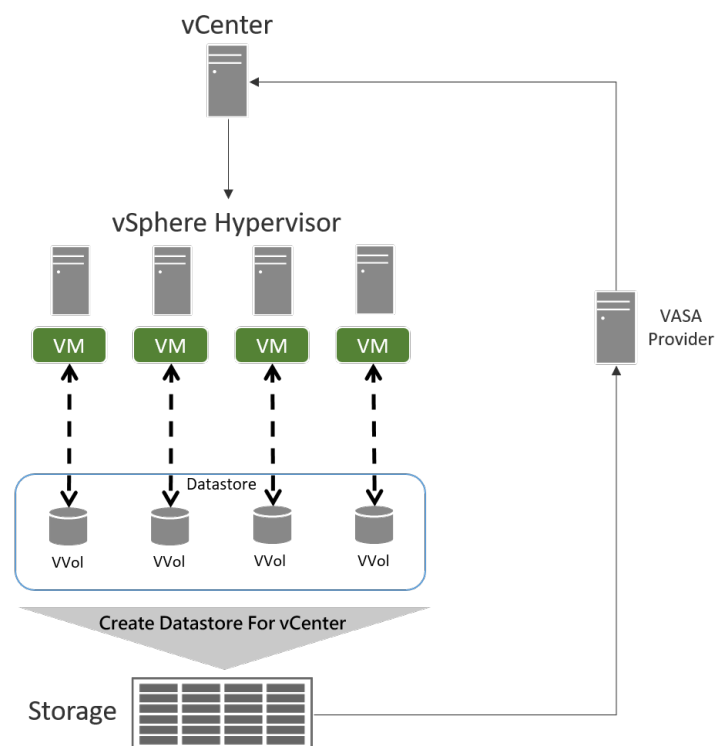


Figure 4-1 QSAN VASA Provider System Architecture

The QSAN VASA Provider integrates APIs between QSAN storage and vCenter, making vSphere vCenter aware of the capabilities of QSAN storage. This support can significantly reduce the workload of virtualization management. In QSM, we support VASA 2.0, which includes:

#### VASA 1.0

Collects data from storage systems and communicates information about storage topology, LUN and volume properties, and events and alerts to vCenter Server.

## **VASA 2.0**

VASA protocol version 2.0 introduces a new set of APIs specific to Virtual Volumes for managing storage containers and Virtual Volumes. It also provides communication between vCenter, hosts, and storage.

# **4.2. Implementation**

## **4.2.1. Preparation**

### **QSAN VASA Provider Software**

Please download and install from QSAN [Download Center](#).

### **Requirement**

Support OS: Windows server 2016, Windows server 2019, Windows 10, and Windows 11.

## **4.2.2. Install VASA Provider**

Here are the steps.

1. Double-click the setup file.
2. Enter the installation and follow the process.

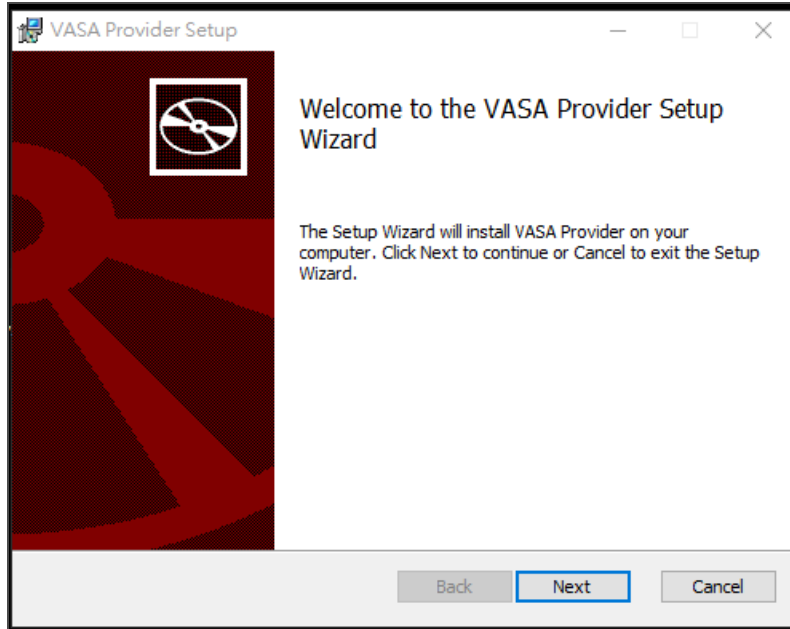


Figure 4-2 Install VASA Provider Step 1

3. Click the **Next** button.
4. Accept the license agreement.
5. Click the **Install** button.

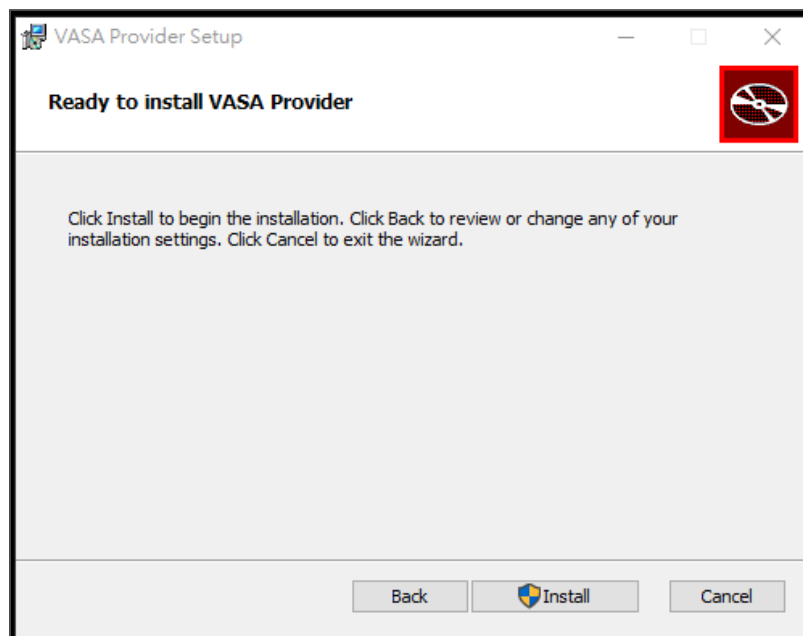


Figure 4-3 Install VASA Provider Step 2

6. Wait for the installation complete and click the **Finish** button.



## INFORMATION

To install VASA Provider, make sure you have the latest JAVA JRE (Java Execution Engine) installed first.

### 4.2.3. Setup VASA Provider

Here are the steps.

1. Once installed, double-click the application.

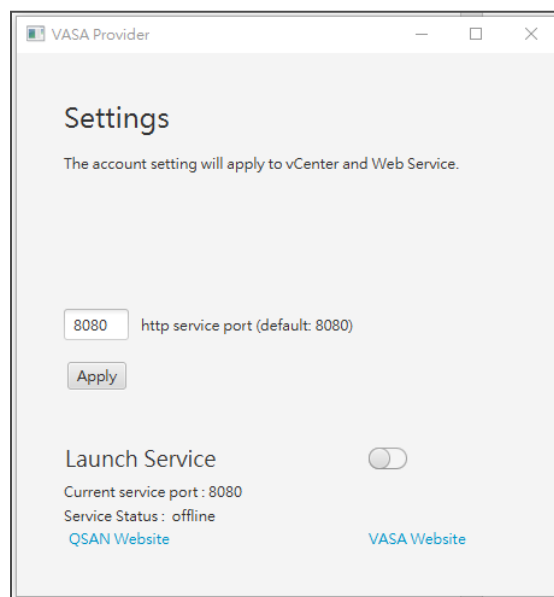


Figure 4-4 Setting the Port Number



## INFORMATION

This application will not run in background, please keep it running.

2. You can set up the port number (default 8080), and then click the **Apply** button.



3. Turn on the launch service button to and press VASA website to enter web UI.
4. Use default username and password to login.
  - Username: **username**
  - Password: **password**




Figure 4-5 Login to QSAN VASA Provider

5. After login, you will see **Settings** and **Registered**.
  - **Settings:** You can change login information / Generate Signed Certificate / URL to join vCenter.
  - **Registered:** Add / manage devices in VASA provider.

The screenshot shows the QSAN configuration interface. At the top, there are icons for 'Settings' and 'Registered', and the 'QSAN' logo with an 'About' link. The main content area is divided into three sections: 1. 'Server Status' with 'Username' and 'Password' input fields and an 'Apply' button. 2. 'Self Signed Certificate' with a 'Generate' button. 3. 'VASA Provider' with a 'VASA URL' dropdown menu and a 'Copy URL to Clipboard' button.

Figure 4-6 Configuring QSAN VASA Provider

6. After setting up and registering the device, you can enter vCenter to add the datastore.



## INFORMATION

To create a datastore, please refer to the [QSM Software Manual](#).

After setting up the VASA Provider and creating the datastore, the final step is to connect them to vCenter and vSphere. You are then free to create any virtual machine on the datastore.

### 4.2.4. Add a VASA Provider

Here are the steps.

1. Go to **vCenter -> Configure -> More -> Storage Provider -> Add** to add a VASA provider.

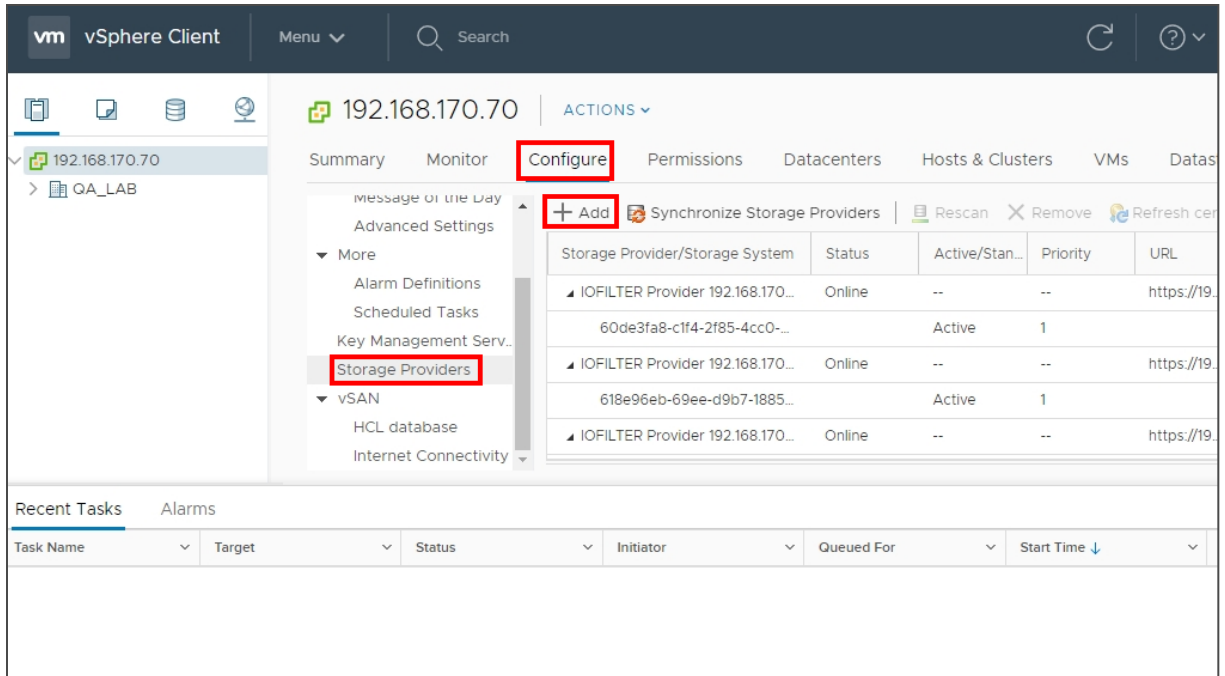


Figure 4-7 Add a VASA Provider

2. Complete the content and click the **OK** button.

- **Name:** You can decide the display name.
- **URL:** VASA Provider URL, you can copy it from VASA Provider setting page.
- **Username and password:** Enter the username and password of the VASA provider.
- **Use storage provider certificate:** This is not the most commonly used item; you can download it from the VASA provider settings page and upload it here to authenticate the device.

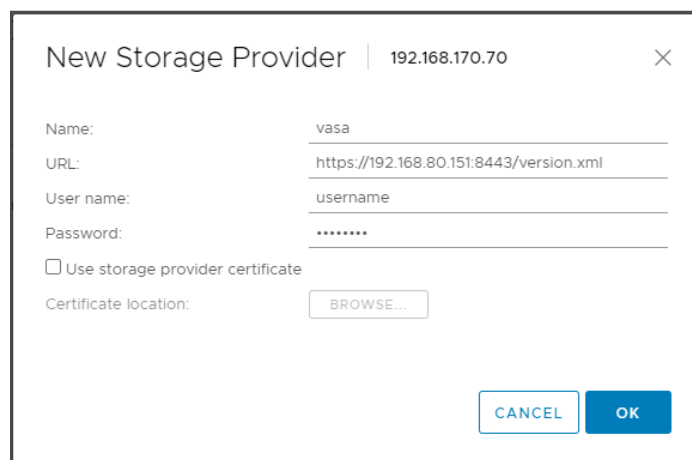


Figure 4-8 Add a New Storage Provider

3. After adding the provider, go to **vSphere** -> **Configure** -> **Storage Adapter**, select an existing storage adapter -> **Static discovery** -> **Add**. If don't have a storage adapter, please add a new one.

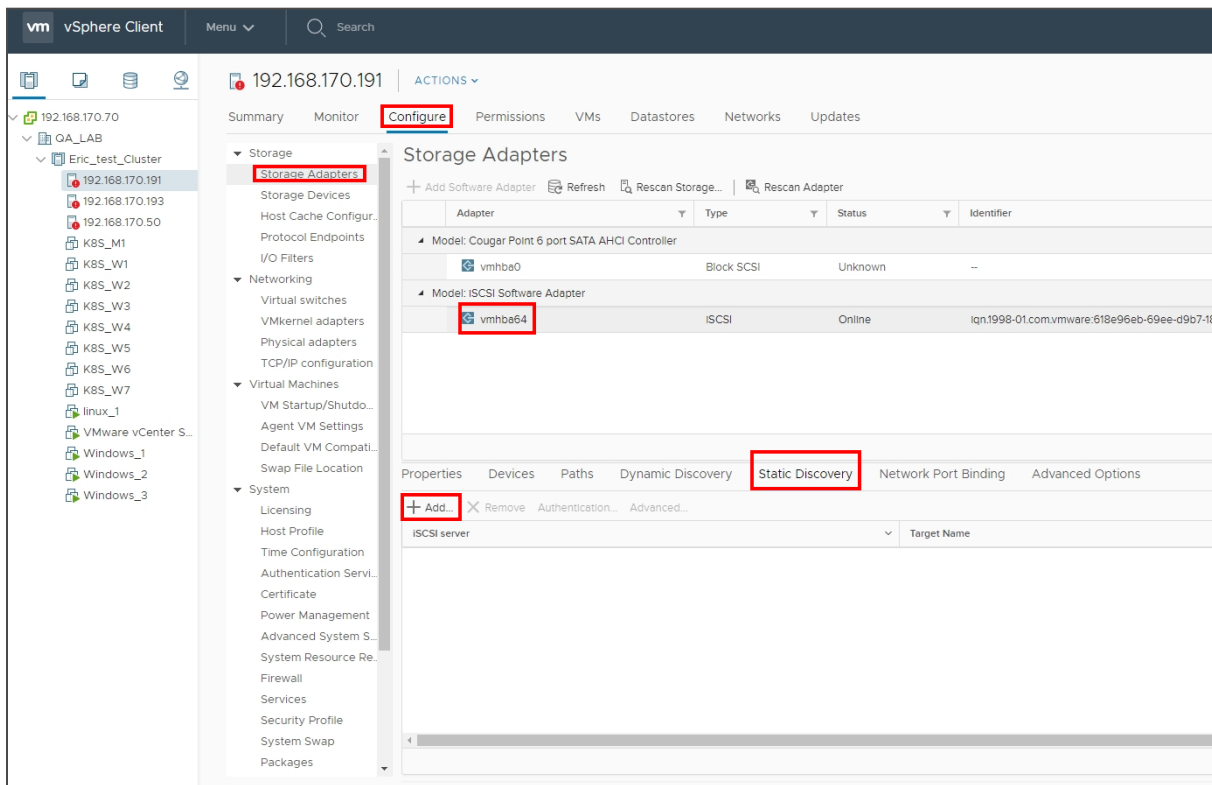


Figure 4-9 Add a Storage Adapter

4. Fill in the content and click the **OK** button.
  - **iSCSI Server:** Enter the storage IP.
  - **Port:** Enter the iSCSI port of the storage.

Add Static Target Server | vmhba64

iSCSI Server: Fully Qualified Domain Name or IP

Port: 3260

iSCSI Target Name:

Inherit authentication settings from parent

CANCEL OK

Figure 4-10 Add an iSCSI Target

5. Go to **Storage** and select **vSphere**, then click **Actions -> Storage -> New Datacenter** to add a datastore in vCenter.

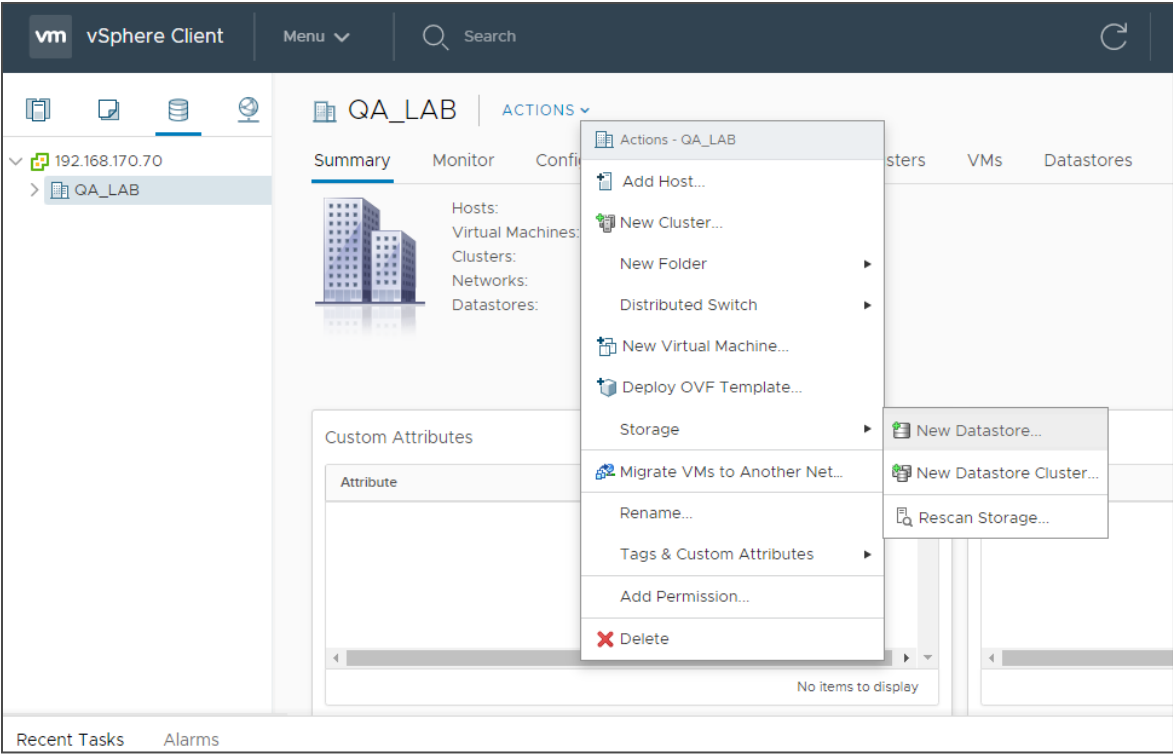


Figure 4-11 Add a Datastore Step 1

6. Select the **VVol** type and complete the process.

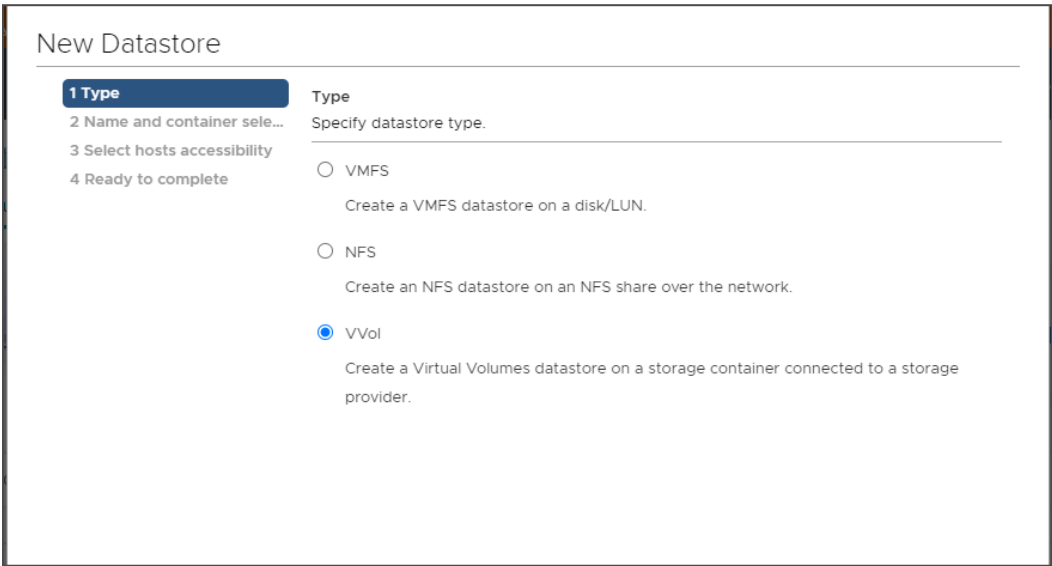


Figure 4-12 Add a Datastore Step 2

7. After completion, you can use the datastore to create a virtual machine. Simply select the new added datastore on the Create Virtual Machine page.

## 4.3. Conclusion

Increased data production and demand in virtualized and cloud environments have led to an increasing need for storage devices that provide sufficient capacity for all applications.

Considering the consumption of limited IT resources, it is necessary to find ways to help IT managers manage them easily and allow IT managers to complete more valuable things.

QSM is now fully compatible with VASA and has released the QSAN VASA Provider. QSAN storage and VMware ESXi provide an efficient and cost-effective solution. It also optimizes IT resources and provides agile solutions for increasing amounts of data.

## 4.4. Appendix

### Apply To

- QSM firmware 3.4.0 and later

### Reference

Document

- [QSM 4 Software Manual](#)

## 5. CONFIGURE VMWARE CLUSTER VMDK

---

This chapter provides technical guidance for setting up WSFC (Windows Server Failover Clustering) in a VMware environment. A comparison of different deployment methods is presented. And point out the advantages of the new feature cluster VMDK. It also provides installation tips to help users not get tripped up.

### 5.1. Introduction to VMware Cluster VMDK

This section provides an overview of WSFC (Windows Server Failover Clustering) and a brief introduction to its setup in a VMware environment. VMware's new feature cluster VMDK is released in ESXi 7. Finally, a comparison is made between the traditional architecture and the new features.

#### 5.1.1. Windows Server Failover Clustering

A failover cluster is a group of independent computers that work together to improve the availability and scalability of the cluster role. To reduce system downtime and ensure high availability for Windows, you can cluster servers (called nodes) so that if one node fails, one or more other nodes automatically take over processing. This is also referred to as Windows clustering.



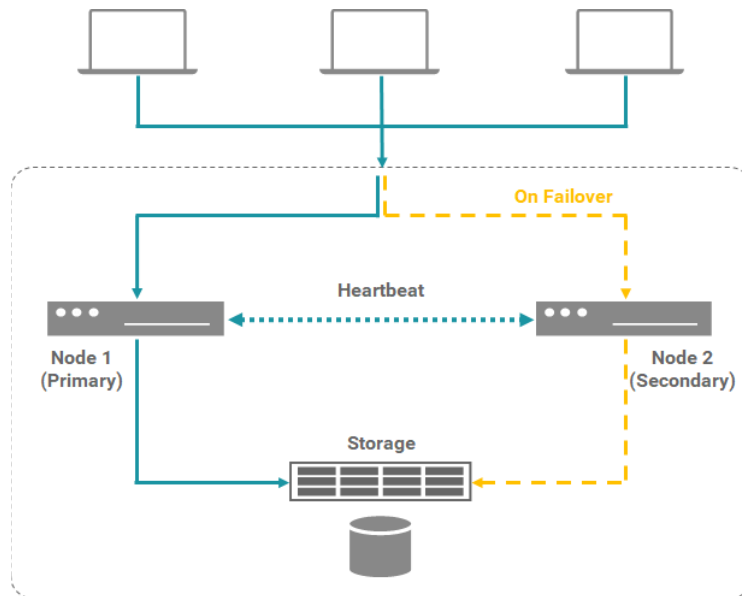


Figure 5-1 Windows Server Failover Clustering Architecture

Cluster servers are connected by physical cables and software. If one or more cluster nodes fail, other nodes begin to provide service (a process known as failover). Additionally, cluster roles are actively monitored to verify that they are functioning properly. If they don't work, they will be restarted or moved to another node.

Failover clustering has many practical applications, including highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines.

### 5.1.2. How Windows Clustering Works

Cluster software is required to monitor the health of the primary node and initiate recovery operations when problems are detected. High availability clustering also need a way to ensure that in the event of a failure, the secondary node is accessing the newest data in storage. In most cases, this is achieved by connecting all nodes of the cluster to the same shared storage.

Failover clustering also provides the CSV (Cluster Shared Volume) feature, which provides a consistent distributed namespace that cluster roles can use to access shared storage from all nodes. With failover clustering, users can minimize service interruptions.

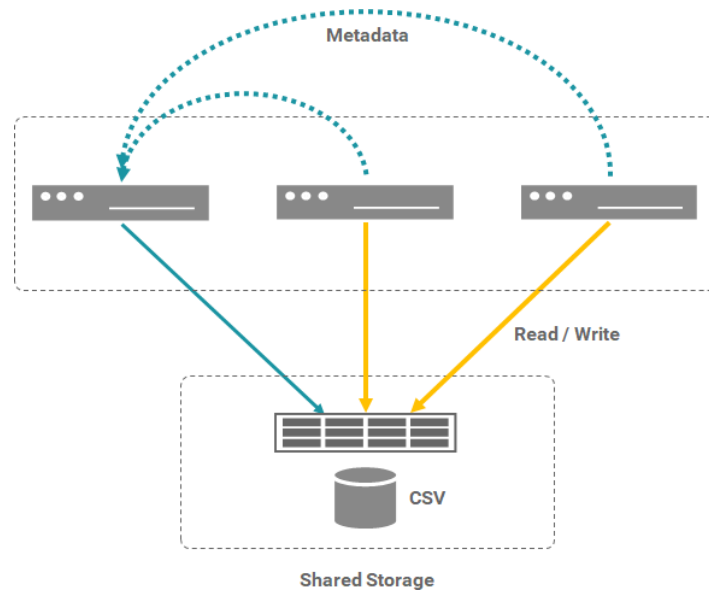


Figure 5-2 Cluster Shared Volume

With CSV feature, all cluster nodes can access the CSV at the same time. Server-side metadata synchronization avoids I/O interruptions. It is recommended that cluster nodes should be geographically separated to protect applications from site-area disasters.

### 5.1.3. VMware RDM

WSFC deployments can be virtualized. VMware vSphere supports Windows clustering using WSFC across virtual machines. Clustering virtual machines can reduce the hardware costs of traditional high-availability Windows clusters.

RDM (Raw Device Mapping) is VMware virtualization technology that allows a VM (Virtual Machine) to directly access LUN (Logical Unit Number). It is a special mapping file in a VMFS volume that manages the metadata of its mapped devices. The mapping file is provided to the management software as a normal disk file and can be used for file system operations. For virtual machines, the storage virtualization layer presents mapped devices as virtual SCSI devices.




Figure 5-3 Raw Device Mapping



## INFORMATION

Regarding RDM, see the following documentation for more details.

- [About Raw Device Mapping](#)
- [Differences between Virtual and Physical RDM](#)

### 5.1.4. Clustered VMDK

In ESXi 7.0, when the VMs hosting the cluster nodes are on different ESXi hosts, clustered VMDKs (Virtual Machine Disk Format) on VMFS (Virtual Machine File System) datastores are supported. VMware has added support for SCSI-3 PR (Persistent Reservations) at the virtual disk level. You can now deploy WSFC using a clustered (shared) VMDK.




Figure 5-4 Enabling Clustered VMDK

Clustered VMDK support can be enabled when creating a new VMFS6 datastore, or enabled on an existing VMFS6 datastore. Before enabling clustered VMDK support, ensure that all hosts connected to the datastore are using ESXi 7.0 or later and managed by vCenter 7.0 or later. All hosts connected to the datastore must be managed by the same vCenter with the cluster VMDK flag disabled or enabled on the datastore. With the cluster VMDK flag enabled or disabled, the host can be managed by any vCenter with version 7.0 or higher.

### 5.1.5. RDM vs. Clustered VMDK

Both RDM and clustered VMDK can help you to setup WSFC, but what we'll highlight here is how RDM compares to clustered VMDK in VMware.

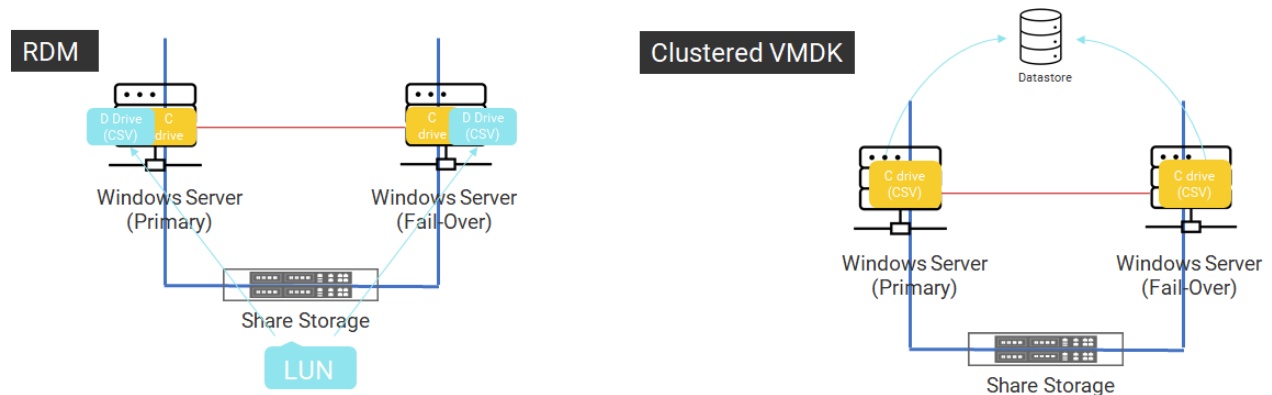


Figure 5-5 RDM vs. Clustered VMDK

Storage in a clustered environment should have a locking mechanism to prevent writes to the same block. It essentially uses this command to lock the volume so only active nodes are allowed to write to it. But since VMFS has its own locking mechanism, these SCSI commands are

intercepted and dropped by traditional virtual disks. Therefore, RDM disks need to be used as mapping devices for physical LUNs.

Clustered VMDK allows SCSI-3 PR commands to be issued to virtual disks, which means you will no longer need a dedicated physical LUN to start a Windows failover cluster.

Table 5-1 RDM vs. Clustered VMDK

	RDM	CLUSTERED VMDK
<b>Compare Items</b>	<ul style="list-style-type: none"> <li>Local C drive for OS only</li> <li>Need to map new LUN for D drive</li> <li>Set D drive into CSV<sup>1</sup> for data saving</li> </ul>	<ul style="list-style-type: none"> <li>Local C drive can be set as CSV<sup>1</sup></li> <li>Extra capacity share the same datastore</li> </ul>
<b>Conclusion</b>	<ul style="list-style-type: none"> <li>Extra LUN mapping</li> <li>Extra drive for data</li> </ul>	<ul style="list-style-type: none"> <li>No extra LUN mapping</li> <li>One drive only</li> <li>Reduce setting configuration process</li> </ul>

<sup>1</sup> CSV (Cluster Shared Volume)

In summary, supporting clustered VMDK simplified the process of VM application environment when setting up WSFC. You can now migrate and delete those RDMs created in your environment to handle failover clusters, allowing these Windows VMs to access VMware's unified and simplified virtual disk management.

## 5.2. Installation Tips

This section provides tips for installing WSFC on VMware using a clustered VMDK. We emphasize tips rather than complete installation steps, as you can find some detailed documentations on the Internet.

### 5.2.1. Prerequisites for Clustered VMDK

Clustered VMDKs come with a bunch of limitations in VMware documents. Some of which are related to arrays.

- Only supported with arrays using FC (Fibre Channel) for connectivity.
- The array must support ATS SCSI commands.
- The array must support SCSI-3 PR (Persistent Reservations), specifically WEAR (Write Exclusive - All Registrants).

There are some prerequisites to VMware and WSFC for using clustered VMDKs.

- CIB (Cluster in a Box) configuration is not supported.
- The datastore must be formatted with VMFS 6 (VMFS 5 is not supported).
- VMDK must be Eager Zero Thick (no thin provisioned VMDKs).
- If you have DRS configured in your environment, you must create an anti-affinity rule so that the VMs can run on separate hosts.
- vCenter server 7.0 and higher.
- Snapshots, cloning, and storage vMotion are not supported (no backup of nodes is possible, because backup software uses snapshots).
- Fault tolerance, hot change to the VMS virtual hardware, and hot expansion of clustered disks are not.
- vMotion is supported, but only for hosts that meet the same requirements.

### 5.2.2. WSFC Topology

The WSFC environment on VMware is shown below.

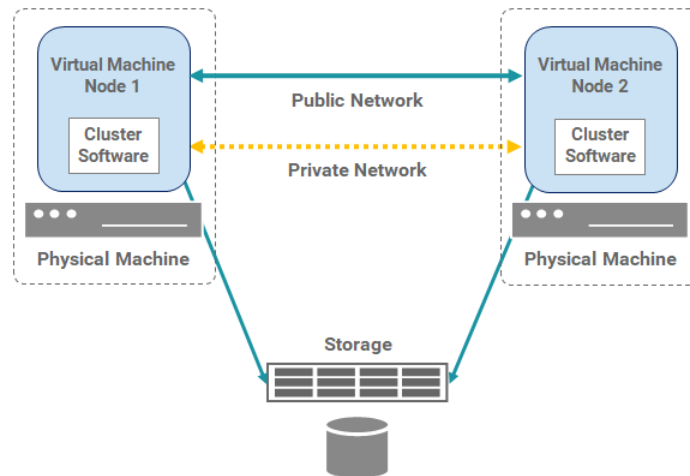


Figure 5-6 Virtual Machines Clustered Across Hosts

Clustering of virtual machines across physical ESXi hosts protects against software and hardware failures on physical ESXi hosts by placing cluster nodes on separate ESXi hosts. This configuration requires shared storage for cluster disk resources. Note that the clustered VMDK supports CAB (Cluster Across Boxes) instead of CIB (Cluster in a Box).

The above figure shows a CAB setup.

- Two virtual machines on two different ESXi hosts running WSFC.
- Virtual machines share private and public network connections for private heartbeats.
- Each virtual machine is connected to shared storage.

### 5.2.3. Setup Clustered VMDK Tips

We provide some tips for setting up WSFC via a clustered VMDK, not full steps, as there is already detailed documentation on the installation.

- When you create a datastore in vSphere 7, there is a new column in the Create Datastore wizard called **Clustered VMDK Supported** that will tell you if the array device is.
- Click the **Enable** button during datastore creation.

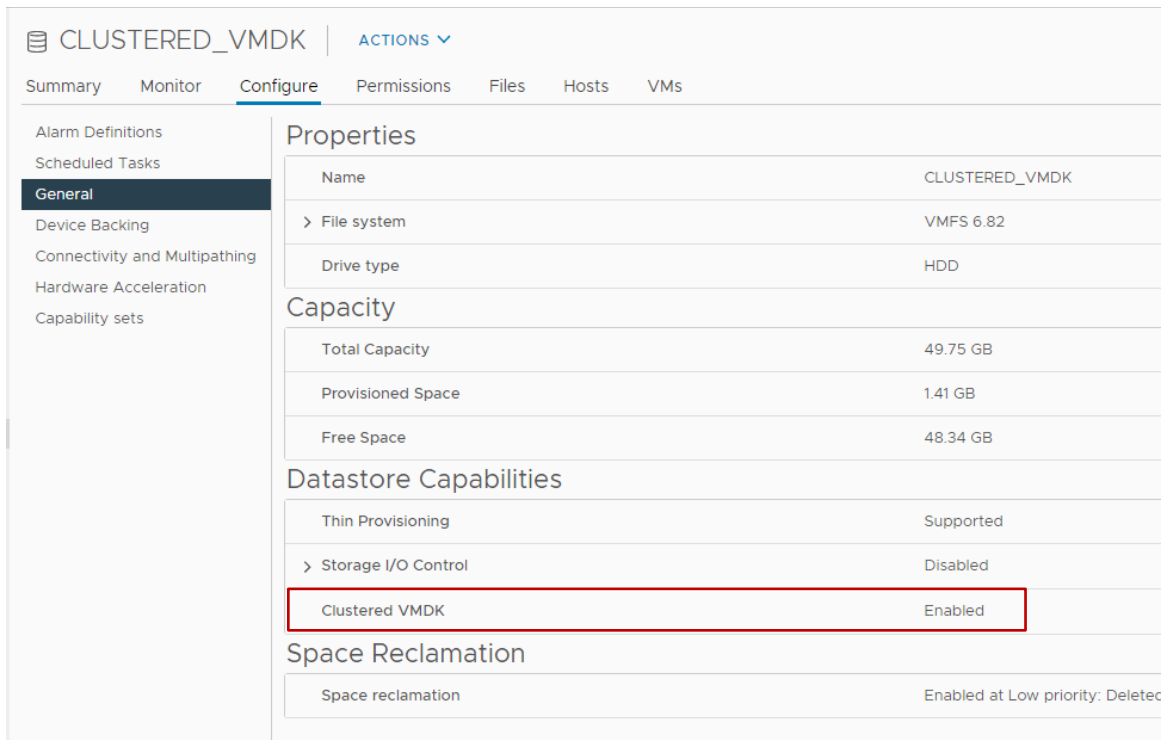


Figure 5-7 Enable Clustered VMDK

- When creating the 1<sup>st</sup> virtual machine, select the **Disk Provisioning to Thick Provision Eager Zeroed** and select the **Virtual Device Node to New SCSI controller**.
- On the 2<sup>nd</sup> virtual machine, click the **Add New Device with Existing Hard Disk**.
- Set the Windows Cluster Parameter **QuorumArbitrationTimeMax** to 60.





## INFORMATION

About **QuorumArbitrationTimeMax**, see the following documentation for more details.

- [QuorumArbitrationTimeMax](#)

### 5.2.4. Not Supported for WSFC Setups

The following environments and functions are not supported for WSFC setups with vSphere 7. Some of which are related to arrays.

- Using VMDKs on NFS datastore as a shared disk resource for WSFC.
- Increasing the size of a shared disk.

The second point means that the datastore cannot be expanded online. Therefore users must plan enough capacity for the application.

## 5.3. Test Results

This section provides performance results for building a WSFC using traditional RDMs and clustered VMDKs. We verified the random IOPS and throughput through the IOmeter benchmark too. This is the environment and results.

### 5.3.1. Environment and Topology

- Storage
  - Model: XS5324D  
Memory: 8 GB per controller  
Firmware: XEVO 2.1.3  
Host Card: 2 x 4-port 16Gb FC (SFP+)

- SAS HDD: 3 x Seagate NL-SAS HDD 12.0 Gb/s 1TB
- Pools: 2 x (3 x HDDs per Pool for RAID 5)
- Volumes: 1 x 100GB in Pool 1 (Ctrl 1)
- Volume Block Size: 512 Bytes
- Server
  - Model: 2 x ASUS RS700
  - 16Gb FC HBA: Marvell QLogic QLE2672
  - OS: VMware ESXi 7.0U2
  - Software
  - VMware vSphere 7.0U2
  - VM OS: 2 x Windows Server 2016 with Failover Cluster feature
  - IOmeter
  - Version: 1.1.0
  - Workers: 1
  - Queue Depth: 128

### 5.3.2. RDM vs. Clustered VMDK Performance Results

The following are the performance results of random 4K and sequential 32K.

Table 5-2 Performance Results of RDM and Clustered VMDK

	RDM	CLUSTERED VMDK
<b>Random 4K</b>	Read IOPS: 117K Write IOPS: 26.7K	Read IOPS: 125K Write IOPS: 22.6K
<b>Sequential 32K</b>	Read Throughput: 838 MB/s Write Throughput: 749 MB/s	Read Throughput: 1,557 MB/s Write Throughput: 1,307 MB/s

The random 4K in the table doesn't make much difference, but in terms of throughput, using clustered VMDK is better than RDM.

## 5.4. Conclusion

In ESXi 7.0, the WSFC configuration supports clustered VMDK on VMFS datastore. It simplifies the process of VM application environment. QSAN storage follows in VMware's footsteps, supports clustered VMDKs, and is well tested. In terms of performance results, using clustered VMDK also has better throughput behavior. QSAN storage is an ideal solution for building Windows cluster architecture. Although there are some limitations for using VMDK, users can avoid them when planning.

## 5.5. Appendix

### Apply To

- XEVO firmware 2.1.3 and above.

### Reference

Setup WSFC Document

- [Setup for Windows Server Failover Clustering](#)

About VMware RDM

- [About Raw Device Mapping](#)
- [Differences between Virtual and Physical RDM](#)

About Clustered VMDK

- [Setup Windows Server Failover Cluster with Clustered VMDKs on vSphere 7 with Hitachi VSP Series](#)
- [VMware Clustered VMDK, SCSI3-PR and WEAR](#)

## 6. DR SOLUTION FOR VMWARE

In this chapter, we provide detailed operations for configuring the DR solution in the VMware environment, and ensure that the replicated data is consistent with the special script implemented in the ESXi server. The procedure is as follows.

1. The prerequisite is to set up an ESXi server.
2. Configure a remote replication task to backup VM files.
3. Create a script in the ESXi server to rotate the snapshots.
4. Roll back replication task for disaster drills.

### 6.1. Configure DR Solution

#### 6.1.1. Setup ESXi server

The environment prepared here is an ESXi 6.5 server, installed with a 10 GbE HBA card, directly connected to QSAN storage, and ensure that the ESXi server is managed by vCenter.




Figure 6-1 ESXi Server Architecture

## 6.1.2. Configure Remote Replication

To configure a remote replication task, you need to set up two QSAN storage systems, and the available space of the target unit must be greater or equal to the source unit. Otherwise, the snapshot replica function may fail due to insufficient storage space. Although the setting method is different, the following sections describe the configuration separately.

### XEVO Configuration

Prepare two XCubeSAN or XCubeFAS storage systems named SAN-a and SAN-b. The following is the procedure.

1. Connect one of 10 GbE ports from SAN -a to SAN -b.
2. In SAN-a, create a pool and a volume. And then set the snapshot space to make the snapshot replica function work normally.
3. In SAN-a, mount the created volume to the prepared ESXi server.
4. Create a VM (Virtual Machine) based on the mounted Datastore in the ESXi server.
5. In SAN-b, repeat steps 2 above to create the same or larger volume size as SAN-a. You may also need to set up snapshot space. Or you may skip this step if you use auto replication to configure the remote replication task.
6. In SAN-a, select the **Protection** tab to create a remote replication task to the replica volume in SAN-b.

#### Protection Volumes

The screenshot shows the 'Protection Volumes' interface with the 'Replication Tasks' tab selected. It displays a table with one replication task and its associated snapshot space details.

Volume Name	The Last Task	Capacity	Target Name	Target LUN	Created	Completed	Speed	Status
Volume_01	QREP163433	100GB	iqn.2004-08.com.qsan.dev0.ctr1	0	Thu Jul 16 17:19:21 2020	<div style="width: 100%;"></div>	20 MB/s	Replicating

Below the table, the 'Provisioned Snapshot Space' is shown as 100GB, with 873 MB/10.00 GB used.

Figure 6-2 Configure a Replication Task

7. Open the console of the VM in the ESXi server, and periodically put some files (such as robocopy utility) to continuously increase the data.



## INFORMATION

Robocopy, for "Robust File Copy", is a command-line directory and/or file replication command for Microsoft Windows. Please see [Robocopy in Wikipedia](#).

8. Create scheduled snapshots in this VM from the vCenter UI, in this example, we take 5 snapshots.




Figure 6-3 Create a Scheduled Snapshot in the VM

9. The preparation work is over here.

## QSM Configuration

Prepare two XCubeNXT or XCubeNAS storage systems named NAS-a and NAS-b. The following is the procedure.

1. Connect one of 10 GbE ports from NAS-a to NAS-b.
2. In NAS-a, create a volume and a shared folder.
3. Access NFS shared folders to assign RW permissions to all connected hosts.
4. In NAS-b, create a volume the same size or larger as the volume in NAS-a.

5. In NAS-a, mount the created shared folder to the prepared ESXi server.
6. Create a VM (Virtual Machine) based on the mounted Datastore in the ESXi server.
7. In NAS-a, select the **Backup Manager** function submenu to create a snapshot replica task to the volume in NAS-b.
8. Open the console of the VM in the ESXi server, and periodically put some files (such as robocopy utility) to continuously increase the data.



## INFORMATION

Robocopy, for "Robust File Copy", is a command-line directory and/or file replication command for Microsoft Windows. Please see [Robocopy in Wikipedia](#).

9. Create scheduled snapshots in this VM from the vCenter UI, in this example, we take 5 snapshots

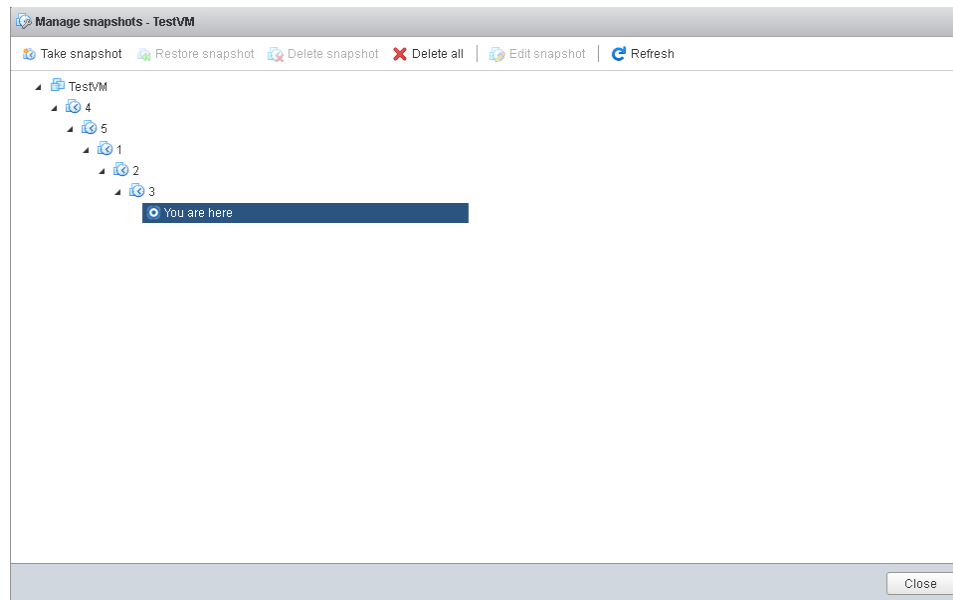


Figure 2-4 Create a Scheduled Snapshot in the VM

10. The preparation work is over here.

### 6.1.3. Create a Script in ESXi server

According to the above operations, we first take a snapshot in the VM from the ESXi server itself, and then replicate the .VMDK file along with the taken snapshots to the remote site. After mounting the volume at the remote site, registering and rolling back the snapshot taken, everything will be consistent with this method.

However, VMware does not automatically delete or rotate snapshots, so it retains a large number of snapshot images, which can cause poor performance for a long time. The script we provide here is to specify a fixed quantity of snapshots. ESXi servers can maintain rotation to prevent too many snapshots from affecting virtual machine performance. Take SAN-a as an example below. FAS and NAS are the same.

1. Create a "Crontabs" folder in the Datastore mounted from SAN-a.
2. Upload the following script "SnapshotAutoDelete.sh" to the "Crontabs" folder.

```
# cat SnapshotAutoDelete.sh

#!/bin/sh

LOG_PATH="/var/log/Schedule_Snapshot.log"
[ -f "$LOG_PATH" ] && rm $LOG_PATH;

QTY=2 # Reserved quantity
for i in `vim-cmd vmsvc/getallvms 2>/dev/null | awk '{print $1}' | grep -e "[0-9]"`
# Grab all Vmid on esxi
do
    SNAPSHOT_COUNT=`vim-cmd vmsvc/snapshot.get $i | egrep -- '--\|-CHILD|^|\-ROOT'
| wc -l`
    GuestName=$(vim-cmd vmsvc/get.summary $i | grep name | awk '{ print $3 }' | cut
-d \ " -f 2)
    if [ $SNAPSHOT_COUNT -gt $QTY ]; then # If the number of snapshots is greater
than the number of reservations
        DELETE_COUNT=$(( $SNAPSHOT_COUNT - $QTY ))
        OLD_SNAPSHOT_ID=`vim-cmd vmsvc/snapshot.get $i | grep Id | head -
$DELETE_COUNT | awk -F: '{print $2}'`
        for n in $OLD_SNAPSHOT_ID
        do
            vim-cmd vmsvc/snapshot.remove $i $n; ret=$?
            sleep 30s
            if [ $ret -eq 0 ];then
                echo "$(date +%F %T)" : $GuestName snapshot $n Delete
Success.." >> $LOG_PATH # Output to log path after deletion
            else
                echo "$(date +%F %T)" : $GuestName snapshot $n Delete
FAILED.." >> $LOG_PATH
            fi
        done
    else
        echo "$(date +%F %T)" : $GuestName snapshot not found." >> $LOG_PATH
    fi
done
```



3. Change the permission of the script to 777, from the SSH session of ESXi server.

```
[root@local:~] cd vmfs/volumes/SANI/Crontabs/
[root@local:~/vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680/Crontabs] chmod 777 SnapshotAutoDelete.sh
[root@local:~/vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680/Crontabs] ls -al
total 1152
drwxr-xr-x 1 root root 73728 Aug 2 16:38 .
drwxr-xr-t 1 root root 73728 Aug 2 16:38 ..
-rwxrwxrwx 1 root root 1088 Aug 2 18:52 SnapshotAutoDelete.sh
[root@local:~/vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680/Crontabs]
```

4. Locate the Datastore via the following command in the SSH session.

```
# esxcli storage filesystem list
```

```
[root@local:~] esxcli storage filesystem list
Mount Point                               Volume Name  UUID                               Mounted  Type      Size      Free
-----
/vmfs/volumes/5bc3fd0f-f996289d-ba94-001018edee60  datastore1  5bc3fd0f-f996289d-ba94-001018edee60  true    VMFS-6   492042190848  442177159168
/vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680  SAN1        5d445d0a-fae8654e-a676-001b21d4d680  true    VMFS-6   166792838144  88226136664
/vmfs/volumes/5ceb8d20-96976e3b-25ef-08606e151c65  5ceb8d20-96976e3b-25ef-08606e151c65  true    vfat     299712512     80486400
/vmfs/volumes/9bfaa77a-a157614d-7923-8cc7a16bcdea  9bfaa77a-a157614d-7923-8cc7a16bcdea  true    vfat     261853184     261844992
/vmfs/volumes/3d40c777-b5b2f4fb-b003-5dfeca8c4b86  3d40c777-b5b2f4fb-b003-5dfeca8c4b86  true    vfat     261853184     113819648
/vmfs/volumes/5ceb8d28-4a26e650-7a8a-08606e151c65  5ceb8d28-4a26e650-7a8a-08606e151c65  true    vfat     4293591040    4264230912
[root@local:~]
```

5. Use the following command to add a cron job to execute the script at 23:30 every day. You can specify the time point according to your environment. This point in time should be earlier than the periodic snapshot task created by vCenter. Or you can edit this file directly.

```
# echo "30 23 * * * sh /vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680/Crontabs/SnapshotAutoDelete.sh" >> /var/spool/cron/crontabs/root
```



## INFORMATION

The **YELLOW** word above is the UUID of the Datastore, please check yours with the above command.

6. Since the configuration will be cleared after the ESXi server restarts, you need to add the above commands to permanently save the configuration. Edit the local cron job file (/etc/rc.local.d/local.sh) of the ESXi server and add the following commands at the end of the configuration file.

```
# vi /etc/rc.local.d/local.sh

...
/bin/echo "30 23 * * * sh /vmfs/volumes/5d445d0a-fae8654e-a676-001b21d4d680/Crontabs/SnapshotAutoDelete.sh" >>/var/spool/cron/crontabs/root
/bin/kill $(cat /var/run/crond.pid)
/usr/lib/vmware/busybox/bin/busybox crond
```

7. Check the quantity of retained snapshots from the ESXi UI and confirm that the snapshots have been retained as the latest two.

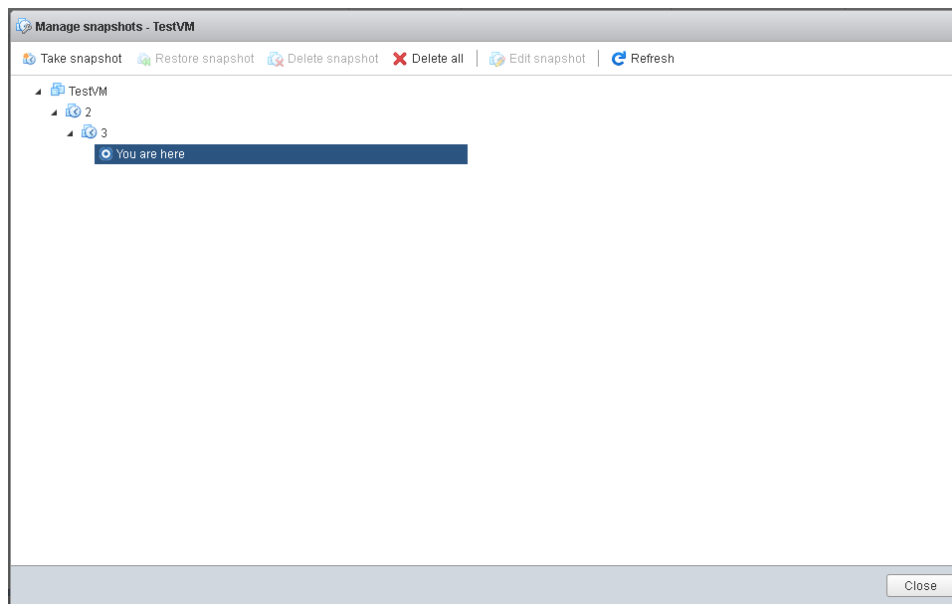


Figure 6-5 List the Snapshots in the VM

8. Use the following command to check the log.

```
# cat /var/log/Schedule_Snapshot.log
```

```
[root@local:~] cat /var/log/Schedule_Snapshot.log
2019-08-05 11:30:38 : 2012R2-SAN1 snapshot 1 Delete Success..
2019-08-05 11:31:12 : 2012R2-SAN1 snapshot 2 Delete Success..
2019-08-05 11:32:10 : 2012R2-SAN1 snapshot 3 Delete Success..
[root@local:~] █
```

9. The ESXi server configuration is complete.

## 6.1.4. Disaster Drill

We provide disaster drills to prove the effectiveness of backups. Similarly, the setting method is different; the following sections describe the configuration separately.

### XEVO Configuration

Continue the previous section, two XCubeSAN or XCubeFAS storage systems named SAN-a and SAN-b. The following is the procedure.

1. In SAN-a, select the **Protection** tab to find the remote replication task.

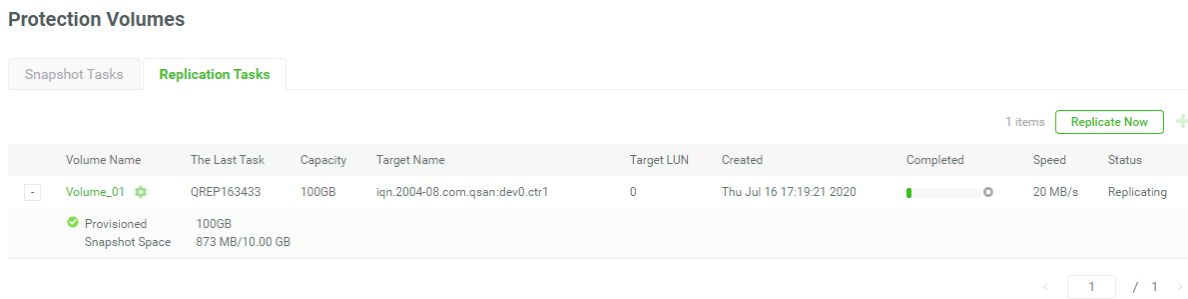


Figure 6-6 Remote Replication Task

2. You may need to unmount the original Datastore (of SAN-a) from the ESXi server to simulate a disaster on SAN-a.
3. In SAN-b, select the **Protection** tab to expose the replicated snapshot as a writable volume, and its exposed snapshot capacity is greater than 0 (GB) by default. This is called the writable snapshot function.

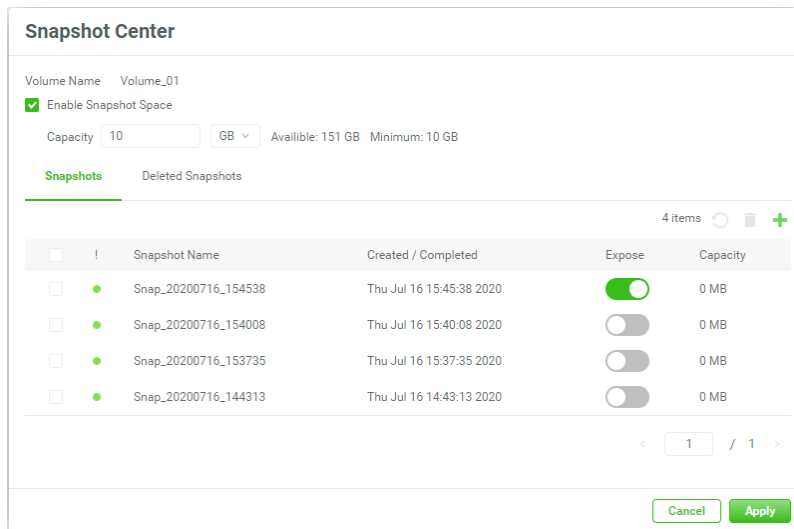


Figure 6-7 Expose the Snapshot

4. Map the volume as a LUN with read-write permission, and then access the vCenter UI (of the ESXi server) to mount the exposed snapshot volume to be a Datastore.
5. During the process of mounting the Datastore, the ESXi system will ask you to assign a New Signature or use an Existing signature. Please choose to use an Existing signature.
6. Right click on the Datastore, you will be able to see the VM replicated from SAN-a, then you can register this VM and try to boot up after the snapshot on the VM is rolled back.

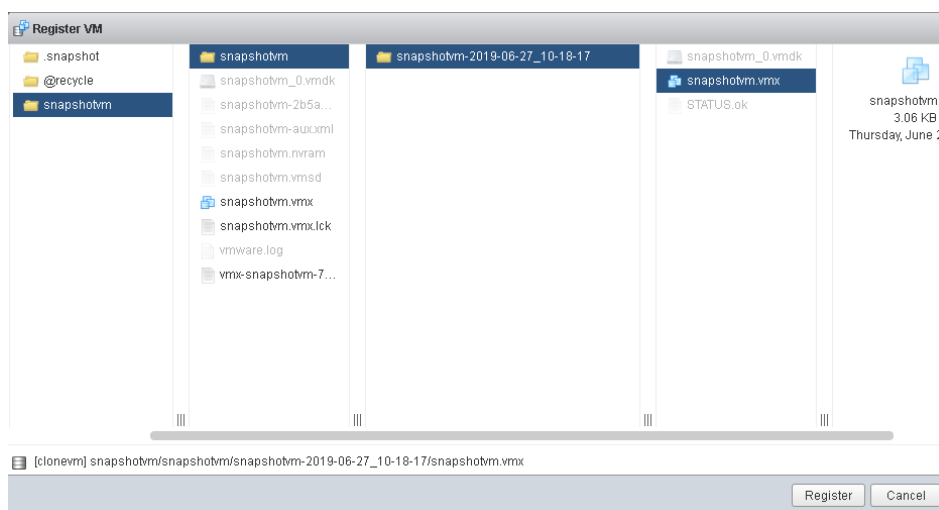


Figure 6-8 Snapshot is rolled back



## TIP

It is necessary to roll back the snapshot of the VM because the .VMDK file may be inconsistent due to the data cached by the ESXi server. Please roll back the last snapshot before powering on the VM to ensure that the VM can be successfully booted up.

7. Done.

## QSM Configuration

Continue the previous section, two XCubeNXT or XCubeNAS storage systems named NAS-a and NAS-b. The following is the procedure.

1. In NAS-a, select the **Protection** tab to find the created snapshot replica task.
2. In NAS-b, select the **Protection** tab, and clone the replicated snapshot into the volume.
3. After the clone is completed, change the permission from RO to RW in shared folder page.
4. Assign the folder with RW permission to the NFS protocol, just like we did in NAS-a.
5. Go to ESXi server, mount the NFS shared folder as a Datastore.
6. Right click on the Datastore, you will be able to see the VM replicated from NAS-a, then you can register this VM and try to boot up after the snapshot on the VM is rolled back.



## TIP

It is necessary to roll back the snapshot of the VM because the .VMDK file may be inconsistent due to the data cached by the ESXi server. Please roll back the last snapshot before powering on the VM to ensure that the VM can be successfully booted up.

7. Done.

## 6.2. Conclusion

This document discusses continuous backup solutions and disaster drills in a VMware environment. Configuring a data protection solution helps prevent unexpected situations. In addition, this is a cost-effective method and does not require any agent to be installed in the environment. The solution we provide can be easily implemented with the simple script and snapshot copies stored in QSAN storage.

## 6.3. Appendix

### Apply To

- XEVO firmware 2.0.0 and later
- QSM firmware 4.0.0 and later

### Reference

#### Documents

- [XEVO Software Manual](#)
- [QSM 4 Software Manual](#)
- [Remote Replication White Paper](#)