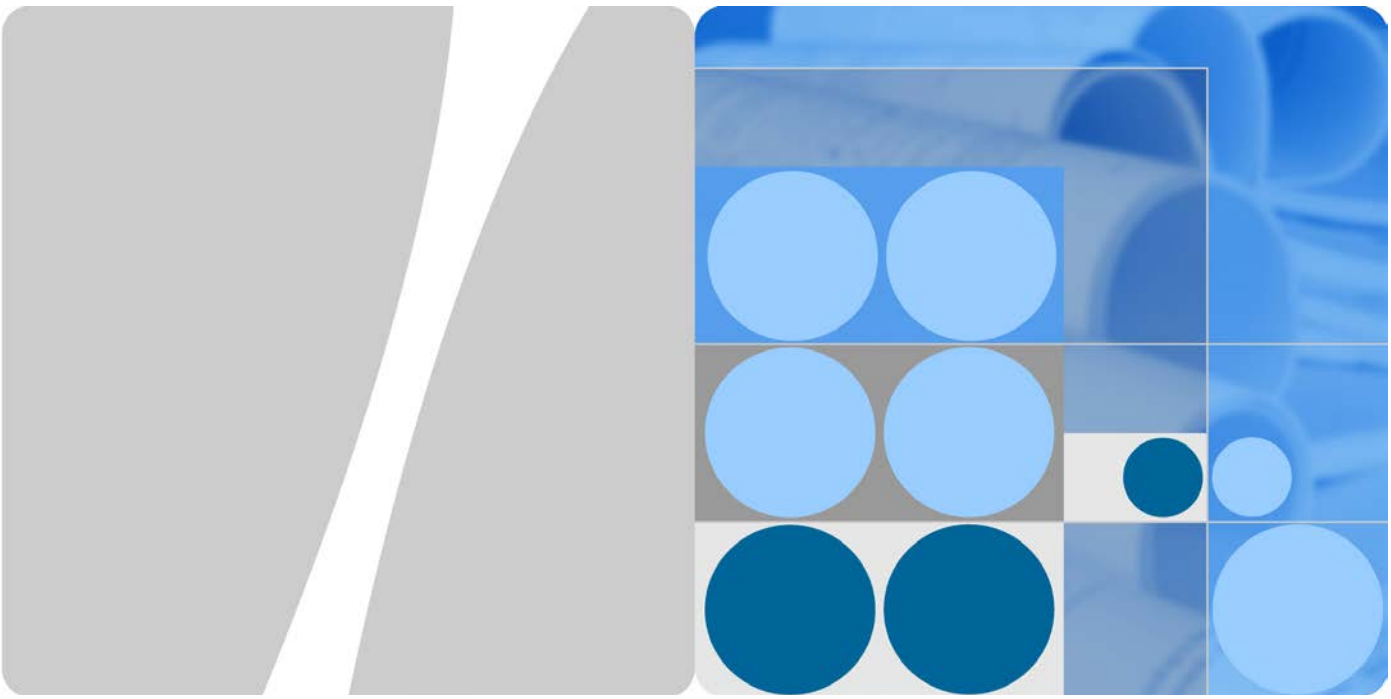

Huawei SAN Storage Host Connectivity Guide for VMware ESXi

Issue 06

Date 2018-04-10

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

Contents

1 About This Document

Purpose

This document details the configuration methods and precautions for connecting Huawei SAN storage devices to VMware ESXi hosts.

Intended Audience

This document is intended for:

- Huawei technical support engineers
- Technical engineers of Huawei's partners
- Other personnel who are involved in interconnecting Huawei SAN and VMware ESXi hosts or who are interested in the interconnection.

Readers of this guide are expected to be familiar with the following topics:

- Huawei OceanStor V3, OceanStor V5, and Dorado V3
- VMware ESXi

Related Documents

For the hosts, host bus adapters (HBAs), and operating systems that are compatible with Huawei storage devices, go to support-open.huawei.com.

For the latest Huawei storage product documentation, go to support.huawei.com.

For VMware ESXi documents or support, go to www.vmware.com/support.

Conventions

Symbol Conventions

Symbol	Description
 DANGER	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.

Symbol	Description
	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
	Indicates a potentially hazardous situation which, if not avoided, may result in minor or moderate injury.
	Indicates a potentially hazardous situation which, if not avoided, could result in equipment damage, data loss, performance deterioration, or unanticipated results. NOTICE is used to address practices not related to personal injury.
	Calls attention to important information, best practices and tips. NOTE is used to address information not related to personal injury, equipment damage, and environment deterioration.

General Conventions

Convention	Description
Times New Roman	Normal paragraphs are in Times New Roman.
Boldface	Names of files, directories, folders, and users are in **boldface**. For example, log in as user **root**.
Italic	Book titles are in *italics*.
Courier New	Examples of information displayed on the screen are in Courier New.

Command Conventions

Format	Description
Boldface	The keywords of a command line are in **boldface**.
Italic	Command arguments are in *italics*.

Where To Get Help

Huawei support and product information can be obtained on the Huawei Online Support site.

Product Information

For documentation, release notes, software updates, and other information about Huawei products and support, go to the Huawei Online Support site (registration required) at

http://support.huawei.com/enterprise/.

Technical Support

Huawei has a global technical support system, able to offer timely onsite and remote technical support service.

For any assistance, contact:

- Your local technical support

  http://e.huawei.com/en/branch-office-query
- Huawei company headquarters.

  Huawei Technologies Co., Ltd.

  Address: Huawei Industrial Base Bantian, Longgang Shenzhen 518129 People's Republic of China

  Website: http://enterprise.huawei.com/

Document Feedback

Huawei welcomes your suggestions for improving our documentation. If you have comments, send your feedback to mailto:infoit@huawei.com.

Change History

Changes between document issues are cumulative. The latest document issue contains all the changes made in earlier issues.

- Issue 01 (2015-12-30)

  This issue is the first official release.
- Issue 02 (2016-03-30)

  This issue is the second official release and includes the following changes:

  Added the product models and versions.
- Issue 03 (2016-11-30)

  This issue is the third official release and includes the following changes:

  Changed the document template.

  Added the description about SAN Boot.
- Issue 04 (2017-03-30)

  This issue is the fourth official release and includes the following changes:

  Added the description that the host must be restarted after rules have been configured for VMware to chapter 7.
- Issue 05 (2018-01-30)

  This issue is the fifth official release and includes the following changes:

  Added the Fibre Channel networking diagram and description.
- Issue 06 (2018-04-10)

This issue is the sixth official release and includes the following changes:

Added the description of the initiator parameters to chapter 7.

Added the HyperMetro working modes to chapter 7.

2 Introduction

2.1 Introduction to VMware ESXi

2.1.1 VMware Infrastructure

Legacy x86 computers are designed merely for running a single operating system or application program. Therefore, most of these computers are under-utilized. To address the under-utilization, virtualization technologies are adopted, enabling an x86 physical machine to host multiple virtual machines (VMs) and allowing the multiple VMs to run different operating systems and applications. That is, virtualization allows resources on this physical machine to be shared among multiple environments, thereby improving x86 hardware utilization.

VMware virtualization technology (VMware Infrastructure) adds a condensed software layer on the computer hardware or in the host operating system. This software layer includes a VM monitor utility that allocates hardware resources in a dynamic and transparent way. Each operating system or application can access desired resources anytime as required.

As an outstanding software solution for x86 virtualization, VMware Infrastructure enables users to manage their virtualized environments in an effective and easy manner. As Figure 2-1 shows, a typical VMware Infrastructure datacenter consists of basic physical building blocks such as x86 computing servers, storage networks and arrays, IP networks, a management server, and desktop clients.

Figure 2-1 VMware Infrastructure datacenter

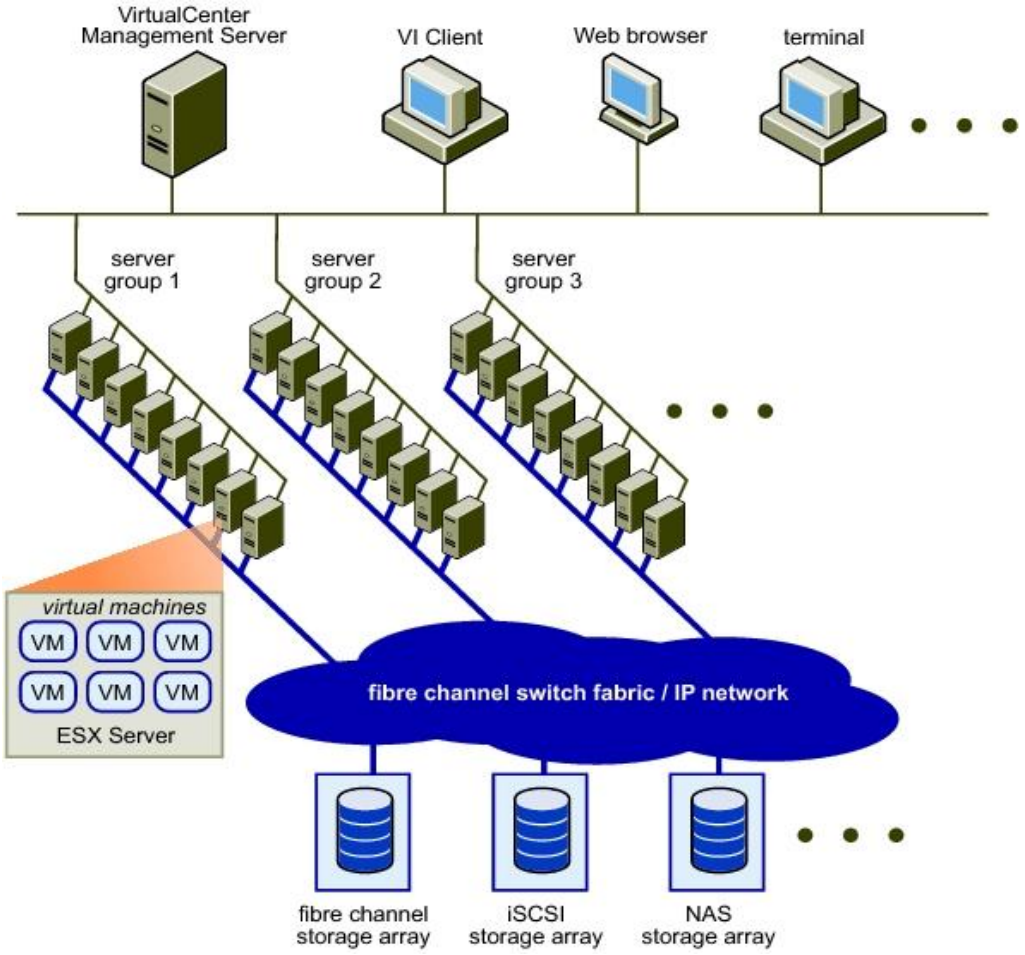

In a VMware Infrastructure storage architecture (as shown in Figure 2-2):

- A Virtual Machine File System (VMFS) volume contains one or more LUNs which belong to one or more storage arrays.
- Multiple ESX servers share one VMFS volume; virtual disks are created on the VMFS volume for use by VMs.

Figure 2-2 Storage architecture in VMware Infrastructure

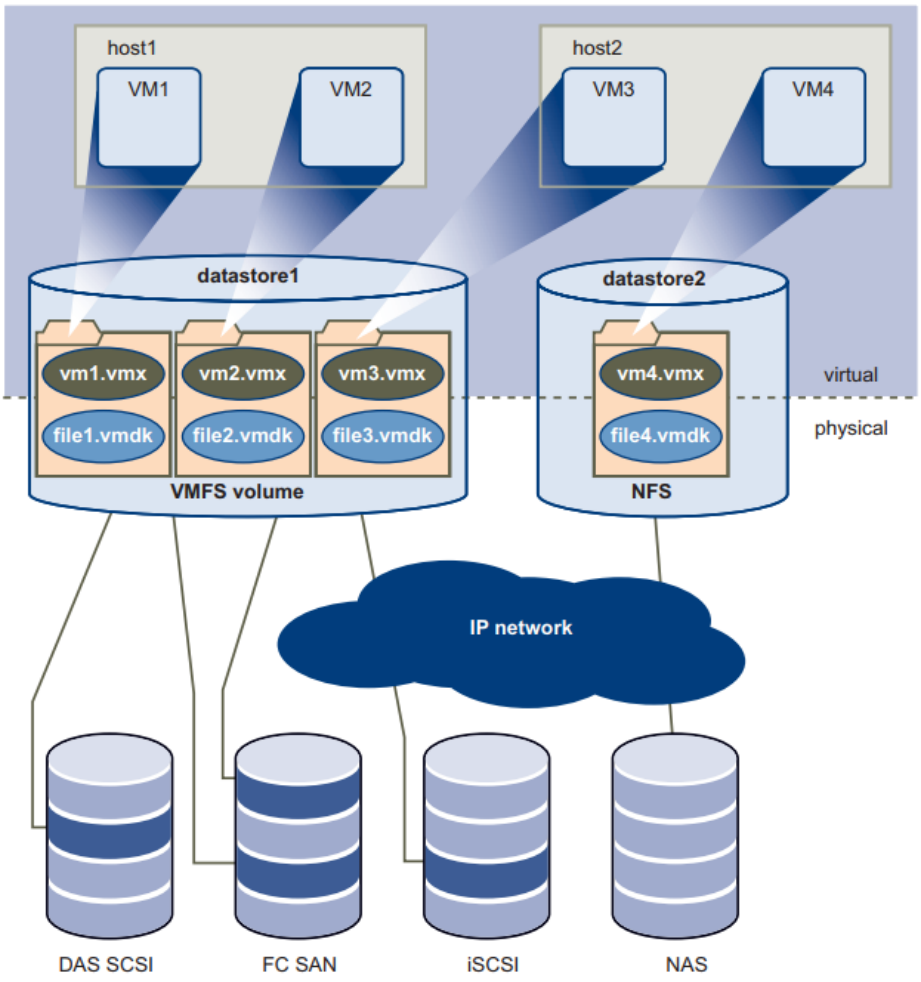

VMware uses VMFS to centrally manage storage systems. VMFS is a shared cluster file system designed for VMs. This file system employs a distributed locking function to enable independent access to disks, ensuring that a VM is accessed by one physical host at a time. Raw Device Mapping (RDM) is a special file acting as the agent for raw devices on a VMFS volume.

2.1.2 File Systems in VMware

VMFS is a clustered file system that leverages shared storage to allow multiple physical hosts to read and write to the same storage simultaneously. It lays a solid foundation for the management of VMware clusters and dynamic resources.

Features of VMFS

- Automated maintenance directory structure
- File locking mechanism
- Distributed logical volume management
- Dynamic capacity expansion
- Clustered file system

- Journal logging
- Optimized VMs' file storage

Advantages of VMFS

- Improved storage utilization
- Simplified storage management
- ESX server clusters of enhanced performance and reliability

Architecture of VMFS

In the VMFS architecture shown in Figure 2-3, a LUN is formatted into a VMFS file system, whose storage space is shared by three ESX servers each carrying two VMs. Each VM has a Virtual Machine Disk (VMDK) file that is stored in a directory (named after a VM) automatically generated by VMFS. VMFS adds a lock for each VMDK to prevent a VMDK from being accessed by two VMs at the same time.

Figure 2-3 VMFS architecture

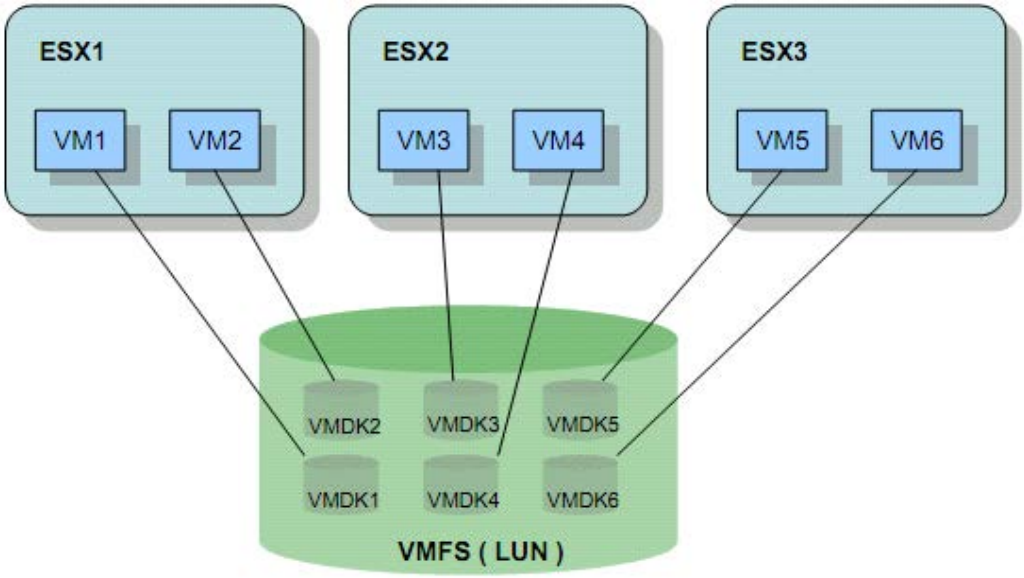

Structure of a VMFS Volume

Figure 2-4 shows the structure of a VMFS volume. A VMFS volume consists of one or more partitions that are arranged in lines. Only after the first partition is used out can the following partitions be used. The identity information about the VMFS volume is recorded in the first partition.

Figure 2-4 Structure of a VMFS volume

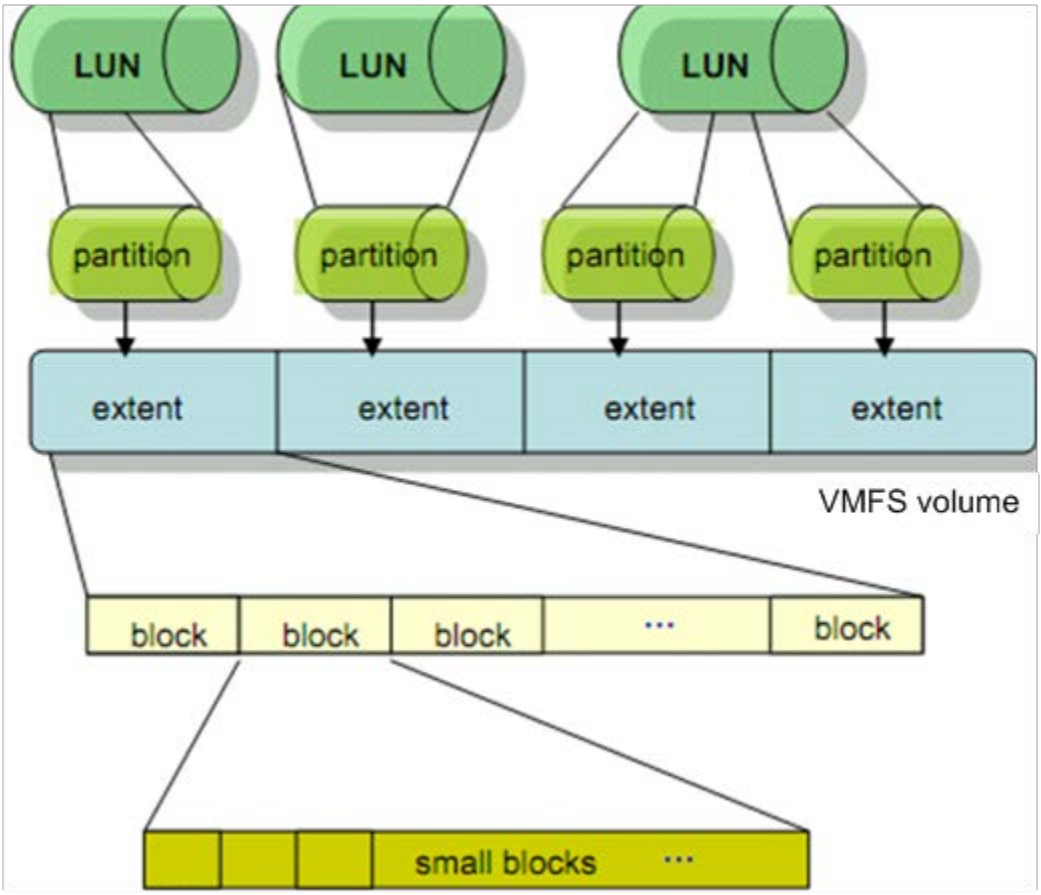

VMFS divides each extent into multiple blocks, each of which is then divided into smaller blocks. This block-based management is typically suitable for VMs. Files stored on VMs can be categorized as large files (such as VMDK files, snapshots, and memory swap files) and small files (such as log files, configuration files, and VM BIOS files). Large and small blocks are allocated to large and small files respectively. In this way, storage space is effectively utilized and the number of fragments in the file system is minimized, improving the storage performance of VMs.

The VMFS-3 file system supports four data block sizes: 1 MB, 2 MB, 4 MB, and 8 MB. Sizes of files and volumes supported by VMFS-3 file systems vary with a file system's block size.

The VMFS-5 file system uses a fixed data block size of 1 MB. Supporting VMFS-5, VMware ESXi 5.0/5.1 supports a maximum VMDK file size of 2 TB, and VMware ESXi 5.5/6.0/6.5 supports a maximum VMDK file size of 62 TB.

The VMFS-6 file system also uses a fixed data block size of 1 MB. Supporting VMFS-6, VMware ESXi 6.5 supports a maximum VMDK file size of 62 TB.

2.1.3 VMware RDM

VMware RDM enables VMs to directly access storage. As shown in Figure below, an RDM disk exists as an address mapping file on the VMFS volume. This mapping file can be considered as a symbolic link that maps a VM's access to an RDM disk to LUNs.

Figure 2-5 RDM mechanism

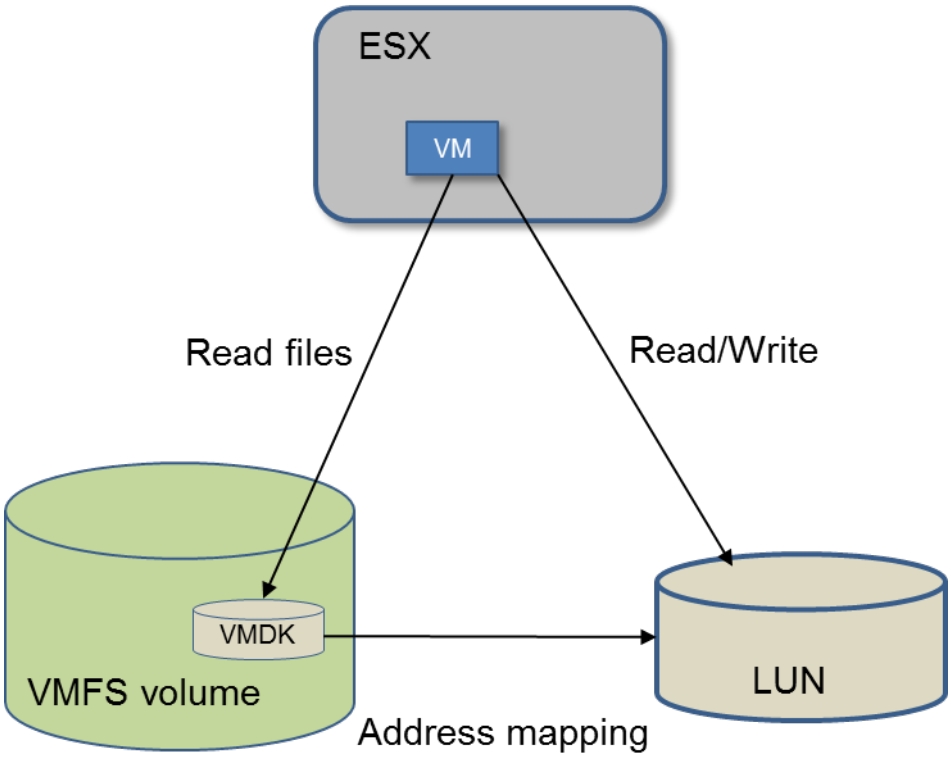

RDM provides two compatible modes, both of which supports vMotion, Distributed Resource Scheduler (DRS), and High Availability (HA)

- Virtual compatibility: fully simulates VMDK files and supports snapshots.
- Physical compatibility: directly accesses SCSI devices and does not support snapshots.

RDMs are applicable in the following scenarios:

- Physical to Virtual (P2V): migrates services from a physical machine to a virtual machine.
- Virtual to Physical (V2P): migrates services from a virtual machine to a physical machine.
- Clustering physical machines and virtual machines.

2.2 Host-SAN Connectivity

2.2.1 FC Connectivity

A Fibre Channel (FC) SAN is a specialized high-speed network that connects host servers to storage systems. The FC SAN components include HBAs in the host servers, switches that help route storage traffic, cables, storage processors (SPs), and storage disk arrays.

To transfer traffic from host servers to shared storage, the FC SAN uses the Fibre Channel protocol that packages SCSI commands into Fibre Channel frames.

- Ports in FC SAN

Each node in the SAN, such as a host, a storage device, or a fabric component has one or more ports that connect it to the SAN. Ports are identified in a number of ways, such as by:

- World Wide Port Name (WWPN)

  A globally unique identifier for a port that allows certain applications to access the port. The FC switches discover the WWPN of a device or host and assign a port address to the device.

- Port_ID (or port address)

  Within a SAN, each port has a unique port ID that serves as the FC address for the port. This unique ID enables routing of data through the SAN to that port. The FC switches assign the port ID when the device logs in to the fabric. The port ID is valid only when the device is logged on.

- Zoning

  Zoning provides access control in the SAN topology. Zoning defines which HBAs can connect to which targets. When you configure a SAN by using zoning, the devices outside a zone are not visible to the devices inside the zone.

  Zoning has the following effects:

  - Reduces the number of targets and LUNs presented to a host.
  - Controls and isolates paths in a fabric.
  - Prevents non-ESXi systems from accessing a specified storage system and protects VMFS data against damage.
  - Separates different environments, for example, a test from a production environment.

VMware ESXi hosts support zone division on a per-initiator basis or on a per-initiator and per-target basis. Per-initiator per-target zone division is preferred, since this stricter zone division has more capabilities of preventing SAN faults or configuration errors.

2.2.2 iSCSI Connectivity

In computing, Internet Small Computer Systems Interface (iSCSI) is an IP-based storage networking standard for linking data storage systems.

By carrying SCSI commands over IP networks, iSCSI is used to access remote block devices in the SAN, providing hosts with the illusion of locally attached devices.

A single discoverable entity on the iSCSI SAN, such as an initiator or a target, represents an iSCSI node.

Each iSCSI node can be identified in a number of ways, such as by:

- IP address

  Each iSCSI node can have an IP address associated with it so that routing and switching equipment on your network can establish the connection between the server and storage. This address is just like the IP address that you assign to your computer to get access to your company's network or the Internet.

- iSCSI name

  A worldwide unique name for identifying the node. iSCSI uses the iSCSI Qualified Name (IQN) and Extended Unique Identifier (EUI).

  By default, VMware ESXi generates unique iSCSI names for your iSCSI initiators, for example, iqn.1998-01.com.vmware:iscsitestox-68158ef2. Usually, you do not have to

change the default value, but if you do, make sure that the new iSCSI name you enter is worldwide unique.

2.2.3 Multipath Connectivity

2.2.3.1 UltraPath

UltraPath is a Huawei-developed multipathing software. It can manage and process disk creation/deletion and I/O delivery of operating systems.

UltraPath provides the following functions:

- Masking of redundant LUNs

  In a redundant storage network, an application server with no multipathing software detects a LUN on each path. Therefore, a LUN mapped through multiple paths is mistaken for two or more different LUNs. UltraPath installed on the application server masks redundant LUNs on the operating system driver layer to provide the application server with only one available LUN, the virtual LUN. In this case, the application server only needs to deliver data read and write operations to UltraPath that masks the redundant LUNs, and properly writes data into LUNs without damaging other data.
- Optimum path selection

  In a multipath environment, the owning controller of the LUN on the storage system mapped to an application server is the prior controller. With UltraPath, an application server accesses the LUN on the storage system through the prior controller, thereby obtaining the highest I/O speed. The path to the prior controller is the optimum path.
- Failover and failback
  - Failover

    When a path fails, UltraPath fails over its services to another functional path.
  - Failback

    UltraPath automatically delivers I/Os to the first path again after the path recovers from the fault. There are two methods to recover a path:
- I/O Load balancing

  UltraPath provides load balancing within a controller and across controllers.
  - For load balancing within a controller, I/Os poll among all the paths of the controller.
  - For load balancing across controllers, I/Os poll among the paths of all these controllers.
- Path test

  UltraPath tests the following paths:
  - Faulty paths

    UltraPath tests faulty paths with a high frequency to detect the path recover as soon as possible.
  - Idle paths

    UltraPath tests idle paths to identify faulty paths in advance, preventing unnecessary I/O retires. The test frequency is kept low to minimize impact on service I/Os.

2.2.3.2 VMware NMP

Overview

VMware ESXi has its own multipathing software Native Multipath Module (NMP), which is available without the need for extra configurations.

This section details the NMP multipathing software.

VMware PSA

Overview

VMware ESXi 4.0 incorporates a new module Pluggable Storage Architecture (PSA) that can be integrated with third-party Multipathing Plugin (MPP) or NMP to provide storage-specific plug-ins such as Storage Array Type Plug-in (SATP) and Path Selection Plugin (PSP), thereby enabling the optimal path selection and I/O performance.

Figure 2-6 VMware pluggable storage architecture

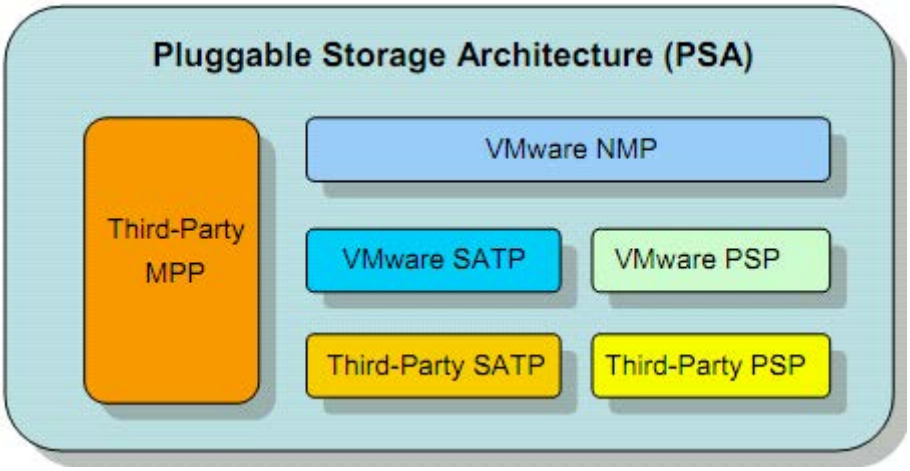

VMware NMP

NMP is the default multipathing module of VMware. This module provides two submodules to implement failover and load balancing.

- SATP: monitors path availability, reports path status to NMP, and implements failover.
- PSP: selects optimal I/O paths.

PSA is compatible with the following third-party multipathing plugins:

- Third-party SATP: Storage vendors can use the VMware API to customize SATPs for their storage features and optimize VMware path selection.
- Third-party PSP: Storage vendors or third-party software vendors can use the VMware API to develop more sophisticated I/O load balancing algorithms and achieve larger throughput from multiple paths.

VMware Path Selection Policy

- Built-in PSP

By default, the PSP of VMware ESXi 5.0 or later supports three I/O policies: Most Recently Use (MRU), Round Robin, and Fixed. VMware ESXi 4.1 supports an additional policy: Fixed AP.

- Third-Party software

Third-party MPP supports comprehensive fault tolerance and performance processing, and runs on the same layer as NMP. For some storage systems, Third-Party MPP can substitute NMP to implement path failover and load balancing.

Functions and Features

To manage storage multipathing, ESX/ESXi uses a special VMkernel layer, Pluggable Storage Architecture (PSA). The PSA is an open modular framework that coordinates the simultaneous operations of MPPs.

The VMkernel multipathing plugin that ESX/ESXi provides, by default, is VMware NMP. NMP is an extensible module that manages subplugins. There are two types of NMP plugins: SATPs and PSPs. Figure 2-7 shows the architecture of VMkernel.

Figure 2-7 VMkernel architecture

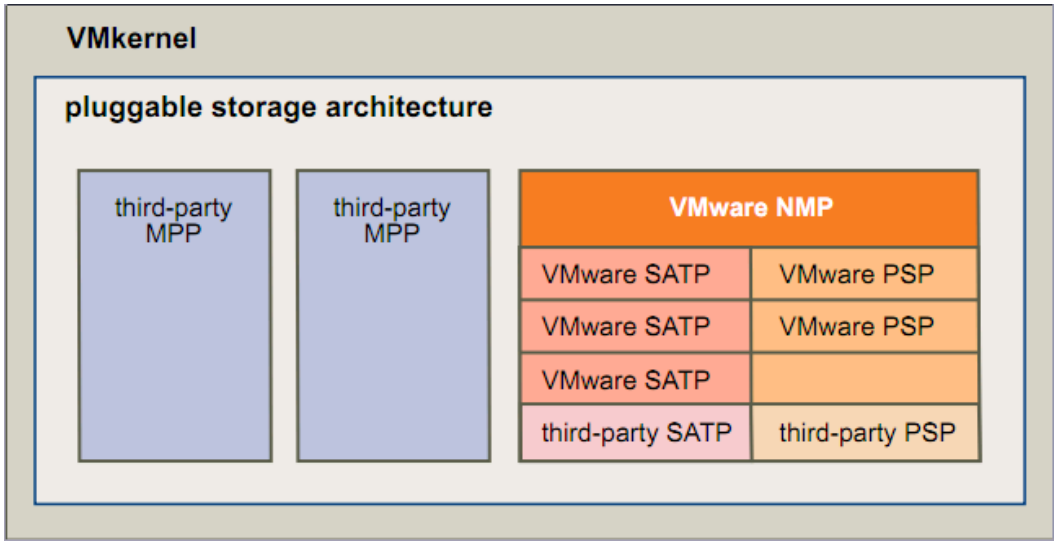

If more multipathing functionality is required, a third party can also provide an MPP to run in addition to, or as a replacement for, the default NMP. When coordinating with the VMware NMP and any installed third-party MPPs, PSA performs the following tasks:

- Loads and unloads multipathing plug-ins.
- Hides virtual machine specifics from a particular plug-in.
- Routes I/O requests for a specific logical device to the MPP managing that device.
- Handles I/O queuing to the logical devices.
- Implements logical device bandwidth sharing between virtual machines.
- Handles I/O queuing to the physical storage HBAs.
- Handles physical path discovery and removal.
- Provides logical device and physical path I/O statistics.

VMware NMP Path Selection Policy

VMware supports the following path selection policies, as described in Table 2-1.

Table 2-1 Path selection policies

Policy/Controller	Active/Active	Active/Passive
Most Recently Used	Administrator action is required to fail back after path failure.	Administrator action is required to fail back after path failure.
Fixed	VMkernel resumes using the preferred path when connectivity is restored.	VMkernel attempts to resume using the preferred path. This can cause path thrashing or failure when another SP now owns the LUN.
Round Robin	The host uses automatic path selection algorithm to ensure that I/Os are delivered to all active paths in turn. It will not switch back even after the faulty path recovers.	The host uses automatic path selection algorithm to always select the next path in the RR scheduling queue, therefore ensuring that I/Os are delivered to all active paths in turn.
Fixed AP	For ALUA arrays, VMkernel picks the path set to be the preferred path. For both A/A, A/P, and ALUA arrays, VMkernel resumes using the preferred path, but only if the path-thrashing avoidance algorithm allows the failback. Fixed AP is available only in VMware ESX/ESXi 4.1.	

The following details each policy.

- **Most Recently Used (VMW_PSP_MRU)**

The host selects the path that is used recently. When the path becomes unavailable, the host selects an alternative path. The host does not revert to the original path when the path becomes available again. There is no preferred path setting with the MRU policy. MRU is the default policy for active-passive storage devices.

Working principle: uses the most recently used path for I/O transfer. When the path fails, I/O is automatically switched to the last used path among the multiple available paths (if any). When the failed path recovers, I/O is not switched back to that path.

- **Round Robin (VMW_PSP_RR)**

The host uses an automatic path selection algorithm rotating through all available active paths to enable load balancing across the paths. Load balancing is a process to distribute host I/Os on all available paths. The purpose of load balancing is to achieve the optimal throughput performance (IPOS, MB/s, and response time).

Working principle: uses all available paths for I/O transfer.

- **Fixed (VMW_PSP_FIXED)**

The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. The default policy for active-active storage devices is Fixed. After the preferred path recovers from fault, VMkernel continues to use the preferred path. This attempt may results in path thrashing or failure because another SP now owns the LUN.

Working principle: uses the fixed path for I/O transfer. When the current path fails, I/O is automatically switched to a random path among the multiple available paths (if any). When the original path recovers, I/O will be switched back to the original path.

- **Fixed AP (VMW_PSP_FIXED_AP)**

This policy is only supported by VMware ESX/ESXi 4.1.x and is incorporated to VMW_PSP_FIXED in later ESX versions.

Fixed AP extends the Fixed functionality to active-passive and ALUA mode arrays.

2.2.3.3 ALUA

- ALUA definition:

  Asymmetric Logical Unit Access (ALUA) is a multi-target port access model. In a multipathing state, the ALUA model provides a way of presenting active/passive LUNs to a host and offers a port status switching interface to switch over the working controller. For example, when a host multipathing program that supports ALUA detects a port status change (the port becomes unavailable) on a faulty controller, the program will automatically switch subsequent I/Os to the other controller.

- Support by Huawei storage:

  Old-version Huawei storage supports ALUA only in two-controller configuration, but not in multi-controller or HyperMetro configuration.

  New-version Huawei storage supports ALUA in two-controller, multi-controller, and HyperMetro configurations.

  Table 2-2 defines old- and new-version Huawei storage.

Table 2-2 Old- and new-version Huawei storage

Storage Type	Version	Remarks
Old-version Huawei storage (namely, storage that does not support multi-controller ALUA or ALUA HyperMetro)	T V1/T V2/18000 V1/V300R001/V300R002/V300R003C00/V300R003C10/V300R005/Dorado V300R001C00	-
New-version Huawei storage (namely, storage that supports multi-controller ALUA and ALUA HyperMetro)	V500R007C00 and later versions V300R003C20/V300R006C00/Dorado V300R001C01	V300R003C20: refers to only V300R003C20SPC200 and later versions. V300R006C00: refers to only V300R006C00SPC100 and later versions. Dorado V300R001C01: refers to only V300R001C01SPC100 and

Storage Type	Version	Remarks
		later versions.

- ALUA impacts

  ALUA is mainly applicable to a storage system that has only one prior LUN controller. All host I/Os can be routed through different controllers to the working controller for execution. ALUA will instruct the hosts to deliver I/Os preferentially from the LUN working controller, thereby reducing the I/O routing-consumed resources on the non-working controllers.

  If all I/O paths of the LUN working controller are disconnected, the host I/Os will be delivered only from a non-working controller and then routed to the working controller for execution.

- Suggestions for using ALUA on Huawei storage

  To prevent I/Os from being delivered to a non-working controller, you are advised to ensure that:

  - LUN home/working controllers are evenly distributed on storage systems so that host service I/Os are delivered to multiple controllers for load balancing.
  - Hosts always try the best to select the optimal path to deliver I/Os even after an I/O path switchover.

2.2.4 SAN Boot

SAN Boot is a network storage management system that stores data (including servers' operating systems) totally on storage systems. Specifically, operating systems are installed on and booted from SAN storage devices. SAN Boot is also called Remote Boot or boot from SAN.

SAN Boot can help to improve system integration, enable centralized management, and facilitate recovery.

- Server integration: Blade servers are used to integrate a large number of servers within a small space. There is no need to configure local disks.
- Centralized management: Boot disks of servers are centrally managed on a storage device. All advanced management functions of the storage device can be fully utilized. For example, the snapshot function can be used for backup. Devices of the same model can be quickly deployed using the snapshot function. In addition, the remote replication function can be used for disaster recovery.
- Quick recovery: Once a server that is booted from SAN fails, its boot volume can be quickly mapped to another server, achieving quick recovery.

2.3 Interoperability Query

When connecting a storage system to a VMware ESXi host, consider the interoperability of upper-layer applications and components (such as storage systems, VMware ESXi systems, HBAs, and switches) in the environment.

You can query the latest compatibility information by performing the following steps:

Step 1 Log in to the website support-open.huawei.com.

Step 2 On the home page, choose **Interoperability Center > Storage Interoperability**.

Figure 2-8 Interoperability query page

Then, the **OceanStor Interoperability Navigator** is displayed.

Step 3 Select the components to query and click **Submit**.

Figure 2-9 Query on OceanStor Interoperability Navigator

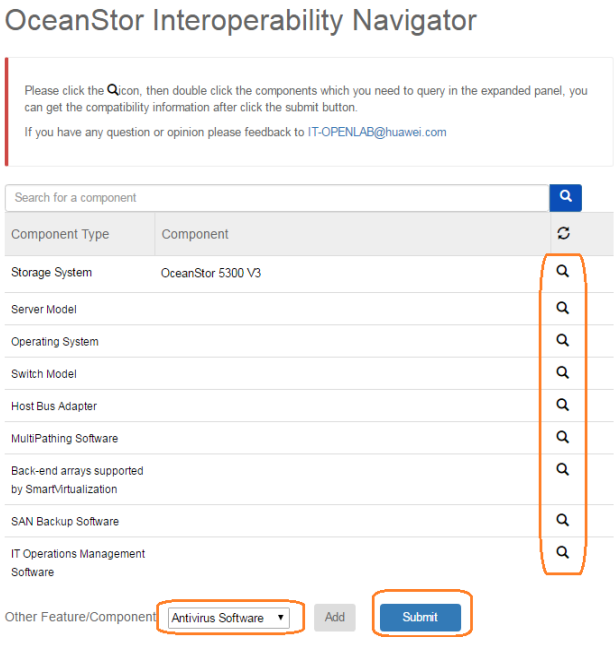

----End

2.4 Specifications

VMware specifications vary with VMware versions. Table 2-3 lists major VMware specifications.

Table 2-3 Major specifications of VMware

Category	Specifications	Max. Value						
		4.0	4.1	5.0	5.1	5.5	6.0	6.5
iSCSI Physical	LUNs per server	256[a]	256	256	256	256	256	512
	Paths to a LUN	8	8	8	8	8	8	8
	Number of total paths on a server	1024	1024	1024	1024	1024	1024	2048
Fibre Channel	LUNs per host	256[a]	256	256	256	256	256	512
	LUN size	2 TB-512B	2 TB-512B	-	64 TB	64 TB	64 TB	64 TB
	LUN ID	255	255	255	255	255	1023	16383
	Number of paths to a LUN	16	32	32	32	32	32	32
	Number of total paths on a server	1024	1024	1024	1024	1024	1024	2048
	Number of HBAs of any type	8	8	8	8	8	8	8
	HBA ports	16	16	16	16	16	16	16
	Targets per HBA	256	256	256	256	256	256	256
FCoE	Software FCoE adapters	-	-	4	4	4	4	4
NFS	Default NFS datastores	8	-	-	-	-	-	-
	NFS datastores	64 (requires changes to advanced settings)	-	-	-	-	-	-
VMFS	RDM size	2 TB-512B	2 TB-512B	-	-	-	-	-
	Volume size	64 TB-16 KB	64 TB	-	-	64 TB	64 TB	64 TB
	Volume per host	256	256	256	256	256	256	512
VMFS-2	Files per volume	256 + (64 x additional	-	-	-	-	-	-

		extents)						
	Block size	256 MB	-	-	-	-	-	-
VMFS-3	VMFS-3 volumes configured per host	256	-	-	-	-	-	-
	Files per volume	~30,720[b]	~30,720[b]	~30,720[b]	~30,720[b]	~30,720[b]	~30,720[b]	-
	Block size	8 MB	8 MB	8 MB	8 MB	8 MB	8 MB	-
	Volume size	-	-	64 TB[c]	64 TB[c]	-	64 TB[c]	-
VMFS-5	Volume size	-	-	64 TB[d]	64 TB[d]	64 TB[d]	64 TB[d]	64 TB[d]
	Block size			1 MB	1 MB	1 MB	1 MB	1 MB
	Files per volume	-	-	~130690	~130690	~130690	~130690	~130690
VMFS-6	Volume size	-	-	-	-	-	-	64 TB
	Block size	-	-	-	-	-	-	1 MB
	Files per volume	-	-	-	-	-	-	~130690

NOTE

a. Local disks are included.

b. The file quantity is sufficient to support the maximum number of VMs.

c. If the block size supported by the file system is 1 MB, the maximum volume size is 50 TB.

d. The volume size is also subject to RAID controllers or adapter drivers.

Table 2-3 lists only part of specifications. For more information, see:

- VMware vSphere Configuration Maximums (4.0)
- VMware vSphere Configuration Maximums (4.1)
- VMware vSphere Configuration Maximums (5.0)
- VMware vSphere Configuration Maximums (5.1)
- VMware vSphere Configuration Maximums (5.5)
- VMware vSphere Configuration Maximums (6.0)
- VMware vSphere Configuration Maximums (6.5)

2.5 Common VMware Commands

This chapter describes the commands commonly used in VMware.

Viewing the Version

Run the following commands to view the VMware version:

```
~ # vmware -l
VMware ESXi 5.1.0 GA
~ # vmware -v
VMware ESXi 5.1.0 build-799733
~ #
```

Viewing Hardware Information

Run the following commands to view hardware information including the ESX hardware and kernel:

```
esxcfg-info -a (Displays all related information.)
esxcfg-info -w (Displays ESX hardware information.)
```

Obtaining Help Documentation

Command syntax varies with host system versions. You can perform the following steps to obtain help documentation for different versions of host systems.

Step 1 Log in to the VMware official website.

http://www.vmware.com/support/developer/vcli/

Step 2 Select a VMware version.

Select the latest version of VMware and click **vSphere Command-Line Interface Reference**, as shown in Figure 2-10.

Figure 2-10 Selecting a VMware version

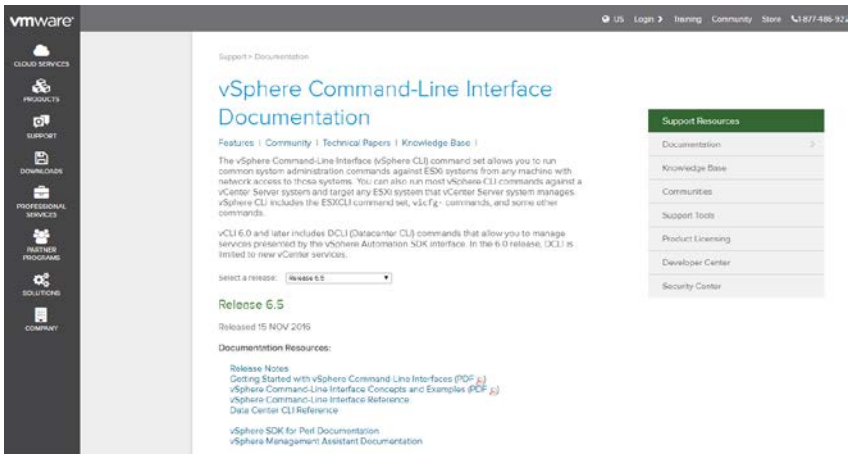

Then, you are navigated to the help page of the selected VMware version.

----End

3 Planning Connectivity

VMware hosts and storage systems can be connected based on different criteria. Table 3-1 describes the typical connection modes.

Table 3-1 Connection modes

Criteria	Connection Mode
Interface module type	Fibre Channel connection/iSCSI connection
Whether switches are used	Direct connection (no switches are used)/Switch-based connection (switches are used)
Whether multiple paths exist	Single-path connection/Multi-path connection
Whether HyperMetro is configured	HyperMetro/Non-HyperMetro

Fibre Channel connections are the most widely used. To ensure service data security, both direct connections and switch-based connections require multiple paths.

The following details Fibre Channel and iSCSI connections in HyperMetro and non-HyperMetro scenarios.

3.1 HyperMetro Scenarios

For details about how to plan connectivity in HyperMetro scenarios, see the *BC&DR Solution Product Documentation (Active-Active Data Center)*.

3.2 Non-HyperMetro Scenarios

3.2.1 Direct FC Connections

Huawei provides two-controller and multi-controller storage systems, which directly connect to VMware ESXi hosts through FC multi-path connections in different ways.

Two-Controller Storage

The following uses Huawei OceanStor 5500 V3 as an example to explain how to directly connect a VMware ESXi host to a two-controller storage system through FC multi-path connections, as shown in Figure 3-2.

Figure 3-2 Direct FC multi-path connections (two-controller storage)

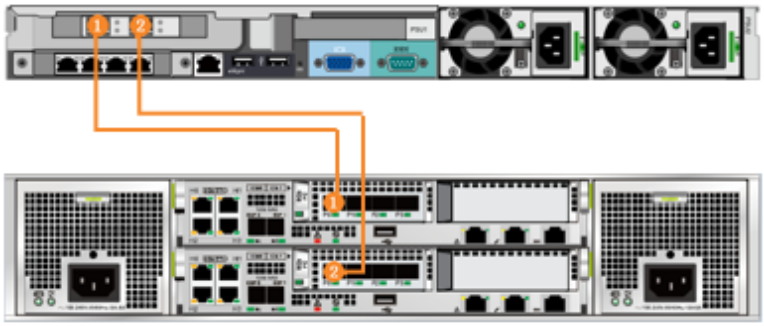

NOTE

In this connection diagram, each of the two controllers is connected to a host HBA port with an optical fiber. The cable connections are detailed in Table 3-2.

Table 3-2 Cable connection description (two-controller storage)

Cable No.	Description
1	Connects Port P0 on the VMware ESXi host to Controller A on the storage system.
2	Connects Port P1 on the VMware ESXi host to Controller B on the storage system.

Multi-Controller Storage

The following uses Huawei OceanStor 18800 V3 (four-controller) as an example to explain how to directly connect a VMware ESXi host to a multi-controller storage system through FC multi-path connections, as shown in Figure 3-2.

Figure 3-3 Direct FC multi-path connections (four-controller storage)

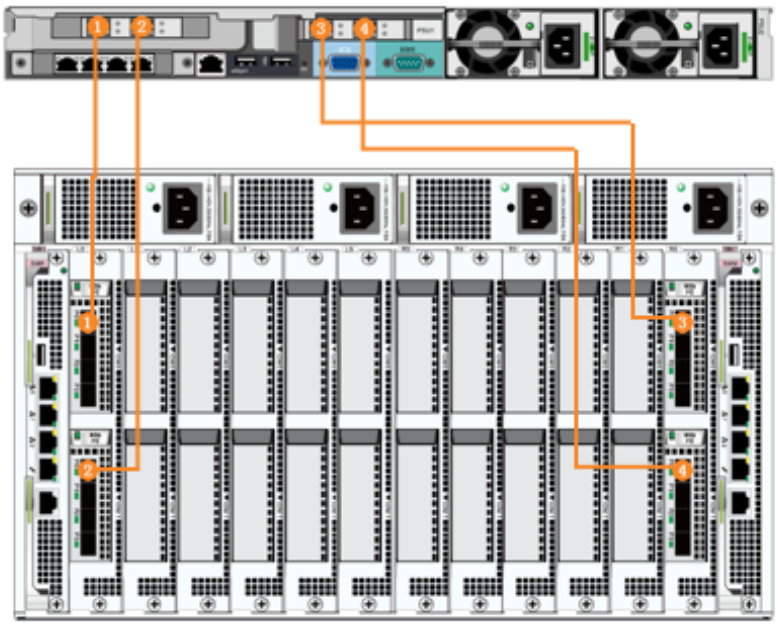

NOTE

In this connection diagram, each of the four controllers is connected to a host HBA port with an optical fiber. The cable connections are detailed in Table 3-3.

Table 3-3 Cable connection description (four-controller storage)

Cable No.	Description
1	Connects Port P0 on the VMware ESXi host to Controller A on the storage system.
2	Connects Port P1 on the VMware ESXi host to Controller B on the storage system.
3	Connects Port P2 on the VMware ESXi host to Controller C on the storage system.
4	Connects Port P3 on the VMware ESXi host to Controller D on the storage system.

3.2.2 Switch-Based FC Connections

Huawei provides two-controller and multi-controller storage systems, which connect to VMware ESXi hosts through FC multi-path connections using a switch in different ways.

Two-Controller Storage

The following uses Huawei OceanStor 5500 V3 as an example to explain how to connect a VMware ESXi host to a two-controller storage system through FC multi-path connections using a switch, as shown in Figure 3-3.

Figure 3-4 Switch-based FC multi-path connections (two-controller storage)

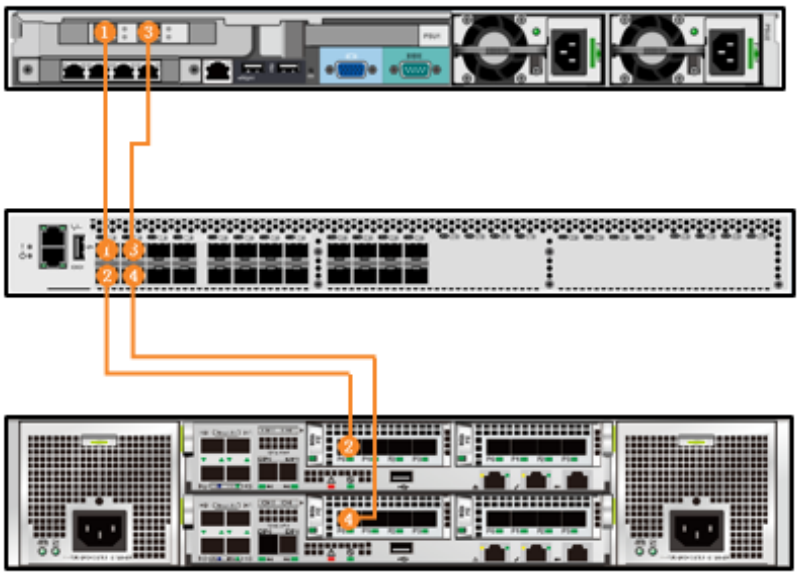

NOTE

In this connection diagram, two controllers of the storage system and two ports of the VMware ESXi host are connected to the FC switch through optical fibers. On the FC switch, the ports connecting to the storage controllers and to the VMware ESXi host are grouped in a zone, ensuring connectivity between the host ports and the storage.

Table 3-4 Zone division on the FC switch (two-controller storage)

Zone Name	Zone Members	Zone Description
Zone001	Ports 1 and 2	Connects Port P0 on the VMware ESXi host to Controller A on the storage system.
Zone002	Ports 3 and 4	Connects Port P1 on the VMware ESXi host to Controller B on the storage system.
Zone003	Ports 1 and 4	Connects Port P0 on the VMware ESXi host to Controller B on the storage system.
Zone004	Ports 3 and 2	Connects Port P1 on the VMware ESXi host to Controller A on the storage system.

NOTE

Zone division in this table is for reference only. Plan zones based on site requirements.

Multi-Controller Storage

The following uses Huawei OceanStor 18800 V3 (four-controller) as an example to explain how to connect a VMware ESXi host to a four-controller storage system through FC multi-path connections using a switch, as shown in Figure 3-4.

Figure 3-5 Switch-based FC multi-path connections (four-controller storage)

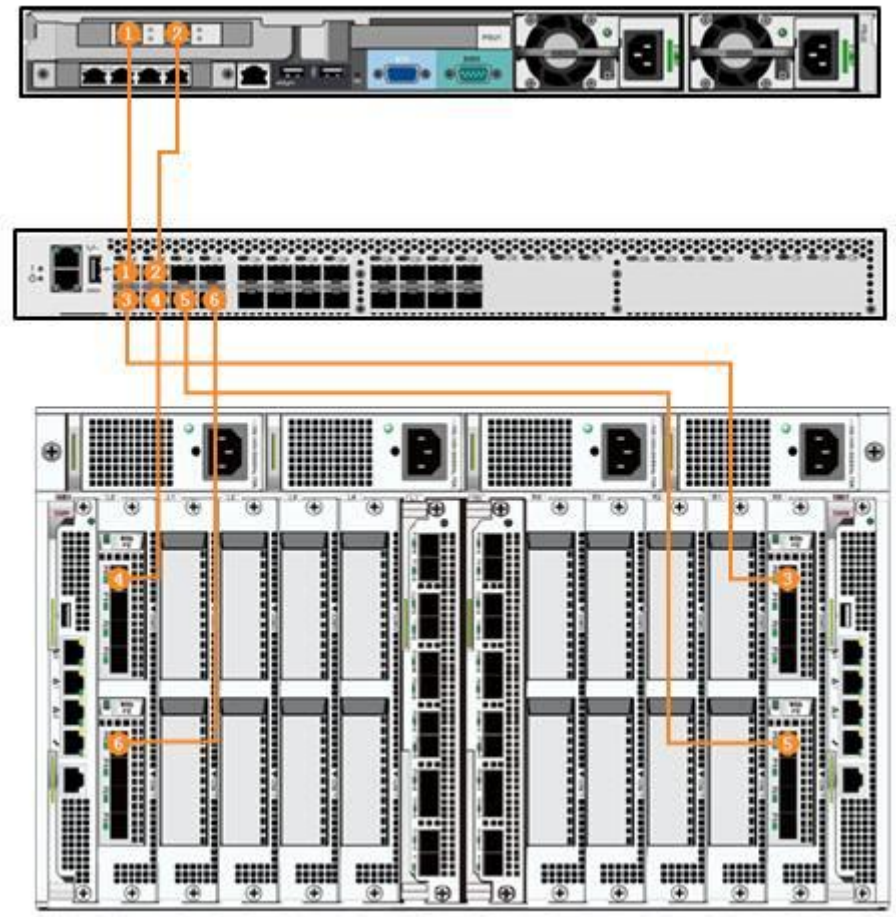

NOTE

In this connection diagram, four controllers of the storage system and two ports of the VMware ESXi host are connected to the FC switch through optical fibers. On the FC switch, the ports connecting to the storage controllers and to the VMware ESXi host are grouped in a zone, ensuring connectivity between the host ports and the storage.

Table 3-5 Zone division on the FC switch (four-controller storage)

Zone Name	Zone Members	Zone Description
Zone001	Ports 1 and 3	Connects Port P0 on the VMware ESXi host to Controller A on the storage system.
Zone002	Ports 1 and 4	Connects Port P0 on the VMware ESXi host to Controller B on the storage system.
Zone003	Ports 1 and 5	Connects Port P0 on the VMware ESXi host to Controller C on the storage system.
Zone004	Ports 1 and 6	Connects Port P0 on the VMware ESXi host to Controller D on the storage system.
Zone005	Ports 2 and 3	Connects Port P1 on the VMware ESXi host to Controller A on the storage system.

Zone Name	Zone Members	Zone Description
Zone006	Ports 2 and 4	Connects Port P1 on the VMware ESXi host to Controller B on the storage system.
Zone007	Ports 2 and 5	Connects Port P1 on the VMware ESXi host to Controller C on the storage system.
Zone008	Ports 2 and 6	Connects Port P1 on the VMware ESXi host to Controller D on the storage system.

NOTE
Zone division in this table is for reference only. Plan zones based on site requirements.

3.2.3 Direct iSCSI Connections

Huawei provides two-controller and multi-controller storage systems, which directly connect to VMware ESXi hosts through iSCSI multi-path connections in different ways.

Two-Controller Storage

The following uses Huawei OceanStor 5500 V3 as an example to explain how to directly connect a VMware ESXi host to a two-controller storage system through iSCSI multi-path connections, as shown in Figure 3-5.

Figure 3-6 Direct iSCSI multi-path connections (two-controller storage)

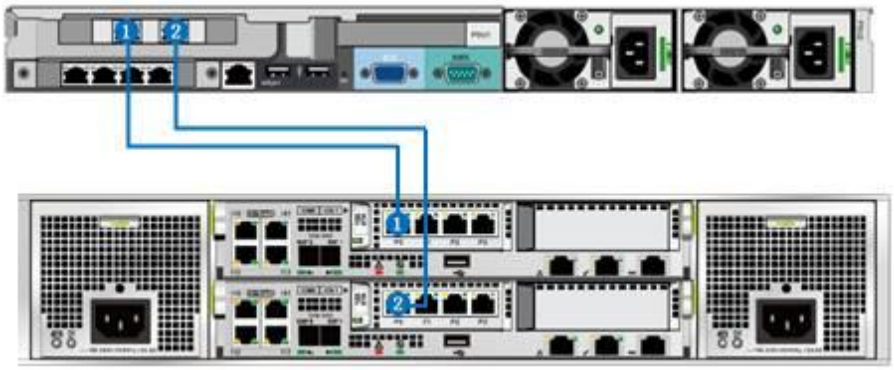

NOTE
In this connection diagram, each of the two controllers is connected to a port on the host network adapter with a network cable. The IP address plan is detailed in Table 3-6.

Table 3-6 IP address plan for direct iSCSI multi-path connections (two-controller storage)

Port Name	Port Description	IP Address	Subnet Mask
Host.P0	Connects the VMware ESXi host to Controller A on the storage system.	192.168.5.5	255.255.255.0

Port Name	Port Description	IP Address	Subnet Mask
Host.P1	Connects the VMware ESXi host to Controller B on the storage system.	192.168.6.5	255.255.255.0
Storage.A.P0	Connects Controller A on the storage system to the VMware ESXi host.	192.168.5.6	255.255.255.0
Storage.B.P0	Connects Controller B on the storage system to the VMware ESXi host.	192.168.6.6	255.255.255.0

NOTE
IP addresses in this table are for reference only. Plan IP addresses based on site requirements.

Multi-Controller Storage

The following uses Huawei OceanStor 18800 V3 (four-controller) as an example to explain how to directly connect a VMware ESXi host to a multi-controller storage system through iSCSI multi-path connections, as shown in Figure 3-6.

Figure 3-7 Direct iSCSI multi-path connections (four-controller storage)

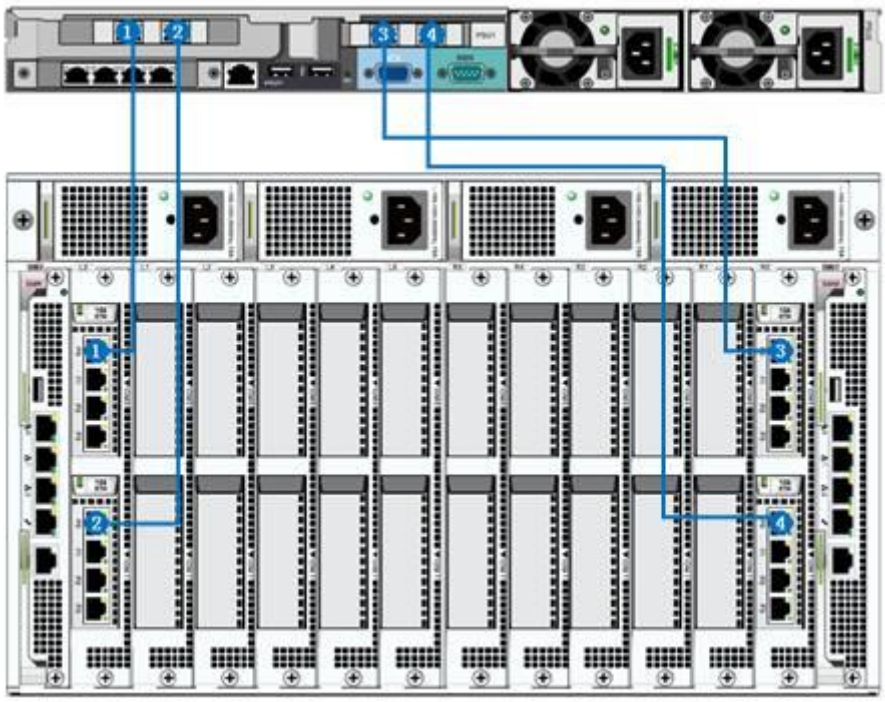

NOTE
In this connection diagram, each of the four controllers is connected to a port on host network adapters with a network cable. The IP address plan is detailed in Table 3-7.

Table 3-7 IP address plan for direct iSCSI multi-path connections (four-controller storage)

Port Name	Port Description	IP Address	Subnet Mask
Host.P0	Connects the VMware ESXi host to Controller A on the storage system.	192.168.5.5	255.255.255.0
Host.P1	Connects the VMware ESXi host to Controller B on the storage system.	192.168.6.5	255.255.255.0
Host.P2	Connects the VMware ESXi host to Controller C on the storage system.	192.168.7.5	255.255.255.0
Host.P3	Connects the VMware ESXi host to Controller D on the storage system.	192.168.8.5	255.255.255.0
Storage.A.P0	Connects Controller A on the storage system to the VMware ESXi host.	192.168.5.6	255.255.255.0
Storage.B.P0	Connects Controller B on the storage system to the VMware ESXi host.	192.168.6.6	255.255.255.0
Storage.C.P0	Connects Controller C on the storage system to the VMware ESXi host.	192.168.7.6	255.255.255.0
Storage.D.P0	Connects Controller D on the storage system to the VMware ESXi host.	192.168.8.6	255.255.255.0

NOTE
IP addresses in this table are for reference only. Plan IP addresses based on site requirements.

3.2.4 Switch-Based iSCSI Connections

Huawei provides two-controller and multi-controller storage systems, which connect to VMware ESXi hosts through Ethernet switches in different ways.

Two-Controller Storage

The following uses Huawei OceanStor 5500 V3 as an example to explain how to connect a VMware ESXi host to a two-controller storage system through iSCSI multi-path connections using an Ethernet switch, as shown in Figure 3-7.

Figure 3-8 Switch-based iSCSI multi-path connections (two-controller storage)

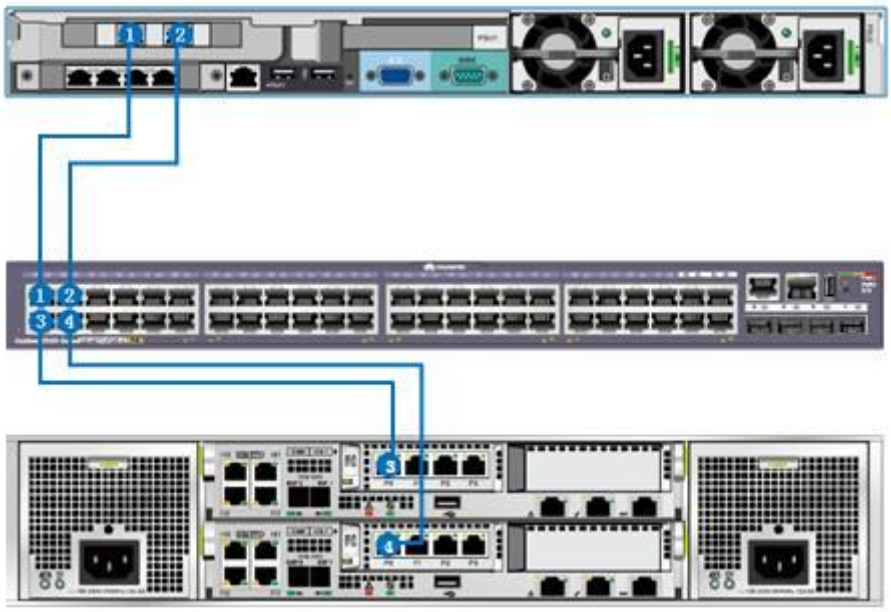

NOTE

In this connection diagram, two controllers of the storage system and two ports of the VMware ESXi host network adapter are connected to the Ethernet switch through network cables. IP addresses of the ports on the storage and host are in the same subnet, ensuring connectivity between the host ports and the storage.

Table 3-8 IP address plan for switch-based iSCSI multi-path connections (two-controller storage)

Port Name	Port Description	IP Address	Subnet Mask
Host.P0	Connects the VMware ESXi host to Controller A on the storage system.	192.168.5.5	255.255.255.0
Host.P1	Connects the VMware ESXi host to Controller B on the storage system.	192.168.6.5	255.255.255.0
Storage.A.P0	Connects Controller A on the storage system to the VMware ESXi host.	192.168.5.6	255.255.255.0
Storage.B.P0	Connects Controller B on the storage system to the VMware ESXi host.	192.168.6.6	255.255.255.0

NOTE

IP addresses in this table are for reference only. Plan IP addresses based on site requirements.

Multi-Controller Storage

The following uses Huawei OceanStor 18800 V3 (four-controller) as an example to explain how to connect a VMware ESXi host to a multi-controller storage system through iSCSI multi-path connections using an Ethernet switch, as shown in Figure 3-8.

Figure 3-9 Switch-based iSCSI multi-path connections (four-controller storage)

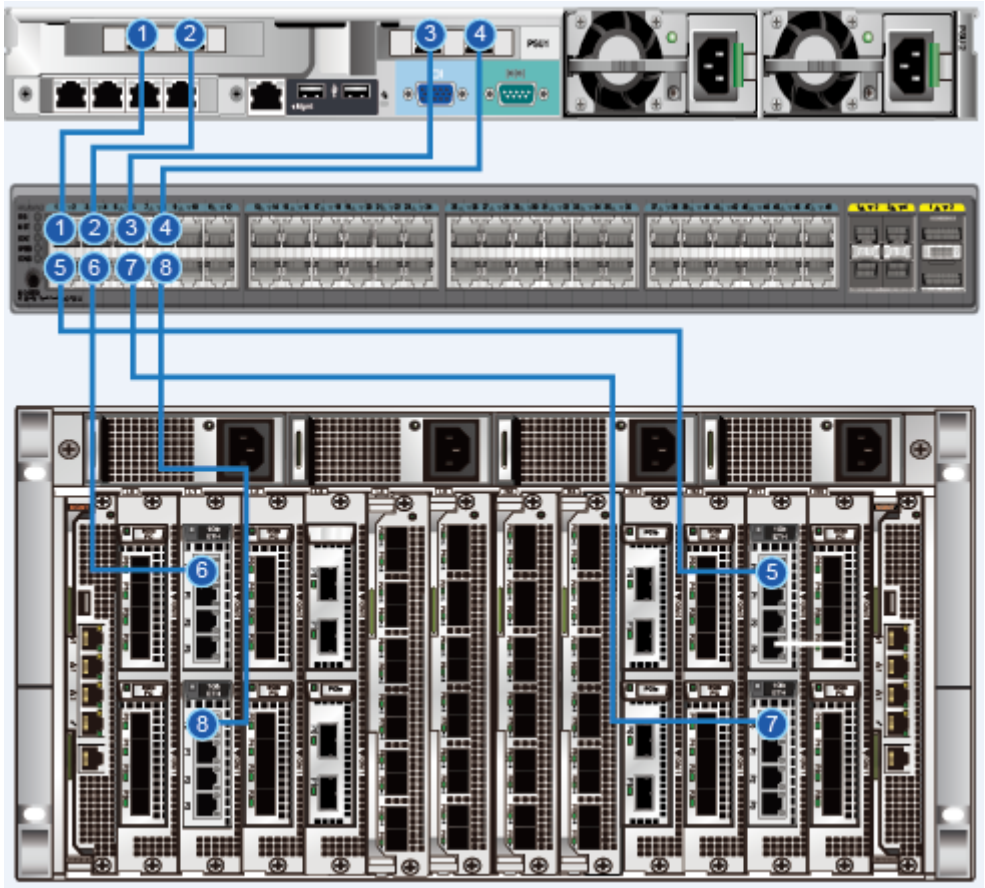

NOTE

In this connection diagram, four controllers of the storage system and four ports of the VMware ESXi host network adapters are connected to the Ethernet switch through network cables. IP addresses of the ports on the storage and host are in the same subnet, ensuring connectivity between the host ports and the storage.

Table 3-9 IP address plan for switch-based iSCSI multi-path connections (four-controller storage)

Port Name	Port Description	IP Address	Subnet Mask
Host.P0	Connects the VMware ESXi host to Controller A on the storage system.	192.168.5.5	255.255.255.0
Host.P1	Connects the VMware ESXi host to Controller B on the storage system.	192.168.6.5	255.255.255.0

Port Name	Port Description	IP Address	Subnet Mask
Host.P2	Connects the VMware ESXi host to Controller C on the storage system.	192.168.7.5	255.255.255.0
Host.P3	Connects the VMware ESXi host to Controller D on the storage system.	192.168.8.5	255.255.255.0
Storage.A.P0	Connects Controller A on the storage system to the VMware ESXi host.	192.168.5.6	255.255.255.0
Storage.B.P0	Connects Controller B on the storage system to the VMware ESXi host.	192.168.6.6	255.255.255.0
Storage.C.P0	Connects Controller C on the storage system to the VMware ESXi host.	192.168.7.6	255.255.255.0
Storage.D.P0	Connects Controller D on the storage system to the VMware ESXi host.	192.168.8.6	255.255.255.0

NOTE

IP addresses in this table are for reference only. Plan IP addresses based on site requirements.

4 Preparations Before Configuration

4.1 Switch

Ensure that the switches are running properly and their ports have the necessary licenses and transmit data normally. Figure 4-1 shows an example of a port failure due to lack of a license.

Figure 4-1 Switch port status

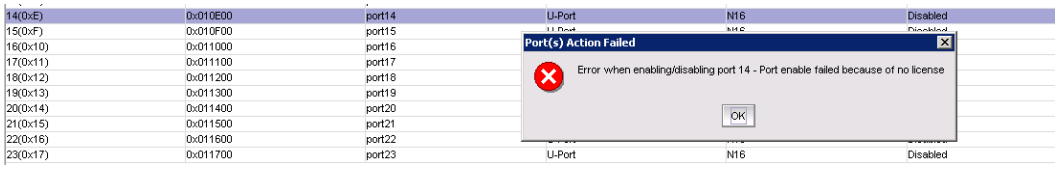

It is recommended that you obtain the product documentation of the switches for reference.

4.2 Storage System

Create disk domains, storage pools, LUNs, hosts, and mapping views on the storage system according to your service requirements. For details about these operations, see the *Basic Storage Service Configuration Guide* corresponding to your storage system.

4.3 Host

Before connecting a host to a storage system, you need to start the ESXi Shell and SSH services, verify that the host HBAs are identified and working correctly, and obtain the WWNs of HBA ports. The WWNs will be used in subsequent configuration on the storage system.

4.3.1 Enabling the Shell and SSH Services for the ESXi Hosts

Enable ESXi Shell and SSH respectively, as shown in Figure 4-2. If you do not need the Shell and SSH services any more, you can disable the services.

Figure 4-2 Enabling the Shell and SSH services

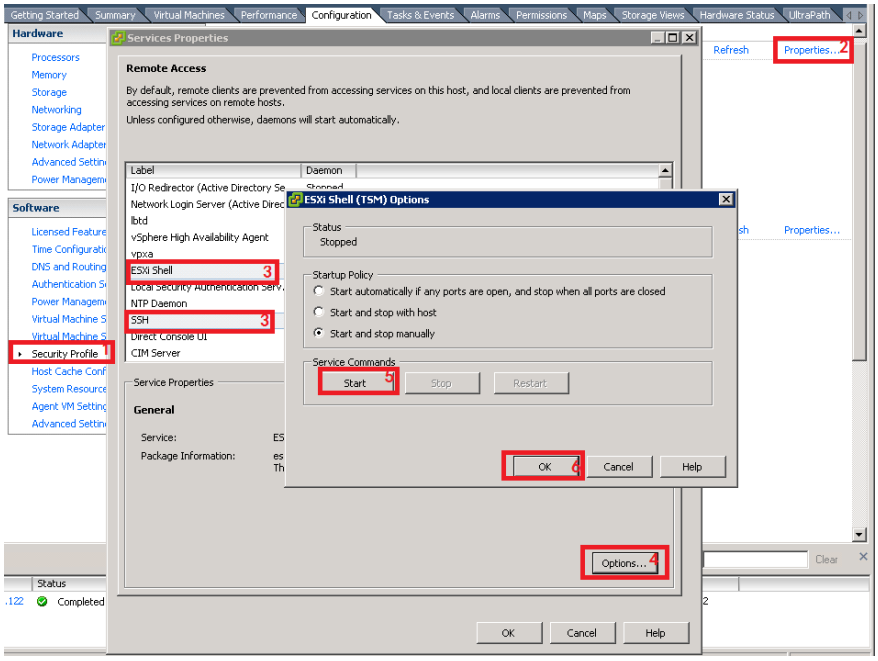

4.3.2 Disabling ATS Heartbeat

VMware ESXi 5.5 Update 2, ESXi 6.0, and later versions support VAAI ATS Heartbeat, which may cause the host to lose connectivity to datastores under certain conditions. For details, see VMware KB ESXi host loses connectivity to a VMFS3 and VMFS5 datastore (2113956). Huawei recommends that you disable ATS Heartbeat on VMware ESXi 5.5 Update 2, ESXi 6.0, and later versions. Versions earlier than ESXi 5.5 Update 2 do not support this function and you can skip this operation.

To disable ATS Heartbeat, perform the following steps:

Step 1 Log in to the ESXi host using SSH and run the following command to disable ATS Heartbeat:

```
#esxcli system settings advanced set -i 0 -o /VMFS3/UseATSForHBOnVMFS5
```

Step 2 Verify the result. If ATS Heartbeat has been disabled, the value of **/VMFS3/UseATSForHBOnVMFS5** is **0**.

```
#esxcli system settings advanced list -o /VMFS3/UseATSForHBonVMFS5
Path: /VMFS3/UseATSForHBOnVMFS5
Type: integerInt Value: 0 <--- check this valueDefault
Int Value: 1
Min Value: 0
Max Value: 1
String Value:Default
String Value:
Valid Characters:Description: Use ATS for HB on ATS supported VMFS5 volumes
```

----End

4.3.3 Identifying HBAs

After an HBA is installed on a host, view information about the HBA on the host. Go to the page for configuration management and choose **Storage Adapters** in the navigation tree. In the function pane, hardware devices on the host are displayed, as shown in Figure 4-3.

Figure 4-3 Viewing the HBA information

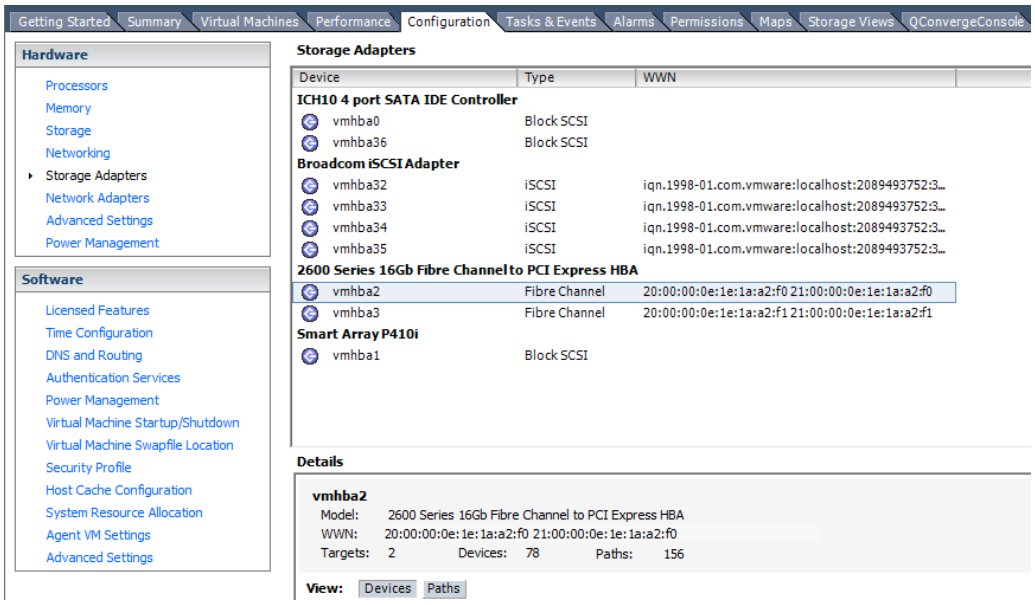

4.3.4 Querying HBA Priorities

After a host identifies a newly installed HBA, you can view properties of the HBA on the host.

The method of querying HBA information varies with operating system versions. The following details how to query HBA information on ESXi 5.5 and versions earlier than ESXi 5.5.

Versions Earlier than VMware ESXi 5.5

The command for viewing the HBA properties varies according to the HBA type. The details are as follows:

- QLogic HBA

  The command syntax is as follows:

  ```
  cat /proc/scsi/qla2xxx/N
  ```

  The command return provides information such as the HBA driver version, topology, WWN, and negotiated rate.
- Emulex HBA

  The command syntax is as follows:

  ```
  cat /proc/scsi/lpfcxxx/N
  ```

  The command return provides information such as HBA model and driver.
- Brocade HBA

```
cat /proc/scsi/bfaxxx/N
```

VMware ESXi 5.5 and Later Versions

Since VMware ESXi 5.5, the **/proc/scsi/** directory contains no content. Run the following commands to query HBA information:

```
~ # esxcli storage core adapter list
HBA Name  Driver     Link State  UID                                   Description
--------  ---------  ----------  ------------------------------------  --------------------------------------------------------
vmhba0    ata_piix   link-n/a    sata.vmhba0                           (0:0:31.2) Intel Corporation ICH10 4 port SATA IDE Controller
vmhba1    hpsa       link-n/a    sas.5001438017531290                  (0:5:0.0) Hewlett-Packard Company Smart Array P410i
vmhba2    qlnativefc link-up     fc.2000000e1e1aa2f0:2100000e1e1aa2f0  (0:11:0.0) QLogic Corp 2600 Series 16Gb Fibre Channel to PCI Express HBA
vmhba3    qlnativefc link-up     fc.2000000e1e1aa2f1:2100000e1e1aa2f1  (0:11:0.1) QLogic Corp 2600 Series 16Gb Fibre Channel to PCI Express HBA
vmhba32 bnx2i      unbound    iscsi.vmhba32                         Broadcom iSCSI Adapter
vmhba33 bnx2i      unbound    iscsi.vmhba33                         Broadcom iSCSI Adapter
vmhba34 bnx2i      unbound    iscsi.vmhba34                         Broadcom iSCSI Adapter
vmhba35 bnx2i      unbound    iscsi.vmhba35                         Broadcom iSCSI Adapter
vmhba36 ata_piix   link-n/a   sata.vmhba36                         (0:0:31.2) Intel Corporation ICH10 4 port SATA IDE Controller
```

You can run the following command to obtain more HBA details:

```
# /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -a
```

For more information, visit:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1031534

For details about how to modify the HBA queue depth, visit:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1267

5 Configuring Connectivity

5.1 Establishing Fibre Channel Connections

5.1.1 Host Configuration

Query the HBA WWN. Figure 5-1 provides an example.

Figure 5-1 Querying the WWN of the host HBA

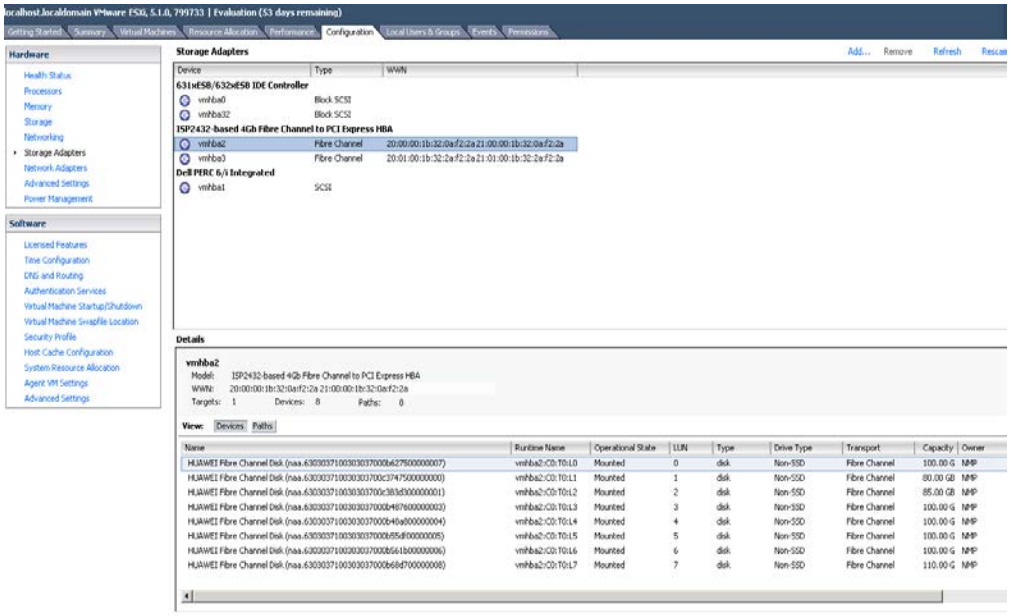

5.1.2 (Optional) Switch Configuration

5.1.2.1 Querying the Switch Model and Version

Perform the following steps to query the switch model and version:

Step 1 Log in to the Brocade switch from a web page.

On the web page, enter the IP address of the Brocade switch. The **Web Tools** switch login dialog box is displayed. Enter the account and password. The default account and password are **admin** and **password**. The switch management page is displayed.

CAUTION

Web Tools works correctly only when Java is installed on the host. Java 1.6 or later is recommended.

Step 2 On the switch management page that is displayed, click **Switch Information**. The switch information is displayed.

Figure 5-2 Switch information

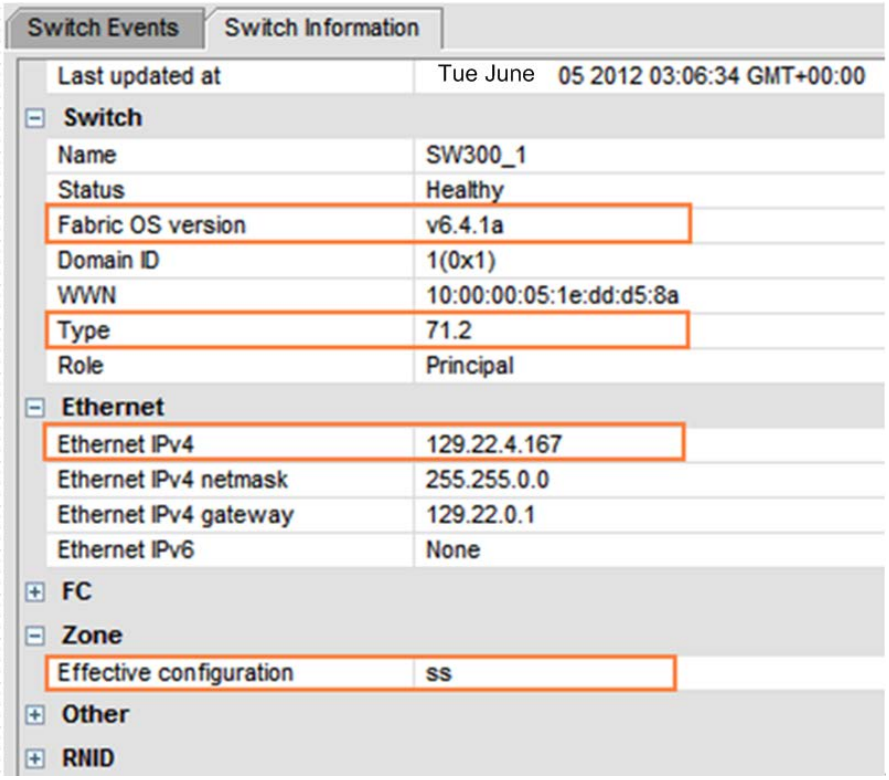

Note the following parameters:

- **Fabric OS version**: indicates the switch version information. The interoperability between switches and storage systems varies with the switch version. Only switches of authenticated versions can interconnect correctly with storage systems.
- **Type**: This parameter is a decimal consisting of an integer and a decimal fraction. The integer indicates the switch model and the decimal fraction indicates the switch template version. You only need to pay attention to the switch model. Table 5-1 describes the switch model mappings.
- **Ethernet IPv4**: indicates the switch IP address.

- **Effective configuration**: indicates the currently effective configurations. This parameter is important and is related to zone configurations. In this example, the currently effective configuration is **ss**.

Table 5-1 Mapping between switch types and names

Switch Type	Switch Name	Switch Type	Switch Name
1	Brocade 1000 Switch	64	Brocade 5300 Switch
2,6	Brocade 2800 Switch	66	Brocade 5100 Switch
3	Brocade 2100, 2400 Switches	67	Brocade Encryption Switch
4	Brocade 20x0, 2010, 2040, 2050 Switches	69	Brocade 5410 Blade
5	Brocade 22x0, 2210, 2240, 2250 Switches	70	Brocade 5410 Embedded Switch
7	Brocade 2000 Switch	71	Brocade 300 Switch
9	Brocade 3800 Switch	72	Brocade 5480 Embedded Switch
10	Brocade 12000 Director	73	Brocade 5470 Embedded Switch
12	Brocade 3900 Switch	75	Brocade M5424 Embedded Switch
16	Brocade 3200 Switch	76	Brocade 8000 Switch
17	Brocade 3800VL	77	Brocade DCX-4S Backbone
18	Brocade 3000 Switch	83	Brocade 7800 Extension Switch
21	Brocade 24000 Director	86	Brocade 5450 Embedded Switch
22	Brocade 3016 Switch	87	Brocade 5460 Embedded Switch
26	Brocade 3850 Switch	90	Brocade 8470 Embedded Switch
27	Brocade 3250 Switch	92	Brocade VA-40FC Switch
29	Brocade 4012 Embedded Switch	95	Brocade VDX 6720-24 Data Center Switch
32	Brocade 4100 Switch	96	Brocade VDX 6730-32 Data Center Switch
33	Brocade 3014 Switch	97	Brocade VDX 6720-60 Data Center Switch
34	Brocade 200E Switch	98	Brocade VDX 6730-76 Data

Switch Type	Switch Name	Switch Type	Switch Name
			Center Switch
37	Brocade 4020 Embedded Switch	108	Dell M8428-k FCoE Embedded Switch
38	Brocade 7420 SAN Router	109	Brocade 6510 Switch
40	Fibre Channel Routing (FCR) Front Domain	116	Brocade VDX 6710 Data Center Switch
41	Fibre Channel Routing, (FCR) Xlate Domain	117	Brocade 6547 Embedded Switch
42	Brocade 48000 Director	118	Brocade 6505 Switch
43	Brocade 4024 Embedded Switch	120	Brocade DCX 8510-8 Backbone
44	Brocade 4900 Switch	121	Brocade DCX 8510-4 Backbone
45	Brocade 4016 Embedded Switch	124	Brocade 5430 Switch
46	Brocade 7500 Switch	125	Brocade 5431 Switch
51	Brocade 4018 Embedded Switch	129	Brocade 6548 Switch
55.2	Brocade 7600 Switch	130	Brocade M6505 Switch
58	Brocade 5000 Switch	133	Brocade 6520 Switch
61	Brocade 4424 Embedded Switch	134	Brocade 5432 Switch
62	Brocade DCX Backbone	148	Brocade 7840 Switch

----End

5.1.2.2 Configuring Zones

Skip this section if you use direct connections.

Zone configuration is important for Fibre Channel switches. The configurations differ with the switch vendor, model, and version. For details, refer to the specific switch's *Configuration Guide*. The following explains the zone configuration procedure by using the Brocade 6510 switch as an example.

Step 1 Log in to the Brocade switch on a web browser.

On the web browser, enter the IP address of the Brocade switch and press **Enter**. The **Web Tools** switch login dialog box is displayed. Enter the account and password (**admin** and **password** by default) to log in.

Step 2 Check the port status on the switch.

In normal conditions, port indicators on the switch are steady green after the corresponding ports have been connected to hosts and storage arrays using optical fibers. This example uses ports 0, 1, 4, and 5, as shown in Figure 5-3.

Figure 5-3 Port status

Step 3 Go to the **Zone Admin** page.

Choose **Configure** > **Zone Admin** from the main menu of **Web Tools**.

Figure 5-4 Zone Admin page

Step 4 Check whether the switch has identified hosts and storage systems.

On the **Zone Admin** page, click the **Zone** tab. In **Member Selection List**, check whether all related ports have been identified, as shown in Figure 5-5.

Figure 5-5 Identified ports

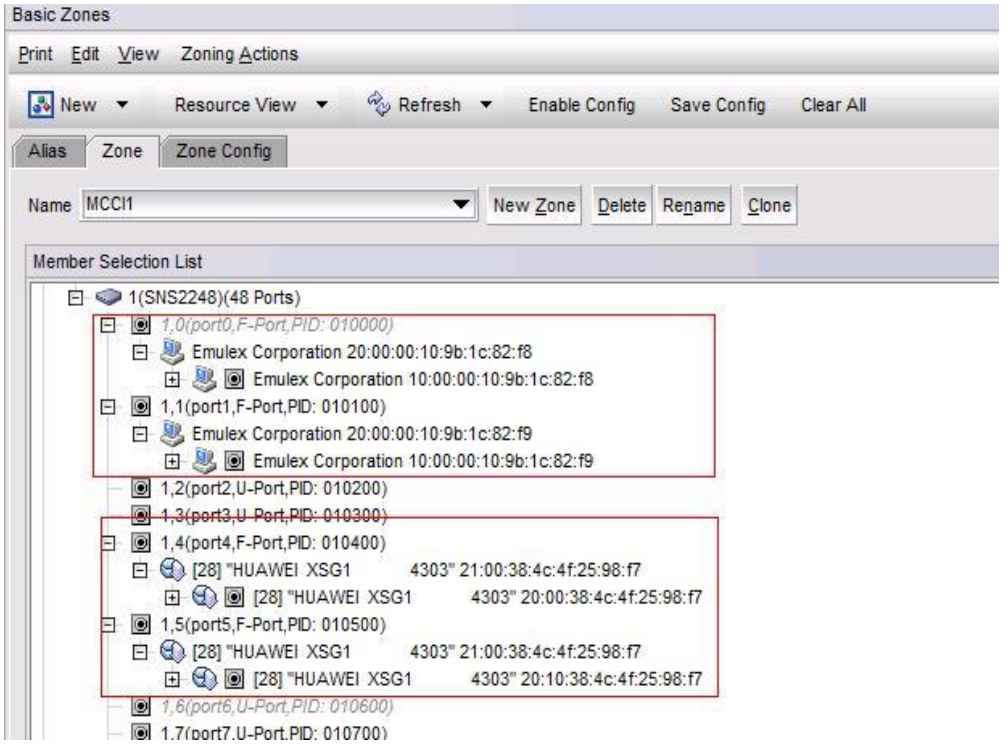

In this example, the hosts use ports 0 and 1, while the storage systems use ports 4 and 5. The display indicates that the switch has correctly identified the devices connected by the four ports.

Step 5 Create a zone.

On the **Zone** tab page, click **New Zone** and enter a name (**Zone001** in this example). Add port 0 (connecting to port P0 of a host) and port 4 (connecting to controller A of a storage system) to this zone, as shown in Figure 5-6.

Figure 5-6 Creating a zone

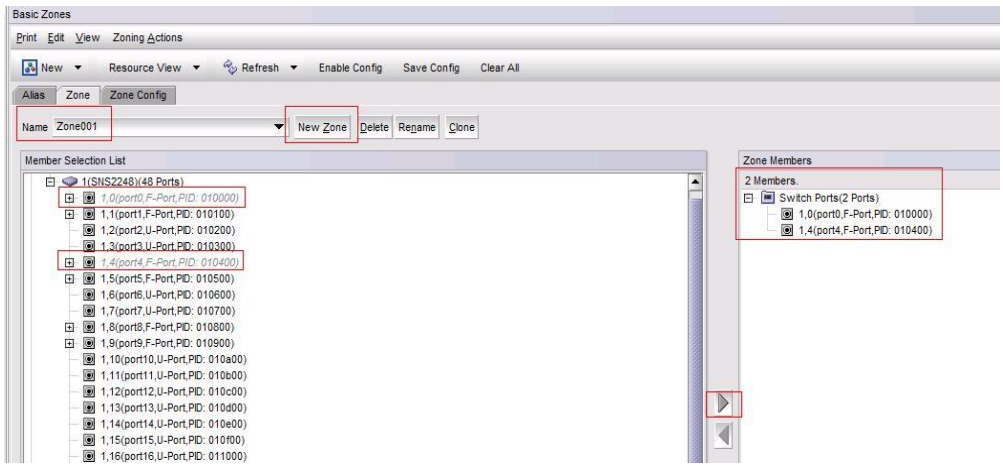

Use the same method to create **Zone002** to **Zone004**. Add ports 1 and 5 to **Zone0002**, ports 0 and 5 to **Zone003**, and ports 1 and 4 to **Zone004**.

Step 6 Add the new zones to the configuration file and activate them.

On the **Switch View** tab page, identify the effective configuration file, as shown in Figure 5-7.

Figure 5-7 Effective configuration file

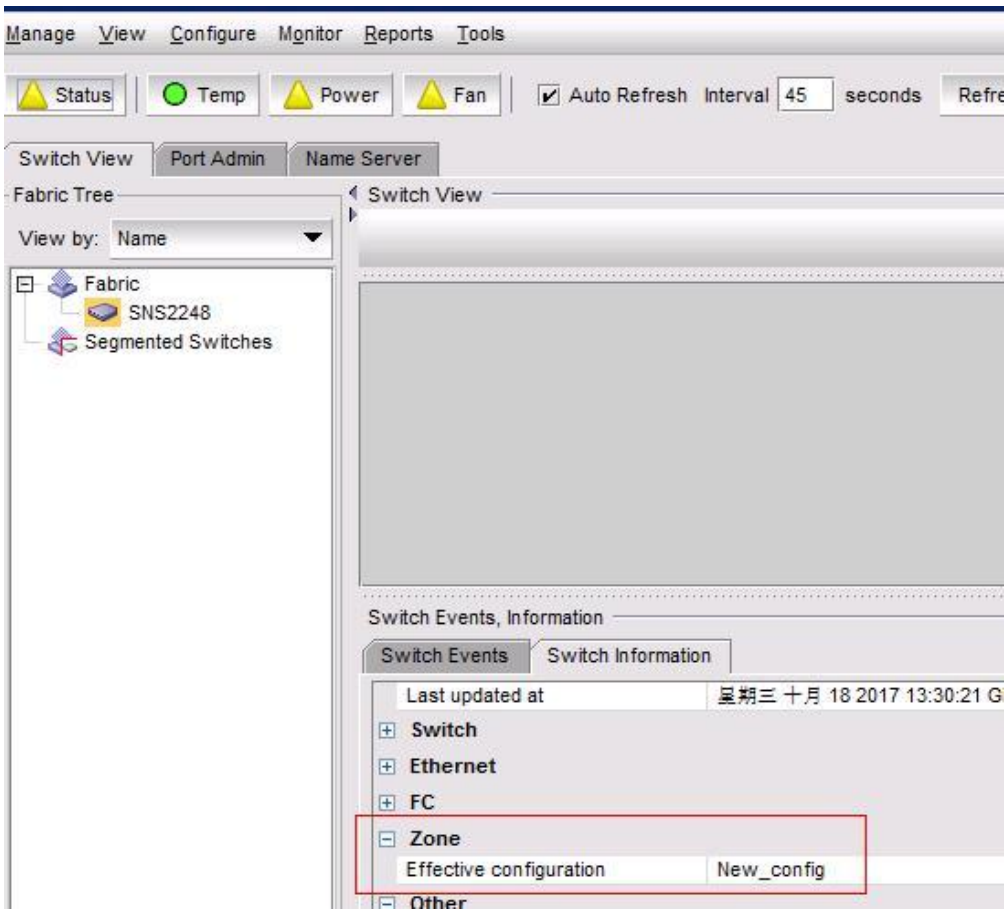

On the **Zone Admin** page, click the **Zone Config** tab. In the **Name** drop-down list, choose the effective configuration file **New_config**.

In **Member Selection List**, select **Zone001** to **Zone004** and add them to the configuration file.

Click **Save Config** to save the configuration and then click **Enable Config** for the configuration to take effect.

Figure 5-8 shows the configuration on the GUI.

Figure 5-8 Adding zones to the configuration file

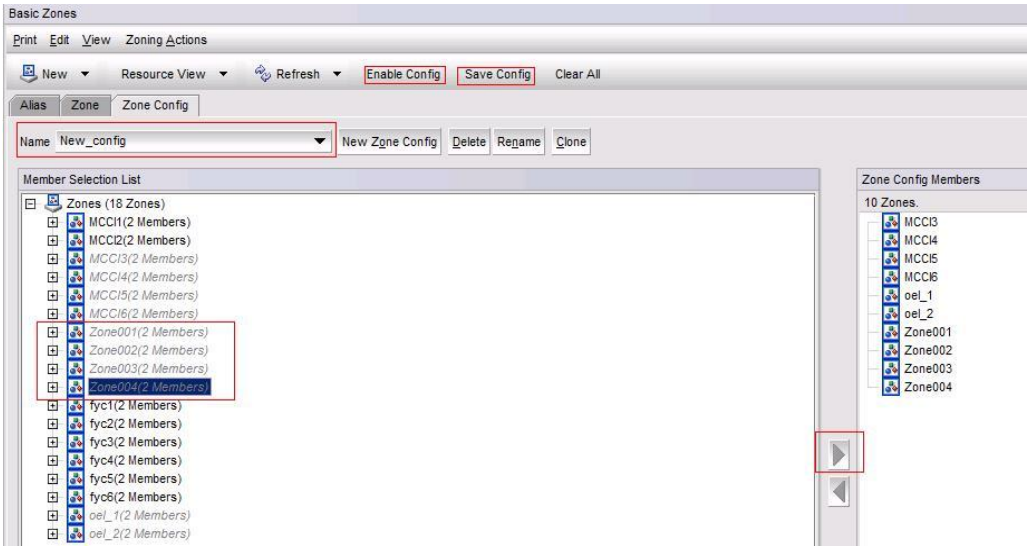

Step 7 Verify that the configuration has taken effect.

On the **Name Server** tab page, verify that the ports have been added to the zones and the zones have taken effect (marked * in the upper right corner), as shown in Figure 5-9.

Figure 5-9 Verifying the configuration

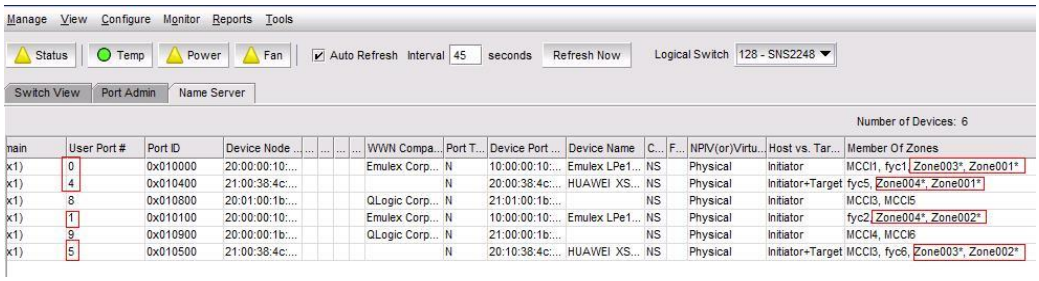

----End

5.1.3 Storage System Configuration

This section details how to add initiators to the hosts on the storage system. For other storage configurations, see the *Basic Storage Service Configuration Guide* corresponding to your storage system.

Step 1 Log in to the storage system on a web browser.

After you have configured the zones on the switch, log in to DeviceManager of the storage system and choose **Provisioning** > **Host** > **Initiator**. On the page that is displayed, select **FC** from the **Initiator Type** drop-down list. Check whether the host initiators have been discovered.

As shown in Figure 5-10, the host initiators have been discovered and are online.

Figure 5-10 Viewing initiators

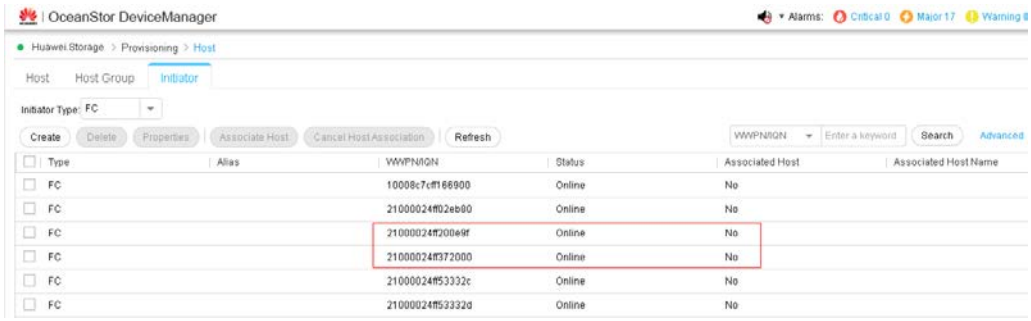

Step 2 Click the **Host** tab, select the host that was created on the storage system, and click **Add Initiator**.

Figure 5-11 Add Initiator dialog box

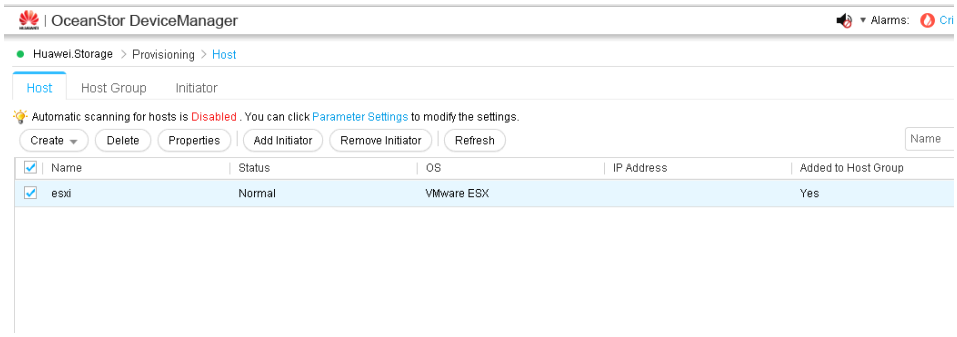

Step 3 Select **FC** from the **Initiator Type** drop-down list and find the host initiators' WWNs.

Figure 5-12 Selecting initiators

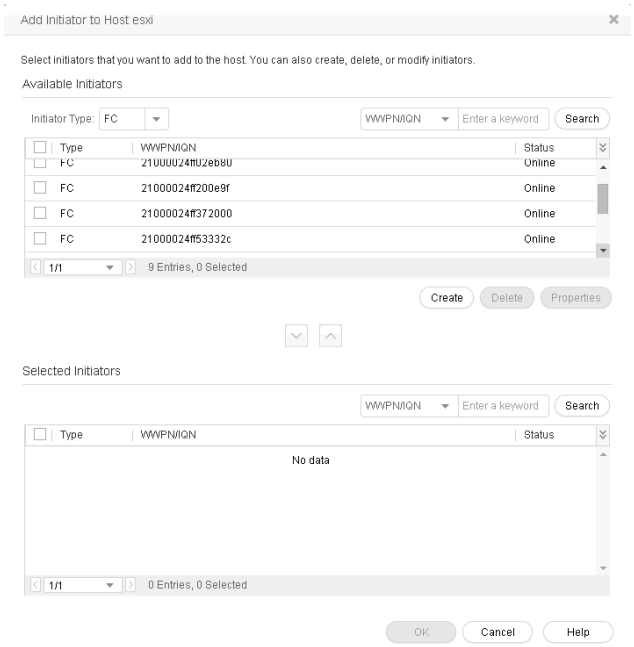

Step 4 Select the host initiators and add them to **Selected Initiators**.

Figure 5-13 Adding initiators

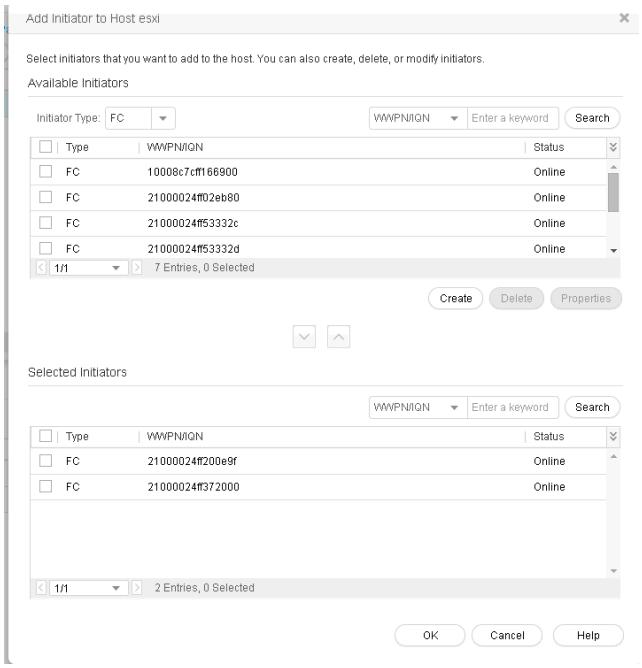

Step 5 Verify that the initiators have been added to the host correctly.

Figure 5-14 Verifying the configuration

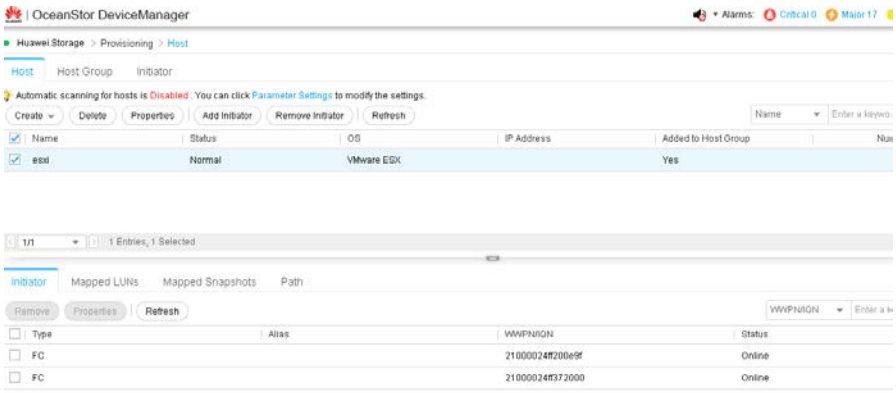

As shown in Figure 5-14, the initiators have been added to the host successfully. The initiator properties depend on the operating system and multipathing software used by the hosts. For details, see the storage-side configuration in the multipathing configuration section. After the initiators have been configured, you can scan for LUNs on the hosts to discover storage resources.

----End

5.2 Establishing iSCSI Connections

Before establishing an iSCSI connection between a host and a storage system, you need to configure the host and storage accordingly.

5.2.1 Host Configuration

5.2.1.1 Configuring Service IP Addresses

On VMware ESXi hosts, you can configure services IP addresses on a VMware host by adding virtual networks.

For vSphere Client

For vSphere Client, perform the following steps to configure services IP addresses:

Step 1 In vSphere Client, choose **Network** > **Add Network**.

Step 2 In **Add Network Wizard** that is displayed, select **VMkernel**, as shown in Figure 5-15

Figure 5-15 Adding VMkernel

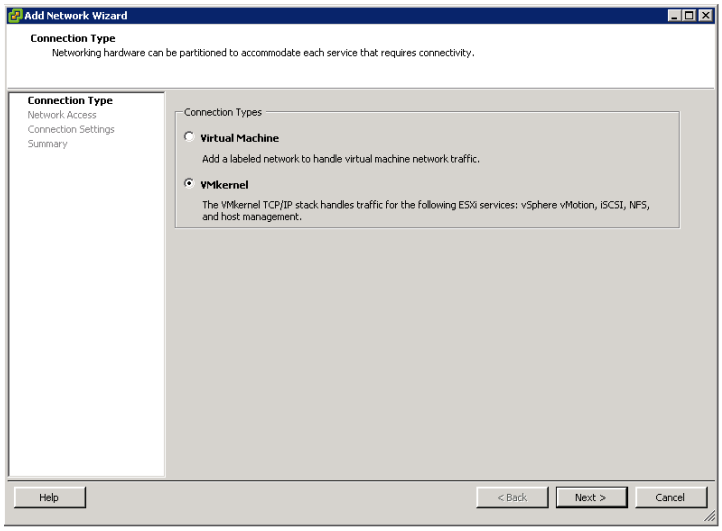

Step 3 Click **Next**.

Step 4 Select the iSCSI service network port, as shown in Figure 5-16

Figure 5-16 Creating a vSphere standard switch

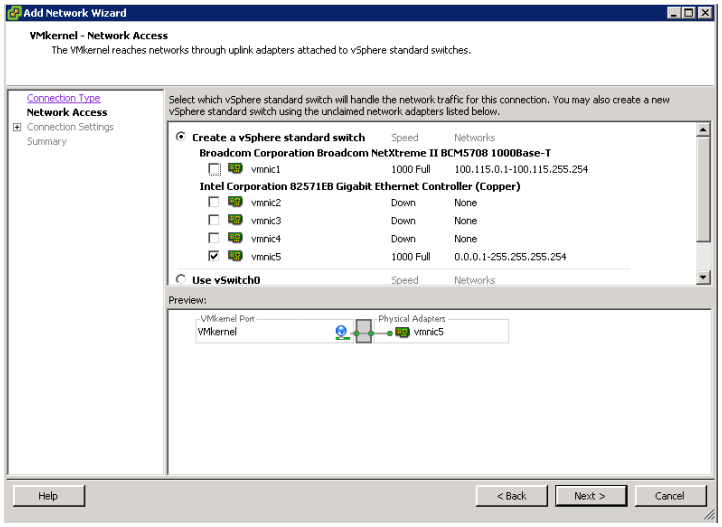

Step 5 Specify the network label, as shown in Figure 5-17

Figure 5-17 Specifying the network label

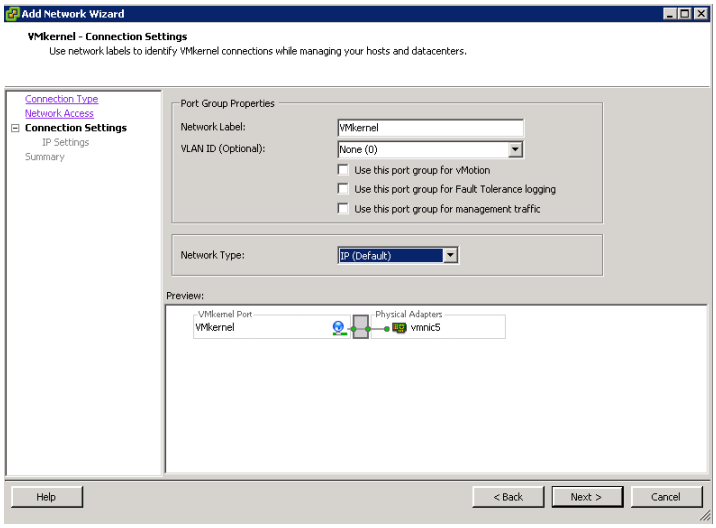

Step 6 Enter the iSCSI service IP address, as shown in Figure 5-18

Figure 5-18 Entering the iSCSI service IP address

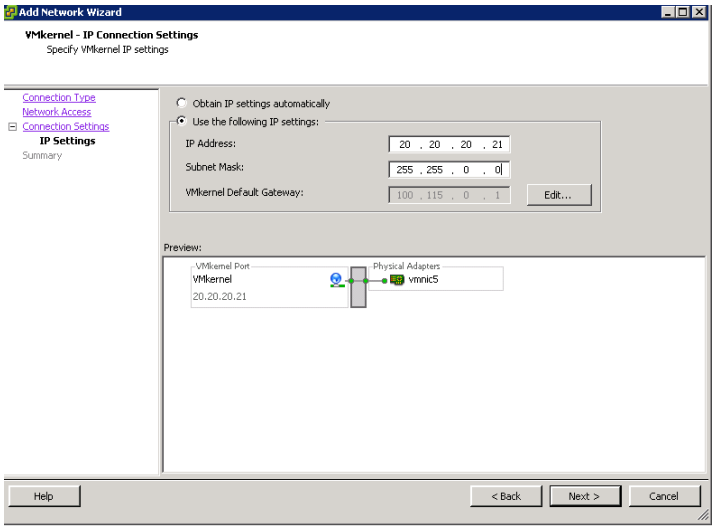

Step 7 Confirm the information that you have configured, as shown in Figure 5-19

Figure 5-19 Information summary

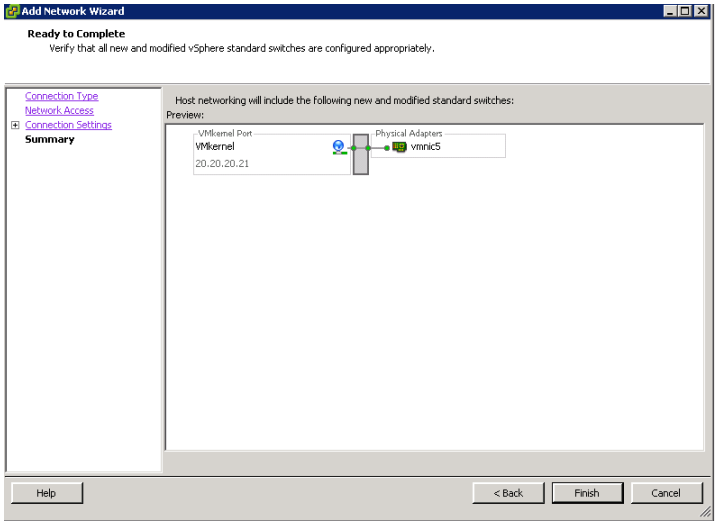

Step 8 For a single-path network, the configuration is completed. For a multi-path network, proceed with the next step.

Step 9 Repeat steps 1 to 6 to create another virtual network.

Figure 5-20 shows the configuration completed for a multi-path network.

Figure 5-20 iSCSI multi-path network with dual adapters

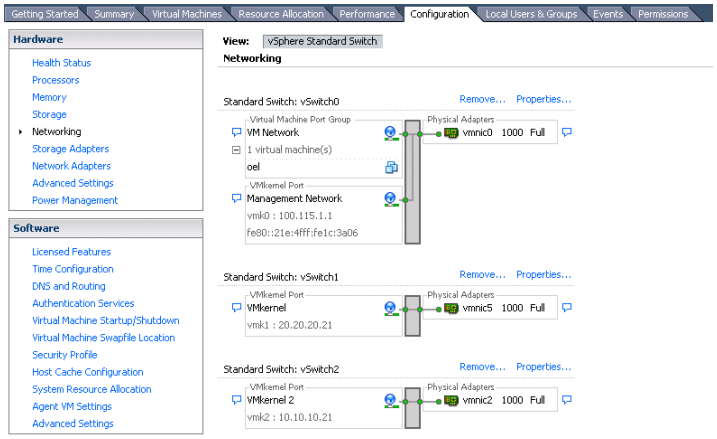

----End

For vSphere Web Client

For vSphere Web Client, perform the following steps to configure services IP addresses:

Step 1 In vSphere Web Client, click **Hosts and Clusters** on the **Home** page.

Figure 5-21 Home page on vSphere Web Client

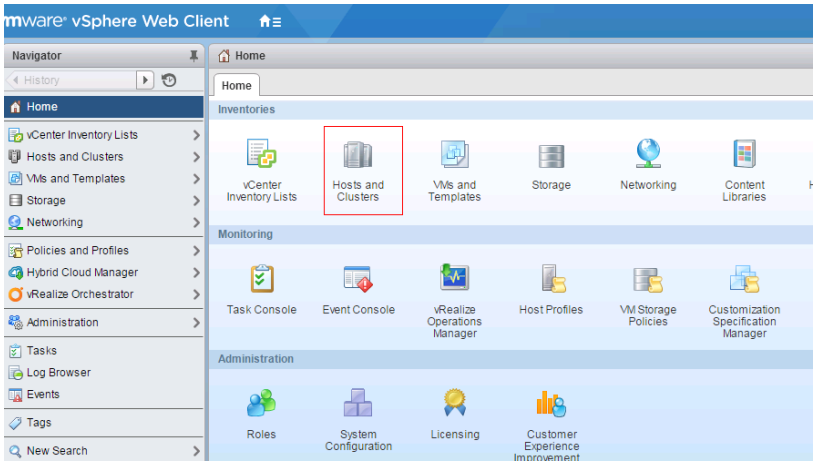

Step 2 Select the target host, click the **Manage** tab, and then click the **Networking** tab.

Figure 5-22 Navigating to the **Networking** tab page

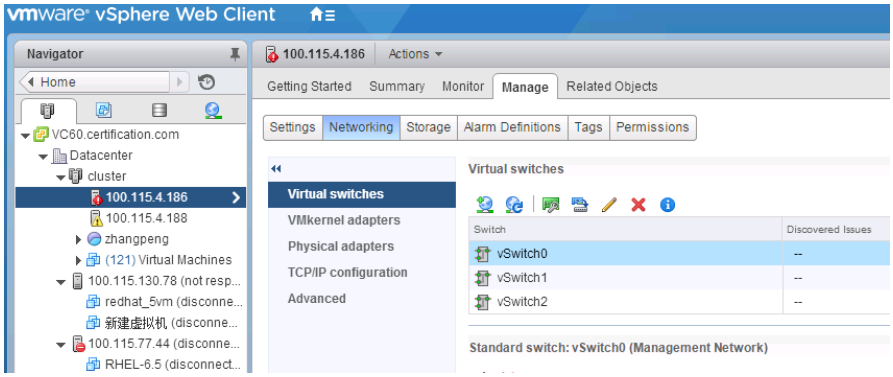

Step 3 Add VMkernel adapters.

Figure 5-23 Adding VMkernel adapters

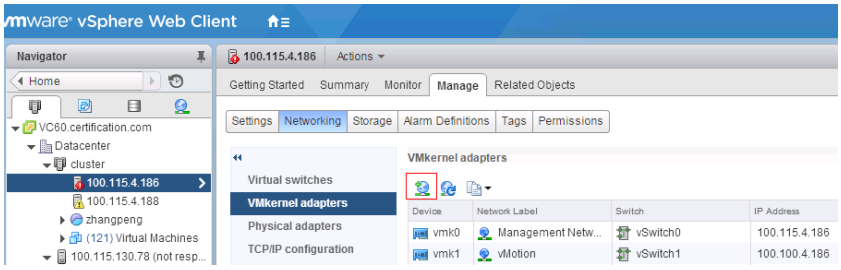

Step 4 In the displayed Add Networking page, select the VMkernel Network Adapter option in the 1 Select connection type.

Figure 5-24 Selecting the connection type

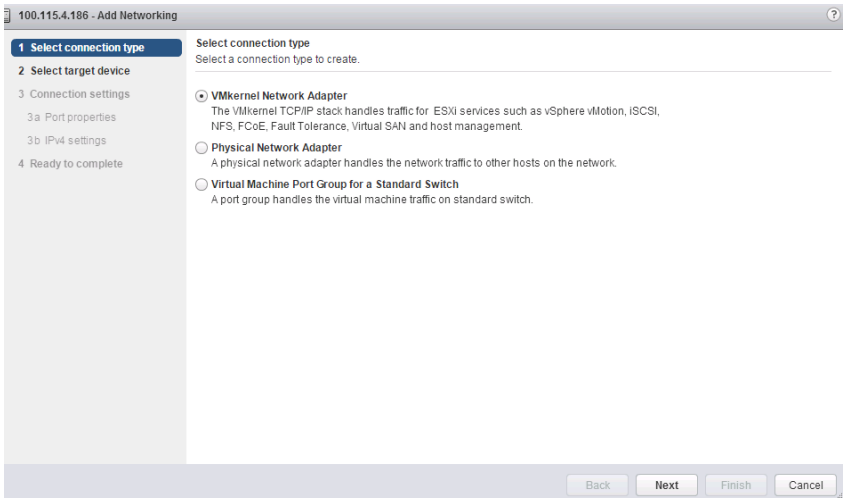

Step 5 In 2 Select target device, select the New standard switch option and click Next.

Figure 5-25 Selecting the target device

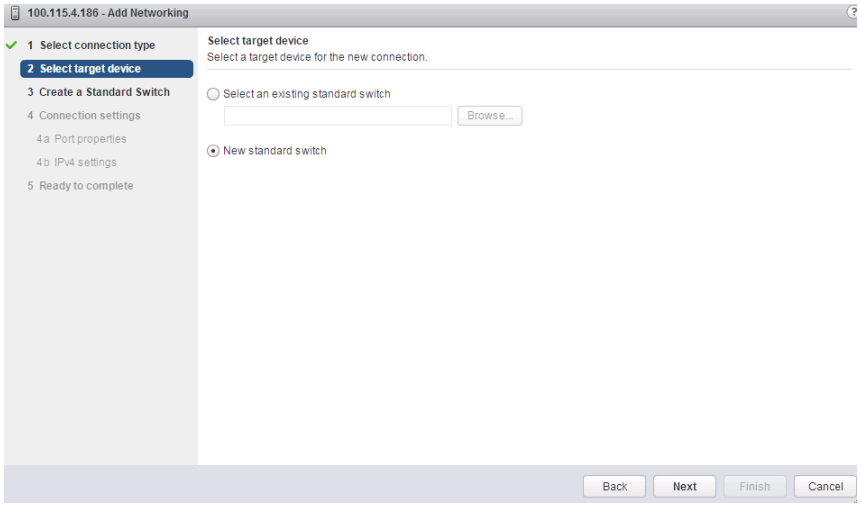

Step 6 In **3 Create a Standard Switch**, add physical adapters and click **Next**.

Figure 5-26 Adding a physical adapter

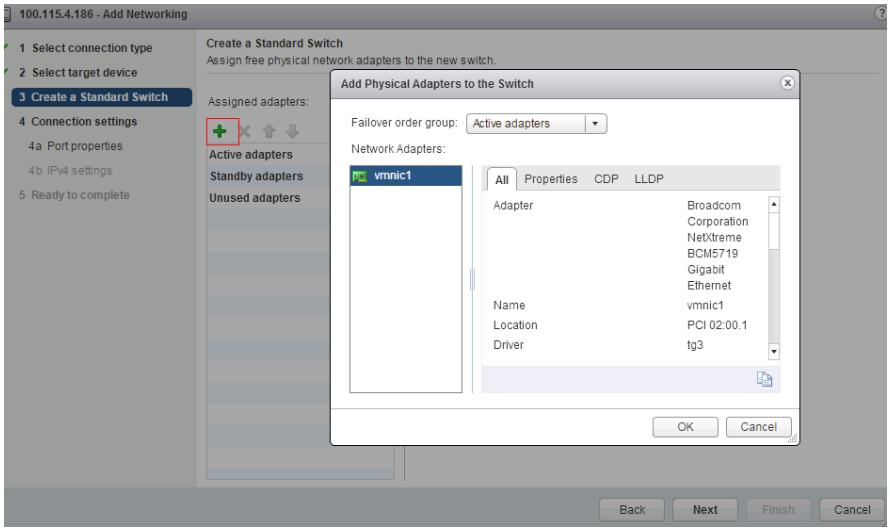

Step 7 Set port properties and click **Next**.

Figure 5-27 Setting port properties

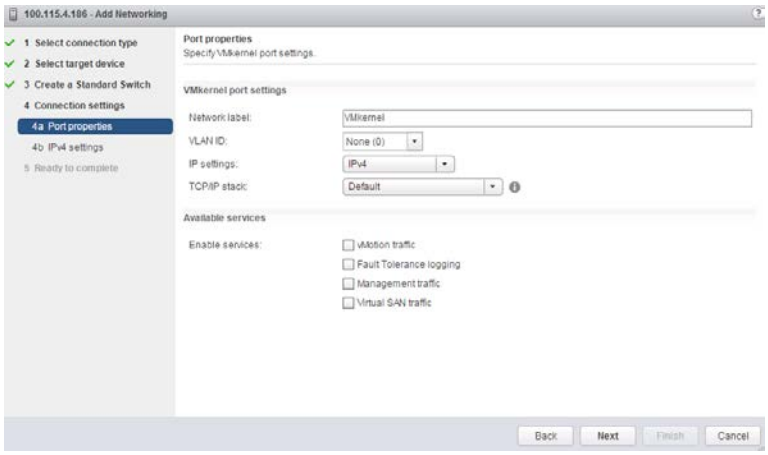

Step 8 Specify the service IP address and click **Next**.

Figure 5-28 Specifying the service IP address

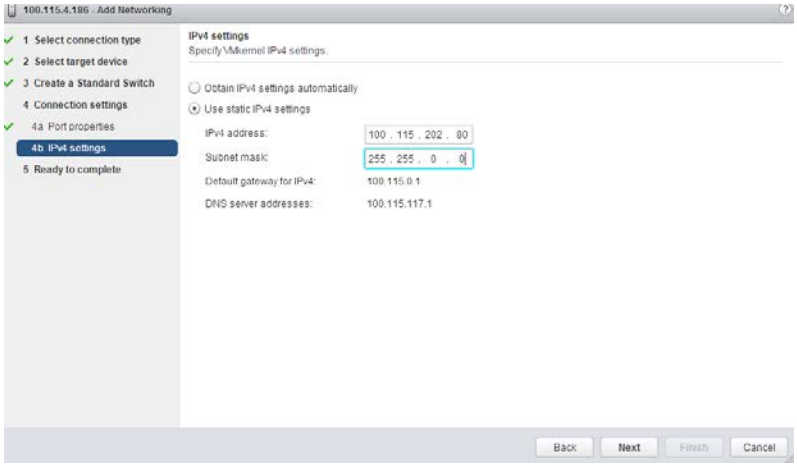

Step 9 Confirm the information and click **Finish**.

Figure 5-29 Checking the settings

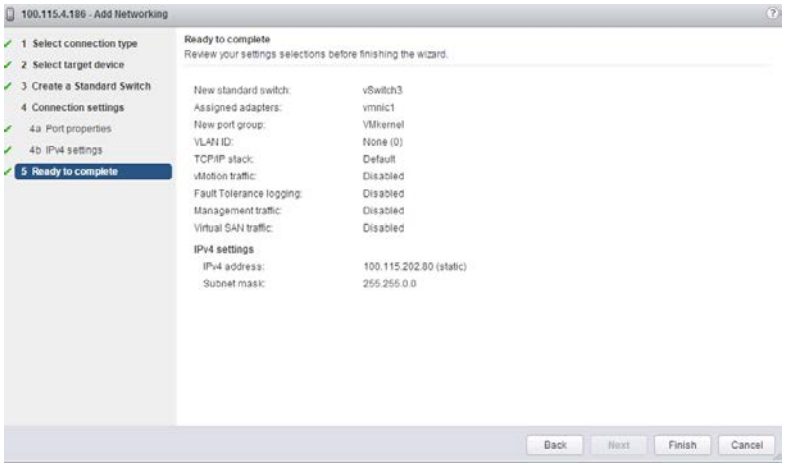

Step 10 Confirm the information that you have configured, as shown in Figure 5-30.

Figure 5-30 Information summary

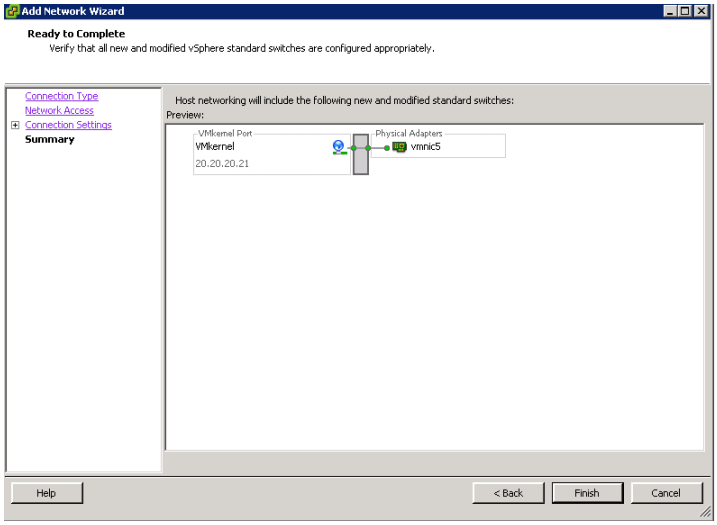

Step 11 If you only need to configure one path, the configuration is complete and you do not need to perform the next step. To configure multiple paths, proceed with the next step.

Step 12 Repeat the preceding steps to create another virtual network. Figure 5-31 shows a multi-path networking configuration.

Figure 5-31 Multi-path networking

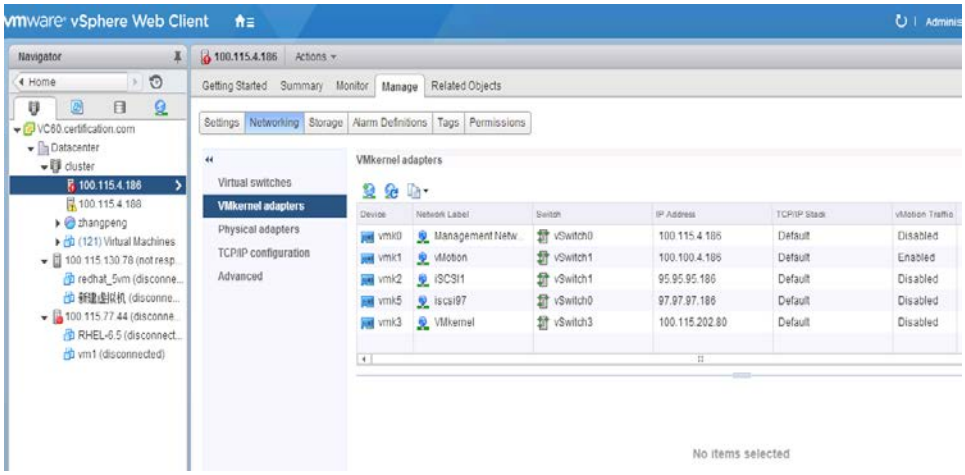

----End

5.2.1.2 Configuring Host Initiators

Host initiator configuration includes creating host initiators, binding initiators to virtual networks created in section 5.2.1.1 Configuring Service IP Addresses, and discovering targets.

In VMware ESX 4.1 and earlier versions, storage adapters have iSCSI adapters. You only need to enable those adapters. In VMware ESXi 5.0 and later versions, you need to manually add iSCSI initiators.

VMware ESXi 5.0

This section uses VMware ESXi 5.0 as an example to explain how to configure host initiators.

Step 1 Choose **Storage Adapters** and right-click the function pane, as shown in Figure 5-32

Figure 5-32 Adding storage adapters

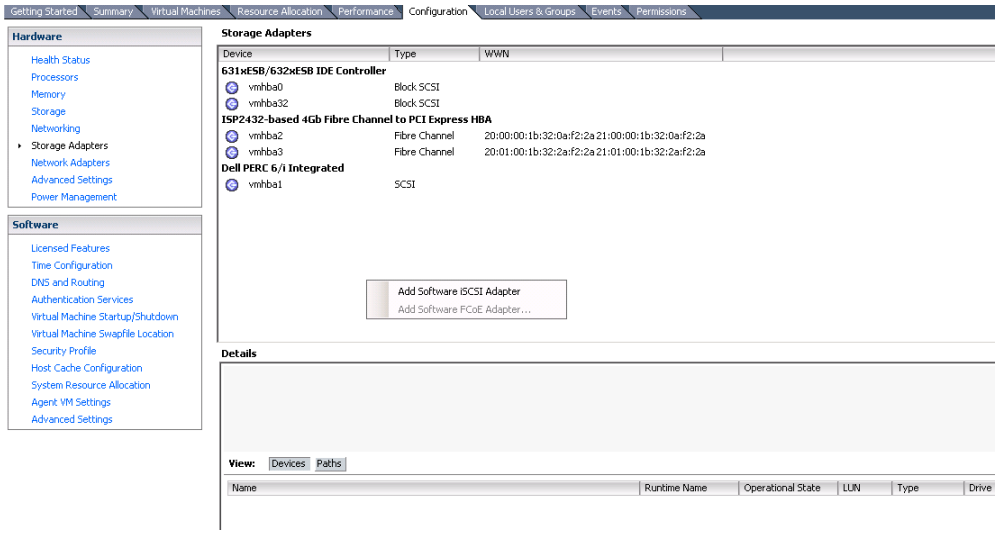

Step 2 Choose **Add Software iSCSI Adapter** from the shortcut menu. On the dialog box that is displayed, click **OK**, as shown in Figure 5-33

Figure 5-33 Adding iSCSI initiators

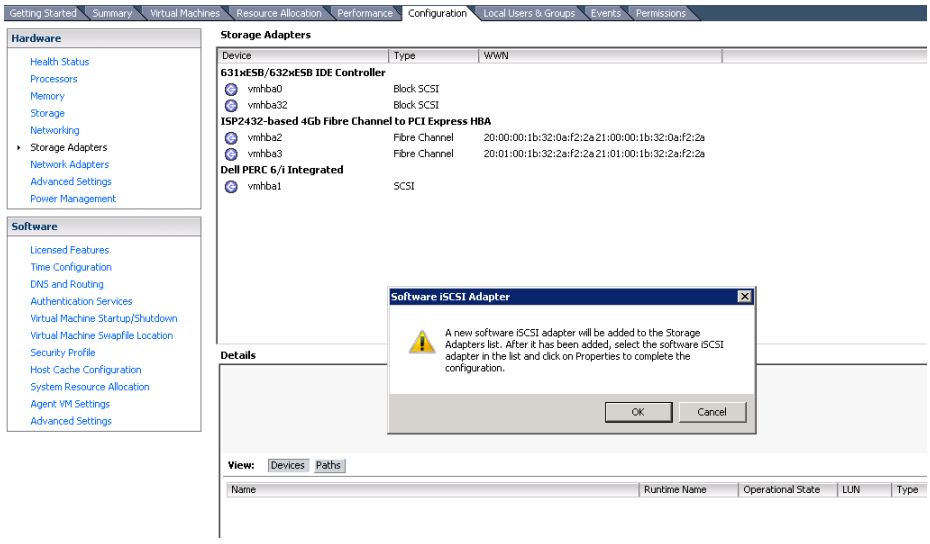

The newly added iSCSI initiators are displayed, as shown in Figure 5-34

Figure 5-34 iSCSI Software Adapter

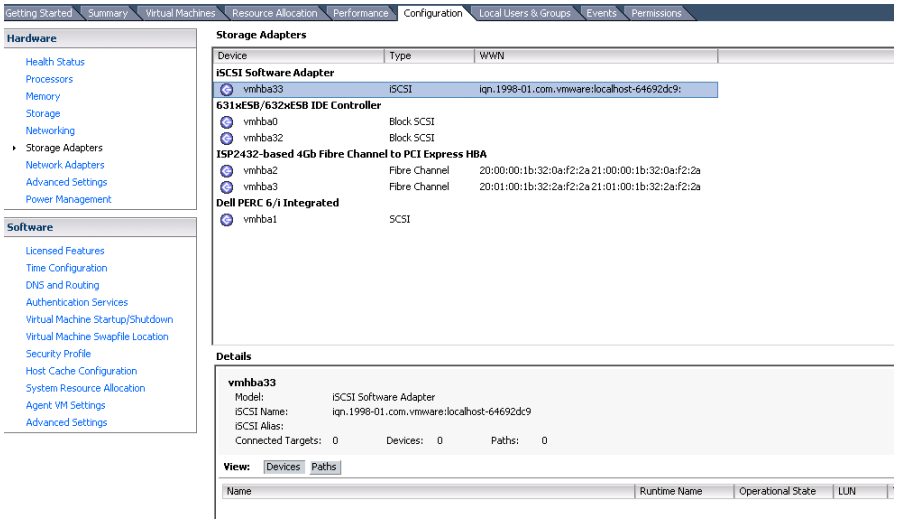

Step 3 Right-click a newly created iSCSI initiator and choose **Properties** from the shortcut menu, as shown in Figure 5-35

Figure 5-35 Initiator properties

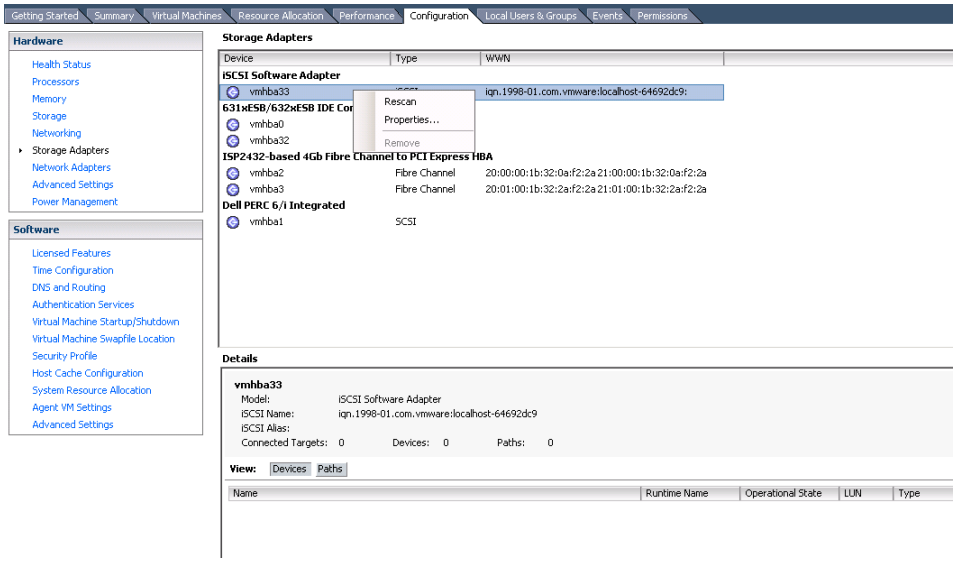

Step 4 On the dialog box that is displayed, click the **Network Configuration** tab and click **Add**, as shown in Figure 5-36

Figure 5-36 iSCSI initiator properties

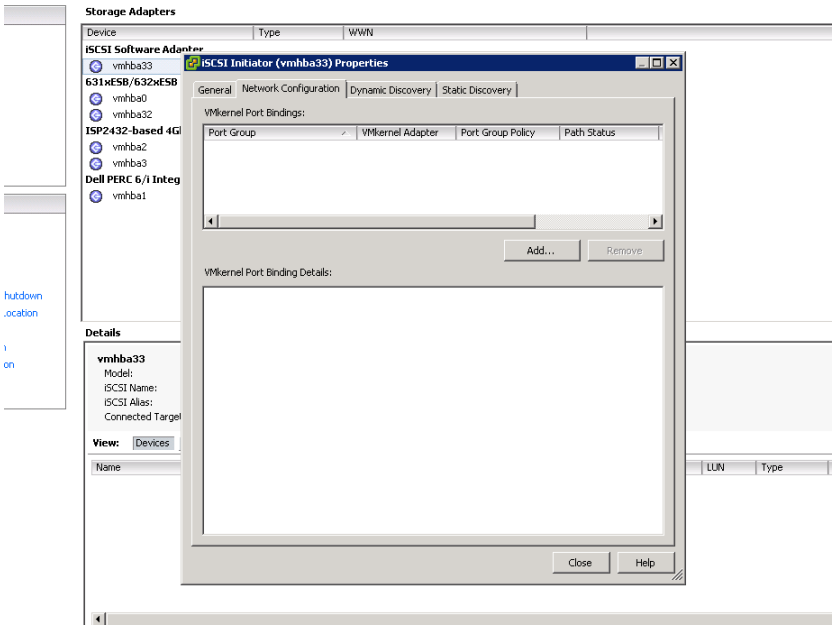

Step 5 Select a virtual network that you have created in section 5.2.1.1 Configuring Service IP Addresses and click **OK**, as shown in Figure 5-37.

Figure 5-37 Binding with a new VMkernel network adapter

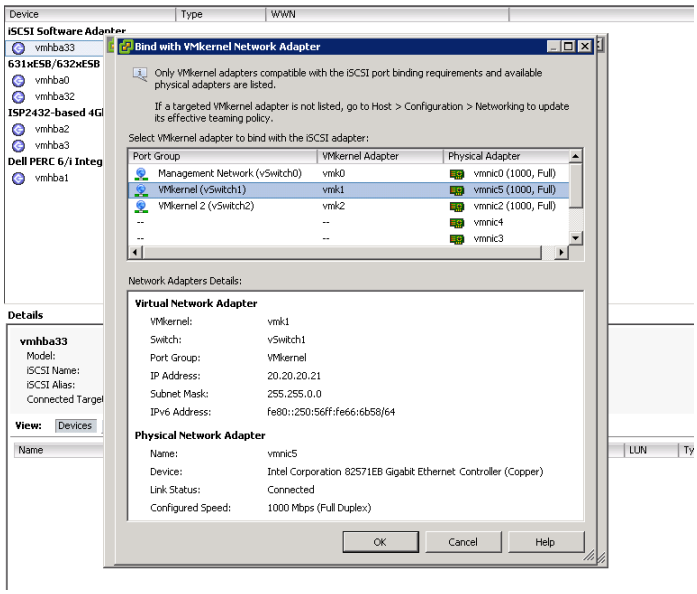

Figure 5-38 shows the properties of an initiator bound to the virtual network.

Figure 5-38 Initiator properties after virtual network binding

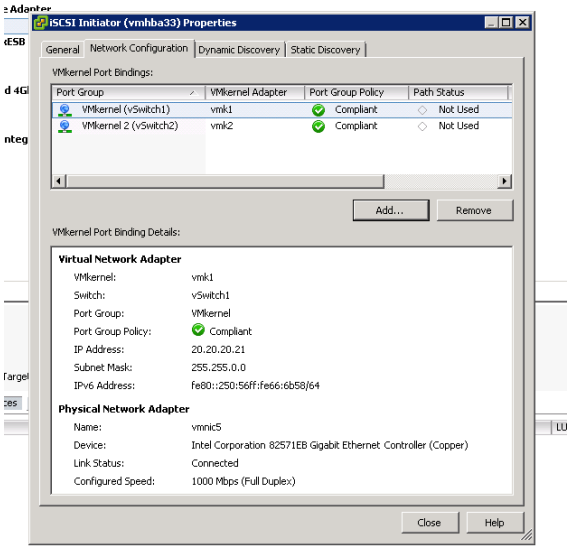

Step 6 In the dialog box for configuring initiator properties, click the **Dynamic Discovery** tab, click **Add**, and enter the target IP address (service IP address of the storage system), as shown in Figure 5-39

Figure 5-39 Adding send target server

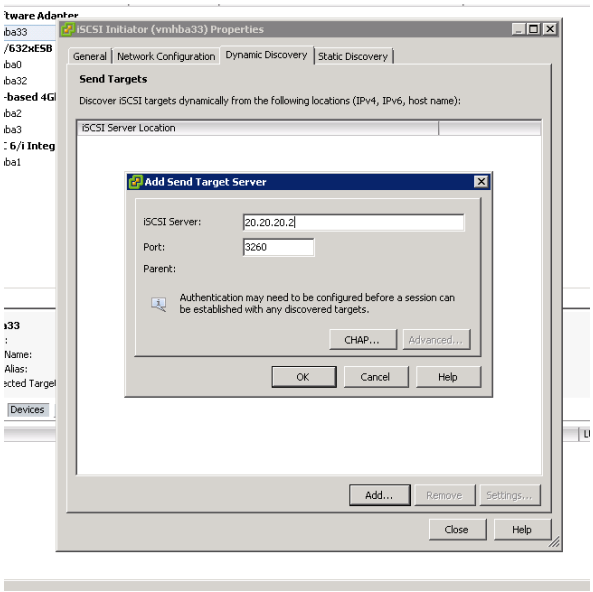

----End

vSphere Web Client

On vSphere Web Client, perform the following steps to configure the host initiator:

Step 1 On vSphere Web Client, click the **Manage** tab and then the **Storage** tab to check the storage adapter.

Figure 5-40 Checking the storage adapter

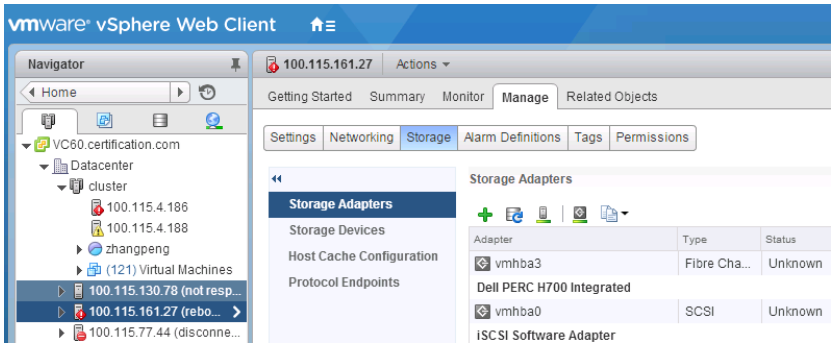

Step 2 Add a storage adapter. In the displayed **Add Software iSCSI Adapter** dialog box, click **OK**.

Figure 5-41 Adding a storage adapter

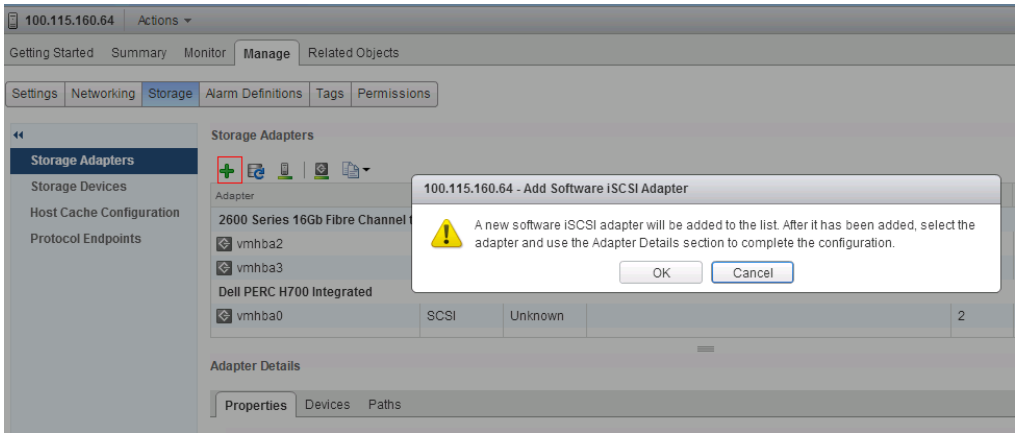

Step 3 Check the created iSCSI adapter.

Figure 5-42 Checking the created iSCSI adapter

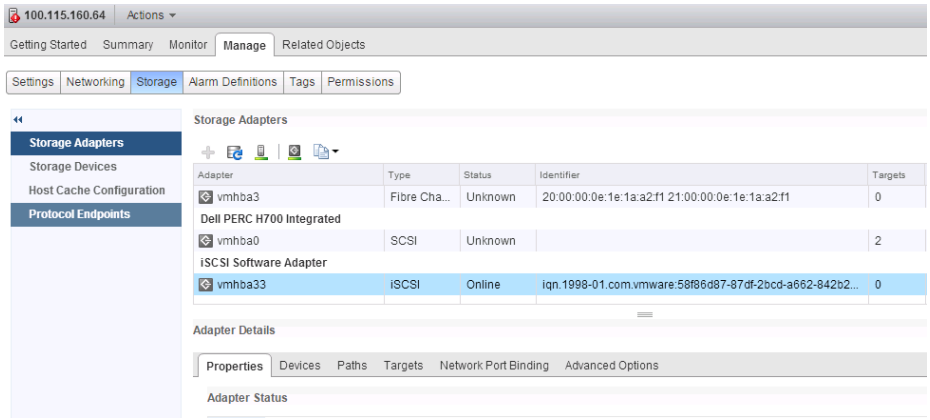

Step 4 In the **Adapter Details** area, click the **Network Port Binding** tab and click the **+** icon.

Figure 5-43 Initiator properties

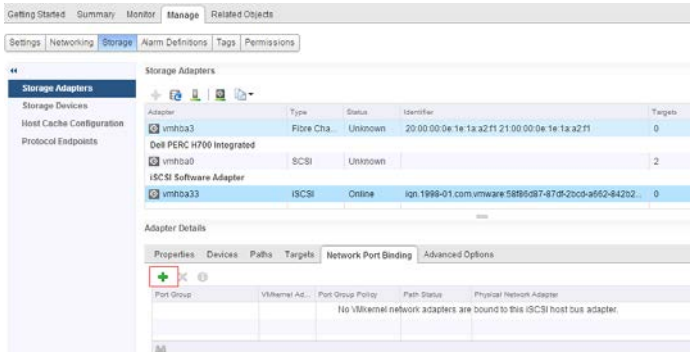

Step 5 Select a virtual network and bind it to the initiator.

Figure 5-44 Binding a virtual network to the initiator

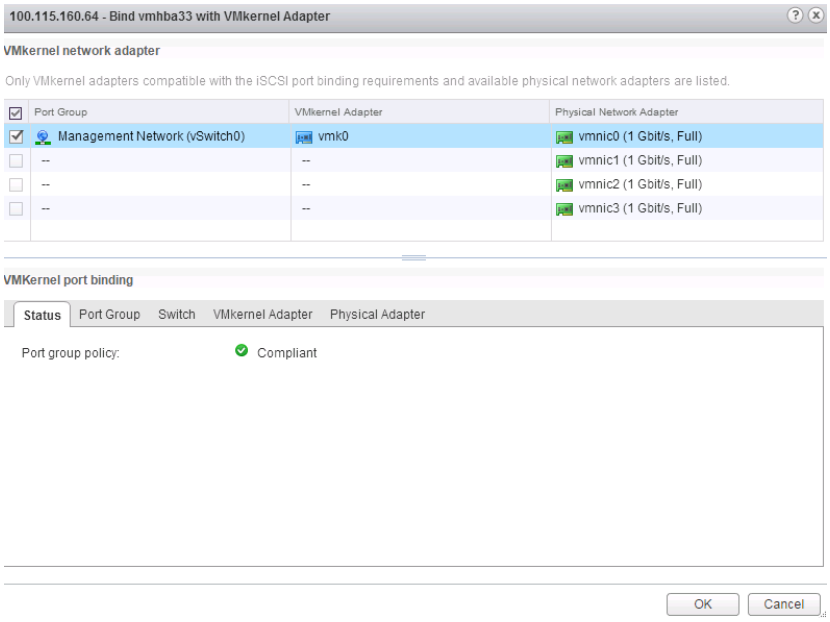

After the binding, the adapter properties are shown as follows:

Figure 5-45 After VMkernal port binding

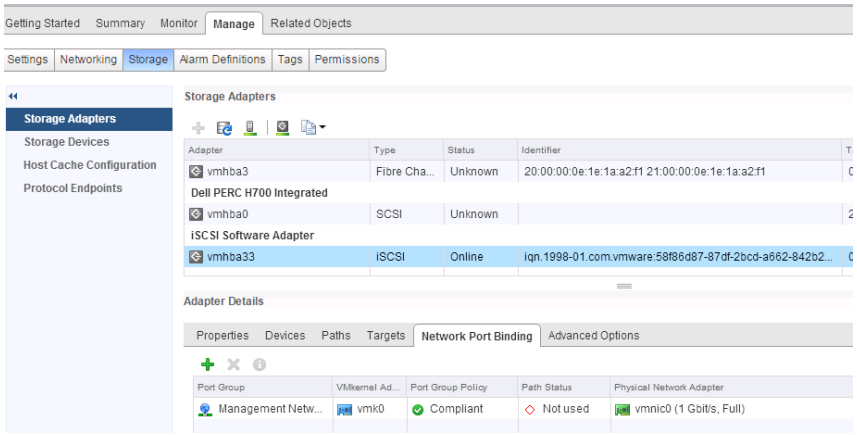

Step 6 In the **Adapter Details** area, click the **Targets** tab. Click the **Dynamic Discovery** button and click **Add**.

Figure 5-46 Dynamic discovery

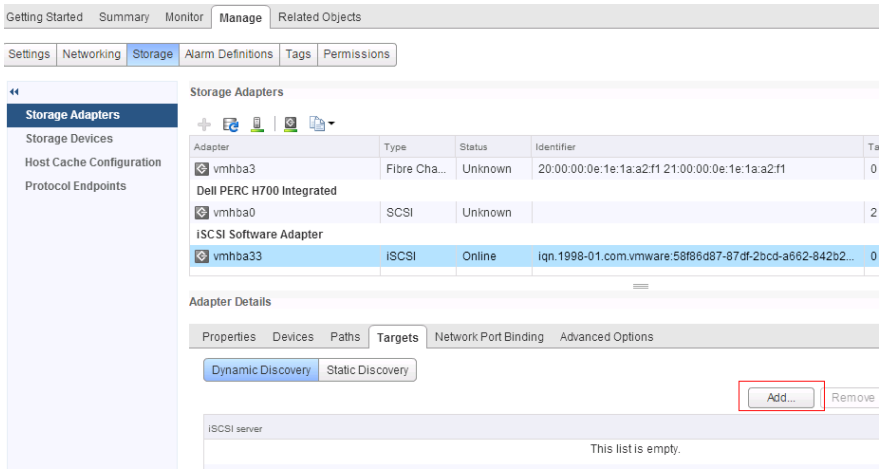

Step 7 Enter the target's IP address (storage's service IP address) and click **OK**.

Figure 5-47 Adding a target

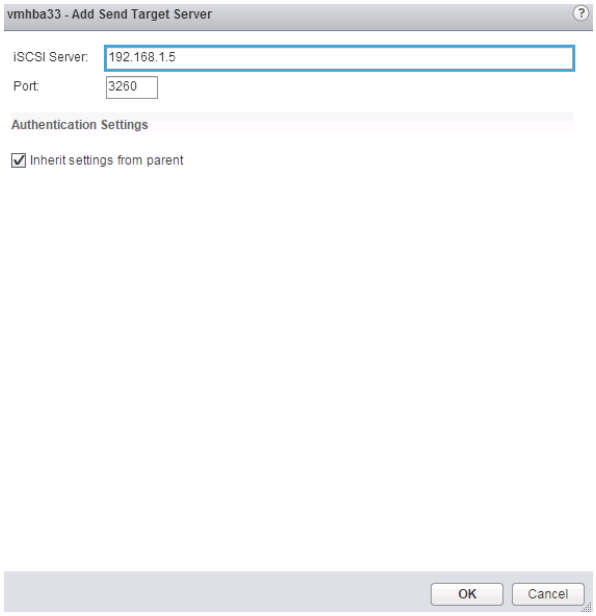

The host initiator configuration is complete.

----End

5.2.1.3 (Optional) Configuring CHAP Authentication

If Challenge Handshake Authentication Protocol (CHAP) authentication is required between a storage system and a host, perform the following steps to configure CHAP authentication.

Prerequisites

- The initiator has been added to the virtual host.
- CHAP has been enabled for the initiator on the storage system.
- No LUN has been mapped to the virtual host that corresponds to the application server.

vSphere Client

On vSphere Client, perform the following steps to configure CHAP authentication:

Step 1 In the dialog box for configuring iSCSI initiator properties, click the **General** tab and click **CHAP…** in the left lower corner, as shown in Figure 5-48.

Figure 5-48 General tab page

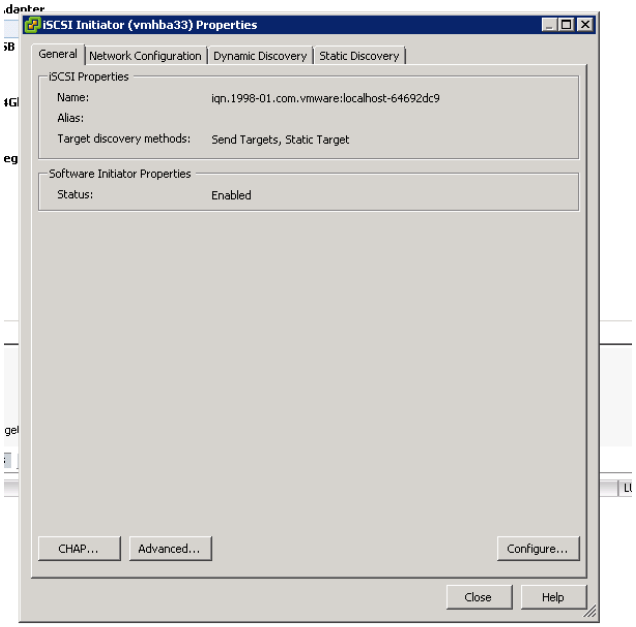

Step 2 In the **CHAP Credentials** dialog box that is displayed, choose **Use CHAP** from the **Select option** drop-down list.

Step 3 Enter the CHAP user name and password configured on the storage system, as shown in Figure 5-49.

Figure 5-49 CHAP credentials dialog box

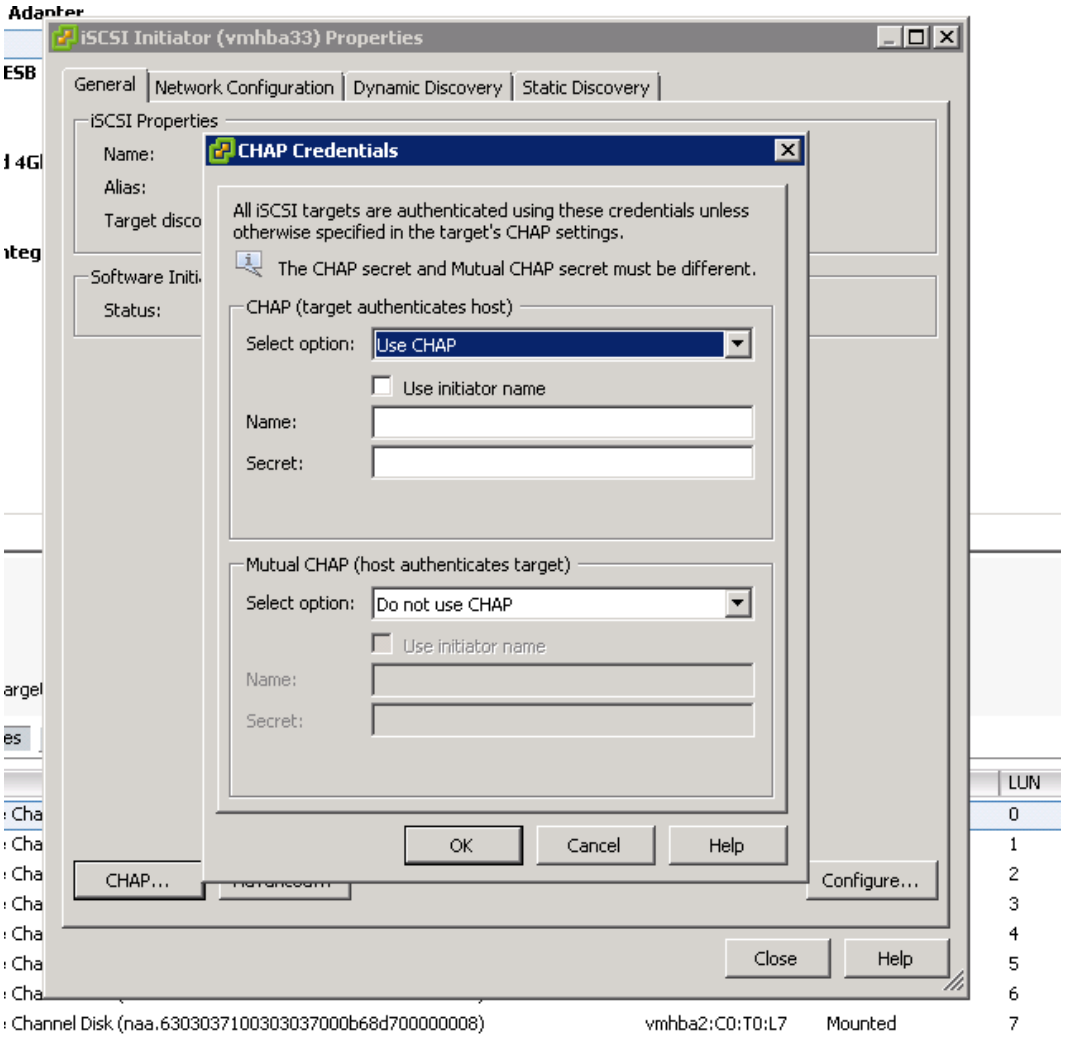

Step 4 Click **OK**.

----End

vSphere Web Client

On vSphere Web Client, perform the following steps to configure CHAP authentication:

Step 1 In the **Adapter Details** area, click the **Properties** tab. On the tab page, click **Edit** following **Authentication**.

Figure 5-50 Editing initiator authentication parameter settings

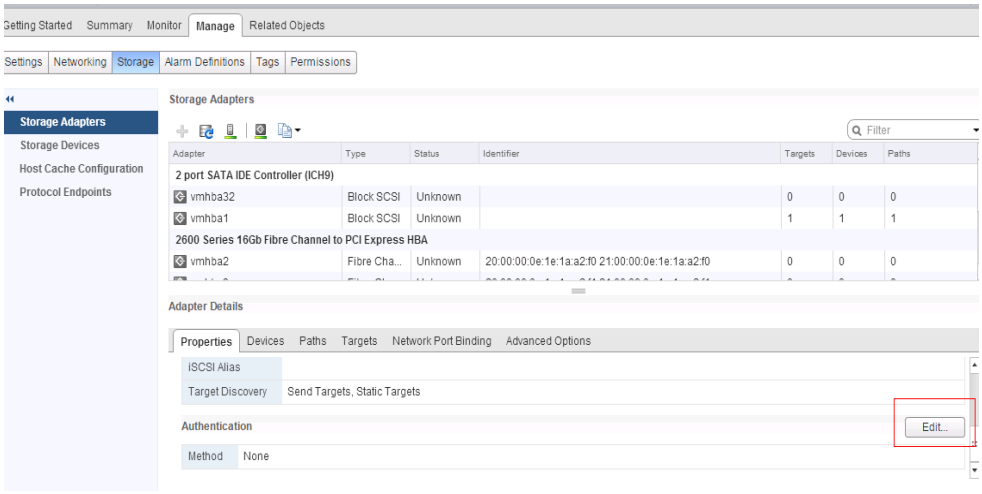

Step 2 In the displayed **Edit Authentication** dialog box, select **Use Unidirectional CHAP** (for example) as **Authentication Method**.

Figure 5-51 Selecting an authentication method

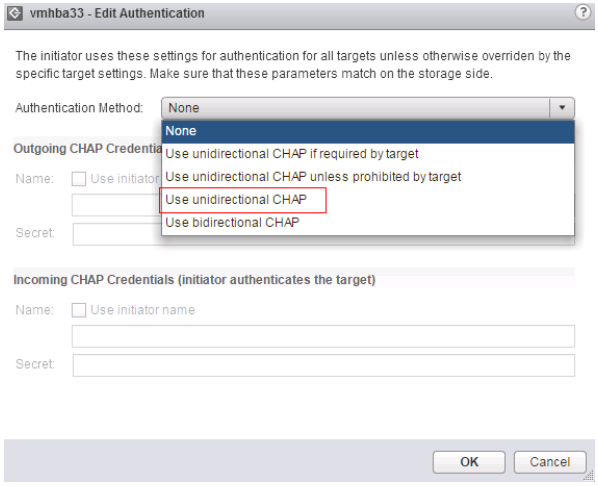

Step 3 Enter the storage system's CHAP name and secret, and click **OK**.

Figure 5-52 Setting CHAP authentication parameters

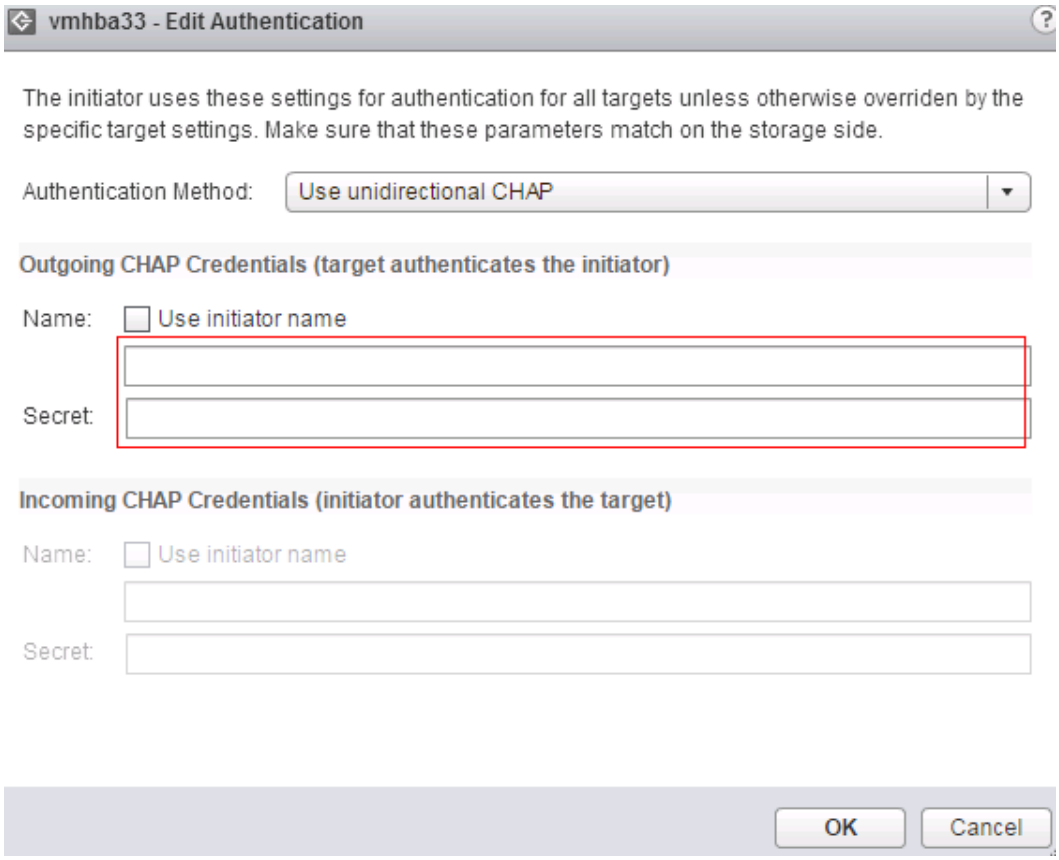

----End

5.2.2 (Optional) Switch Configuration

This section describes how to configure Ethernet switches, including configuring VLANs and binding ports. Skip this section if you use direct connections.

Configuring VLANs

On an Ethernet network to which many hosts are connected, a large number of broadcast packets are generated during the host communication. Broadcast packets sent from one host will be received by all other hosts on the network, consuming more bandwidth. Moreover, all hosts on the network can access each other, resulting data security risks.

To save bandwidth and prevent security risks, hosts on an Ethernet network are divided into multiple logical groups. Each logical group is a VLAN. The following uses Huawei Quidway 2700 Ethernet switch as an example to explain how to configure VLANs.

In the following example, two VLANs (VLAN 1000 and VLAN 2000) are created. VLAN 1000 contains ports GE 1/0/1 to 1/0/16. VLAN 2000 contains ports GE 1/0/20 to 1/0/24.

Step 1 Go to the system view.

```
<Quidway>system-view
System View: return to User View with Ctrl+Z.
```

Step 2 Create VLAN 1000 and add ports to it.

```
[Quidway]VLAN 1000
[Quidway-vlan1000]port GigabitEthernet 1/0/1 to GigabitEthernet 1/0/16
```

Step 3 Configure an IP address for VLAN 1000.

```
[Quidway-vlan1000]interface VLAN 1000
[Quidway-Vlan-interface1000]ip address 1.0.0.1 255.255.255.0
```

Step 4 Create VLAN 2000, add ports, and configure an IP address.

```
[Quidway]VLAN 2000
[Quidway-vlan2000]port GigabitEthernet 1/0/20 to GigabitEthernet 1/0/24
[Quidway-vlan2000]interface VLAN 2000
[Quidway-Vlan-interface2000]ip address 2.0.0.1 255.255.255.0
```

----End

Binding Ports

When storage systems and hosts are connected in point-to-point mode, existing bandwidth may be insufficient for storage data transmission. Moreover, devices cannot be redundantly connected in point-to-point mode. To address these problems, ports are bound (link aggregation) to improve bandwidth and balance load among multiple links.

Three Ethernet link aggregation modes are available:

- Manual aggregation

  Ports are added to an aggregation group by running a command manually. Ports added to the aggregation group must have the same link type.
- Static aggregation

  Ports are added to an aggregation group by running a command manually. Ports added to the aggregation group must have the same link type and LACP enabled.
- Dynamic aggregation

  The protocol dynamically adds ports to an aggregation group. Ports added in this way must have LACP enabled and the same speed, duplex mode, and link type.

Table 5-2 compares these aggregation modes.

Table 5-2 Comparison among link aggregation modes

Link Aggregation Mode	Packet Exchange	Port Detection	CPU Usage
Manual aggregation	No	No	Low
Static aggregation	Yes	Yes	High
Dynamic aggregation	Yes	Yes	High

Huawei OceanStor storage devices support 802.3ad link aggregation (dynamic aggregation). In this link aggregation mode, multiple network ports are in an active aggregation group and work in duplex mode and at the same speed. After binding iSCSI host ports on a storage

device, enable aggregation for their peer ports on the switch. Otherwise, links are unavailable between the storage device and the switch.

This section uses switch ports GE 1/0/1 and GE 1/0/2 and the storage system's ports P2 and P3 as an example to explain how to bind ports.

The port binding method differs with the OceanStor system version. For details, refer to the specific storage product documentation. The following steps use OceanStor V3 V300R003 as an example.

Step 1 Log in to DeviceManager and choose **Provisioning** > **Port**.

Step 2 Bind ports.

1. Select the ports that you want to bind and choose **More** > **Bond Ports**.

   The **Bond Port** dialog box is displayed.
2. Enter a **Bond Name**, select the target ports, and click **OK**.
3. In the security alert dialog box that is displayed, select **I have read and understand the consequences associated with performing this operation** and click **OK**.

After the storage system ports are bound, configure link aggregation on the switch using the following command:

```
<Quidway>system-view
System View: return to User View with Ctrl+Z.
[Quidway-Switch]interface GigabitEthernet 1/0/1
[Quidway-Switch-GigabitEthernet1/0/1]lacp enable
LACP is already enabled on the port!
[Quidway-Switch-GigabitEthernet1/0/1]quit
[Quidway-Switch]interface GigabitEthernet 1/0/2
[Quidway-Switch-GigabitEthernet1/0/2]lacp enable
LACP is already enabled on the port!
[Quidway-Switch-GigabitEthernet1/0/2]quit
```

After the command is executed, LACP is enabled for ports GE 1/0/1 and GE 1/0/2. Then the ports can be automatically detected and added to an aggregation group.

----End

5.2.3 Storage System Configuration

The supported IP protocol stack differs with the storage system version. You need to select the desired IP protocol based on your storage system version and project situations.

Configuring Port IP Addresses

Configure Ethernet port parameters to ensure proper communication between the storage system and application server.

Note the following items when setting the properties of an Ethernet port:

- The default internal heartbeat IP addresses of a two-controller storage system are **127.127.127.10** and **127.127.127.11**, and those of a four-controller storage system are **127.127.127.10**, **127.127.127.11**, **127.127.127.12**, and **127.127.127.13**. Therefore, the IP address of the router must not be in the 127.127.127.XXX segment and the gateway address must not be **127.127.127.10**, **127.127.127.11**, **127.127.127.12**, or **127.127.127.13**. Otherwise, routing will fail. Internal heartbeat links are established between controllers

for these controllers to detect each other's working status. You do not need to separately connect cables. In addition, internal heartbeat IP addresses have been assigned before delivery, and you cannot change these IP addresses.

- The IP address of the Ethernet port cannot be in the same network segment as that of a management network port.
- The IP address of the Ethernet port cannot be in the same network segment as that of a maintenance network port.
- If the Ethernet port connects to an application server, the IP address of the Ethernet port must be in the same network segment as that of the service network port on the application server. If the Ethernet port connects to another storage device, the IP address of the Ethernet port must be in the same network segment as that of the Ethernet port on the other storage device. Add routes if available IP addresses in the desired segment are insufficient.

CAUTION

Changing the IP address of a host port may interrupt services on this host port.

Configurations vary on different storage systems. The following uses OceanStor V3/Dorado V3 as examples to describe how to configure an IP address for an Ethernet port:

Step 1 Go to the **Ethernet Port** dialog box.

1. In the basic information area of the function pane, click the device icon.
2. In the middle function pane, click the cabinet whose Ethernet ports you want to view.
3. Click the controller enclosure where the desired Ethernet ports reside. The controller enclosure view is displayed.
4. Click  to switch to the rear view.
5. Click the Ethernet port whose information you want to modify.

   The **Ethernet Port** dialog box is displayed.
6. Click **Modify**.

Step 2 Modify the Ethernet port, as shown in Figure 5-53.

1. In **IPv4 Address** or **IPv6 Address**, enter an IP address for the Ethernet port.
2. In **Subnet Mask** or **Prefix**, enter a subnet mask or prefix for the Ethernet port.
3. In **MTU (Byte)**, enter the maximum size of data packet that can be transferred between the Ethernet port and the host. The value is an integer ranging from 1500 to 9216.

Figure 5-53 Configuring an IP address

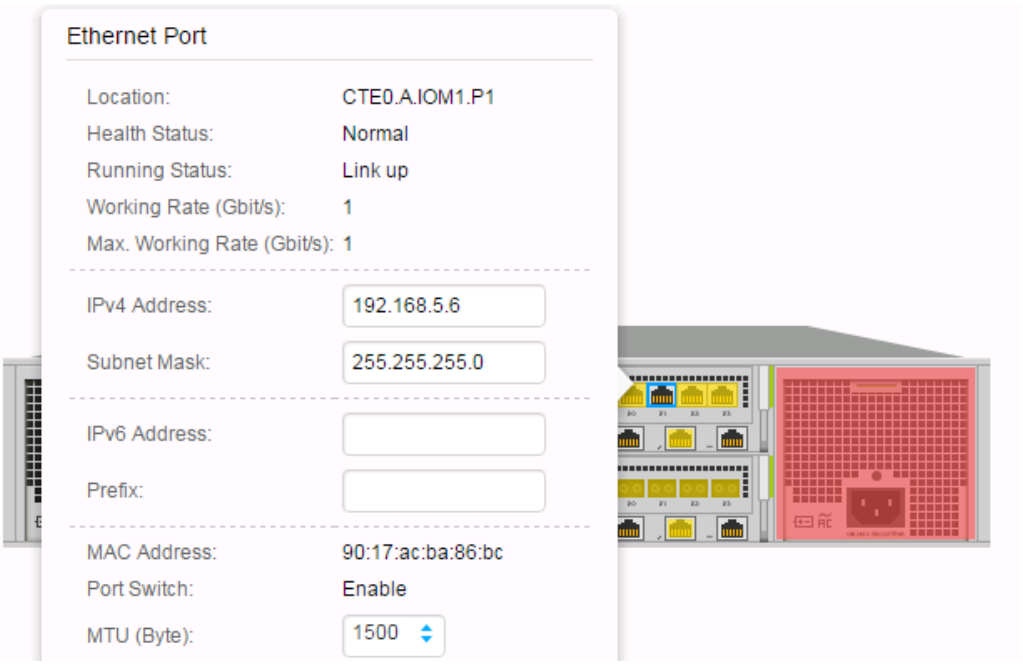

Step 3 Confirm the Ethernet port modification.

1. Click **Apply**.

   The **Danger** dialog box is displayed.
2. Confirm the information in the dialog box and select **I have read and understand the consequences associated with performing this operation.**
3. Click **OK**.

   The **Success** dialog box is displayed, indicating that the operation is successful.
4. Click **OK**.

----End

(Optional) Adding Routes

If iSCSI networking is used and data needs to be transmitted across network segments, you need to configure routes.

Step 1 Log in to DeviceManager.

Step 2 Choose **Provisioning** > **Port** > **Ethernet Ports**.

Step 3 Select the Ethernet port for which you want to add a route and click **Route Management**.

The **Route Management** dialog box is displayed.

Step 4 Configure the route information for the Ethernet port.

1. In **IP Address**, select the IP address of the Ethernet port.
2. Click **Add**.

   The **Add Route** dialog box is displayed.

3. In **Type**, select the type of the route to be added.

   There are three route options:

   - Default route

     Data is forwarded through this route by default if no preferred route is available. The destination address field and the target mask field (IPv4) or prefix (IPv6) of the default route are automatically set to 0. To use this option, you only need to add a gateway.

   - Host route

     A route to an individual host. The destination mask (IPv4: 255.255.255.255) or prefix (IPv6: 128) of the host route is automatically set. To use this option, add the destination address and a gateway.

   - Network segment route

     A route to a network segment. You need to add the destination address, destination address mask (IPv4) or prefix (IPv6), and gateway. For example, the destination address is 172.17.0.0, destination address mask is 255.255.0.0, and gateway is 172.16.0.1.

4. Set **Destination Address**.

   Set **Destination Address** to the IPv4 or IPv6 (depending on which one you use) address or network segment of the application server's service network port or that of the other storage system's Ethernet port.

5. Set **Destination Mask** (IPv4) or **Prefix** (IPv6).

   - If an IPv4 address is used, this parameter specifies the subnet mask of the IP address for the service network port on the application server or the other storage device.
   - If an IPv6 address is used, this parameter specifies the prefix of the IPv6 address for the application server's service network port or that of the other storage system's Ethernet port.

6. In **Gateway**, enter the gateway of the local storage system's Ethernet port IP address.

Step 5 Click **OK**. The route information is added to the route list.

A security alert dialog box is displayed.

Step 6 Confirm the information in the dialog box and select **I have read and understand the consequences associated with performing this operation.**

Step 7 Click **OK**.

The **Success** dialog box is displayed, indicating that the operation is successful.

NOTE
To remove a route, select it and click **Remove**.

Step 8 Click **Close**.

----End

Adding an Initiator to a Host

This section details how to add initiators to the hosts on the storage system. For other storage configurations, see the *Basic Storage Service Configuration Guide* corresponding to your storage system.

Step 1 Log in to DeviceManager.

Step 2 Choose **Provisioning** > **Host**.

Step 3 Select the target host and click **Add Initiator**.

Figure 5-54 Selecting a host

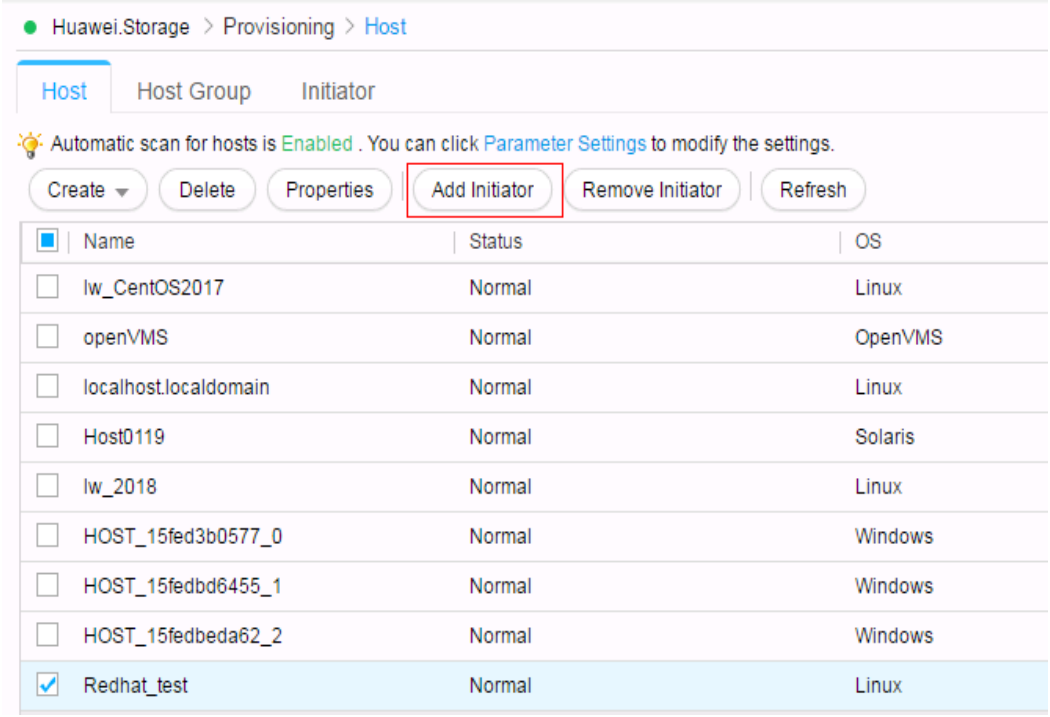

Step 4 Select the initiator and click to add it to **Selected Initiators**.

Figure 5-55 Adding an initiator

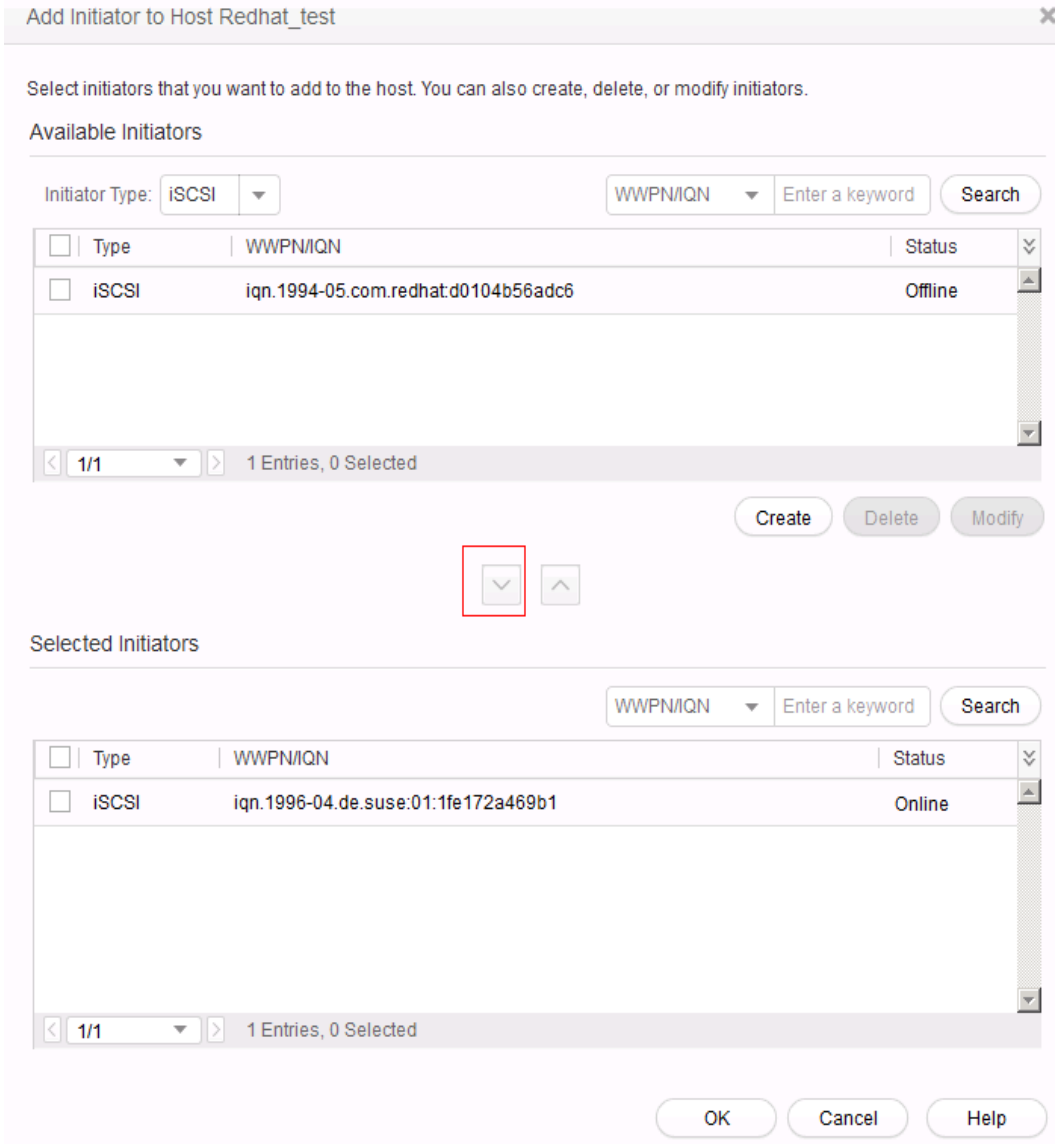

Step 5 Click **OK**.

----End

The initiator properties depend on the operating system and multipathing software used by the hosts. For details, see the storage-side configuration in the multipathing configuration section. After the initiators have been configured, you can scan for LUNs on the hosts to discover storage resources.

6 Mapping and Using LUNs

6.1 Scanning for LUNs on a Host

After LUNs are mapped on a storage system, scan for the mapped LUNs on the host.

- vSphere Client

Figure 6-1 Scanning for the mapped LUNs

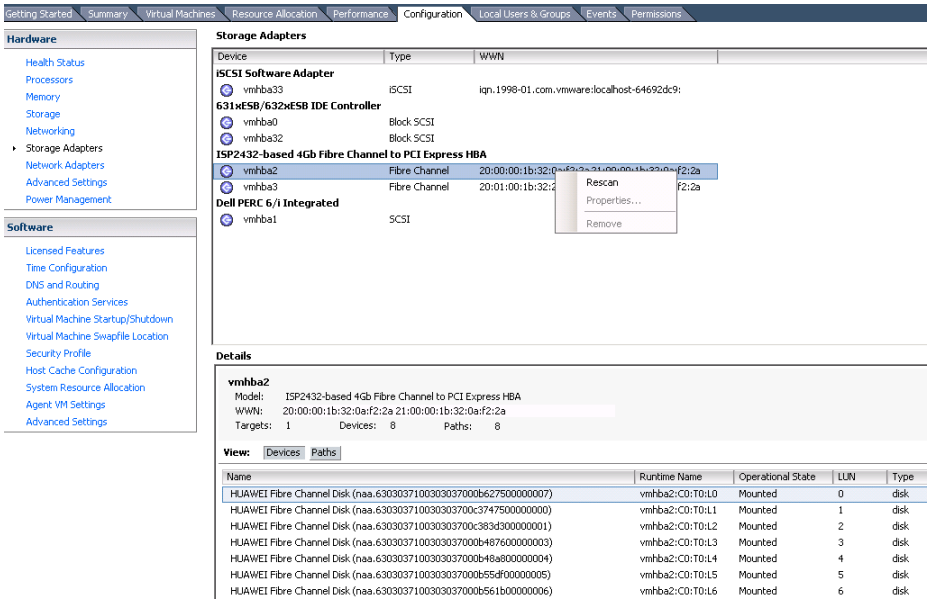

- vSphere Web Client

Figure 6-2 Scanning for the mapped LUNs (on vSphere Web Client)

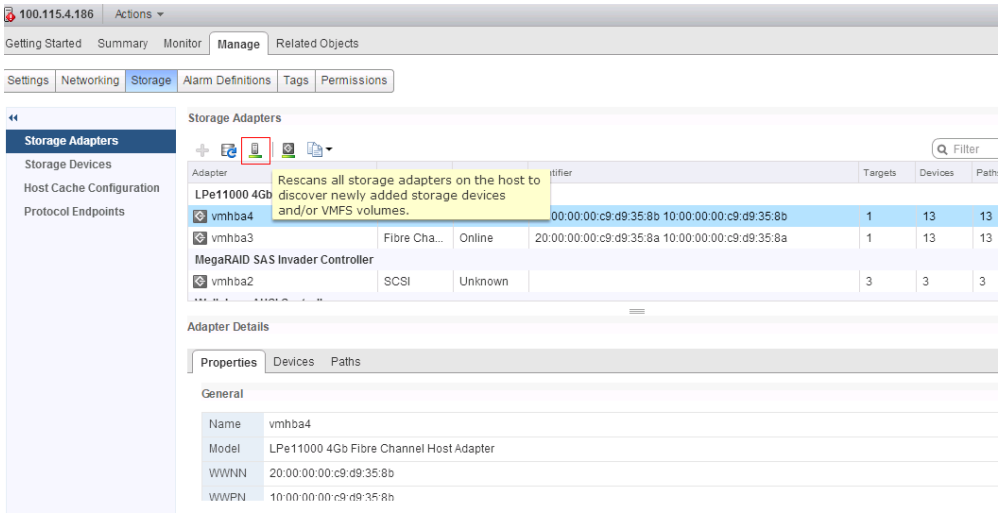

6.2 Using the Mapped LUNs

After the mapped LUNs are detected on a host, you can directly use the raw devices to configure services or use the LUNs after creating a file system.

6.2.1 Raw Device Mapping (RDM)

RDM uses raw devices as disks for VMs. Perform the following steps to map raw devices.

vSphere Client

On vSphere Client, perform the following steps to configure RDM:

Step 1 Right-click a VM and choose **Edit Settings** from the shortcut menu, as shown in Figure 6-3.

Figure 6-3 Editing host settings

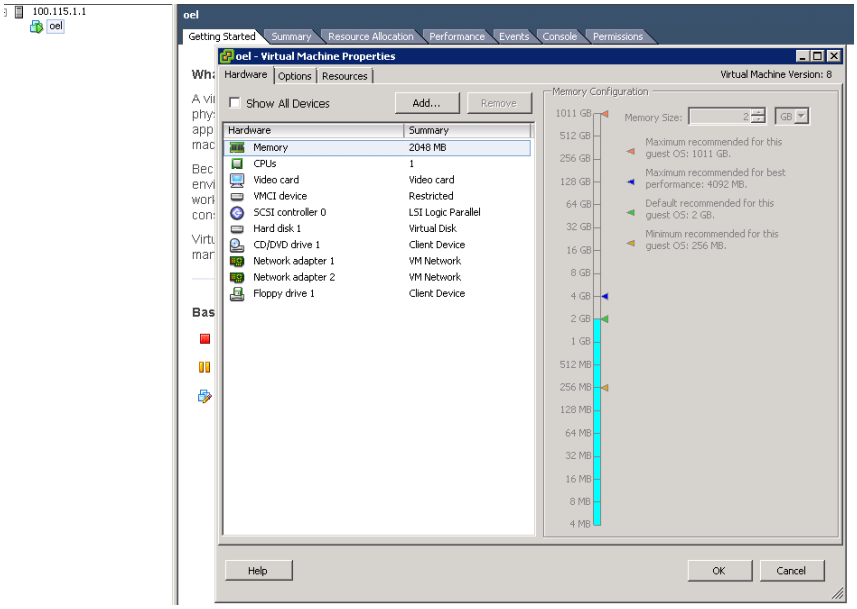

Step 2 On the **Hardware** tab page, click **Add**. In the **Add Hardware** dialog box that is displayed, choose **Hard Disk** in **Device Type** and click **Next**, as shown in Figure 6-4.

Figure 6-4 Adding disks

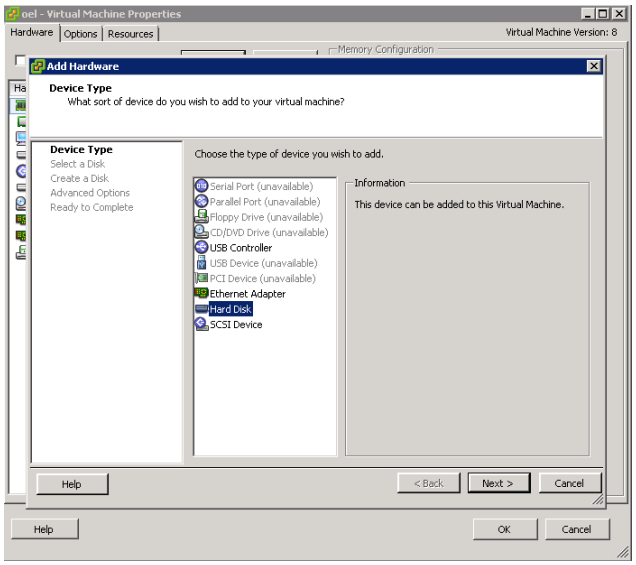

Step 3 Select disks.

You can create a new virtual disk, use an existing virtual disk, or use raw disk mappings, as shown in Figure 6-5.

Figure 6-5 Selecting disks

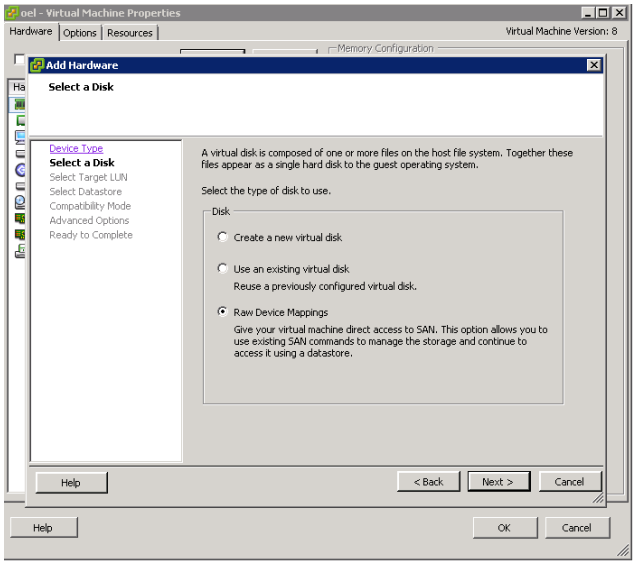

Step 4 Select **Raw Device Mappings** and click **Next**.

Step 5 Select a target LUN and click **Next**, as shown in Figure 6-6.

Figure 6-6 Selecting a target LUN

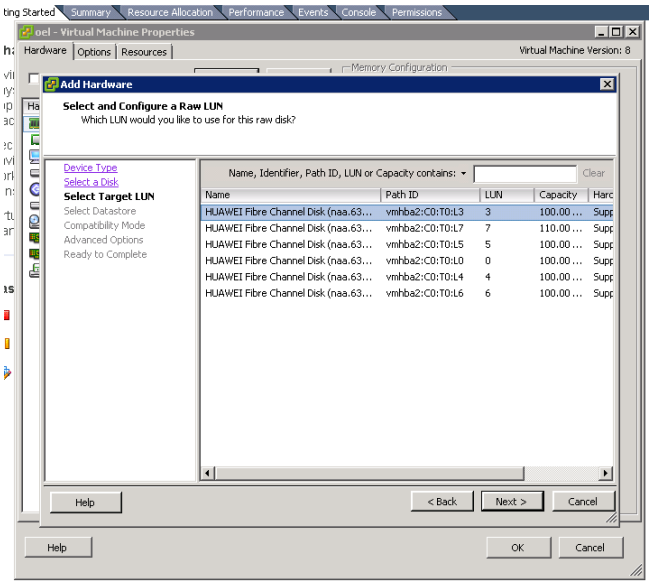

Step 6 Select a datastore. The default datastore is under the same directory as the VM storage. Click **Next**, as shown in Figure 6-7.

Figure 6-7 Selecting a datastore

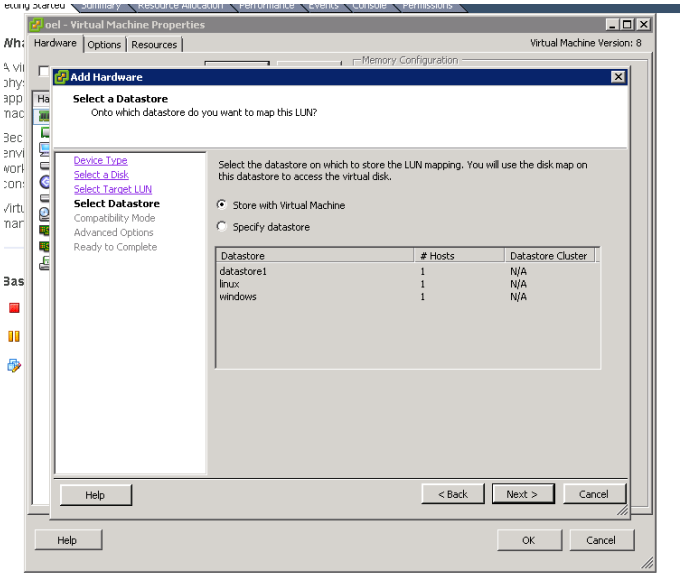

Step 7 Select a compatibility mode.

Select a compatibility mode based on site requirements and click **Next**, as shown in Figure 6-8.

Figure 6-8 Selecting a compatibility mode

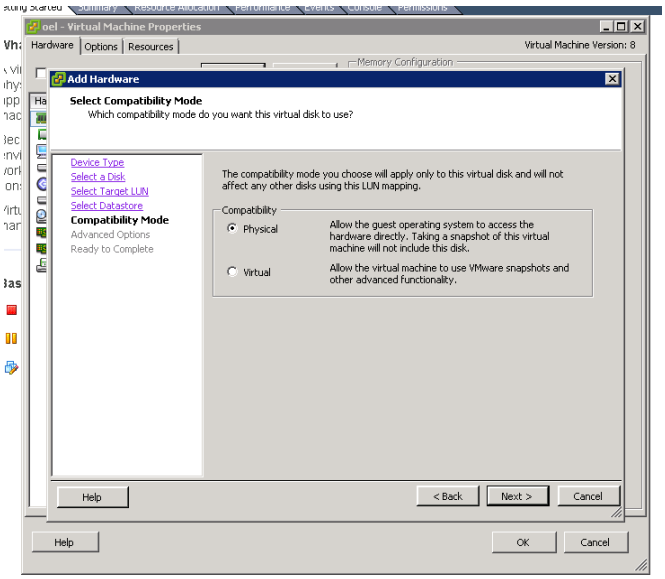

Step 8 In **Advanced Options**, keep the default virtual device node unchanged, as shown in Figure 6-9.

Figure 6-9 Selecting a virtual device node

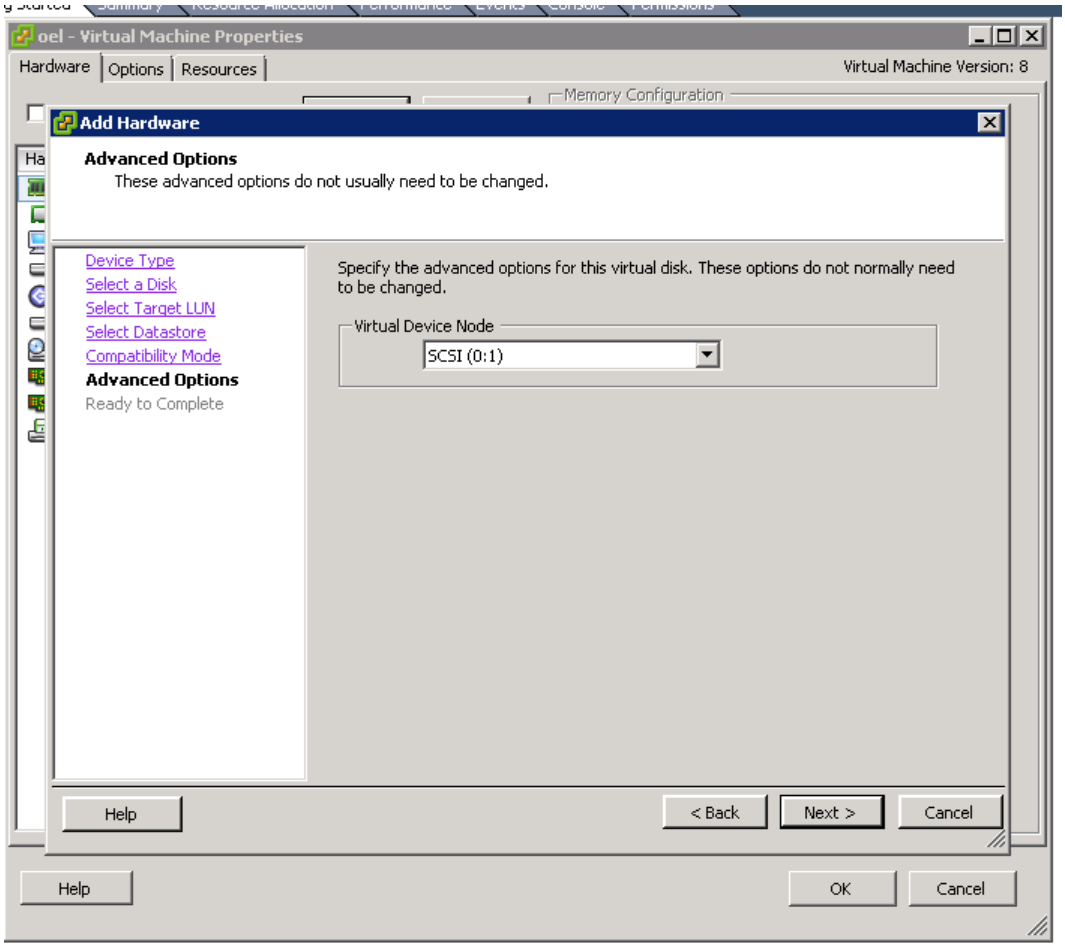

Step 9 In **Ready to Complete**, confirm the information about the disk to be added, as shown in Figure 6-10.

Figure 6-10 Confirming the information about the disk to be added

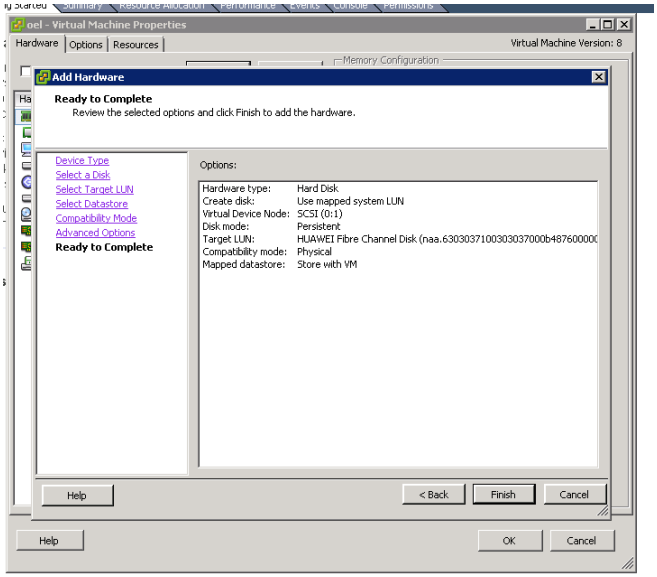

Step 10 Click **Finish**. The system starts to add disks, as shown in Figure 6-11.

----End

Figure 6-11 Adding raw disk mappings

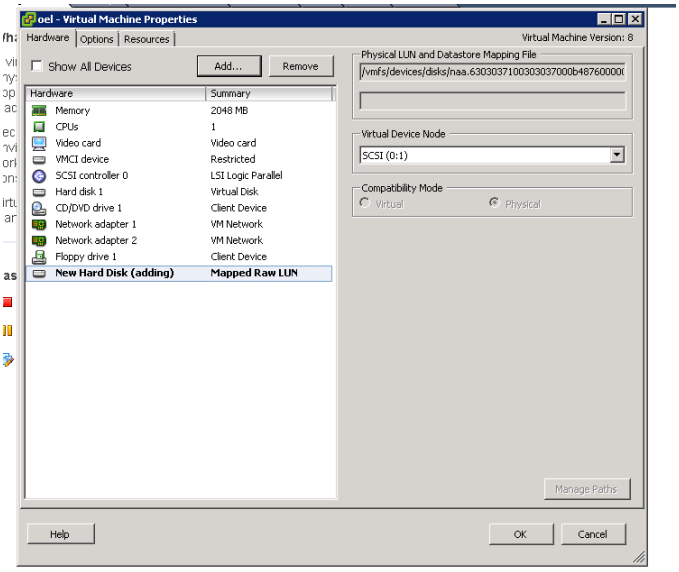

Step 1 After a raw disk is mapped, the type of the newly created disk is **Mapped Raw LUN**.

----End

vSphere Web Client

On vSphere Web Client, perform the following steps to configure RDM:

Step 1 On the **Related Objects** tab page, click the **Virtual Machines** tab. On the left pane of the tab page, select the to-be-added host, right-click it, and choose **Edit Settings** from the shortcut menu.

Figure 6-12 Editing host settings

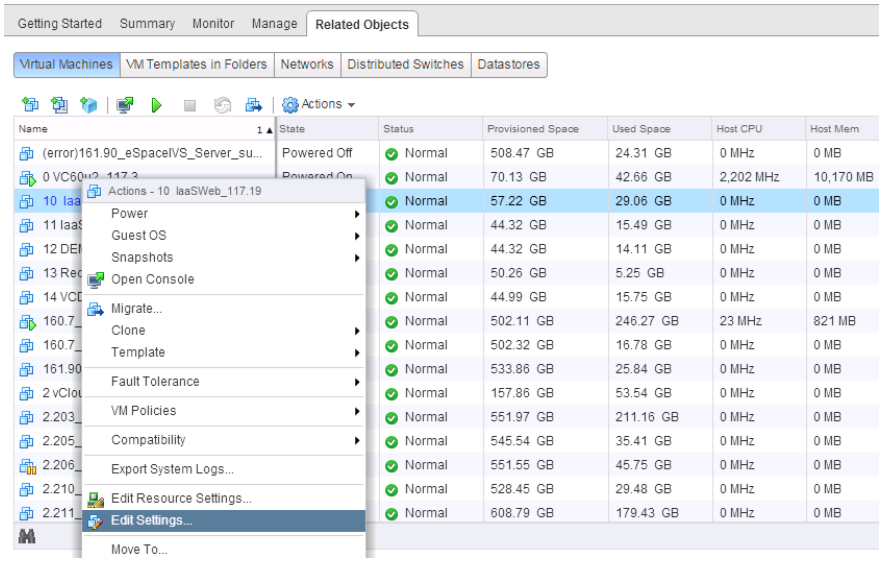

Step 2 In the displayed **Edit Settings** dialog box, click the Virtual Hardware tab. On the tab page, select **RDM Disk** from the **New Device** option list at the bottom.

Figure 6-13 Adding RDM disks

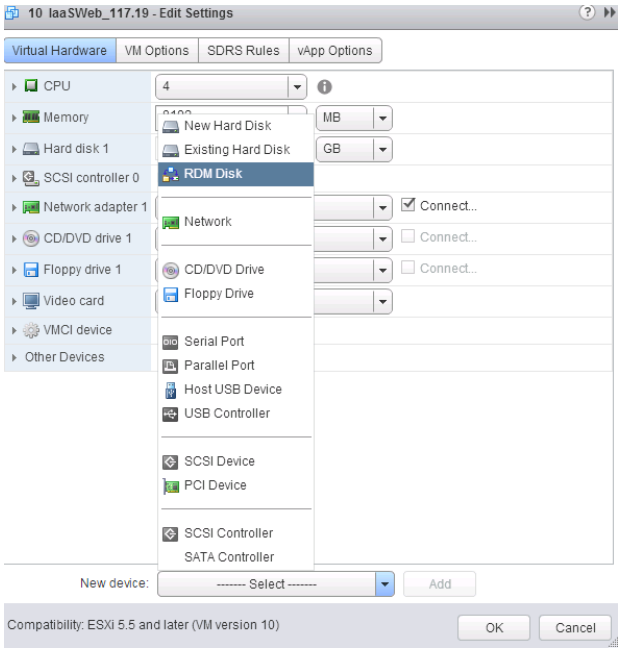

Step 3 Click **Add** to add the target disk.

Figure 6-14 Selecting disks to add

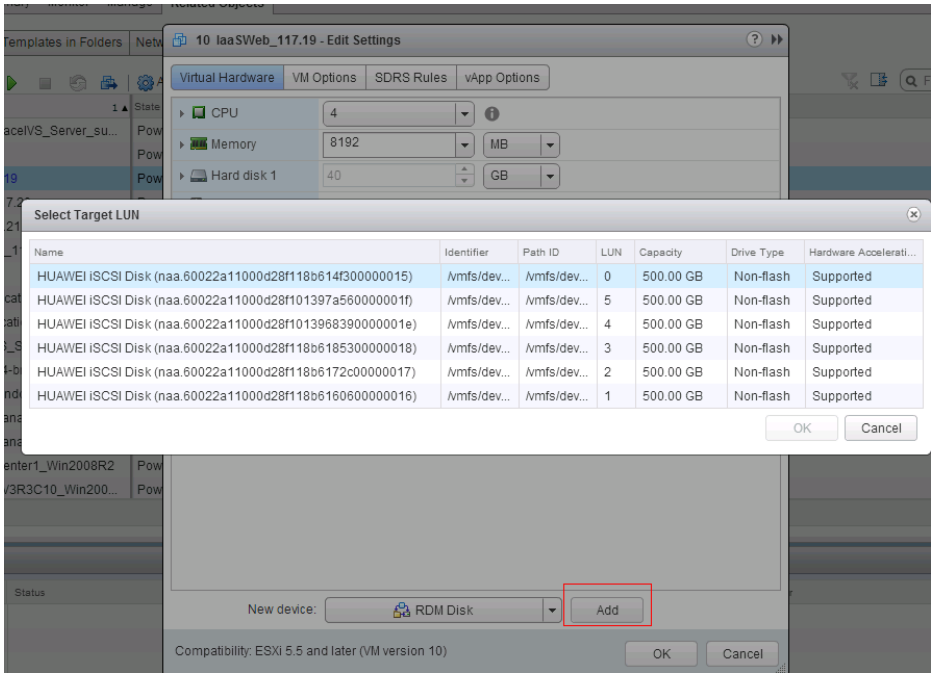

Step 4 Verify the disk information and click **OK**.

Figure 6-15 Completing the disk addition operation

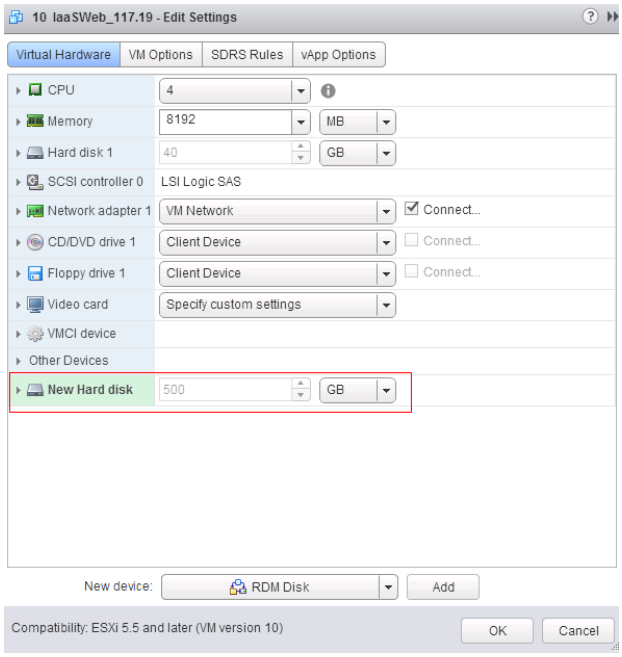

Step 5 Navigate to the Edit Settings tab page again to check whether the target disk is added successfully.

Figure 6-16 Checking whether the disk is successfully added

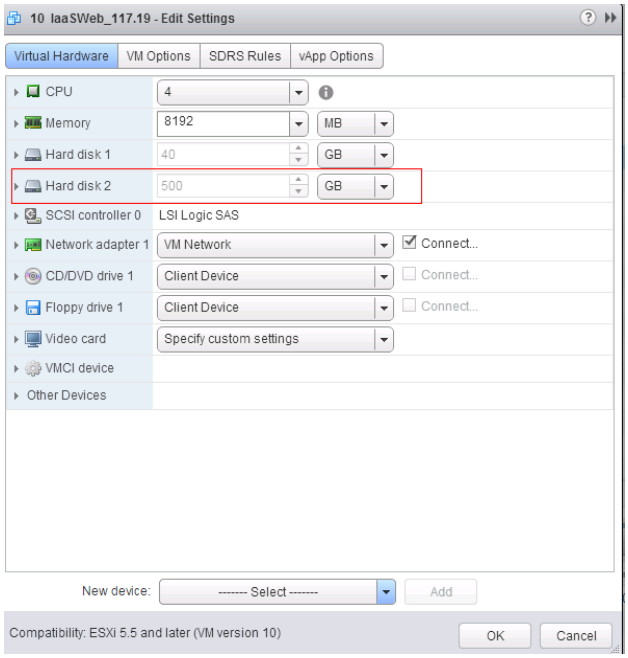

----End

6.2.2 Creating Datastores

Create a file system before creating a virtual disk. A file system can be created using the file system disks in datastores.

This section details how to create a datastore.

vSphere Client

On vSphere Client, perform the following steps to create a datastore:

Step 1 On the **Configuration** tab page, choose **Storage** in the navigation tree. On the **Datastores** tab page that is displayed, click **Add Storage**, as shown in Figure 6-17.

Figure 6-17 Adding storage

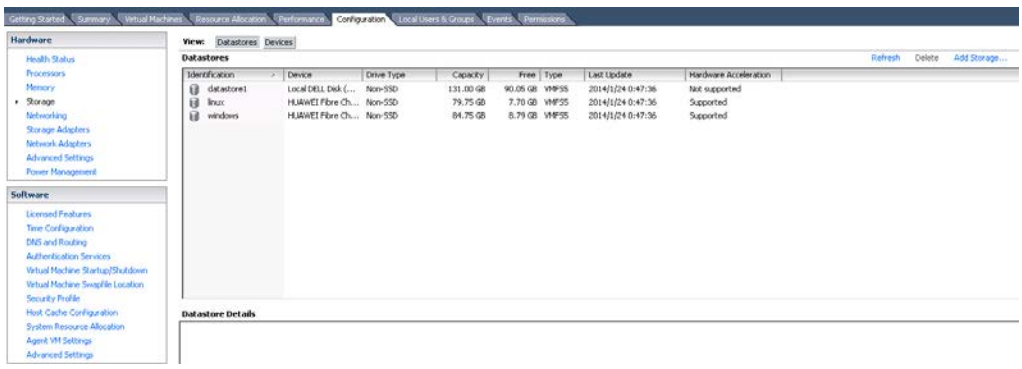

Step 2 Select a storage type and click **Next**, as shown in Figure 6-18.

Figure 6-18 Selecting a storage type

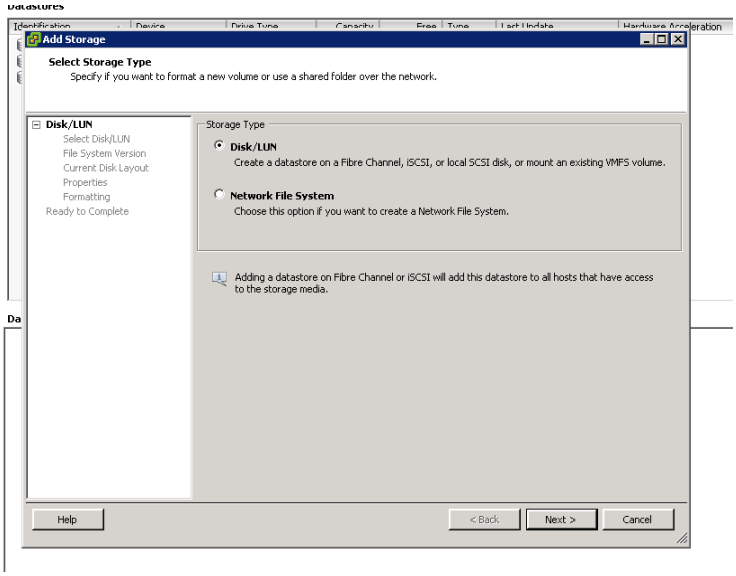

Step 3 On the **Select Disk/LUN** page that is displayed, select a desired disk and click **Next**, as shown in Figure 6-19.

Figure 6-19 Select a disk/LUN

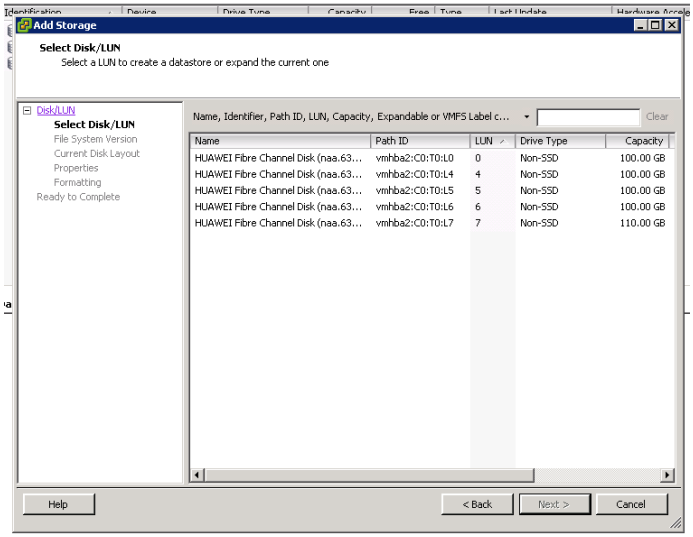

Step 4 Select a file system version. VMFS-5 is selected in this example, as shown in Figure 6-20.

Figure 6-20 Selecting a file system version

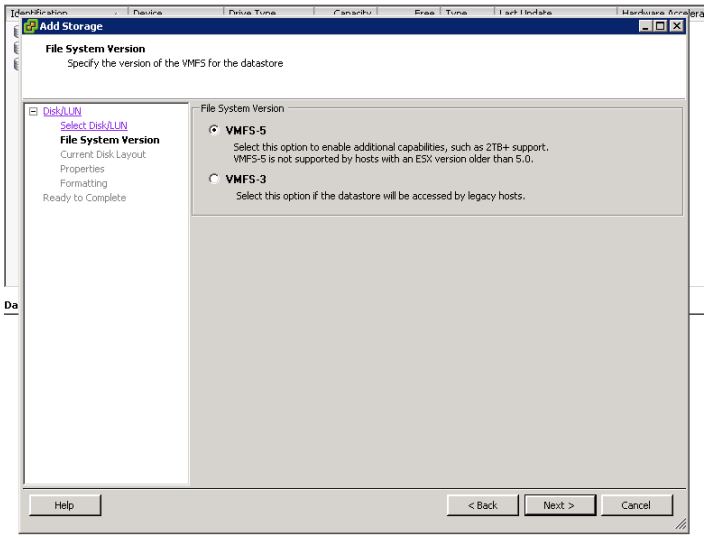

Step 5 View the current disk layout and device information, as shown in Figure 6-21.

Figure 6-21 Viewing the current disk layout

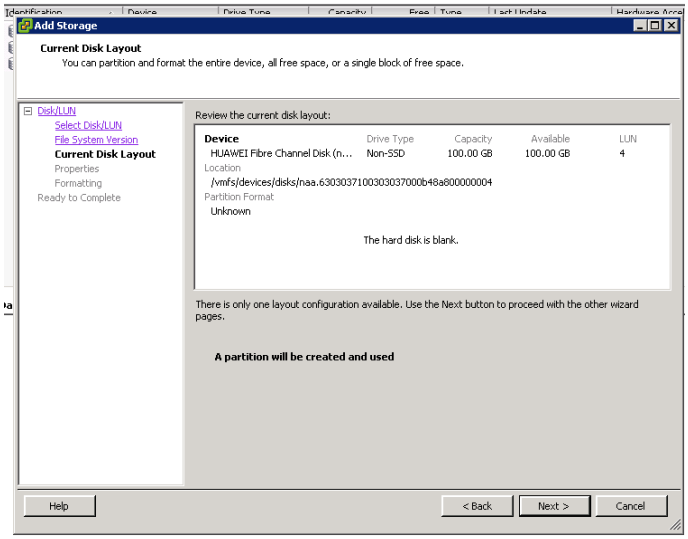

Step 6 Enter the name of a datastore, as shown in Figure 6-22.

Figure 6-22 Entering a datastore name

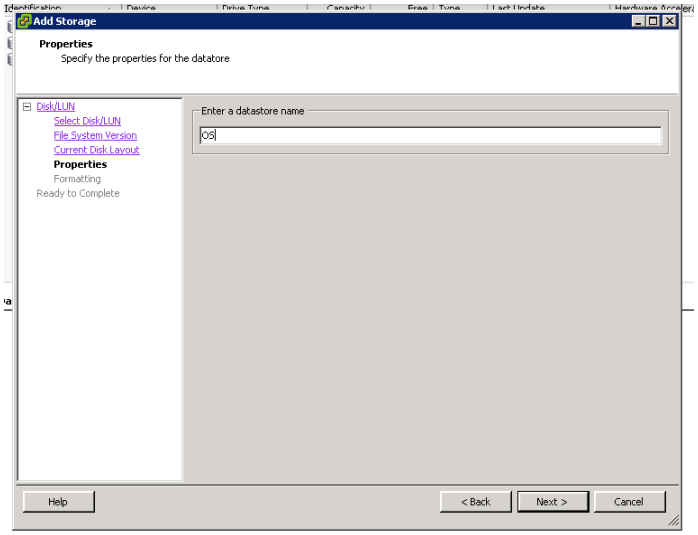

Step 7 Specify a disk capacity. Normally, **Maximum available space** is selected. If you want to test LUN expansion, customize a capacity, as shown in Figure 6-23.

Figure 6-23 Specifying a capacity

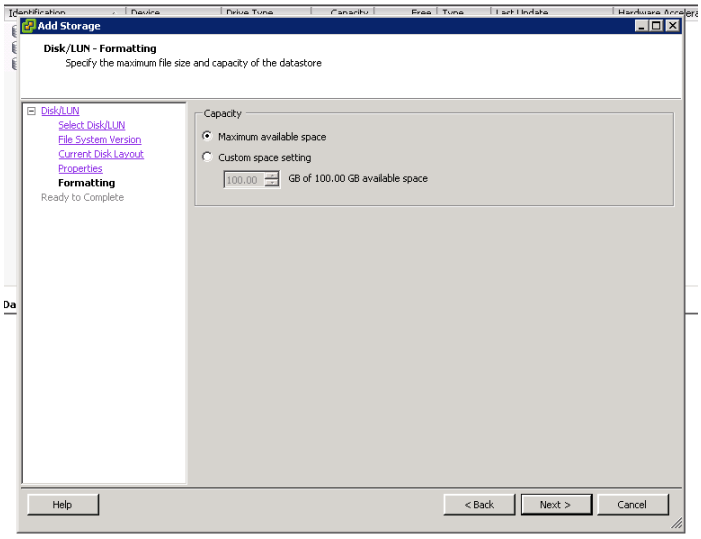

Step 8 Confirm the disk layout. If the disk layout is correct, click **Finish**, as shown in Figure 6-24.

Figure 6-24 Confirming the disk layout

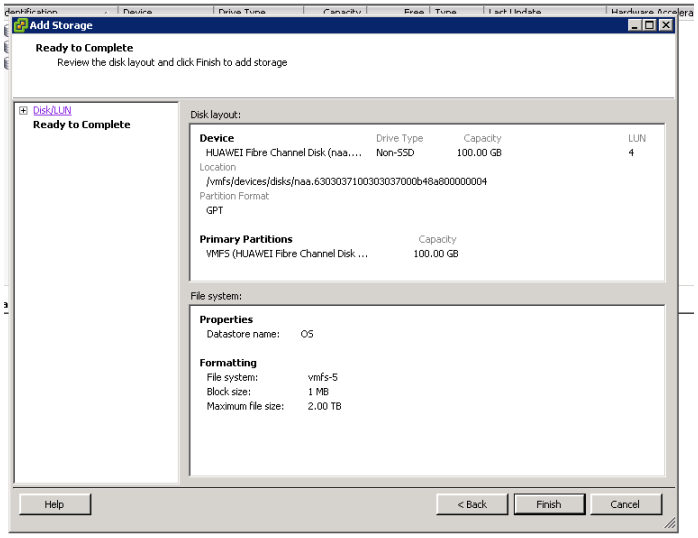

----End

vSphere Web Client

On vSphere Web Client, perform the following steps to create a datastore:

Step 1 On the **Related Objects** tab page, click the **Datastores** tab.

Figure 6-25 Checking the datastores

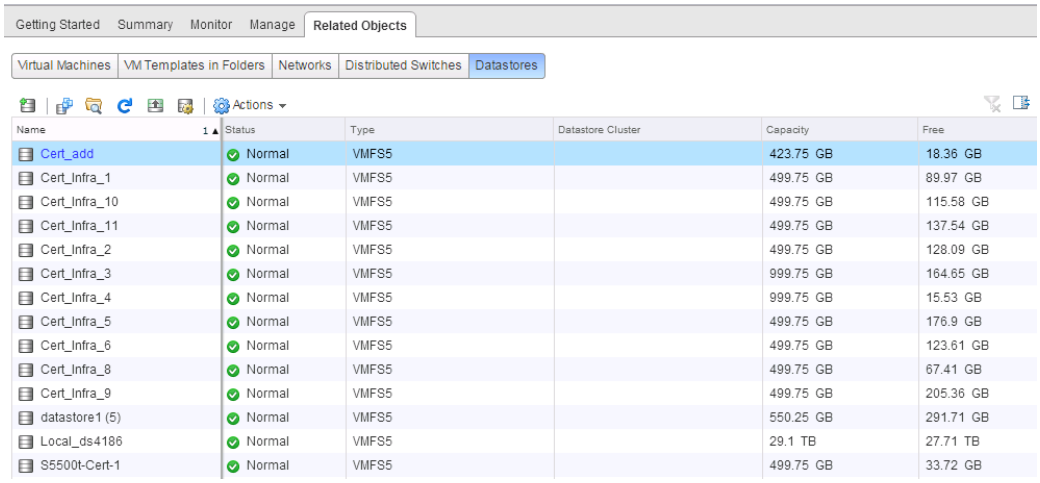

Step 2 Click  to open the **New Datastore** page. On this page, select **VMFS** as **Type**, and click **Next**.

Figure 6-26 Creating the datastore type

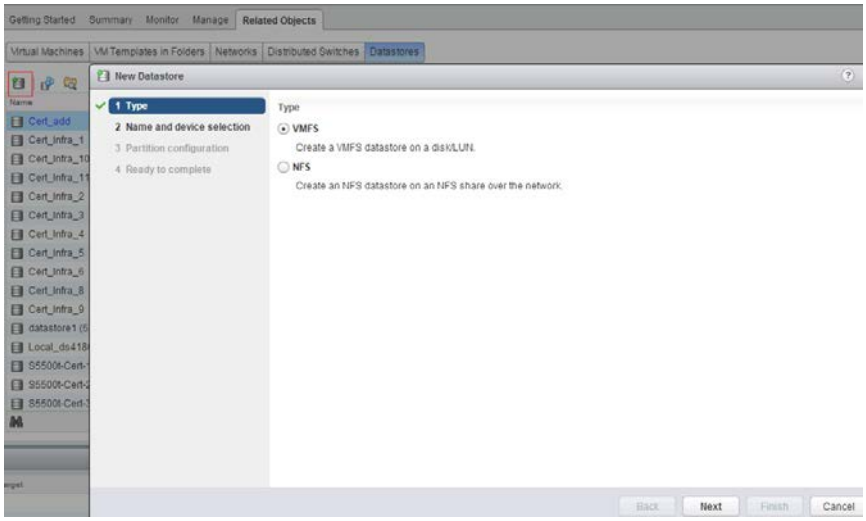

Step 3 Specify the datastore name, select the disks, and click **Next**.

Figure 6-27 Specifying the datastore name and selecting disks

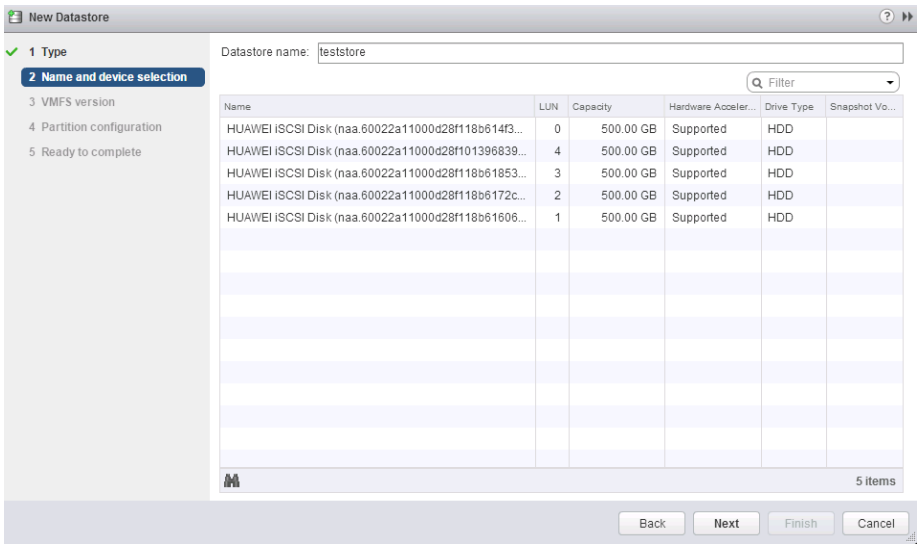

Step 4 Select the file system version (**VMFS 5**, for example), and click **Next**.

Figure 6-28 Selecting the file system version

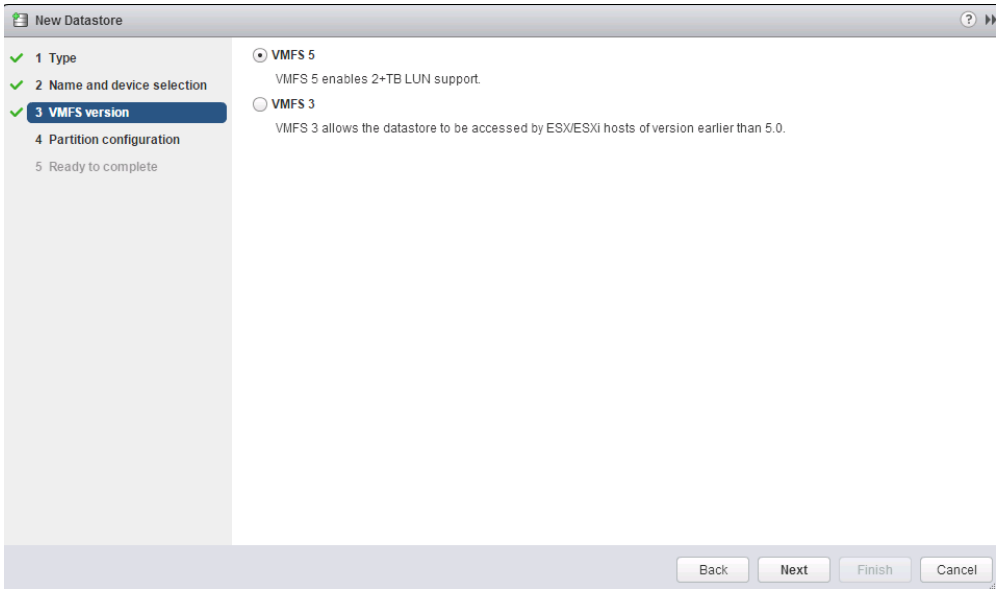

Step 5 Configure datastore partition and click **Next**.

Figure 6-29 Configuring the partition layout

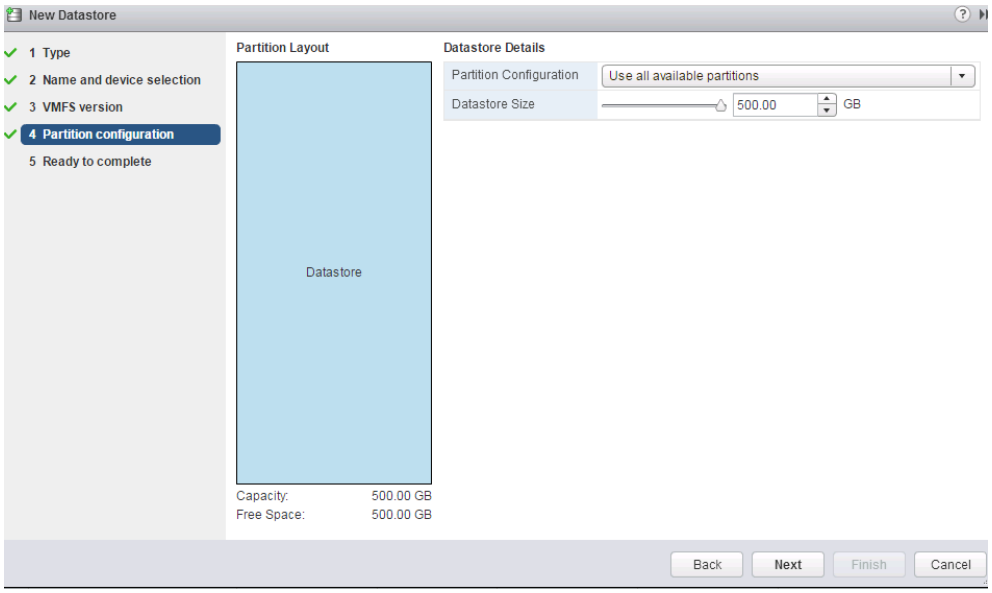

Step 6 Verify the datastore configurations and click **Finish**.

Figure 6-30 Verifying the datastore configurations

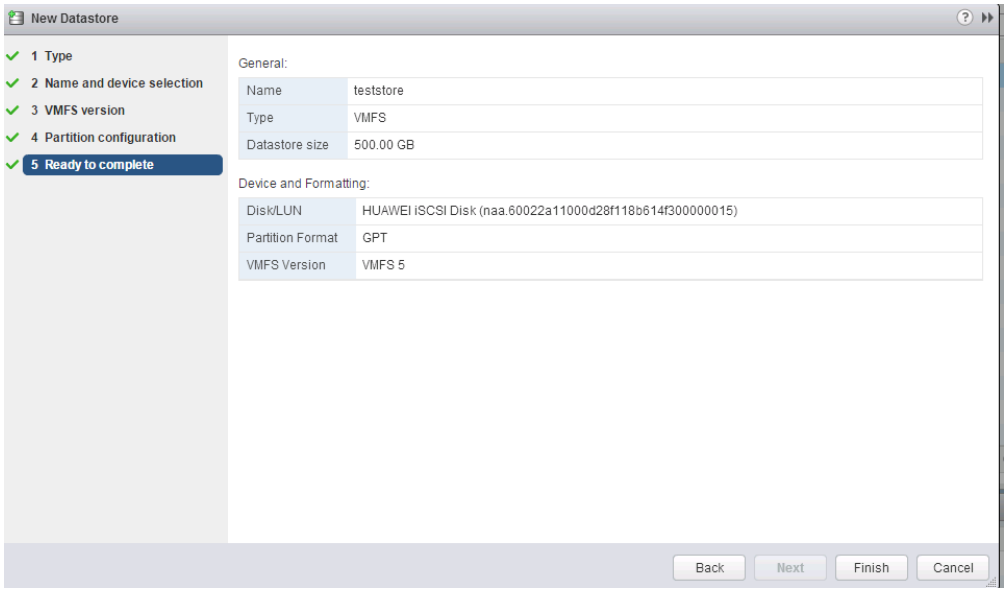

Step 7 Check whether the datastore is successfully created.

Figure 6-31 Checking for the datastore

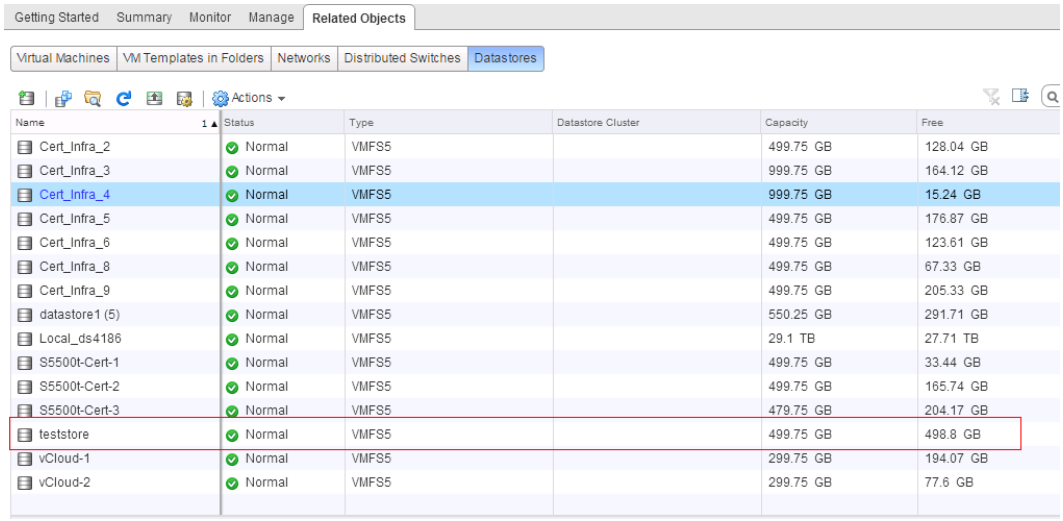

----End

6.2.3 Creating Virtual Disks

This section describes how to add LUNs to VMs as virtual disks.

vSphere Client

On vSphere Client, perform the following steps to create virtual disks:

Step 1 Right-click a VM and choose **Edit Settings** from the shortcut menu, as shown in Figure 6-32.

Figure 6-32 Editing VM settings

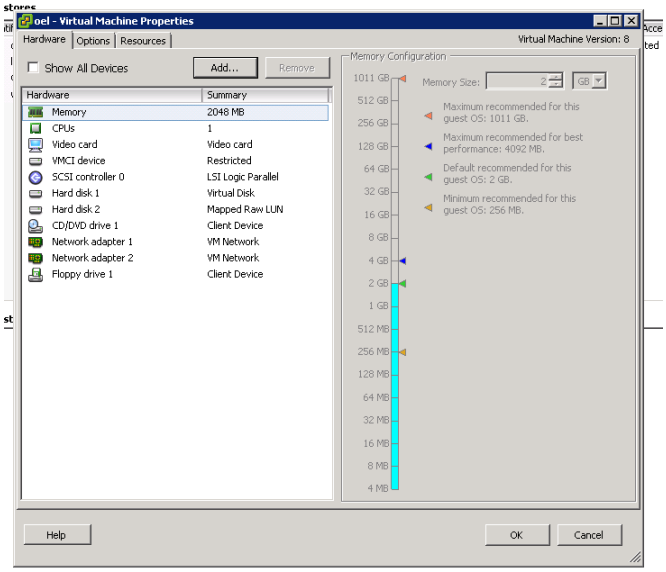

Step 2 Click **Add**, select **Hard Disk** and click **Next**, as shown in Figure 6-33.

Figure 6-33 Adding disks

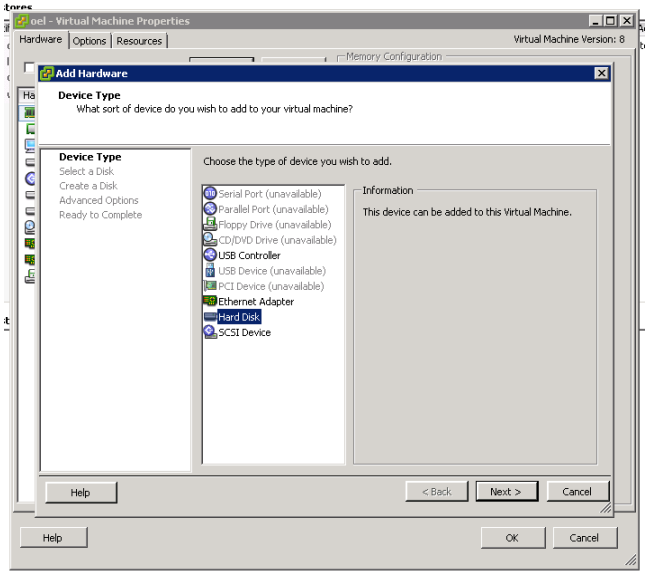

Step 3 In **Select a Disk**, select **Create a new virtual disk**, as shown in Figure 6-34.

Figure 6-34 Creating a new virtual disk

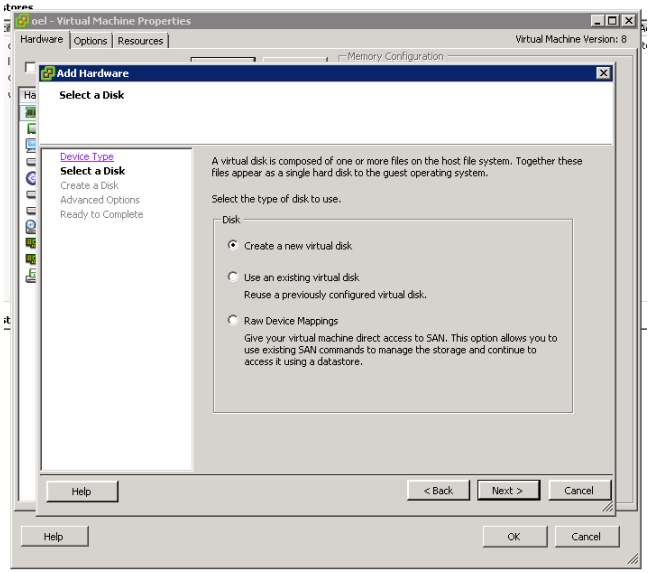

Step 4 Specify the disk capacity based on site requirements, as shown in Figure 6-35.

Figure 6-35 Specifying the disk capacity

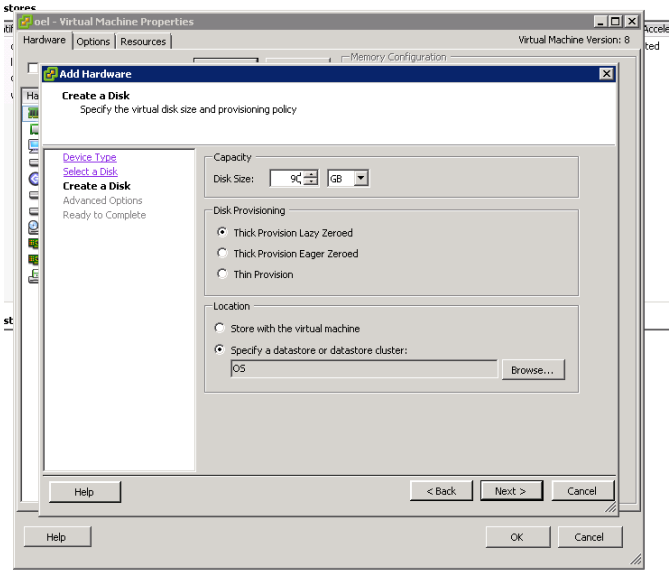

Step 5 Select a datastore. In this example, the datastore is **disk1** and the file system type is **VMFS-5**, as shown in Figure 6-36.

Figure 6-36 Selecting a datastore

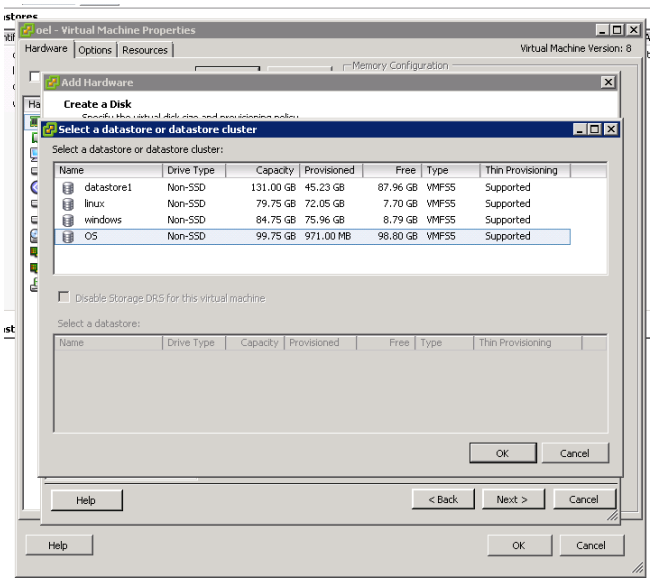

Step 6 Select a virtual device node. If there are no special requirements, keep the default virtual device node unchanged, as shown in Figure 6-37.

Figure 6-37 Selecting a virtual device node

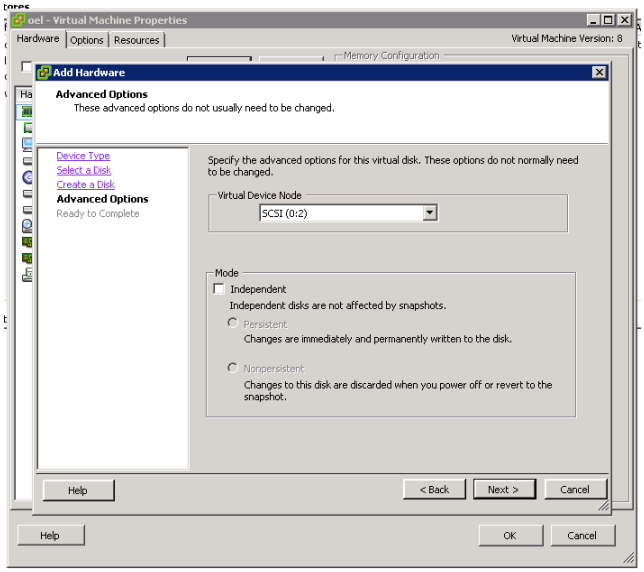

Step 7 View the basic information about the virtual disk, as shown in Figure 6-38.

Figure 6-38 Viewing virtual disk information

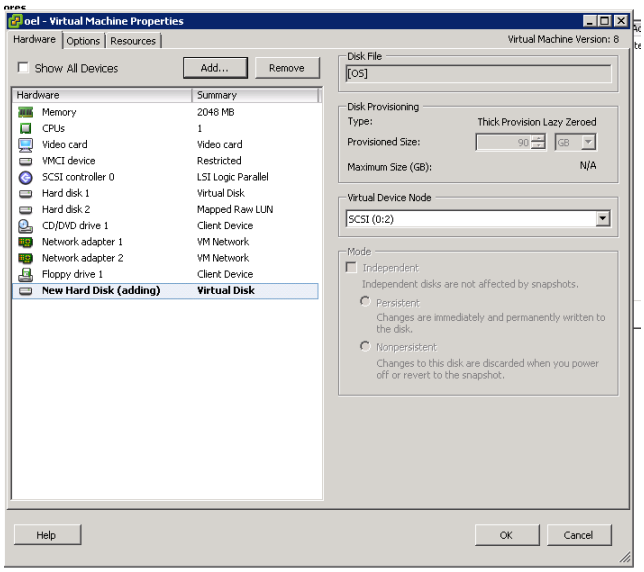

As shown in the preceding figure, hard disk 1 that you have added is a virtual disk.

----End

vSphere Web Client

On vSphere Web Client, perform the following steps to create virtual disks:

Step 1 On the **Related Objects** tab page, click the **Virtual Machines** tab. On the tab page, select the host for which you need to create virtual disks, right-click the host, and choose **Edit Settings** from the shortcut menu.

Figure 6-39 Editing the host settings

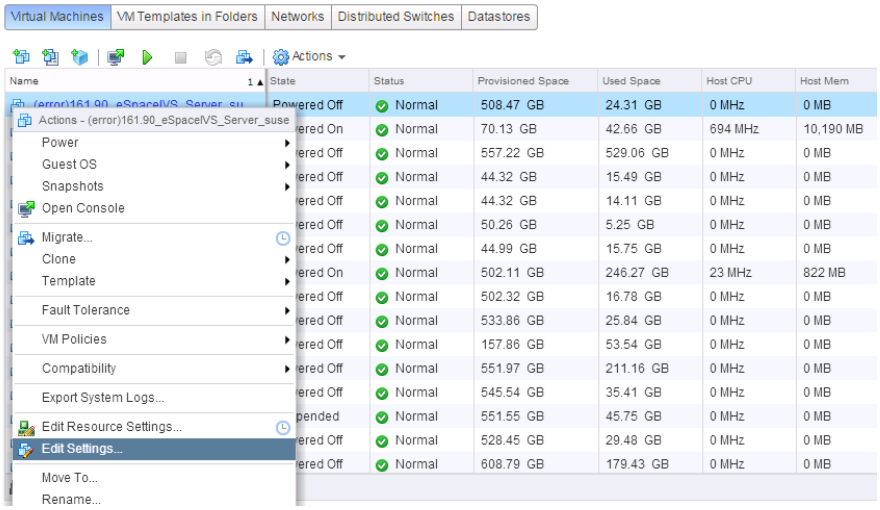

Step 2 On the **Virtual Hardware** tab page, select **New Hard Disk** from the **New Device** option list.

Figure 6-40 Selecting to add a hard disk

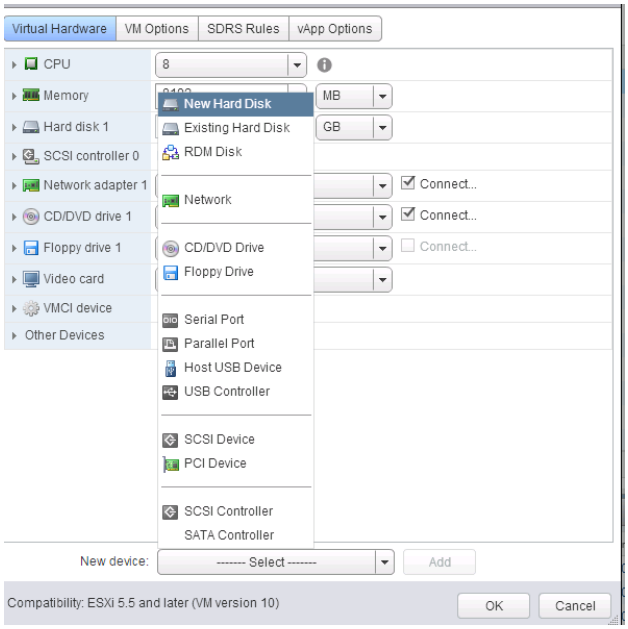

Step 3 Click **Add** and check the information of the added disk.

Figure 6-41 Checking the information of the added disk

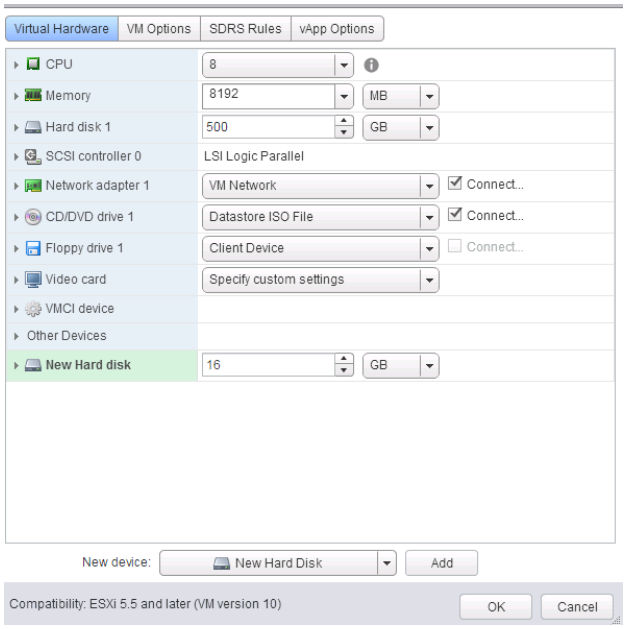

Step 4 To modify any disk properties, expand **New Hard disk** by clicking the arrow icon on its left.

Figure 6-42 Modifying disk properties

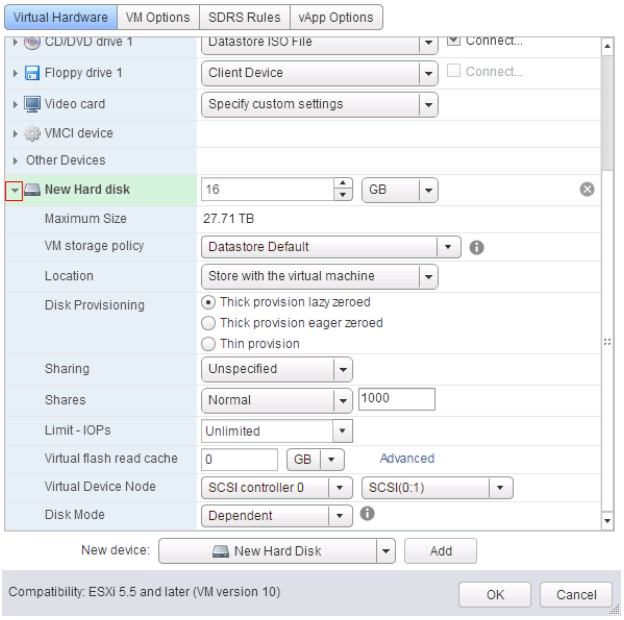

----End

7 Configuring Multipathing in HyperMetro Scenarios

This chapter describes the multipathing software configurations on the hosts and storage systems. For details about how to configure HyperMetro services, see the *HyperMetro Feature Guide*.

7.1 UltraPath

7.1.1 Storage System Configuration

If you use UltraPath, retain the default initiator settings. Do not select **Uses third-party multipath software**.

Figure 7-1 Initiator setting when UltraPath is used

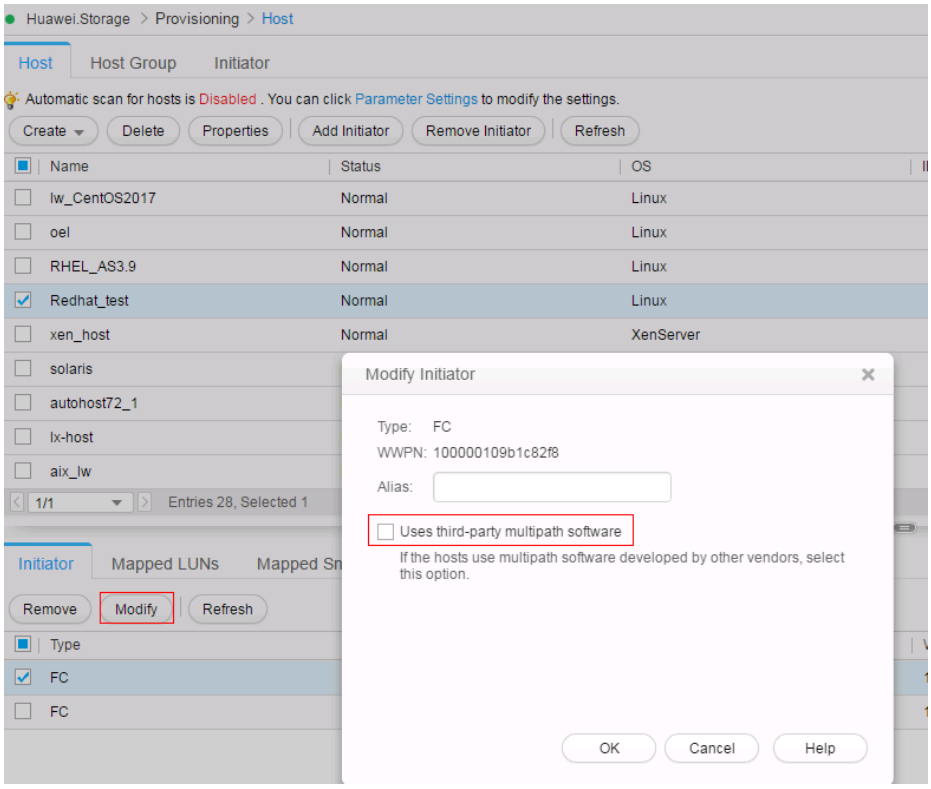

7.1.2 Host Configuration

7.1.2.1 Installing UltraPath

In active-active data center solutions, you can configure UltraPath to increase the I/O performance and decrease the read/write delay.

For details on how to install UltraPath, refer to instructions in the *OceanStor UltraPath for Linux User Guide*.

7.1.2.2 Configuring UltraPath

Prerequisites

- In UltraPath, set the HyperMetro working mode to preferred storage array mode. In this mode, the local storage array is preferred in processing host services. The remote storage array is used only when the local array is faulty. This improves the service response speed and reduces the access latency.
- This configuration must be performed separately on all hosts.
- When UltraPath is configured, retain the **Uses the third-party multipath software** option unselected.

NOTE

For details on how to disable the third-party multipath software function, refer to section 7.2.2 Storage System Configuration.

Procedure

Step 1 Run a command to set the HyperMetro working mode.

Table 7-1 lists the command for setting the HyperMetro working mode.

Table 7-1 Setting the HyperMetro working mode

Operating System	Command	Example
vSphere	**set hypermetro workingmode -m** *mode* **-p primary_array_id**	**esxcli upadm set hypermetro workingmode -m priority -p 0**
NOTE In vSphere, **esxcli upadm** is added in this command for navigating to the OceanSto UltraPath CLI.		

Table 7-2 describes the parameters in the **set hypermetro workingmode** command.

Table 7-2 Parameter description

Parameter	Description	Default Value
-m *mode*	HyperMetro working mode. • **priority**: preferred storage array mode • **balance**: load balancing mode **NOTE** If you set the HyperMetro working mode for a specific virtual LUN first and then the global HyperMetro working mode for the storage system, the working mode for the virtual LUN remains unchanged.	**priority** **priority** is recommended. **balance** is applicable when two active-active data centers are in the same building.
-p primary_array_id	ID of the preferred storage array. The ID is allocated by UltraPath. The storage array that is in the same data center as the application hosts reside is preferred. Run the **esxcli upadm show diskarray** command to obtain the storage array ID. **NOTE** • In **priority** mode, this parameter indicates the storage array to which I/Os are preferentially delivered. • In **balance** mode, this parameter indicates the storage array where the first slice section resides.	None **NOTE** Mapping relationship between application hosts and storage arrays: • Storage array A is the preferred array for all application hosts in data center A. • Storage array B is the preferred array for all application hosts in data center B.

Step 2 Run the **esxcli upadm show upconfig** command to query whether the working mode setting is successful.

NOTE

In vSphere, you can add **esxcli upadm** before a command to enter the OceanStor UltraPath CLI.

If the command returns information as follows (vSphere, for example), it indicates that the working mode setting is successful.

```
HyperMetro WorkingMode : read write within primary array
```

Figure 7-2 Querying the UltraPath settings

```
[root@localhost:~] esxcli upadm show upconfig
=====================================================
UltraPath Configuration
=====================================================
Basic Configuration
    Working Mode : load balancing within controller
    LoadBalance Mode : min-queue-depth
    Loadbanlance io threshold : 1
    LUN Trespass : on

Advanced Configuration
    Io Retry Times : 10
    Io Retry Delay : 0
    Faulty path check interval : 10
    Idle path check interval : 60
    Failback Delay Time : 600
    Max io retry timeout : 1800

Path reliability configuration
    Timeout degraded statistical time : 600
    Timeout degraded threshold : 1
    Timeout degraded path recovery time : 1800
    Intermittent IO error degraded statistical time : 300
    Min. I/Os for intermittent IO error degraded statistical : 5000
    Intermittent IO error degraded threshold : 20
    Intermittent IO error degraded path recovery time : 1800
    Intermittent fault degraded statistical time : 1800
    Intermittent fault degraded threshold : 3
    Intermittent fault degraded path recovery time : 3600
    High latency degraded statistical time : 300
    High latency degraded threshold : 1000
    High latency degraded path recovery time : 3600
    Sensitive delayed degraded threshold : 30000
    Sensitive delayed degraded recovery time : 120

APDtoPDL configuration
    APD to PDL Mode : off
    APD to PDL Timeout : 10

HyperMetro configuration
    HyperMetro Primary Array SN : 210235982510EA000016
    HyperMetro WorkingMode : read write within primary array
    HyperMetro Split Size : 128MB
```

Step 3 Set timeout parameters.

For FC networking, there is no need for extra configurations.

For iSCSI networking, run the following commands:

```
esxcli iscsi adapter param set -A vmhba35 -k NoopOutInterval -v 3
esxcli iscsi adapter param set -A vmhba35 -k NoopOutTimeout -v 10
esxcli iscsi adapter param set -A vmhba35 -k RecoveryTimeout -v 3
```

NOTE

- The preceding commands can be used only in VMware 5.0 and later versions. For the detailed HyperMetro-supported VMware versions, see http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf.
- The field in italic and bold (vmhba35, in this example) indicates the iSCSI adapter. You need to change this value according to your site.
- The settings will not take effect until you restart the host.
- Without the preceding command executions, the default ESXi settings might result in an up-to-35s path switchover time. The preceding command executions can shorten the path switchover time to about 16s.

Step 4 (Optional) Enable APD to PDL for the VMware ESXi hosts.

NOTE

If VMware ESXi hosts are deployed in a cluster and their connected OceanStor V3 is V300R006C00SPC100 or later, you do not need to enable the APD to PDL function; however, you need to enable the APD to PDL function if OceanStor V3 is earlier than V300R006C00SPC100.

1. Run the **esxcli upadm set apdtopdl -m on** command.
2. Run the **esxcli show upconfig** command to query the configuration result.

   If **APD to PDL Mode** is **on**, the APD to PDL function is successfully configured.

```
vSphere Security documentation for more information
~ # esxcli upadm show upconfig
============================================================ UltraPath
Configuration
============================================================ Basic Configuration
 Working Mode : load balancing within controller
 LoadBanlance Mode : min-queue-depth
 Loadbanlance io threshold : 1
 LUN Trespass : on
Advanced Configuration
 Io Retry Times : 10
 Io Retry Delay : 0
 Faulty path check interval : 10
 Idle path check interval : 60
 Failback Delay Time : 600
 Max io retry timeout : 1800
Path reliability configuration
 Timeout degraded statistical time : 600
 Timeout degraded threshold : 1
 Timeout degraded path recovery time : 1800
 Intermittent IO error degraded statistical time : 300
 Min. I/Os for intermittent IO error degraded statistical : 5000 Intermittent IO error
derraded threshold : 20
 Intermittent IO error derraded path recovery time : 1800
 Intermittent fault degraded statistical time : 1800
 Intermittent fault degraded threshold : 3
 Intermittent fault degraded path recovery time : 3600
```

```
 High latency degraded statistical time : 300
 High latency degraded threshold : 1000
 High latency degraded path recovery time : 3600
APDtoPDL configuration
APD to PDL Mode : on
APD to PDL Timeout : 10
```

----End

Configuring a VMware Cluster

If you want to configure VMware clusters, see the BC&DR Solution Product Documentation (Active-Active Data Center). See section "Virtualization Platform Configuration" in this document. The contents in this section are as follows.

Mandatory Configuration Items:

- Deploy ESXi hosts across data centers in a HA cluster and configure the cluster with HA advanced parameter **das.maskCleanShutdownEnabled = True** for VMware vSphere 5.0 u1, 5.1, and 5.5 versions.
- A VM service network requires L2 interworking between data centers for VM migration between data centers without affecting VM services.
- For VMware vSphere 5.0 u1, later 5.0 versions, and 5.1 versions, log in to the CLI of each ESXi host using SSH and add **Disk.terminateVMOnPDLDefault = True** in the **/etc/vmware/settings** file.
- For VMware vSphere 5.5, 6.0 u1, and versions between them, log in to the host using the vSphere Client.
  - Set **VMkernel.Boot.terminateVMOnPDL = True**. The parameter forcibly powers off VMs on a datastore when the datastore enters the PDL state.

    Figure 7-3 Boot parameter settings

    

  - Set **Disk.AutoremoveOnPDL = 0**. The parameter forcibly removes datastores in PDL state.

    Figure 7-4 Disk parameter settings

    

- For VMware vSphere 6.0 u2 and later versions:
  - After connecting to vCenter through the Web Client, enter the cluster HA configuration. The configuration requirements are as follows.

Figure 7-5 vSphere 6.0 cluster configuration

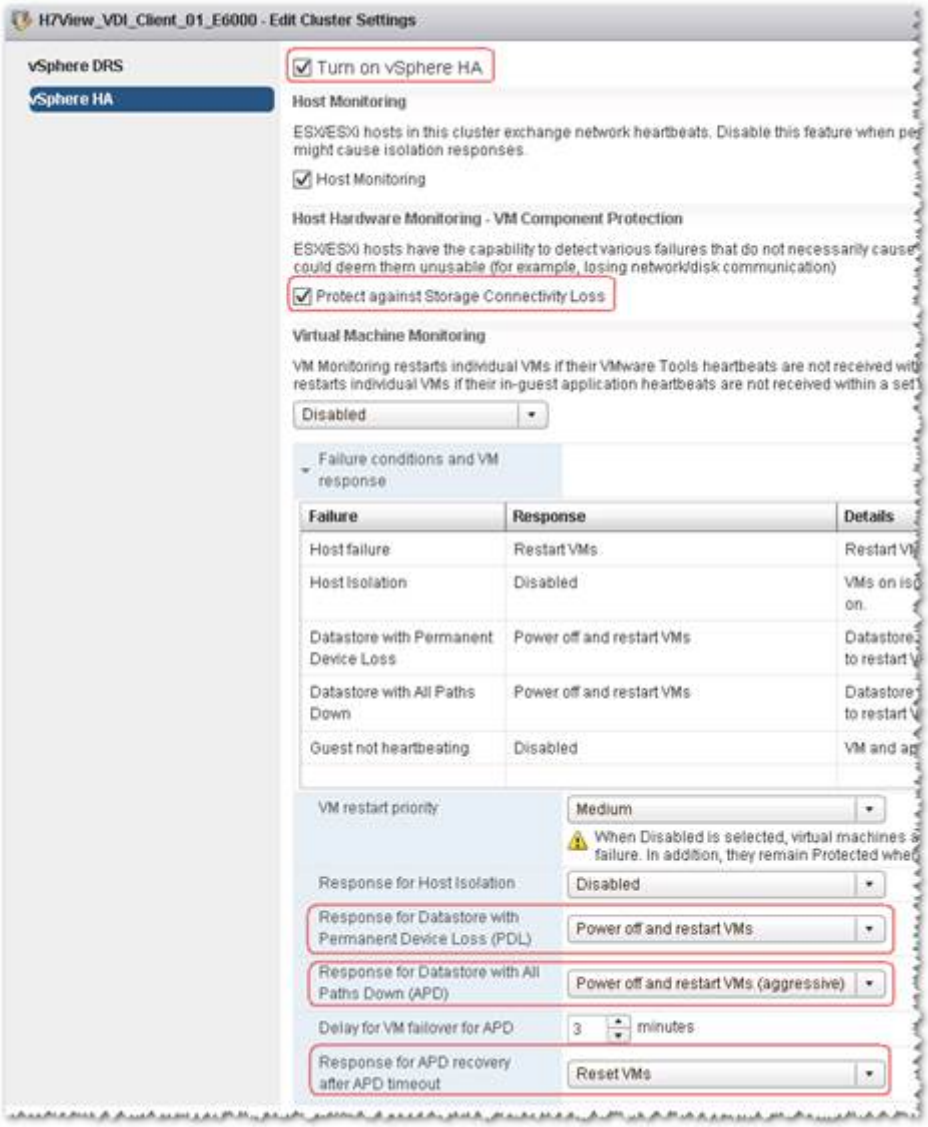

- For VMware vSphere 6.5:

  When using HyperMetro with VMware ESXi, note the following precautions:
  - When using HyperMetro with VMware ESXi, note the following precautions: The two LUNs in a HyperMetro pair must use the same LUN ID when being mapped to a VMware ESXi host (ESXi 6.5.0 GA build 4564106 or a follow-up version earleir than ESXi 6.5 U1 build 5969303). You can query the host LUN ID mapped to the ESXi host in the Mapping View of OceanStor DeviceManager, as shown in Figure 7-6.
  - Before modifying the **Host LUN ID**, read the following warnings carefully since misoperations may cause service interruption. To modify the host LUN ID for a

LUN, right-click the LUN and choose **Change host LUN ID** from the shortcut menu. In the displayed dialog box, set the same **Host LUN ID** value for the two storage devices in the HyperMetro pair and then click **OK**. For details, see Figure 7-7.

WARNING

Changing the host LUN ID with an incorrect procedure may cause service interruption.

- If no datastore has been created on either LUN in the HyperMetro pair, you can directly change the host LUN ID for the LUNs. Wait for about 5 to 15 minutes after the modification is complete, and then run the **Rescan** command in the ESXi host CLI to check whether the LUNs in the HyperMetro pair have been restored and been online.
- If a datastore has been created on either LUN in the HyperMetro pair and a service has been deployed in the datastore, change the host LUN ID using only the following two methods (otherwise, changing the host LUN ID for either LUN will cause the LUN to enter the PDL state and consequently interrupt services):
- Method 1: You do not need to restart the ESXi host. Migrate all VMs in the datastore deployed on the LUNs in the HyperMetro pair to another datastore, and then change the host LUN ID on the OceanStor DeviceManager. Wait for about 5 to 15 minutes after the modification is complete, and then run the **Rescan** command in the ESXi host CLI to verify that the LUNs in the HyperMetro pair have been restored and been online. Then, migrate the VMs back to the datastore deployed on the LUNs in the HyperMetro pair.
- Method 2: You need to restart the ESXi host. Power off all VMs in the datastore deployed on the LUNs in the HyperMetro pair to ensure that no service is running on the LUNs. Then, modify the host LUN ID on the OceanStor DeviceManager. Then, restart the ESXi host to make the modification take effect. After restarting the ESXi host, check whether the LUNs in the HyperMetro pair have been restored and been online.

Figure 7-6 Changing the host LUN ID (1)

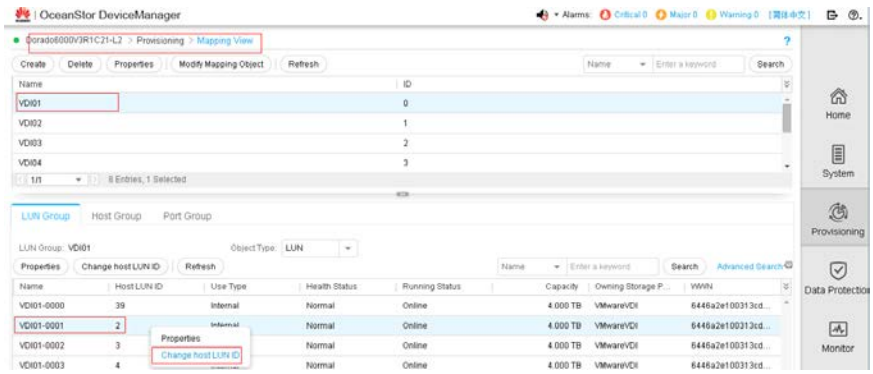

Figure 7-7 Changing the host LUN ID (2)

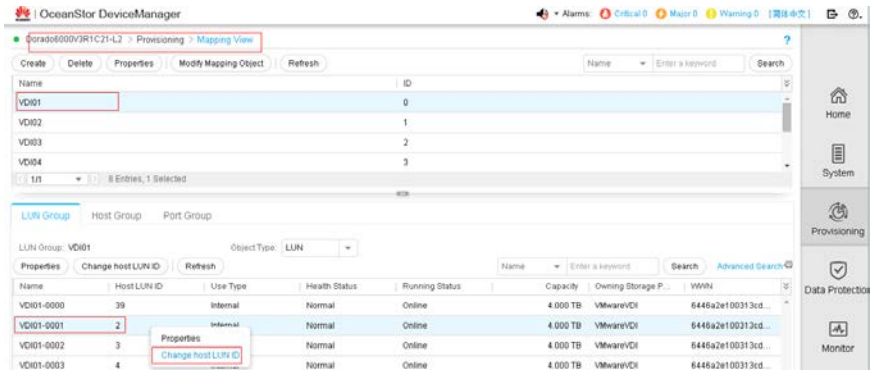

For OceanStor V3 V300R003C20SPC200, a single array with ALUA enabled can have a maximum of 8 controllers; two active-active arrays with ALUA enabled also cannot have more than 8 controllers.

VMware ESXi 6.0 U2 and later versions support HyperMetro configuration. Versions earlier than VMware ESXi 6.0 U2 have their defects.

Dorado V3 must be V300R001C01SPC100 and later versions, supporting multi-controller and HyperMetro ALUA.

- HyperMetro pairs' LUN mappings on two active-active storage arrays must be consistent. That is, the two LUNs in a HyperMetro pair must use the same LUN ID when being mapped to a host. On the storage arrays, you can run the **show host lun host_id=xx** command to query all LUNs mapped to the host. In this command, **xx** indicates the host ID. If you want to modify the ID information, run the **change mapping_view mapping_view_id=*x* host_lun_id_list=*Lun ID:Host Lun ID*** command.

- After connecting to vCenter through the Web Client, enter the cluster HA configuration. The configuration requirements are as follows.

Figure 7-8 vSphere 6.5 cluster configuration-1

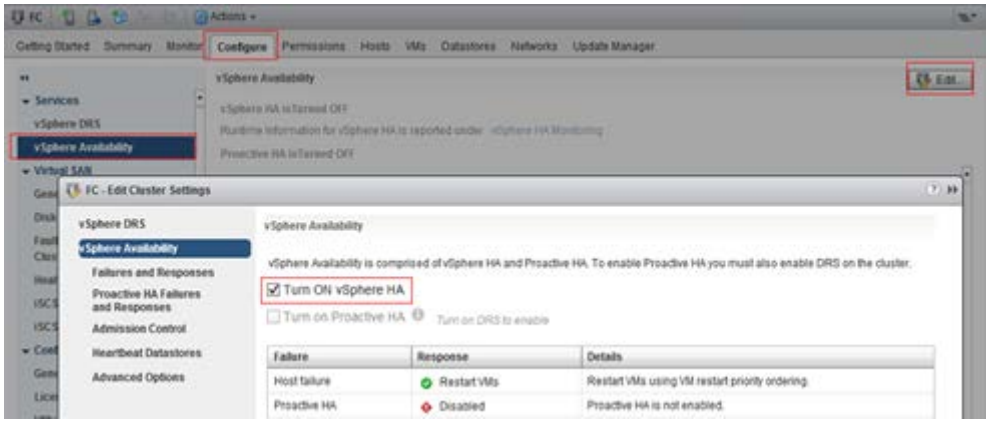

Figure 7-9 vSphere 6.5 cluster configuration-2

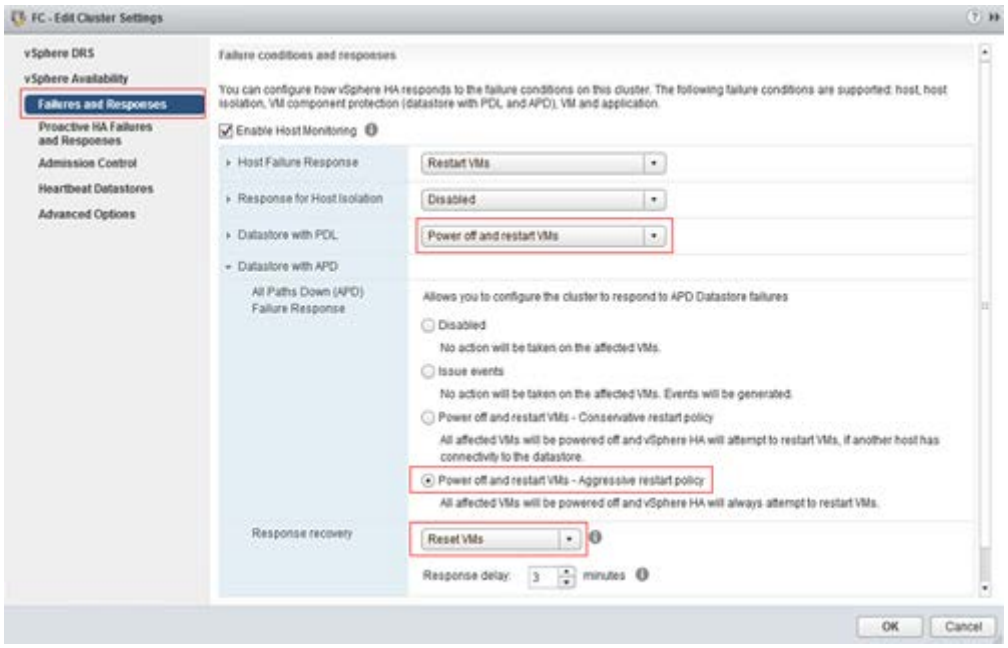

CAUTION

For VMware vSphere 5.1 to 5.5 versions, restart hosts for the configuration to take effect.

For VMware vSphere 6.0 U2 to 6.5 versions, re-enable the HA cluster to make the configuration take effect without restarting hosts.

Recommended Configuration Items

- The vMotion network, service network, and management network are configured with different VLAN IDs to avoid network interference.
- The management network includes the vCenter Server management node and ESXi hosts that are not accessible to external applications.
- The service network is divided into VLANs based on service requirements to ensure logical isolation and control broadcast domains.
- In a single cluster, the number of hosts does not exceed 16. If a cluster has more than 16 hosts, you are advised to use the hosts to create multiple clusters across data centers.
- A DRS group is configured to ensure that VMs can be recovered first in the local data center in the event of the breakdown of a single host.

7.1.3 Verification

On vSphere, run the **esxcli upadm show upconfig** command.

NOTE

In vSphere, **esxcli upadm** in this command is added for navigating to the OceanStor UltraPath CLI.

If the command output contains the following information, the configuration is successful.

```
HyperMetro WorkingMode : read write within primary array
```

Figure 7-10 provides an example.

Figure 7-10 Verifying the HyperMetro working mode

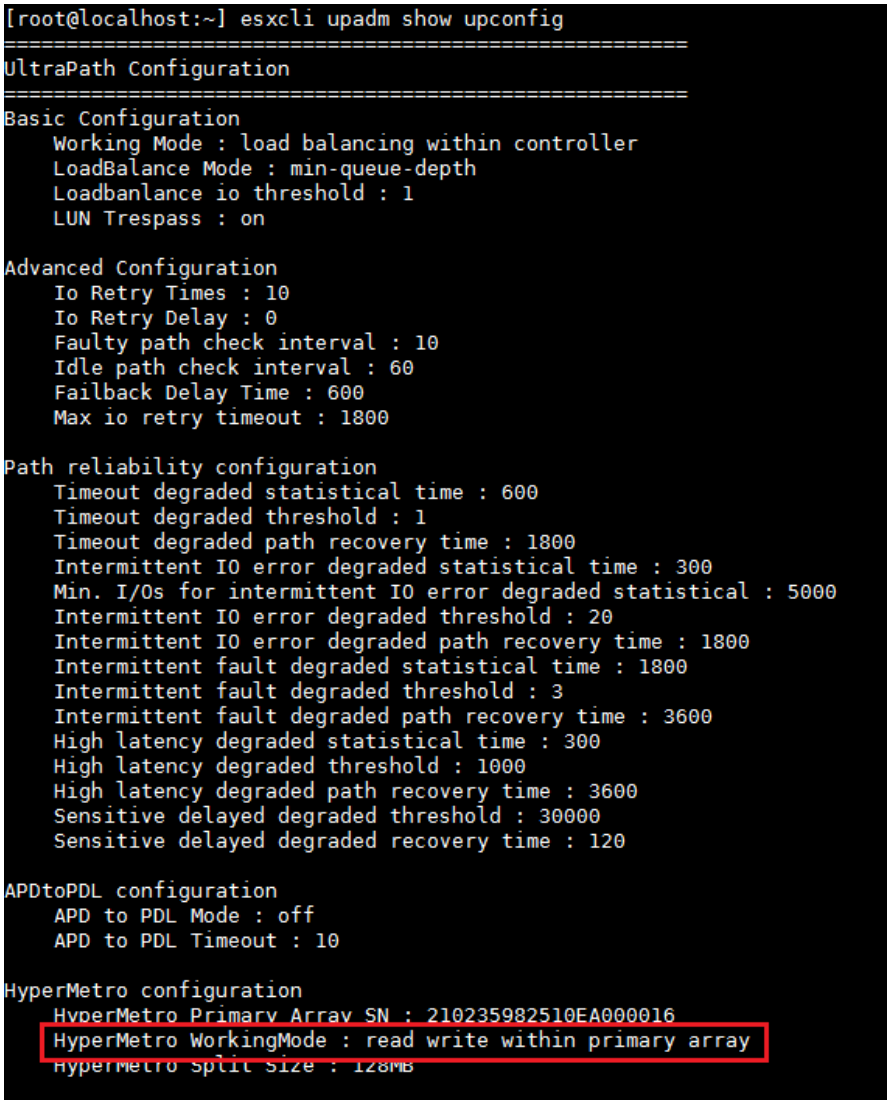

7.2 OS Native Multipathing Software

This section describes the concepts that may be used in configuring OS native multipathing software.

7.2.1 HyperMetro Working Modes

Introduction to HyperMetro Working Modes

Typically, HyperMetro works in load balancing mode or local preferred mode. The typical working modes are valid only when both the storage system and host use ALUA. It is advised to set the host's path selection policy to round-robin. If HyperMetro works in load balancing mode, the host's path selection policy must be round-robin.

HyperMetro storage arrays can be classified into a local and a remote array by their distance to the host. The one closer to the host is the local array and the other one is the remote array.

Table 7-3 describes the configuration methods and application scenarios of the typical working modes.

Table 7-3 HyperMetro working modes

Working Mode	Configuration Method	Application Scenario
Load balancing mode	Enable ALUA on the host and set the path selection policy to round-robin. Configure a switchover mode that supports ALUA for both HyperMetro storage arrays' initiators that are added to the host. Set the path type for both storage arrays' initiators to the optimal path.	The distance between both HyperMetro storage arrays is less than 1 km. For example, they are in the same equipment room or on the same floor.
Local preferred mode	Enable ALUA on the host. It is advised to set the path selection policy to round-robin. Configure a switchover mode that supports ALUA for both HyperMetro storage arrays' initiators that are added to the host. Set the path type for the local storage array's initiators to the optimal path and that for the remote storage array's initiators to the non-optimal path.	The distance between both HyperMetro storage arrays is greater than 1 km. For example, they are in different locations or data centers.

Working Principles and Failover

When ALUA works, the host multipathing software divides the physical paths to disks into Active Optimized (AO) and Active Non-optimized (AN) paths. The host delivers services to the storage system via the AO paths preferentially.

- An AO path is the optimal I/O access path and is between the host and a working controller.
- An AN path is the suboptimal I/O access path and is between the host and a non-working controller.

When HyperMetro works in load balancing mode, the host multipathing software selects the paths to the working controllers on both HyperMetro storage arrays as the AO paths, and those to the other controllers as the AN paths. The host accesses the storage arrays via the AO

paths. If an AO path fails, the host delivers I/Os to another AO path. If the working controller of a storage array fails, the system switches the other controller to the working mode and maintains load balancing.

Figure 7-11 Load balancing mode

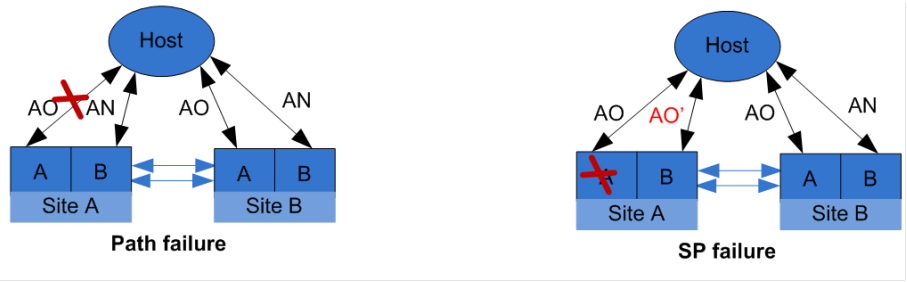

When HyperMetro works in local preferred mode, the host multipathing software selects the paths to the working controller on the local storage array as the AO paths. This ensures that the host delivers I/Os only to the working controller on the local storage array, reducing link consumption. If all AO paths fail, the host delivers I/Os to the AN paths on the non-working controller. If the working controller of the local storage array fails, the system switches the other controller to the working mode and maintains the local preferred mode.

Figure 7-12 Local preferred mode

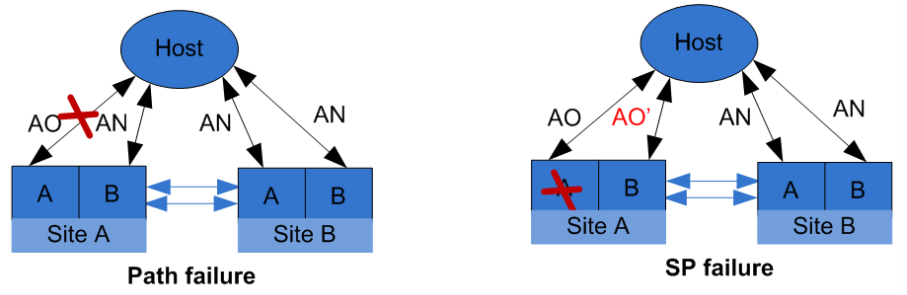

7.2.2 Storage System Configuration

Initiator Modes

Table 7-4 describes the initiator parameters.

Table 7-4 Initiator parameters

Parameter	Description	Example
Uses third-party multipath software	This parameter is displayed only after an initiator has been added to the host. If LUNs have been mapped to the host before you enable or disable this parameter, restart the host after you configure this parameter. You do not need to enable this parameter on a host	Enabled

Parameter	Description	Example
	with UltraPath.	
Switchover Mode	Path switchover mode The system supports the following modes: • **Early-version ALUA**: default value of **Switchover Mode** for an upgrade from an earlier version to the current version. Detailed requirements are as follows: – The storage system must be upgraded from V300R003C10 or an earlier version to V300R003C20, V300R006C00SPC100, or a version later than V300R006C00SPC100; from V300R005 to V300R006C00SPC100 or a later version; from Dorado V300R001C00 to Dorado V300R001C01SPC100 or a later version. – Before the upgrade, the storage system must have a single or dual controllers and has enabled ALUA. • **Common ALUA**: Detailed requirements are as follows: – The storage system version must be V500R007C00 and later, V300R003C20 and later, V300R006C00SPC100 and later, or Dorado V300R001C01SPC100 and later. – The OS of the host that connects to the storage system must be VMware ESXi, Red Hat 6.X, Windows Server 2012 (using Emulex HBAs), Windows Server 2008 (using Emulex HBAs), or HP-UX 11i V3. • **ALUA not used**: does not support ALUA or HyperMetro. This mode is used when a host such as HP-UX 11i V2 does not support ALUA or ALUA is not needed. • **Special mode**: supports ALUA and has multiple values. It is used by host operating systems that are not supported by the **common ALUA** mode. Detailed requirements are as follows: – The storage system version must be V500R007C00 and later, V300R003C20 and later, V300R006C00SPC100 and later, or Dorado V300R001C01SPC100 and later. – The OS of the host that connects to the storage system must be VMware, AIX, Red Hat 7.X, Windows Server 2012 (using QLogic HBAs), or Windows Server 2008 (using QLogic HBAs).	
Special mode	Special modes support ALUA and apply to	Mode 0

Parameter	Description	Example
type	V500R007C00 and later, V300R003C20 and later, V300R006C00SPC100 and later, or Dorado V300R001C01SPC100 and later. The detailed requirements are as follows: • **Mode 0**: – The host and storage system must be connected using a Fibre Channel network. – The OS of the host that connects to the storage system is Red Hat 7.X, Windows Server 2012 (using QLogic HBAs), or Windows Server 2008 (using QLogic HBAs). • **Mode 1**: – The OS of the host that connects to the storage system is AIX or VMware. – HyperMetro works in load balancing mode. • **Mode 2**: – The OS of the host that connects to the storage system is AIX or VMware. – HyperMetro works in local preferred mode.	
Path Type	The value can be either **Optimal Path** or **Non-Optimal Path**. • When HyperMetro works in load balancing mode, set the **Path Type** for the initiators of both the local and remote storage arrays to **Optimal Path**. Enable ALUA on both the host and storage arrays. If the host uses the round-robin multipathing policy, it delivers I/Os to both storage arrays in round-robin mode. • When HyperMetro works in local preferred mode, set the **Path Type** for the initiator of the local storage array to **Optimal Path**, and that of the remote storage array to **Non-Optimal Path**. Enable ALUA on both the host and storage arrays. The host delivers I/Os to the local storage array preferentially.	Optimal Path

Configure the initiators according to the requirements of each OS. The initiators that are added to the same host must be configured with the same switchover mode. Otherwise, host services may be interrupted.

NOTICE

If a LUN has been mapped to the host, you must restart the host for the configuration to take effect after you modify the initiator parameters. If you configure the initiator for the first time, restart is not needed.

Configuring the Initiators

If you want to configure the initiator mode, perform the following operations.

Step 1 Go to the host configuration page.

Open OceanStor DeviceManager. In the right navigation tree, click **Provisioning** and then click **Host**, as shown in the following figure.

Figure 7-13 Going to the host configuration page

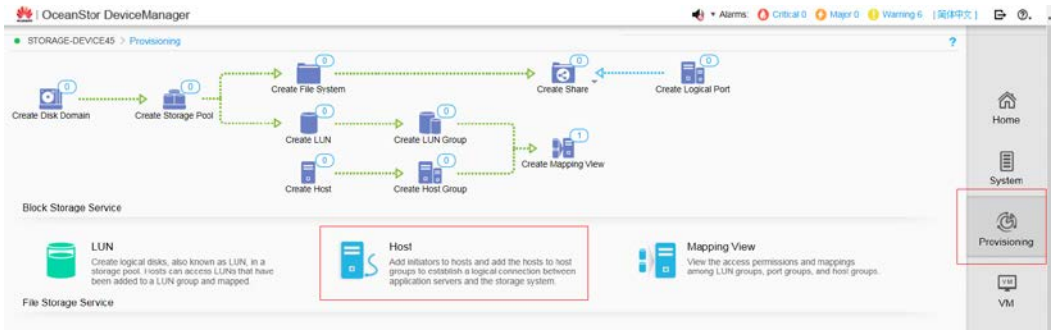

Step 2 Select an initiator of which information you want to modify.

On the **Host** tab page, select a host you want to modify. Then select the initiator (on the host) you want to modify. Click **Modify**.

Figure 7-14 Selecting an initiator of which information you want to modify

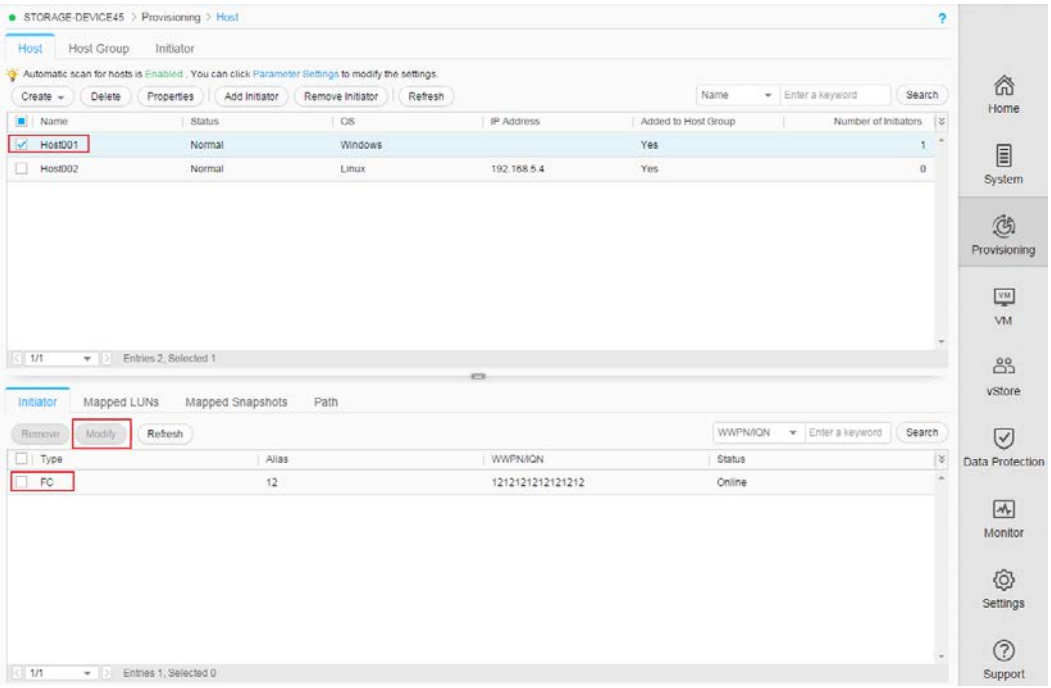

Step 3 Modify the initiator information.

In the **Modify Initiator** dialog box that is displayed, modify the initiator information based on the requirements of your operating system. The following figure shows the initiator information modification page.

Figure 7-15 Modifying initiator information

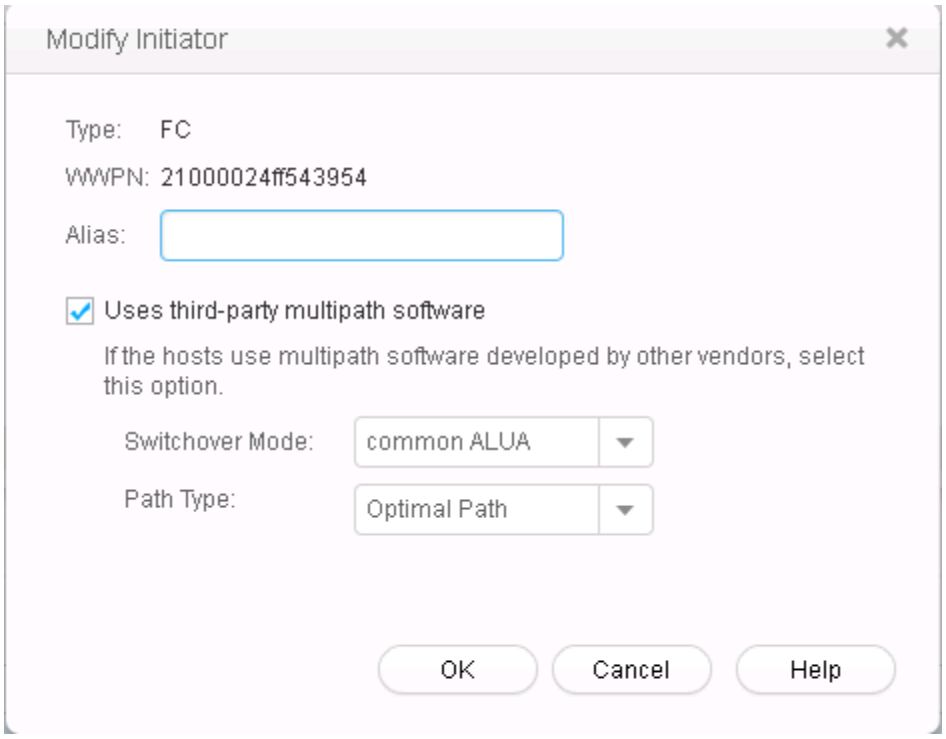

Step 4 Repeat the preceding operations to modify the information about other initiators on the host.

Step 5 Restart the host to enable the configuration to take effect.

----End

Recommended VMware NMP Configuration

This section provides recommended VMware NMP configurations for HyperMetro configuration with different ESXi versions' VMware NMP.

Table 7-5 lists the storage array configurations.

Table 7-5 Configuration on HyperMetro OceanStor V3/OceanStor V5/Dorado V3 storage when interconnected with VMware ESXi

OS	Storage Array Configuration						
	HyperMetro Working Mode	Storage	OS	OS Native Multipathing Software	Switchover Mode	Special Mode Type	Path Type
VMware ESXi	Load balancing	Local storage array	VMware ESXi	Enabled	Special mode	Mode 1	Optimal path
		Remote storage array	VMware ESXi	Enabled	Special mode	Mode 1	Optimal path

	Local preferred	Local storage array	VMware ESXi	Enabled	Special mode	Mode 2	Optimal path
		Remote storage array	VMware ESXi	Enabled	Special mode	Mode 2	Non-optimal path

For details about the VMware ESXi versions, see the compatibility list:

http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf

NOTICE

If a LUN has been mapped to the host, you must restart the host for the configuration to take effect after you modify the initiator parameters. If you configure the initiator for the first time, restart is not needed.

In OceanStor V3 V300R003C20, mode 1 and mode 2 are disabled by default. For details about how to enable them, see the *OceanStor 5300 V3&5500 V3&5600 V3&5800 V3&6800 V3 Storage System V300R003C20 Restricted Command Reference* or *OceanStor 18500 V3&18800 V3 Storage System V300R003C20 Restricted Command Reference*. Contact Huawei technical support engineers to obtain the documents.

In OceanStor V5 V500R007C00, OceanStor V3 V300R006C00SPC100, Dorado V3 V300R001C01SPC100, and later versions, you can configure mode 1 and mode 2 on DeviceManager directly.

Figure 7-16 Querying the special mode type

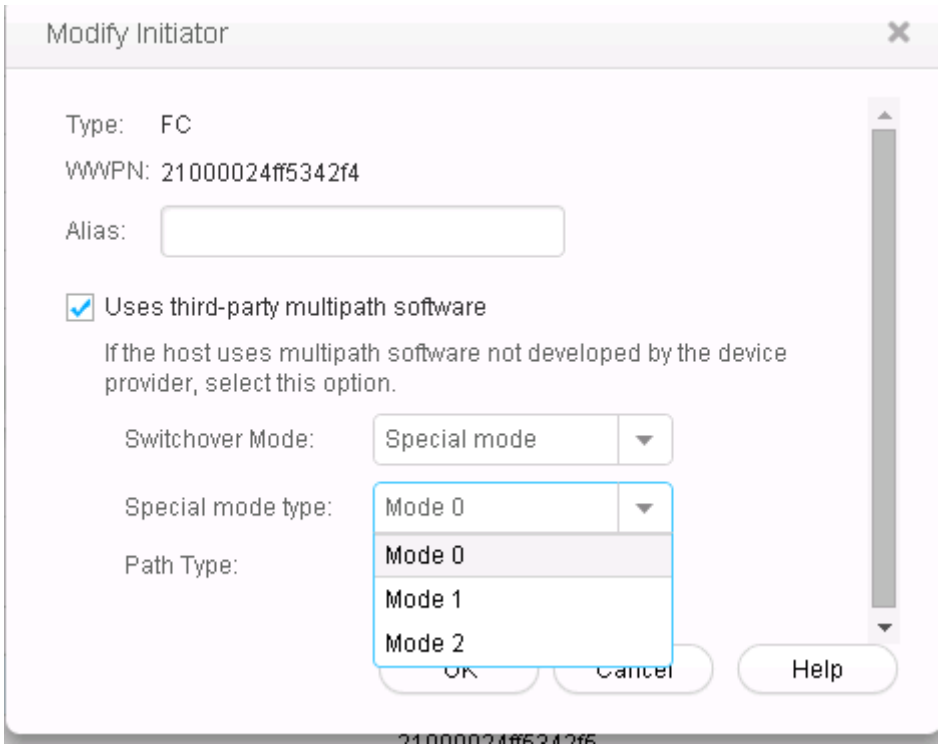

7.2.3 Host Configuration

7.2.3.1 Installing Multipathing Software

VMware NMP is ESXi native multipathing software and therefore does not need to be installed separately.

7.2.3.2 Configuring Multipathing

Recommended VMware NMP Configuration

Table 7-6 Recommended VMware NMP configuration for OceanStor V3/OceanStor V5/Dorado V3 HyperMetro configuration

Storage Device	Number of Controllers	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
VMware ESXi						
Dorado V3 Series, OceanStor V3/OceanStor V5/18000 V3 series V300R003C20	N/A	Y	N/A	VMW_SATP_ALUA	VMW_PSP_RR	See notes 1, 2, and 3

and later						

NOTE

1. For the MSCS and WSFC clusters deployed on VMware ESXi 5.1 or earlier VMs, you cannot set the RDM LUN to Round Robin, but can set it to FIXED. For details, see section 9.1 How To Query and Modify the Path Selection Policy? or *VMware KB1036189*.

2. When using All-Flash Array, it is advised to set **IO Operation Limit** to **1** on ESXi. For ESXi5.x, 6.x, the command is as follows:

esxcli storage nmp psp roundrobin deviceconfig set --device=**device_NAA**** --iops=1 --type iops

You need to change the preceding information in bold based on your actual situation.

3. Dorado V3 systems must be V300R001C01SPC100 or later versions, with multi-controller ALUA and HyperMetro ALUA supported.

For supported ESXi versions, see:

http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf

Precautions

When using HyperMetro with VMware ESXi, note the following precautions:

- The two LUNs in a HyperMetro pair must use the same LUN ID when being mapped to a VMware ESXi host (ESXi 6.5.0 GA build 4564106 or a follow-up version earleir than ESXi 6.5 U1 build 5969303).
  - You can query the host LUN ID mapped to the ESXi host in the Mapping View of OceanStor DeviceManager, as shown in Figure 7-17.
  - Before modifying the **Host LUN ID**, read the following warnings carefully since misoperations may cause service interruption. To modify the host LUN ID for a LUN, right-click the LUN and choose **Change host LUN ID** from the shortcut menu. In the displayed dialog box, set the same **Host LUN ID** value for the two storage devices in the HyperMetro pair and then click **OK**. For details, see Figure 7-18.

WARNING

Changing the host LUN ID with an incorrect procedure may cause service interruption.

If no datastore has been created on either LUN in the HyperMetro pair, you can directly change the host LUN ID for the LUNs. Wait for about 5 to 15 minutes after the modification is complete, and then run the **Rescan** command in the ESXi host CLI to check whether the LUNs in the HyperMetro pair have been restored and been online.

If a datastore has been created on either LUN in the HyperMetro pair and a service has been deployed in the datastore, change the host LUN ID using only the following two methods (otherwise, changing the host LUN ID for either LUN will cause the LUN to enter the PDL state and consequently interrupt services):

- Method 1: You do not need to restart the ESXi host. Migrate all VMs in the datastore deployed on the LUNs in the HyperMetro pair to another datastore, and then change the host LUN ID on the OceanStor DeviceManager. Wait for about 5 to 15 minutes after the modification is complete, and then run the **Rescan** command in the ESXi host CLI to verify that the LUNs in the HyperMetro pair have been restored and been online. Then, migrate the VMs back to the datastore deployed on the LUNs in the HyperMetro pair.
- Method 2: You need to restart the ESXi host. Power off all VMs in the datastore deployed on the LUNs in the HyperMetro pair to ensure that no service is running on the LUNs. Then, modify the host LUN ID on the OceanStor DeviceManager. Then, restart the ESXi host to make the modification take effect. After restarting the ESXi host, check whether the LUNs in the HyperMetro pair have been restored and been online.

Figure 7-17 Changing the host LUN ID (1)

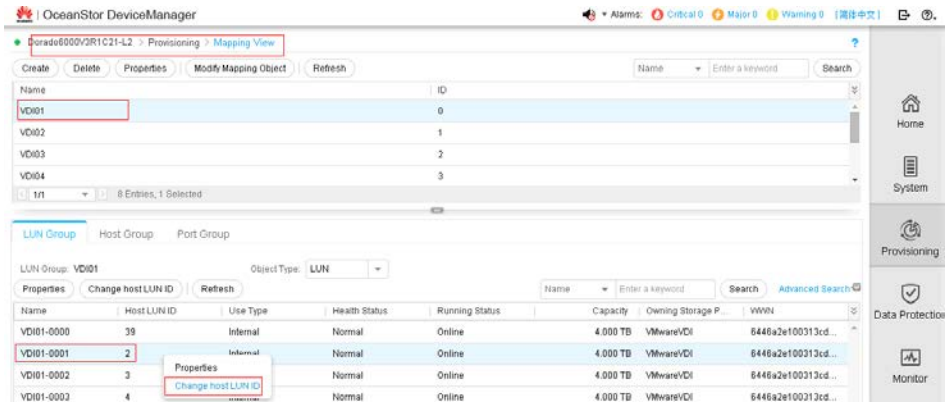

Figure 7-18 Changing the host LUN ID (2)

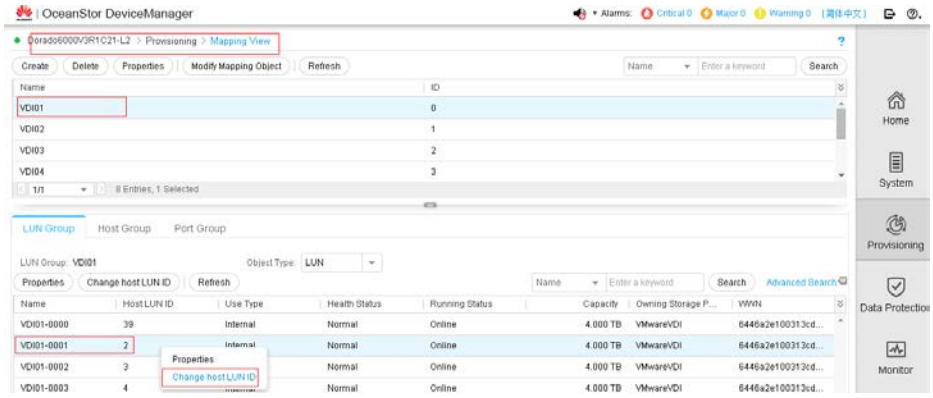

- For OceanStor V3 V300R003C20SPC200, a single array with ALUA enabled can have a maximum of 8 controllers; two active-active arrays with ALUA enabled also cannot have more than 8 controllers.
- VMware ESXi 6.0 U2 and later versions support HyperMetro configuration. Versions earlier than VMware ESXi 6.0 U2 have their defects.
- Dorado V3 must be V300R001C01SPC100 and later versions, supporting multi-controller and HyperMetro ALUA.

CAUTION

Before deploying HyperMetro solution based on VMware ESXi NMP, you need to consider the compatibility between components (such as storage system, operating system, HBAs, and switches) and the application software.

Check the interoperability matrix before deployment: http://support-open.huawei.com/ready/.

This document provides the configuration methods only for HyperMetro interoperability-related components. For specific interoperability configuration scenarios, you must check the corresponding HyperMetro interoperability matrix.

Configuring the Host

For HyperMetro storage, perform the following configuration steps.

Setting the VMware NMP Multipathing Rules

Run the following command on the host:

```
esxcli storage nmp satp rule add -V HUAWEI-M XSG1 -s VMW_SATP_ALUA -P VMW_PSP_RR -c tpgs_on
```

In these commands, HUAWEI is an example of storage Vendor and XSG1 is an example of storage Model. You need to change the two values based on your actual storage configurations. Table 7-7 provides the vendor and model information of Huawei mainstream storage devices.

Table 7-7 Huawei Storage vendor and model information

Storage Device	Vendor	Model
S2200T/S2600T/S5500T/S5600T/S5800T/S6800T	HUAWEI/SYMANTEC/HUASY	S2200T/S2600T/S5500T/S5600T/S5800T/S6800T
Dorado2100 G2	HUAWEI/SYMANTEC/HUASY	Dorado2100\ G2
Dorado5100	HUAWEI/SYMANTEC/HUASY	Dorado5100
18500	HUAWEI	HVS85T
18800/18800F	HUAWEI	HVS88T
V5 series V3 series 18000 V3 series 18000 V5 series Dorado V3 Series	HUAWEI	XSG1

NOTE

- After the command is executed, the new rule will immediately take effect for the newly mapped LUN, but will not take effect for previously mapped LUNs unless ESXi is restarted.

- Restart the host for the configuration to take effect. For details, see: https://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.storage.doc/GUID-D10F7E66-9DF1-4CB7-AAE8-6F3F1F450B42.html.

Set FC timeout parameters as follows:

- For FC networking, no configuration is required.
- For iSCSI networking, execute the following commands on ESXi hosts:

esxcli iscsi adapter param set -A *vmhba35* -k NoopOutInterval -v 3

esxcli iscsi adapter param set -A *vmhba35* -k NoopOutTimeout -v 10

esxcli iscsi adapter param set -A *vmhba35* -k RecoveryTimeout -v 3

NOTE

1. All the preceding commands are available only in VMware 5.0 and later versions. For HyperMetro-supported VMware versions, see:

   http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf

2. The field in italic (vmhba35 in this example) indicates the iSCSI initiator. You need to change it according to your own hosts.

3. The configurations will not take effect unless the host restarts.

4. If you do not perform the preceding configurations but only retain the ESXi host settings, the path switchover time may reach up to 35s. However, the preceding configurations can shorten the path switchover time to about 16s.

Configuring a VMware Cluster

If you want to configure VMware clusters, see the BC&DR Solution Product Documentation (Active-Active Data Center). See section "Virtualization Platform Configuration" in this document. The contents in this section are as follows.

Mandatory Configuration Items:

- Deploy ESXi hosts across data centers in a HA cluster and configure the cluster with HA advanced parameter **das.maskCleanShutdownEnabled = True** for VMware vSphere 5.0 u1, 5.1, and 5.5 versions.
- A VM service network requires L2 interworking between data centers for VM migration between data centers without affecting VM services.
- For VMware vSphere 5.0 u1, later 5.0 versions, and 5.1 versions, log in to the CLI of each ESXi host using SSH and add **Disk.terminateVMOnPDLDefault = True** in the **/etc/vmware/settings** file.
- For VMware vSphere 5.5 and its update versions, log in to the host using the vSphere Client.
  - Set **VMkernel.Boot.terminateVMOnPDL = True**. The parameter forcibly powers off VMs on a datastore when the datastore enters the PDL state.

Figure 7-19 Boot parameter settings

- Set **Disk.AutoremoveOnPDL = 0**. This setting ensures that datastores in the PDL state will not be automatically removed.

Figure 7-20 Disk parameter settings

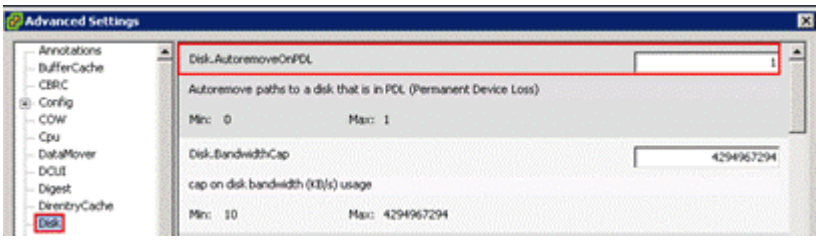

- For VMware vSphere 6.0 u2 and later versions:
  - After connecting to vCenter through the Web Client, enter the cluster HA configuration. The configuration requirements are as follows.

Figure 7-21 vSphere 6.0 cluster configuration

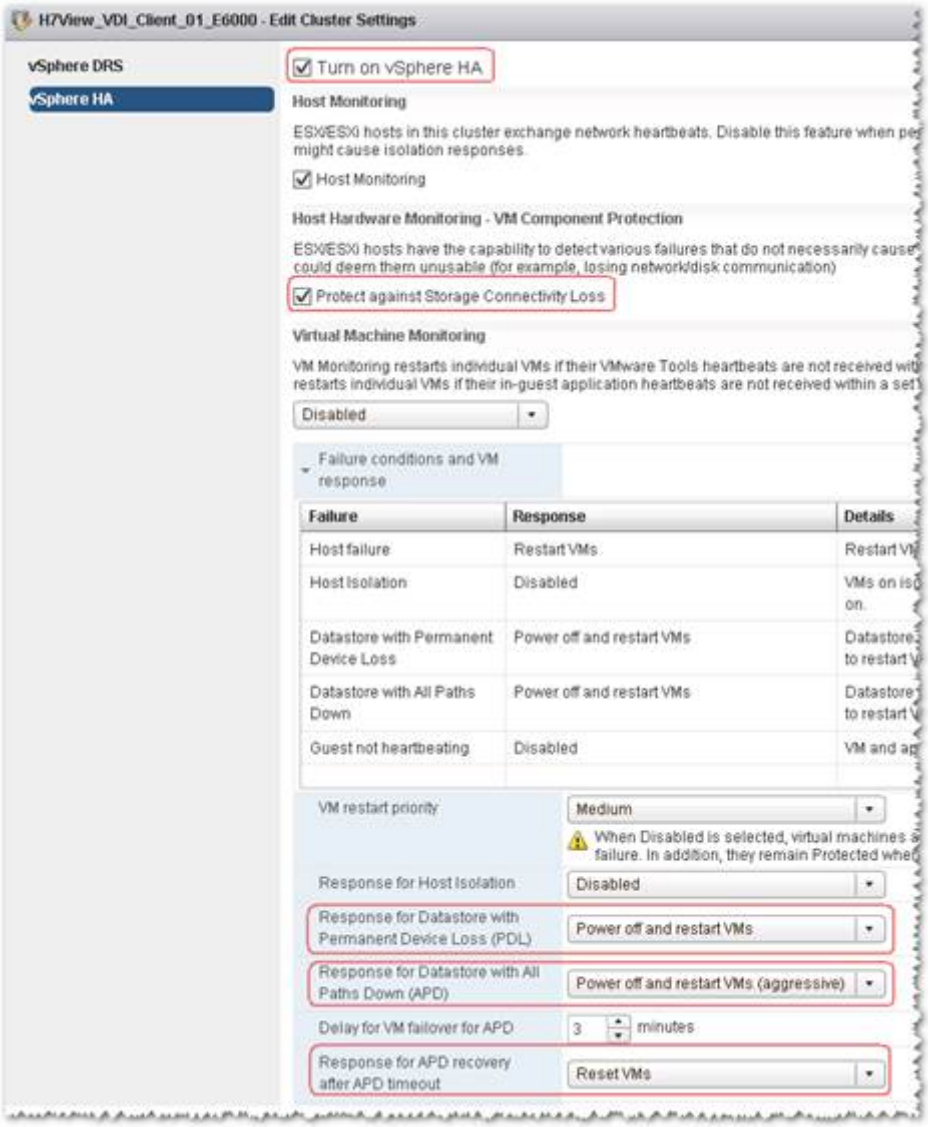

- For VMware vSphere 6.5:
  - After connecting to vCenter through the Web Client, enter the cluster HA configuration. The configuration requirements are as follows.

Figure 7-22 vSphere 6.5 cluster configuration-1

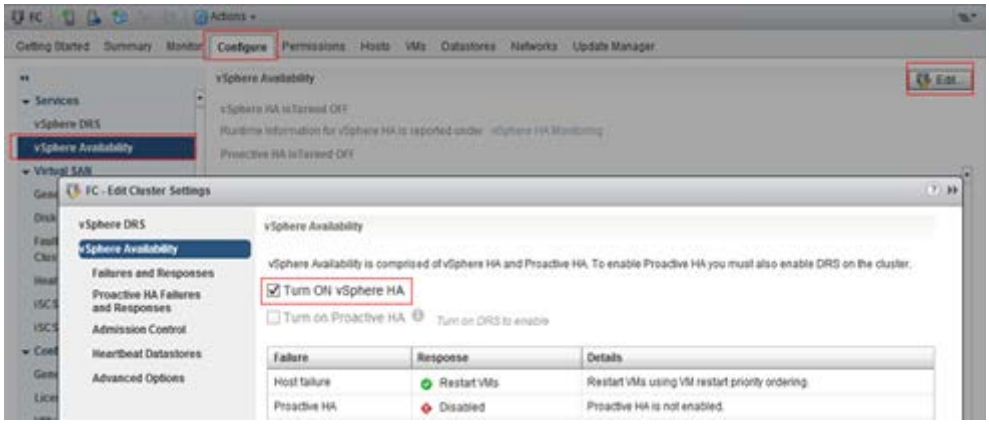

Figure 7-23 vSphere 6.5 cluster configuration-2

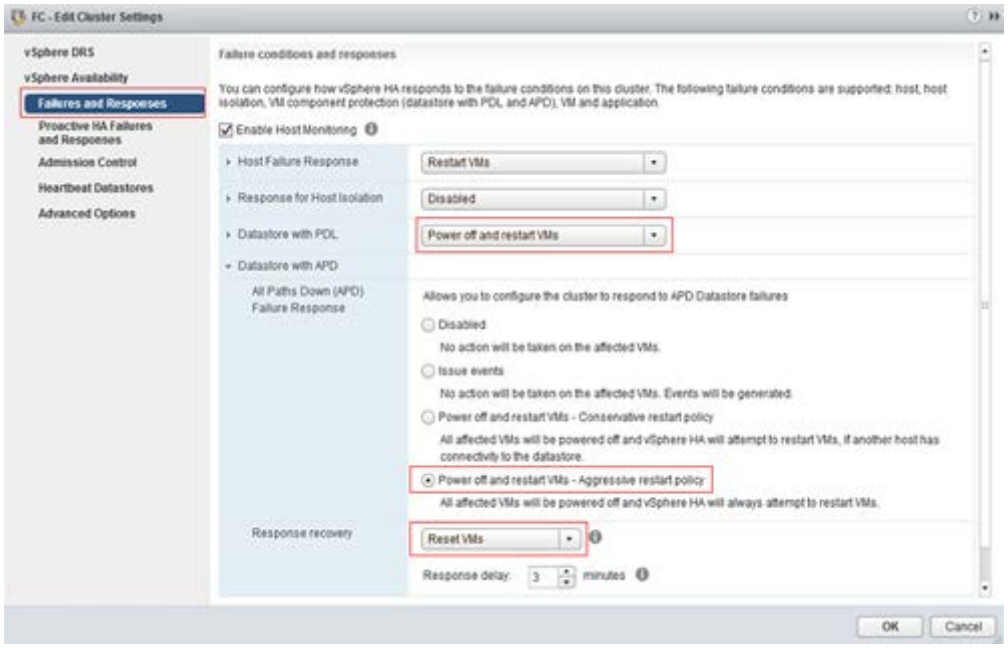

CAUTION

For VMware vSphere 5.1 to 5.5 versions, restart hosts for the configuration to take effect.

For VMware vSphere 6.0 U2 to 6.5 versions, re-enable the HA cluster to make the configuration take effect without restarting hosts.

Recommended Configuration Items

- The vMotion network, service network, and management network are configured with different VLAN IDs to avoid network interference.
- The management network includes the vCenter Server management node and ESXi hosts that are not accessible to external applications.
- The service network is divided into VLANs based on service requirements to ensure logical isolation and control broadcast domains.
- In a single cluster, the number of hosts does not exceed 16. If a cluster has more than 16 hosts, you are advised to use the hosts to create multiple clusters across data centers.
- A DRS group is configured to ensure that VMs can be recovered first in the local data center in the event of the breakdown of a single host.

7.2.4 Verification

Log in to the vSphere Client or vSphere Web Client to check whether the LUNs' VMware NMP configurations have taken effect.

Figure 7-24 Verifying the multipathing configuration

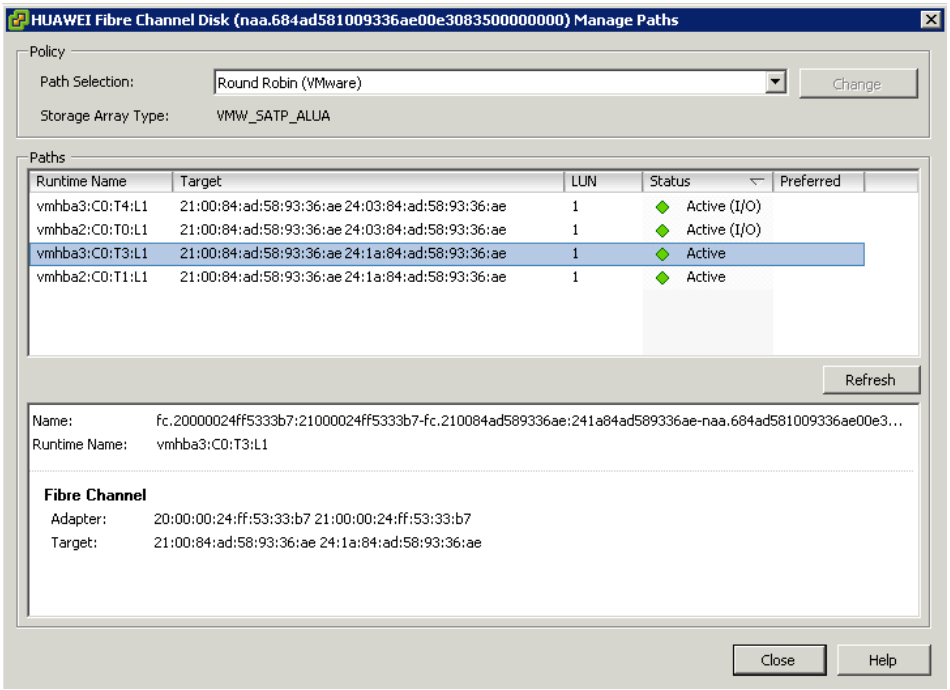

8 Configuring Multipathing in Non-HyperMetro Scenarios

8.1 UltraPath

8.1.1 Storage System Configuration

If you use UltraPath, retain the default initiator settings. Do not select **Uses third-party multipath software**.

Figure 8-1 Initiator setting when UltraPath is used

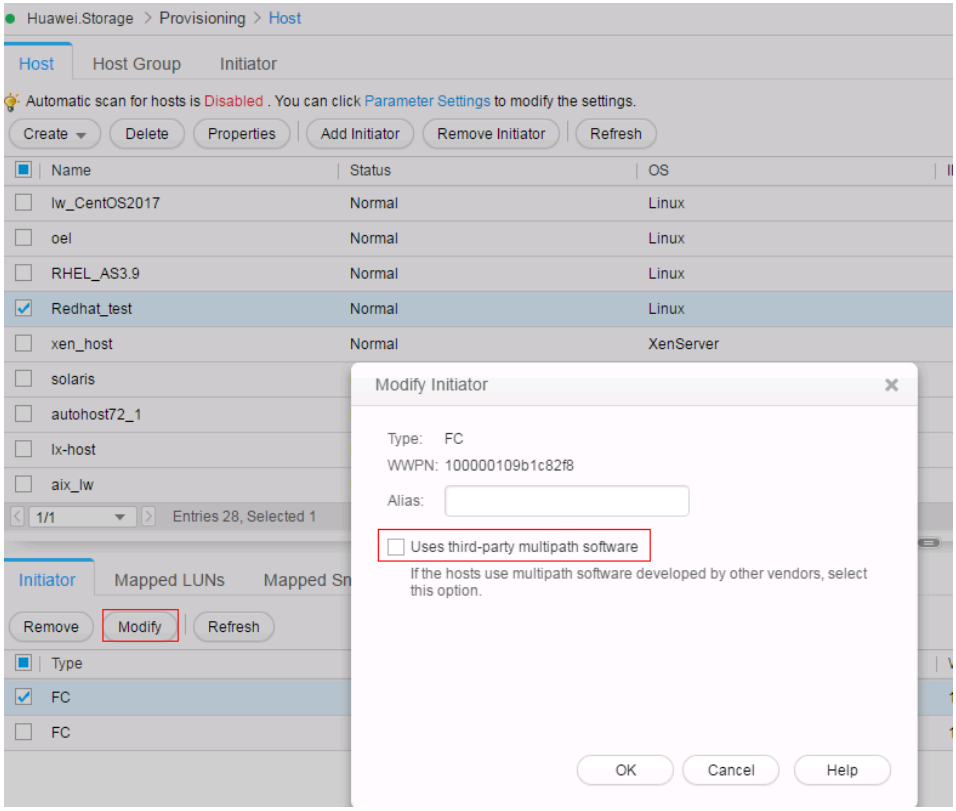

8.1.2 Host Configuration

Install and configure UltraPath by following instructions in the *OceanStor UltraPath for Linux User Guide*.

8.2 OS Native Multipathing Software

Different OS versions support different VMware NMP policies.

This chapter describes the VMware NMP policy recommended by HUAWEI for establishing connections between VMware ESXi and HUAWEI storage systems:

New-version Huawei storage (namely, storage that supports multi-controller ALUA and ALUA HyperMetro): OceanStor V5 series V500R07C00 and later, OceanStor V3/18000 V3 series V300R003C20 (V300R003C20SPC200 and later)/V300R006C00 (V300R006C00SPC100 and later), Dorado V3 V300R001C01 (V300R001C01SPC100 and later)

Old-version Huawei storage (namely, storage that does not support multi-controller ALUA or ALUA HyperMetro): OceanStor T V1/T V2/18000 V1/V300R001/V300R002/V300R003C00/V300R003C10/V300R005, Dorado V300R001C01

8.2.1 New-Version Huawei Storage

8.2.1.1 Recommended NMP Configurations

Table 8-1 provides the recommended NMP configurations when different ESX/ESXi versions interconnect with HUAWEI storage.

CAUTION

The recommended NMP configuration is a universal configuration, but may be not the best configuration in your storage environments.

For example, VMW_PSP_RR has better performance than VMW_PSP_FIXED, but VMW_PSP_RR has some use restrictions: for the MSCS and WSFC clusters deployed on VMs, you can set the RDM LUN to PSP_RR only in VMware ESXi 5.5 and later versions. For details, see VMware KB 2147662.

If you want to configure an optimal path policy, contact local Huawei support.

Table 8-1 Recommended NMP configurations when different ESXi versions interconnect with OceanStor V3/OceanStor V5/Dorado V3

Storage Device	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
ESXi 5.0.*					
OceanStor	Yes	N/A	**VMW_S**	**VMW_P**	See Notes 1, 2, 3, and 6.

Storage Device	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
V5/18000 V5 OceanStor V3/18000 V3 V300R003C20 and later versions Dorado V3			**ATP_ALUA**	**SP_FIXED**	
ESXi 5.1.*					
OceanStor V5/18000 V5 OceanStor V3/18000 V3 V300R003C20 and later versions Dorado V3	Yes	Yes	**VMW_SATP_ALUA**	**VMW_PSP_FIXED**	See Notes 1, 2, 3, and 6.
		No	**VMW_SATP_ALUA**	**VMW_PSP_RR**	See Notes 1, 2, 3, 4, 5, and 6.
ESXi 5.5.*, 6.0.*, 6.5.*					
OceanStor V5/18000 V5 OceanStor V3/18000 V3 V300R003C20 and later versions Dorado V3	Y	N/A	**VMW_SATP_ALUA**	**VMW_PSP_RR**	See Notes 1, 2, 3, 4, 5, and 6.

NOTE

1. Failback is supported upon recovery from a path fault.

2. On the VMware command line interface, run the following commands to add rules:

esxcli storage nmp satp rule add -V **HUAWEI** -M **XSG1** -s **VMW_SATP_ALUA** -P **VMW_PSP_RR** -c tpgs_on.

You need to change the preceding information in bold based on your actual situation. For details, see 7.2.3.2 Configuring Multipathing.

After the command is executed, the new rule will immediately take effect for the newly mapped LUN, but will not take effect for previously mapped LUNs unless ESXi is restarted.

3. This configuration is recommended for ALUA-enabled storage.

4. For the MSCS and WSFC clusters deployed on VMware ESXi 5.1 and earlier versions, you cannot set the RDM LUN to Round Robin, but can set it to FIXED. For details, see section 9.1 How To Query and Modify the Path Selection Policy? or *VMware KB1036189*.

5. When using All-Flash Array, it is advised to set *IO Operation Limit* to 1 on ESXi. For ESXi5.x, 6.x, run the command:

esxcli storage nmp psp roundrobin deviceconfig set --device=**device_NAA**** --iops=1 --type iops

You need to change the preceding information in bold based on your actual situation.

6. Dorado V3 storage must be V300R001C01SPC100 or later versions, supporting multi-controller ALUA and HyperMetro ALUA.

For supported ESXi versions, see:

http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf

8.2.1.2 Storage System Configuration

Configuring the ALUA Mode

For non-HyperMetro configuration, use the configuration listed in Table 8-2.

Table 8-2 Configuration on non-HyperMetro OceanStor V5/OceanStor V3/Dorado V3 storage when interconnected with VMware ESXi

Operating System	Configuration on the Storage Array					
	Storage	Operating System	Third-Party Multipath Software	Switchover Mode	Special Mode	Path Type
ESXi 5.0.x, ESXi 5.1.x, ESXi 5.5.x, ESXi 6.0.x, ESXi 6.5.x	Two-controller, multi-controller	VMware ESX	Enable	Special mode	Mode 1	Optimal Path

NOTE
For supported ESXi versions, see:
http://support-open.huawei.com/ready/pages/user/compatibility/support-matrix.jsf.

NOTICE
If a LUN has been mapped to the host, you must restart the host for the configuration to take effect after you modify the initiator parameters. If you configure the initiator for the first time, restart is not needed.

In OceanStor V3 V300R003C20, mode 1 and mode 2 are disabled by default. For details about how to enable them, see the *OceanStor 5300 V3&5500 V3&5600 V3&5800 V3&6800 V3 Storage System V300R003C20 Restricted Command Reference* or *OceanStor 18500 V3&18800 V3 Storage System V300R003C20 Restricted Command Reference*. Contact Huawei technical support engineers to obtain the documents.

In OceanStor V5 V500R007C00, OceanStor V3 V300R006C00SPC100, Dorado V3 V300R001C01SPC100, and later versions, you can configure mode 1 and mode 2 on DeviceManager directly.

Figure 8-2 Querying the special mode type

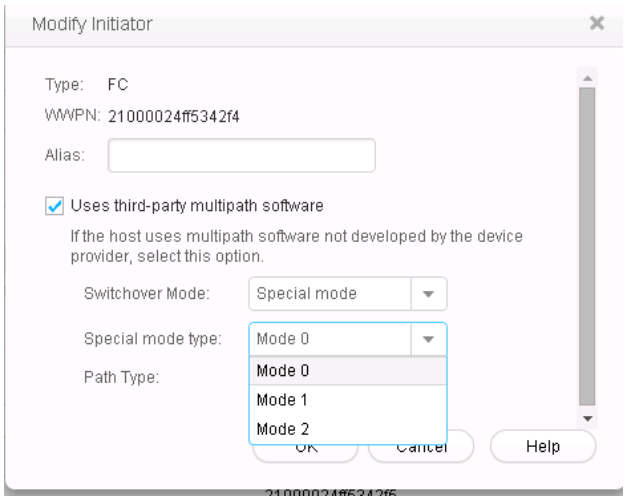

Configuring the AA Mode

Configure the AA mode for the host initiator on the storage as follows:

- Select **Uses third-party multipath software**.
- Set **Switchover Mode** to **ALUA not used**.
- Set **Path Type** to **Optimal Path**.

Figure 8-3 Configuring the AA mode for the host initiator

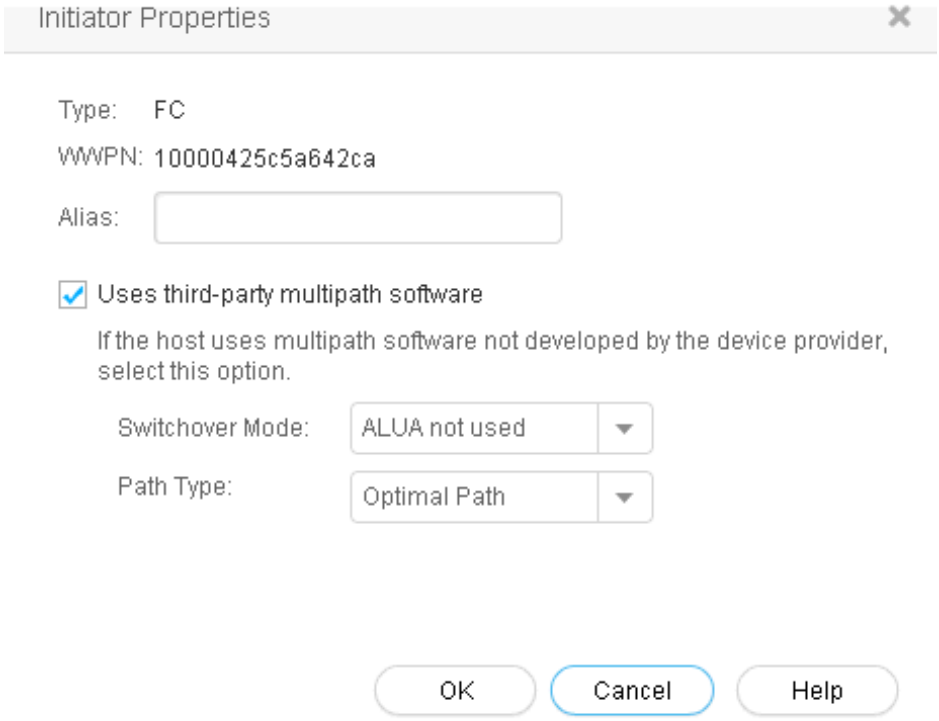

8.2.1.3 Host Configuration

Configuring the ALUA Mode

For non-HyperMetro configuration, perform the following steps to configure VMware NMP.

After enabling ALUA on Huawei storage, perform the following steps to add multipathing rule on the ESXi hosts:

Step 1 Check the vendor and model information of the storage systems.

Use the SSH tool to log in to the ESXi Shell, and run the **esxcli storage core device list** to view **Vendor** and **Model** information of the storage system.

```
[root@localhost:~] esxcli storage core device list
naa.630d17e100b3020708d125f600000026
 Display Name: HUAWEI Fibre Channel Disk (naa.630d17e100b3020708d125f600000026)
 Has Settable Display Name: true
 Size: 15360
 Device Type: Direct-Access
 Multipath Plugin: NMP
 Devfs Path: /vmfs/devices/disks/naa.630d17e100b3020708d125f600000026
 Vendor: HUAWEI
 Model: XSG1
 Revision: 4303
 SCSI Level: 6
 Is Pseudo: false
 Status: on
```

Step 2 Add multipathing rules.

Run different configuration commands for the two different multipathing modes:

- **VMW_PSP_Fixed**:

  esxcli storage nmp satp rule add -V *HUA*WEI -M **XSG1** -s VMW_SATP_ALUA -P **VMW_PSP_FIXED** -c tpgs_on

- **VMW_PSP_RR**:

  esxcli storage nmp satp rule add -V HUAWEI -M **XSG1** -s VMW_SATP_ALUA -P **VMW_PSP_RR** -c tpgs_on

In these commands, HUAWEI is an example of storage Vendor and XSG1 is an example of storage Model. You need to change the two values based on your actual storage configurations. Table 8-3 provides the vendor and model information of Huawei mainstream storage devices.

Table 8-3 Huawei Storage vendor and model information

Storage Device	Vendor	Model
S2200T/S2600T/S5500T/S5600T/S5800T/S6800T	HUAWEI/SYMANTEC/HUASY	S2200T/S2600T/S5500T/S5600T/S5800T/S6800T
Dorado2100 G2	HUAWEI/SYMANTEC/HUASY	Dorado2100\ G2
Dorado5100	HUAWEI/SYMANTEC/HUASY	Dorado5100

Storage Device	Vendor	Model
18500	HUAWEI	HVS85T
18800/18800F	HUAWEI	HVS88T
V5 series 18000 V5 series V3 series 18000 V3 series Dorado V3 Series	HUAWEI	XSG1

NOTE

To delete existing multipathing configuration rules, replace [path policy] with the configured path mode (for example, VMW_PSP_Fixed) and then run the following command:

esxcli storage nmp satp rule remove -V HUAWEI -M XSG1 -s VMW_SATP_ALUA -P [path policy] -c tpgs_on

Step 3 Confirm that the rule is added successfully:

esxcli storage nmp satp rule list | grep *HUAWEI*

----End

WARNING

After the command is executed, the new rule will immediately take effect for the newly mapped LUN, but will not take effect for previously mapped LUNs unless ESXi is restarted.

Configuring the AA Mode

Select one path of the storage system as its preferred path, as shown in Figure 8-4. For a storage device that has a preferred path selected, first set another path as the preferred path and then set original path to the preferred path again.

Figure 8-4 Setting the preferred path for a storage device

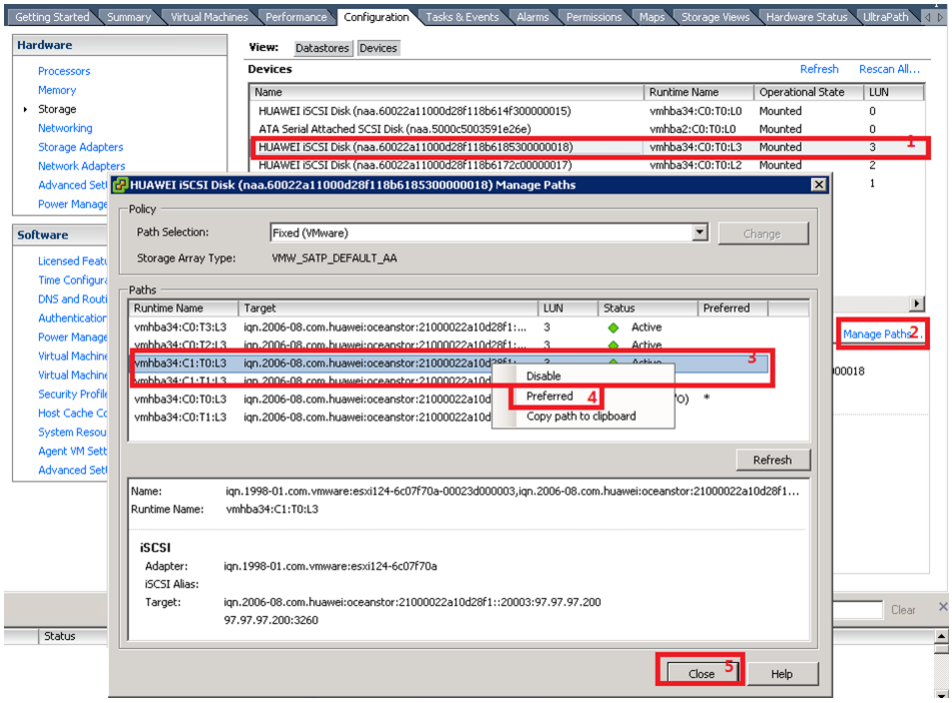

8.2.1.4 Verification

After the configuration, run the following command to confirm that the multipathing software configuration is correct.

```
esxcli storage nmp satp rule list|grep -i huawei
```

The following figure shows the command output.

Figure 8-5 Command output

Run the following command to check whether the path information takes effect.

```
esxcli storage nmp device list -d=naa.6xxxxxxx.
```

The following figure shows the command output.

Figure 8-6 VMware path information

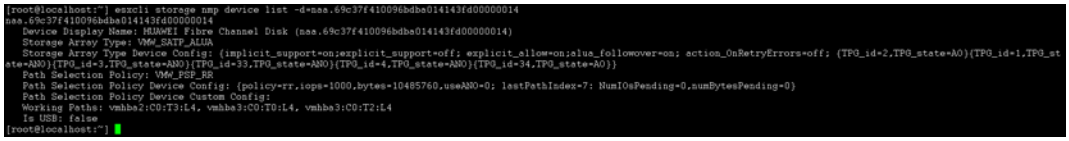

The path information is displayed in the unit of port group rather than in the unit of a single path.

8.2.2 Old-Version Huawei Storage

8.2.2.1 Recommended NMP Configurations

Table 8-4 provides the recommended NMP configurations when different ESX/ESXi versions interconnect with HUAWEI storage.

CAUTION

The recommended NMP configuration is a universal configuration, but may be not the best configuration in your storage environments.

For example, VMW_PSP_RR has better performance than VMW_PSP_FIXED, but VMW_PSP_RR has some use restrictions: for the MSCS and WSFC clusters deployed on VMs, you can set the RDM LUN to PSP_RR only in VMware ESXi 5.5 and later versions. For details, see VMware KB 2147662.

If you want to configure an optimal path policy, contact local Huawei support.

Table 8-4 Recommended NMP configurations when different ESXi versions interconnect with Huawei old-version storage

Storage Device	Number of Controllers	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
ESX 4.0.*						
S2600, S5000 series T V1 series, Dorado5100 Dorado2100 G2, T V2 series 18000 V1 series V3 series	2	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.
	4 or more	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.
ESXi 4.1.*						
S2600, S5000 series T V1 series, Dorado5100 Dorado2100	2	Y	N/A	VMW_SATP_ALUA	VMW_PSP_FIXED_AP	See notes 2, 3, and 4.
	4 or more	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.

Storage Device	Number of Controllers	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
G2 T V2 series 18000 V1 series V3 series						
ESXi 5.0.*						
S2600, S5000 series	2	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.
T V1 series, Dorado5100 Dorado2100 G2 T V2 series 18000 V1 series V3 series	2	Y	N/A	VMW_SATP_ALUA	VMW_PSP_FIXED	See notes 2, 3, and 4.
	4 or more	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, 5, and 7.
ESXi 5.1.*						
S2600, S5000 series	2	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.
T V1 series, Dorado5100 Dorado2100 G2 T V2 series 18000 V1 series V3 series	2	Y	Y	VMW_SATP_ALUA	VMW_PSP_FIXED	See notes 2, 3, and 4.
			N	VMW_SATP_ALUA	VMW_PSP_RR	See notes 2, 3, 5, and 6.
	4 or more	N	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, 5, and 7.	
ESXi 5.5.*, 6.0.*, 6.5.*						
S2600, S5000 series	2	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, and 5.
T V1 series, Dorado5100 Dorado2100 G2, T V2 series	2	Y	N/A	VMW_SATP_ALUA	VMW_PSP_RR	See notes 2, 3, and 4.
	4 or more	N	N/A	VMW_SATP_DEFAULT_AA	VMW_PSP_FIXED	See notes 1, 2, 5, and 7.

Storage Device	Number of Controllers	ALUA Enabled or Not	VM Cluster	Recommended SATP Type	Recommended PSP Type	Remarks
18000 V1 series V3 series						

NOTE

- 1. You need to manually set the primary path for each LUN on the vSphere Client. For the default preferred LUN, you can set a non-preferred path and then set the preferred path.
- 2. A switchback is supported upon recovery from a path fault.
- 3. On the VMware command line interface, run the following commands to add rules:

  For ESX/ESXi 4.x:

  esxcli nmp satp addrule -V **HUAWEI** -M **XSG1** -s **VMW_SATP_ALUA** -o tpgs_on

  esxcli nmp satp setdefaultpsp -s **VMW_SATP_ALUA** -P **VMW_PSP_FIXED**

  For ESXi 5.0 and later:

  esxcli storage nmp satp rule add -V **HUAWEI** -M **XSG1** -s **VMW_SATP_ALUA** -P **VMW_PSP_RR** -c tpgs_on

  You need to change the preceding information in bold based on your actual situation.

  After the command is executed, the new rule will immediately take effect for the newly mapped LUN, but will not take effect for previously mapped LUNs unless ESXi is restarted.
- 4. This configuration is recommended for ALUA-enabled storage.
- 5. This configuration is recommended for ALUA-disabled storage.
- 6. For the MSCS and WSFC clusters deployed on VMware ESXi 5.1 and earlier versions, you cannot set the RDM LUN to Round Robin, but can set it to FIXED. For details, see section 9.1 How To Query and Modify the Path Selection Policy? or *VMware KB1036189*.
- 7. For any future controller expansion purpose, you are advised to disable ALUA and configure VMW_SATP_DEFAULT_AA.

CAUTION

- To avoid the Ping-Pong effect in VMware ESX 4.0 clusters, you are advised to disable ALUA.
- If a path policy or preferred path is set on the VMware page before or after rules are added, this setting prevails. The newly added rule will not take effect to any LUN that has been configured with a path policy or preferred path.
- For a LUN already configured with a preferred path, first switch to the non-preferred path and then set back to the preferred path, thereby ensuring normal switchback upon recovery from the fault.

OceanStor 18000/T V2/V3 supports two or more controllers.

When the storage systems have two controllers, they support ALUA and A/A.

When the storage systems have more than two controllers, they support only A/A but not ALUA (as of the release of this document).

To facilitate future capacity expansion, you are advised to disable ALUA on the OceanStor 18000/T V2/V3 and its host.

8.2.2.2 Storage System Configuration

Configuring the ALUA Mode

- **T Series V100R005/Dorado2100/Dorado5100/Dorado2100 G2**

  Use the Huawei OceanStor storage management system to enable ALUA for all the host initiators, as shown in Figure 8-7.

Figure 8-7 Enabling ALUA for T series V100R005/Dorado2100/Dorado5100/Dorado2100 G2

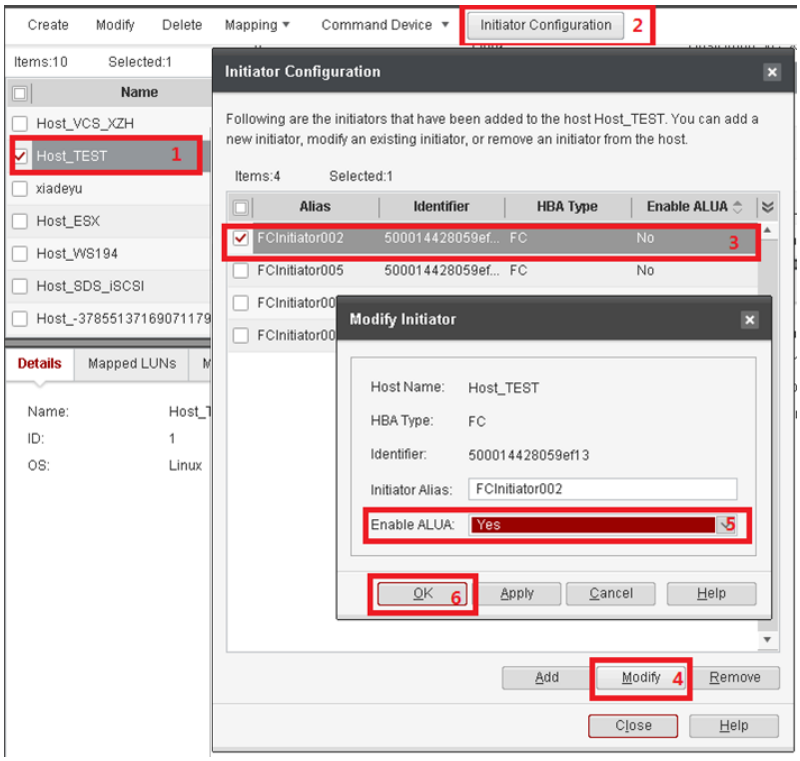

- **T Series V200R002/18000 Series/V3 Series/18000 V3 Series**

  Use the Huawei OceanStor storage management system to enable ALUA for all the host initiators, as shown in Figure 8-8.

Figure 8-8 Enabling ALUA for T series V200R002/18000 series/V3 series/18000 V3 series

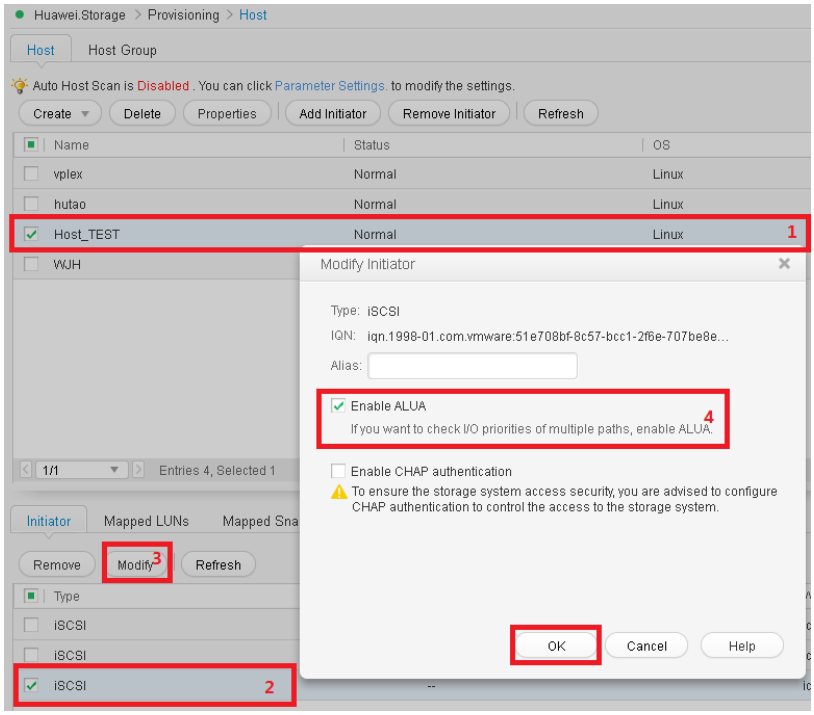

NOTE

If there are more than two controllers and ALUA is disabled by default, the ALUA status cannot be changed.

Configuring the AA Mode

On the storage, the host initiator default setting is the AA mode, and therefore no manual configuration is required.

NOTICE

If a LUN has been mapped to the host, you must restart the host for the configuration to take effect after you modify the initiator parameters. If you configure the initiator for the first time, restart is not needed.

8.2.2.3 Host Configuration

Configuring the ALUA Mode

The same as "Configuring the ALUA Mode" under section 8.2.1.3 Host Configuration.

Configuring the AA Mode

Select one path of the storage system as its preferred path, as shown in Figure 8-9. For a storage device that has a preferred path selected, first set another path as the preferred path and then set original path to the preferred path again.

Figure 8-9 Setting the preferred path for a storage device

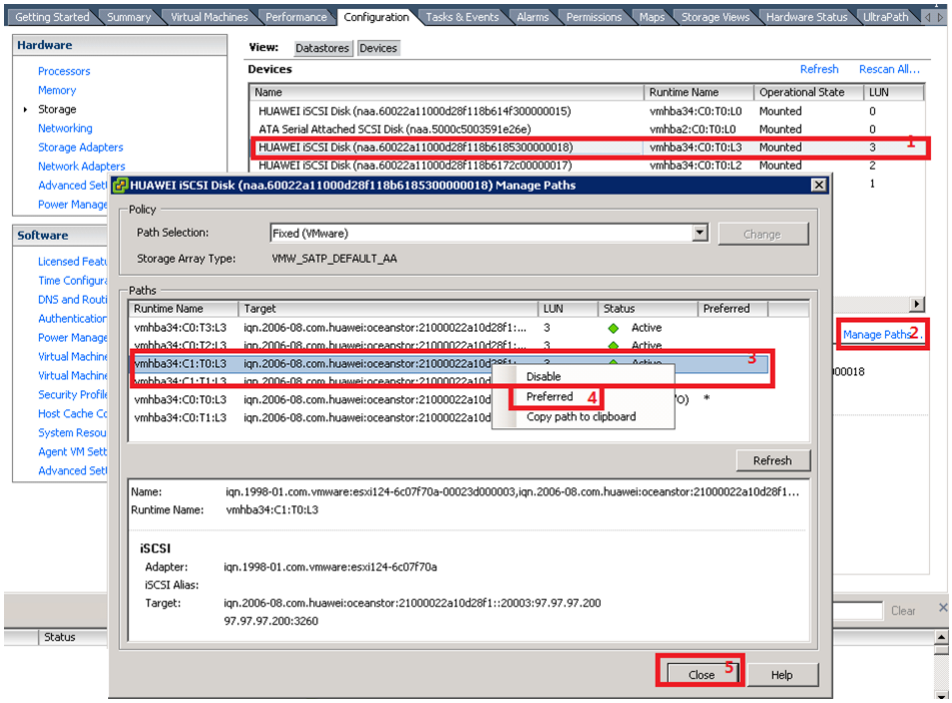

8.2.2.4 Verification

After the configuration, run the following command to confirm that the multipathing software configuration is correct.

```
esxcli storage nmp satp rule list|grep -i huawei
```

The following figure shows the command output.

Figure 8-10 Command output

Run the following command to check whether the path information takes effect.

```
esxcli storage nmp device list -d=naa.6xxxxxxx.
```

The following figure shows the command output.

Figure 8-11 VMware path information

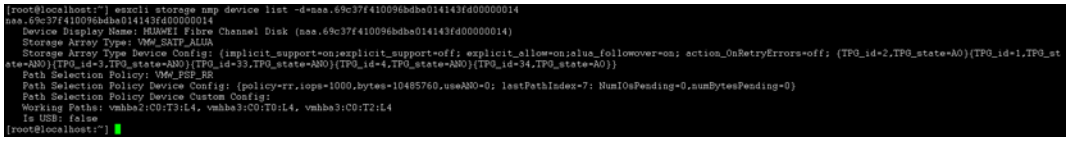

The path information is displayed in the unit of port group rather than in the unit of a single path.

9 FAQs

9.1 How To Query and Modify the Path Selection Policy?

You can check and modify a path policy by running command lines.

This section describes how to use commands to check and modify the path policy.

Querying the Path Policy of a Single LUN

- ESX/ESXi 4.0

The following is an example command for querying a path policy:

```
[root@e4 ~]# esxcli nmp device list -d naa.6666666100666650092f53300000045
naa.6666666100666650092f53300000045
   Device Display Name: HUASY iSCSI Disk (naa.6666666100666650092f53300000045)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config:
   Path Selection Policy: VMW_PSP_FIXED
   Path Selection Policy Device Config:
{preferred=vmhba33:C0:T1:L0;current=vmhba33:C0:T0:L0}
   Working Paths: vmhba33:C0:T0:L0
[root@e4 ~]#
```

- ESX/ESXi 4.1

The following is an example command for querying a path policy:

```
[root@localhost ~]# esxcli corestorage device list
naa.60022a11000416611b2a9d180000000a
   Display Name: HUASY Fibre Channel Disk (naa.60022a11000416611b2a9d180000000a)
   Size: 56320
   Device Type: Direct-Access
   Multipath Plugin: NMP
   Devfs Path: /vmfs/devices/disks/naa.60022a11000416611b2a9d180000000a
   Vendor: HUASY
   Model: S5600T
   Revision: 2105
   SCSI Level: 4
   Is Pseudo: false
   Status: on
   Is RDM Capable: true
   Is Local: false
```

```
   Is Removable: false
   Attached Filters:
   VAAI Status: unknown
   Other UIDs: vml.020001000060022a11000416611b2a9d180000000a533536303054
 [root@localhost ~]# esxcli nmp device list
naa.60022a11000416611b2a9d180000000a
   Device Display Name: HUASY Fibre Channel Disk
(naa.60022a11000416611b2a9d180000000a)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;explicit_support=on;
explicit_allow=on;alua_followover=on;{TPG_id=2,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_FIXED
   Path Selection Policy Device Config: Current Path=vmhba1:C0:T0:L1
   Working Paths: vmhba1:C0:T0:L1
```

esxcli corestorage device list is used to display existing disks.

esxcli nmp device list is used to display disk paths.

- ESXi 5.0 and later

The following is an example command for querying a path policy:

```
~ # esxcli storage nmp device list
naa.666666610066666502b85d9200000014
 Device Display Name: HUASY iSCSI Disk (naa.666666610066666502b85d9200000014)
 Storage Array Type: VMW_SATP_ALUA
 Storage Array Type Device Config: {implicit_support=on;explicit_support=on;
explicit_allow=on;alua_followover=on;{TPG_id=1,TPG_state=AO}{TPG_id=2,TPG_state=AN
O}}
 Path Selection Policy: VMW_PSP_RR
 Path Selection Policy Device Config: Current Path=vmhba39:C0:T0:L2
 Path Selection Policy Device Custom Config:
 Working Paths: vmhba39:C0:T0:L2
```

Modifying the Path Policy for a Single LUN

You can run the following command to modify the PSP path policy of a LUN:

- VMware ESXi/ESX 4.1

```
# esxcli nmp device setpolicy -d naa.6006016055711d00cff95e65664ee011
--psp=VMW_PSP_FIXED
```

- VMware ESXi 5.0 and later

```
# esxcli storage nmp device set -d naa.6006016055711d00cff95e65664ee011
--psp=VMW_PSP_FIXED
```

Run the following command to check the modification result:

- VMware ESXi/ESX 4.1

```
# esxcli nmp device list -d naa.6006016055711d00cff95e65664ee011 |grep PSP
naa.6006016055711d00cff95e65664ee011
```

- VMware ESXi 5.0 and later

```
# esxcli storage nmp device list -d naa.6006016055711d00cff95e65664ee011 |grep PSP
naa.6006016055711d00cff95e65664ee011
```

9.2 VMware APD and PDL

For details about All-Paths-Down (APD) and Permanent Device Loss (PDL) in the VMware ESXi system, see the VMware Knowledge Base:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004684

9.3 How Can I Select a Fixed Preferred Path for a Storage Device with Active-Active Controllers?

Determine which path is preferred based on the performance optimization principle and load balance principle.

- Select the path connected to the working controller of the LUN.
- If there are multipath paths connected to the working controller of the LUN, distribute the preferred paths evenly to the multiple paths.

9.4 How Can I Determine Which Controller a Path Is Connected to?

A path can be located using the initiator (host port) and the target (storage device port). Figure 9-1 illustrates how to obtain the path information about a storage device.

- Determine a path connected to a Fibre Channel storage device by the initiator name's former part (for example **vmhba4:C0**) and the target name. Then based on the target name's latter part (for example **20:08:f8:4a:bf:57:af:b7**), determine the storage device's Fibre Channel port that corresponds to the WWPN. The previous information can be combined to determine which controller the path connects to.
- Determine the IPv4 address of an iSCSI storage device's Ethernet port by the target name's latter part (for example **97.97.97.201:3260**). Then you can determine which controller the path connects to.

Figure 9-1 Port information about a path

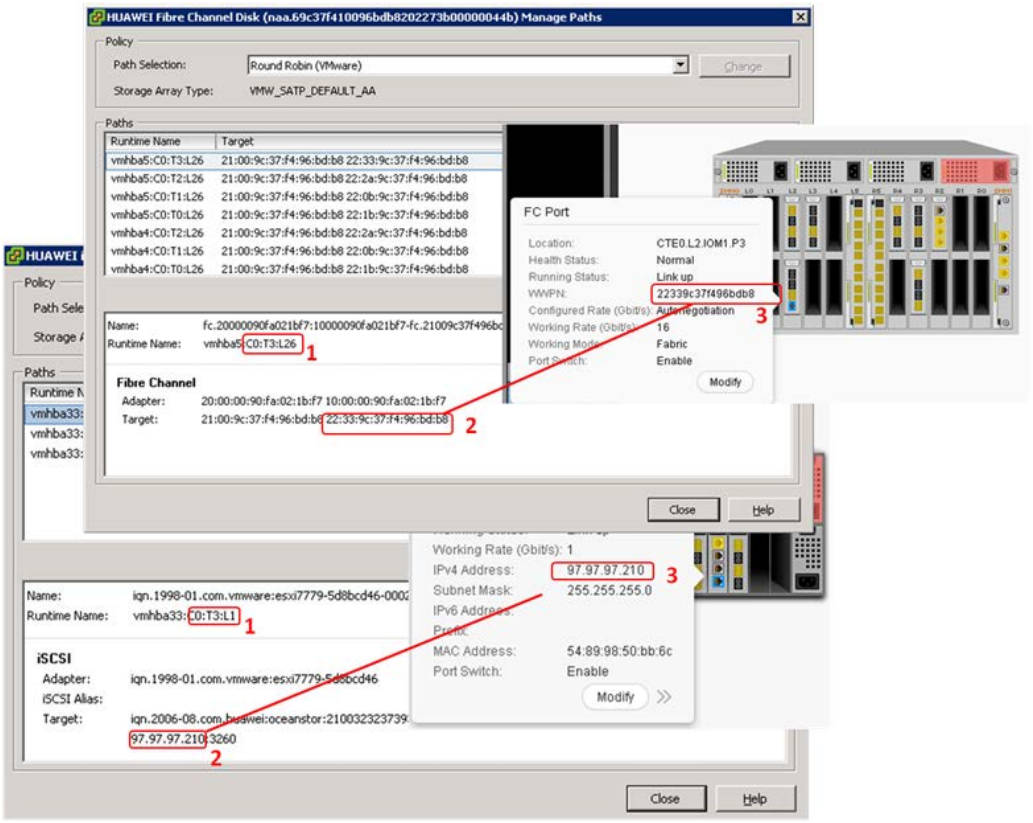

9.5 Precautions for Configuring iSCSI Multipathing

9.5.1 When Is Port Binding Needed?

Port binding is needed when multiple VMkernel ports are available on a VMware ESXi host and all these ports are in the same IP subnet and broadcast domain.

An example is shown as follows:

Figure 9-2 Example wherein port binding is required

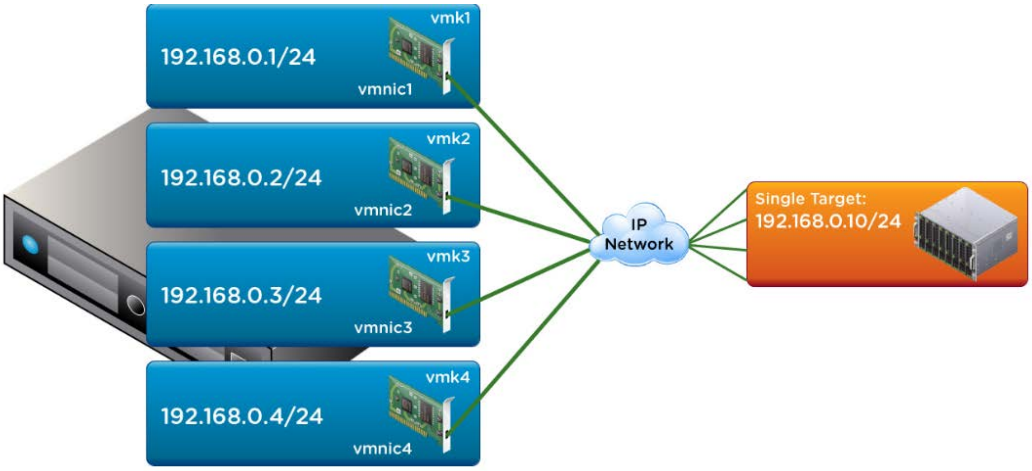

In this example, the networking configurations on the ESXi host are as follows:

Figure 9-3 Example network configurations where port binding is required

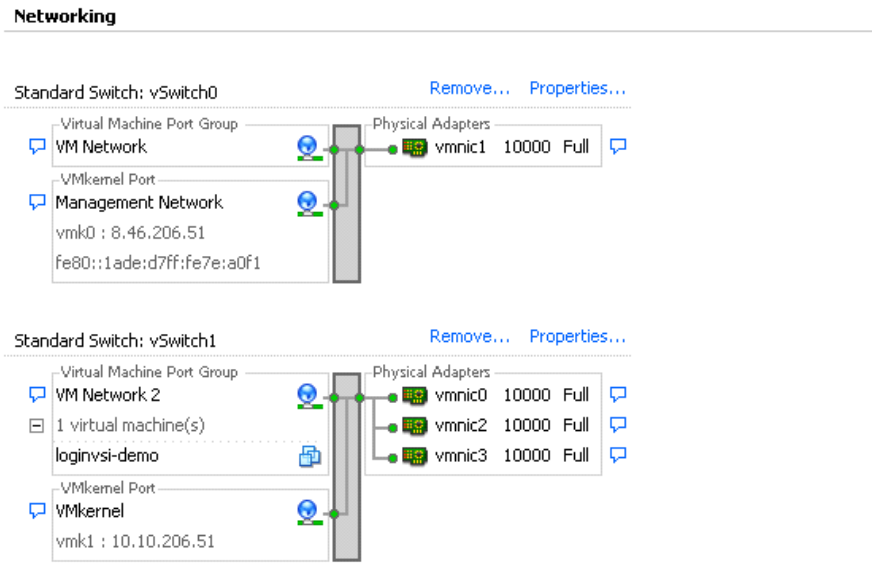

For port binding details, see the VMware knowledge base: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2045040.

9.5.2 When Is Port Binding Not Needed?

Port binding is not needed if:

- In the array, the target iSCSI ports are in different broadcast domains or IP subnets.
- The VM kernel ports for iSCSI connections are in different broadcast domains, IP subnets, or vSwitches.

- Connections can only be routed to the iSCSI array.

An example is shown as follows.

Figure 9-4 Example wherein port binding is not required

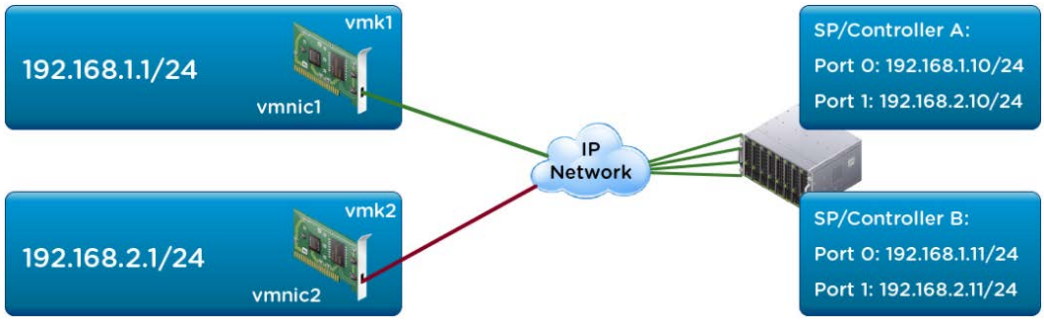

In this example, the networking configurations on the ESXi host are as follows:

Figure 9-5 Example network configurations where port binding is not required

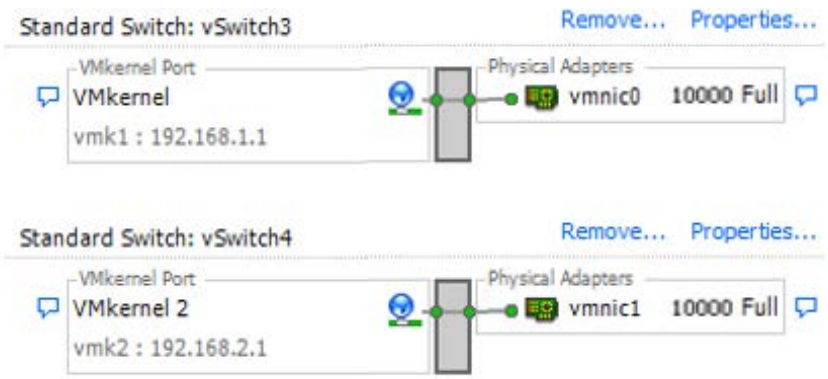

For port binding details, see the VMware knowledge base: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2045040.

10 Acronyms and Abbreviations

Acronyms and Abbreviations	Full Spelling
A	
ALUA	Asymmetric logical unit number
AN	Active Optimized
AO	Active Non-optimized
APD	All-Paths-Down
C	
CHAP	Challenge Handshake Authentication Protocol
CLI	Command Line Interface
D	
DRS	Distributed Resource Scheduler
E	
EQI	Extended Unique Identifier
F	
FC	Fiber Channel
FCR	Fibrer Channel Routing
G	

Acronyms and Abbreviations	Full Spelling
GE	Gigabit Ethernet
H	
HA	High Availability
HBA	Host Bus Adapter
I	
IP	Internet Protocol
IQN	iSCSI Qualified Name
iSCSI	Internet Small Computer Systems Interface
L	
LACP	Link Aggregation Control Protocol
LUN	Logical Unit Number
M	
MB	Megabyte
MRU	Most Resently Use
N	
NFS	Network File System
NMP	Native Multipath Module
P	
PDC	Permanent Device Loss
PSA	Pluggable Storage Architecture
PSP	Path Selection Plugin
P2V	Physical to Virtual
R	

Acronyms and Abbreviations	Full Spelling
RAID	Redundant Array of Independent Disks
RDM	Raw Device Mapping
S	
SAN	Storage Area Network
SATP	Storage Array Type Plugin
SP	Storage Processor
V	
VLAN	Virtual Local Area Network
VM	Virtual Machine
VMDK	Virtual Machine Disk
VMFS	Virtual Machine File System
V2P	Virtual to Physical
W	
WWN	World Wide Name
WWPN	World Wide Port Name