



Hewlett Packard
Enterprise

VMware Integration Guide

Legal Notices

Copyright © 2020 - 2021 by Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Hewlett Packard Enterprise believes in being unconditionally inclusive. If terms in this document are recognized as offensive or noninclusive, they are used only for consistency within the product. When the product is updated to remove the terms, this document will be updated.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Publication Date

Tuesday December 27, 2022 10:51:55

Document ID

sjn1591843375879

Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: support@nimblestorage.com

For all other general support contact information, go to <https://www.hpe.com/us/en/services/nimble-storage.html>.



Contents

How to Find the Information You Need.....	9
HPE Storage Integration with VMware.....	11
VMware Integration Features.....	11
Key Points About Setting Up an Integrated Environment.....	13
Planning Your Installation and Setup.....	15
HPE Storage Connection Manager for VMware.....	19
Understand How HPE Storage Connection Manager Supports Groups and Pools.....	19
HPE Storage Connection Manager for VMware Installation Options.....	22
Requirements for Installing HPE Storage Connection Manager.....	22
Important Information About Updating the ESXi Host and HPE Storage Connection Manager 7.x.....	23
Manually Download the HPE Storage Connection Manager Software Package.....	23
Manually Copy HPE Storage Connection Manager to the ESXi Host.....	24
Requirements for Installing HPE Storage Connection Manager.....	25
HPE Storage Connection Manager Installation Using vSphere 7.0 Lifecycle Manager When Connected to the Internet.....	25
HPE Storage Connection Manager Installation Using vSphere 6.5 and 6.7 Update Manager When Connected to the Internet.....	27
HPE Storage Connection Manager Installation Using vSphere 7.0 Lifecycle Manager With No Internet Connection.....	28
HPE Storage Connection Manager Installation Using vSphere 6.5 and 6.7 Update Manager With No Internet Connection.....	30
ESXCLI Installation of HPE Storage Connection Manager on ESXi 7.x Using an Online Bundle.....	32
ESXCLI Installation of HPE Storage Connection Manager on ESXi 7.x Using an Offline Bundle.....	33
ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Online Bundle.....	34
ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Offline Bundle	35
Updating HPE Storage Connection Manager for ESXi 7.x.....	37
Update HPE Storage Connection Manager for ESXi 6.x.....	37
Verify the HPE Storage Connection Manager Installation.....	38
Verify Settings for CompareLUNNumber and FailDiskRegistration.....	39
Configure HPE Storage Connection Manager on the ESXi Host.....	40
Using an HPE Dual 8GB MicroSD Enterprise Midline USB	41
View HPE Storage Connection Manager Logs.....	42
Configure Custom Log Locations.....	42
Uninstall HPE Storage Connection Manager 7.x from an ESXi 7.x Host.....	43

Uninstall HPE Storage Connection Manager 7.x from an ESXi 6.x Host.....	45
Uninstall HPE Legacy NCM Versions (7.0.1 and lower) from an ESXi 6.x Host.....	47
Verify the HPE Storage Connection Manager Uninstall	48

VMware iSCSI Configuration.....50

High-Level Steps to Set Up the ESXi iSCSI Network Configuration	50
Configure ESXi iSCSI Networking with Multiple vSwitches	51
Configure ESXi iSCSI Networking with a Single vSwitch.....	52
Configure the ESXi iSCSI Software Adapter	53
Bind VMK Ports to ESXi iSCSI Software Adapter	53
Use HPE Storage Connection Manager to Set iSCSI Path Selection Policy.....	54
iSCSI Initiator Groups and the Array.....	54
Create an iSCSI Initiator Group Using the Array OS GUI.....	54
Create an iSCSI Initiator Group Using the Array OS CLI.....	55
Assign Volumes to an iSCSI Initiator Group Using the Array OS GUI.....	56
Assign Volumes to an iSCSI Initiator Group Using the Array OS CLI.....	56
Unassign Volumes from an iSCSI Initiator Group Using the Array OS GUI.....	57
Unassign Volumes from an iSCSI Initiator Group Using the Array OS CLI.....	57
Edit an iSCSI Initiator Group Using the Array OS GUI.....	57
Edit an iSCSI Initiator Group Using the Array OS CLI.....	58
Delete an iSCSI Initiator Group Using the Array OS GUI.....	59
Delete an iSCSI Initiator Group Using the Array OS CLI.....	60
Use a CHAP Account.....	60
iSCSI Digest Considerations.....	60
Enable iSCSI Digest.....	61
Disable iSCSI Digest.....	61
Additional Information for Working with iSCSI.....	62

VMware Fibre Channel Configuration.....63

Fibre Channel Target Limits.....	63
----------------------------------	----

The HPE Nimble Storage vCenter Plugin.....64

Clients Supported in the vCenter Plugin	64
---	----

Set Up the vCenter Plugin.....65

Registration Requirements for Using vCenter Plugin.....	65
LUN Restrictions When Using the vCenter Plugin.....	65
Register or Add a vCenter Plugin Using the Array OS GUI.....	65
Register or Add a vCenter Plugin Using the Array OS CLI.....	66
Display a List of Registered Plugins Using the Array OS CLI.....	66
Edit the Registered vCenter Using the Array GUI.....	67
Unregister the vCenter Plugin.....	67

Unregister the vCenter Plugin from the vCenter Server.....	67
Use the GUI to Unregister the vCenter Plugin.....	68
Use the CLI to Unregister the vCenter Plugin	68
Removing Stale Entries Following Update to vSphere 7.....	68

RBAC and the vCenter Plugin.....70

Privileges Required for Registering the vCenter Plugin and VASA Provider.....	70
Privileges Required for Working with Datastores	71
Privileges Required for Working with VMs.....	72
Privileges Required for Creating VMware Synchronized Snapshots.....	74
Storage Array Roles and vVols.....	75
Create RBAC Roles	75

Using the vCenter Plugin with VMFS Datastores.....76

Create a VMFS Datastore from the vCenter Plugin.....	76
Mount an Existing VMFS Datastore Using the vCenter Plugin.....	78
Clone a VMFS Datastore from the vCenter Plugin.....	79
Grow a VMFS Datastore from the vCenter Plugin.....	80
Take a Snapshot of an HPE Nimble Storage-Based Datastore Using the vCenter Plugin.....	80
Edit a Datastore Mapped to an Array Using the vCenter Plugin.....	80
View VMFS Datastore Details from the vCenter Plugin.....	82
Delete a Datastore Mapped to a Volume Using the vCenter Plugin.....	82

Working with VMware Virtual Volumes.....84

How VASA Provider Works with vVols.....	84
Protocol Endpoints.....	85
Support for vVols Space Reclamation Using UNMAP	85
vVols and Stale Bindings.....	86
Supported Features.....	86
vVols and HPE Storage Connection Manager.....	86
Encryption.....	86
Windows Toolkit.....	86
Storage Policy Based Management.....	86
Group and Pool Merges.....	87
Using VASA Provider to Provide Disaster Recovery for vVols.....	87
Managing vVols.....	88
Configuring vVols Using the vCenter Plugin.....	88
Create a vVols Datastore Using the vCenter Plugin.....	88
Edit a vVol Datasore Using the vCenter Plugin.....	89
Grow a vVol Datastore Using the vCenter Plugin.....	90
Delete a vVol Datastore Using the vCenter Plugin.....	90
Run Configuration Checks Against the vCenter Server, Array.....	90
Configuring vVols and VMs from the VMware GUI.....	91

Overview of the vVols Workflow.....	91
Create a vVols Datastore.....	91
Create a VM.....	92
VSS and vVols.....	92
Overview of the VSS for vVols Setup Process.....	93
VSS for vVols Limitations and Troubleshooting.....	94
Setting Up the Guest Operating System for VSS for vVols.....	95
Synchronize the VM with the Guest OS.....	95
Working with VM Storage Policies.....	96
Virtual Machine Backups.....	96
Create a VM Storage Policy.....	97
Change the Assigned Storage Policy for a VM.....	98
Verifying Policy Compliance.....	99
Using Replication Partner VMs with the vCenter Plugin Dashboard.....	99
Specifying a Replication Partner VM	100
Claim a VM.....	100
Options for Restoring VMs from the vCenter Plugin.....	100
Restore an Entire VM from the vCenter Plugin.....	101
Restore One or More Disks From a VM.....	102
Clone a VM.....	102
Clone One or More Disks.....	103
Options for Deleting vVol VMs from the vCenter Plugin.....	104
Delete a VM.....	104
Undelete a VM.....	105
Purge a VM Using the vCenter Plugin.....	105
Troubleshooting Tips.....	106
Registration Error - Invalid Provider Certificate.....	106
Failure to Add VP - Time Mismatch.....	106
Datastore Inaccessible.....	107
VSS Snapshots Fail with the Message: "No Volume Connected to the Host"	107

VMware Synchronized Snapshots and VMFS Datastores.....108

How HPE Storage Synchronization Works with VMware.....	108
Snapshot Exclusion and Inclusion Options for VMs and Datastores	108
Volume Collections and VMware Objects.....	110
Bringing a VSS Snapshot Online After a Restore Operation.....	111

SRM and Storage Integration.....112

How SRA Works with SRM.....	112
Overview of SRA Setup Process.....	112
SRA for SRM Prerequisites.....	114
Download the SRA for SRM Installation Package.....	114
Install SRA for SRM for Windows.....	115

Install SRA for SRM for Photon OS SRM or VAs.....	115
Update SRA for SRM for Windows.....	117
Update SRA for SRM for Photon.....	117
Uninstall SRA for SRM for Windows.....	117
Uninstall SRA for SRM for Photon OS SRM or VAs.....	118
Working with Array Manager in SRM.....	118
Configure SRM 6.0, 6.1, and 6.5 for HPE Storage Arrays.....	118
Configure SRM 8.x for Storage Arrays.....	119
Using SRM to Restore vVols.....	120
Initiate a Recovery Plan.....	124
Test the Recovery Plan.....	124
Using SRM Test Failover Workflows and Replication.....	124
Implement the Recovery Plan.....	125
SRA and Microsoft Volume Shadow Service	126

VAAI Integration.....127

What is VAAI?.....	127
VAAI Requirements.....	127
Enable the VMware VAAI Provider to Use Storage Volumes.....	128

Using InfoSight Virtualization Data to Evaluate Performance.....129

Enable InfoSight to Collect VMware Information.....	129
Using the Virtualization Dashboard to View VMware Data.....	130
Key Data Provided by the Virtualization Section of InfoSight	130

Common Tasks and Best Practices.....132

VMware Partition Alignment.....	132
Register or Add VM to Inventory.....	133
vCenter Server.....	133
ESXi Host.....	133
Restore a VM from a Datastore.....	133
Restore Entire Datastore to an Existing Snapshot.....	134
Recover a Virtual Machine from a Cloned Snapshot Using the Array OS GUI.....	135
Change Access Information Using the Array OS GUI.....	136
Change Access Information Using the CLI.....	136
iSCSI Best Practices in VMware Environments.....	136
Set Up 1:1 Mapping for vSphere Switches.....	137
Host Disk Timeout Values.....	137
Managing Target Subnets.....	138
Guest Disk Timeout Values.....	139
Using VMware RDM Disks	139
Denylist RDMS Using SCSI 3 Persistent Reservations.....	140
Configure ESXi iSCSI Networking with Distributed Virtual Switches.....	140

Storage LUNs in the Device List.....	141
Locate vCenter Log Files.....	141

Helpful Information.....142

Recommendations for Environments That Do Not Use HPE Storage Connection Manager.....	142
Configure iSCSI Discovery.....	143
iSCSI Host Connection Methods.....	144
Set the iSCSI Host Connection Method to Manual.....	147
Configure Jumbo Frames.....	147
Validate the MTU Settings.....	148
Change NIC Frame Size.....	148
Configure an ESX Datastore.....	148
Enable Application-Consistent Quiescing on Windows Server 2008 VM.....	149
Mount and Unmount a Datastore Outside the Plugin.....	149

How to Find the Information You Need

Creating an environment that supports VMware with the storage array involves working with both VMware and HPE Storage features. HPE storage provides several tools, such as an HPE vCenter Plugin and an HPE Connection Manager toolkit for VMware, to help you set up your environment. You also have the option of performing the necessary steps manually.

This guide focuses on what you need to do to set up the VMware side of your environment. Other documents provide detailed information about using an array and other HPE Storage features.

All documentation provided by HPE is located in the Documentation portal of InfoSight: <https://infosight.hpe.com>. Select **Resources > Documentation**.

The following is a high-level summary of the documentation that you might find helpful.

Review This Documentation ...	To Learn About ...
<p>The Validated Configuration Matrix, which is available at HPE InfoSight.</p> <p>You can access it manually by logging onto InfoSight and choosing:</p> <p>Resources > Validated Configuration Matrix</p>	<p>Supported configurations for using array OS with VMware.</p>
<p>This guide (<i>VMware Integration Guide</i>)</p>	<p>Using VMware features with an HPE Storage array, including:</p> <ul style="list-style-type: none"> • Setting up your environment and registering with the HPE Storage vCenter Plugin with the VMware vCenter Server • Using the HPE Storage Connection Manager and how it can help with setting up iSCSI • Using the vCenter Plugin Web Client <p>Note: Release 5.1.1.0 and later does not support the VMware Desktop Client (also called a C# or Thick Client).</p> <ul style="list-style-type: none"> • Working with virtual machines (VMs) • Working with traditional datastores • Working with VMware virtual volumes (vVols) • Using Storage-Based Policy Management (SBPM) • Using VMware VAAI with HPE Storage volumes • Using disaster recovery features <ul style="list-style-type: none"> • HPE Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) for datastore-based recovery • VASA Provider for vVol-based disaster recovery • Virtual machine (VM) replication partners to recover VMs • Working with role-based access control (RBAC) features and the HPE Storage arrays • Taking snapshots



Review This Documentation ...	To Learn About ...
	<ul style="list-style-type: none"> • HPE Storage Volume Shadow Copy Service with vVols • VMware synchronized snapshots
<i>HPE Storage Windows Integration Guide</i>	Setting up VSS for vVols
HPE Storage Hardware and Administration Guides and Technical Reports	Working with the HPE Storage array and setting up your hardware
The technical report VMware VVols on HPE Storage Implementation Guide	Implementing vVols on HPE Storage
Array OS Release Notes	Learning about new features and late-breaking issues involving array OS
VMware documentation	Getting more details about VMware features



HPE Storage Integration with VMware

You can use HPE Storage products and features to integrate your HPE Storage environment with VMware, or you can perform manual steps.

The HPE vCenter Plugin and a VASA Provider ship pre-installed with the storage array. The plugin simplifies using VMware features with the array. The VASA Provider enables you to use VMware Virtual Volumes (vVols).

There are also several HPE Storage products that you can use to assist you in setting up your environment. These include the HPE Storage Connection Manager and the HPE Storage Windows Toolkit.

If you do not want to use the plugin, you can perform manual steps to integrate your environment with HPE Storage. You can also use a combination of manual steps and HPE Storage products.

You can decide whether to install HPE Storage products based on the features you want to use. The Storage array supports numerous features, including role-based access control (RBAC), disaster recovery, vVols, Storage-Based Policy Management (SBPM), HPE Storage Volume Shadow Copy Service (VSS) snapshots, and VMware synchronized snapshots.



Important: For best results, you should review the key points about setting up an HPE Storage environment that works with VMware ([Key Points About Setting Up an Integrated Environment](#) on page 13).

VMware Integration Features

HPE Storage supports numerous integration features that can assist you in setting up your VMware environment and using a storage array. Some of these integration features are pre-installed as part of the array OS software that ships on the array. There are other features that you can download and install based on the needs of your environment.

Note: The HPE InfoSight Software Downloads page contains links to HPE Storage products that are not included in the array OS software. See <https://infosight.hpe.com/resources/nimble/software>.

The integration features include the following:

HPE Storage Connection Manager

Creates the optimal number of iSCSI sessions for each storage volume and manages the selection of paths to the storage array in a VMware environment using iSCSI. HPE Storage Connection Manager is an optional product that you install on the ESXi host.

Note: To achieve optional I/O performance on HPE Storage devices, HPE recommends using HPE Storage Connection Manager for automatic iSCSI session management and for optimal path selection in both FC and iSCSI environments. HPE Storage Connection Manager is required in environments that use striped pools as well as environments that use the volume migration feature to ensure traffic is directed to the optimal array

Components: Connection Service and HPE Storage Path Selection Plugin (PSP)

Location: HPE Storage Connection Manager is a separate download

HPE Storage Replication Adapter (SRA)

Enables integration and interoperability with VMware Site Recovery Manager (SRM). You can use these features to set up disaster recovery plans.

HPE Nimble Storage Replication Adapter supports Windows SRM and Photon OS SRM while HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM supports only Photon OS SRM.

Location: SRA is a separate download

HPE Alletra 6000 and HPE Nimble Storage vCenter Plugin



Allows you to create and manage VMware VMFS datastores and certain other VMware features on the storage array. You must register the plugin with the vCenter Server.

Location: Pre-installed in array OS

HPE Storage Volume Shadow Copy Service (VSS) for vVols

Enables application-consistent snapshots for a Microsoft Exchange Server and a Microsoft SQL Server when the application data is hosted on a VMware Virtual Volume (vVol) and you are running release 5.0.1.0 or later with HPE Storage Windows Toolkit 5.0.0 or later installed on the guest operating system.

Location: vCenter Plugin, which is part of array OS, and Windows Toolkit, which is a separate download

HPE Storage Windows Toolkit

Supports HPE Storage VSS for vVols. Windows Toolkit is an optional product that you install on the Windows guest operating system.

Location: Windows Toolkit is a separate download

VAAI

Enables WRITE SAME, UNMAP, THIN PROVISION STUN, ATS, and XCOPY APIs.

Located in: Pre-installed in array OS

VASA Provider

Enables management of vVols by providing information about vVols and Storage-Based Policy Management (SBPM). VASA 3.0 provides disaster recovery at a vVols level. You must register VASA Provider with the vCenter Server.

Location: Pre-installed in array OS

VMware synchronized snapshots

Enables application consistent snapshots within VMFS datastore environments.

Location: Pre-installed in array OS



Key Points About Setting Up an Integrated Environment

There are a few things to know before setting up an integrated environment that allows you to use VMware with the storage array.

System Requirements

- You must verify that your configuration is supported. Check the Validated Configuration Matrix, which is online at <https://infosight.hpe.com/resources/nimble/validated-configuration-matrix>.
- You must ensure that all firewall ports dealing with communication between the vCenter Server Appliance or vCenter Server on Windows and the storage array are open. These ports include:

Port	Source	Destination	Use
443	Array: Group IP and Support IP	vCenter server	VASA and vCenter Plugin registration
443	vCenter server	Array: Group IP	vCenter Plugin communication
4210	Windows host with toolkit	Array: Group IP	VSS for vVols integration
4311	Array: Group IP and Support IP	Windows Host with toolkit	VSS for vVols integration
5392	Windows host with toolkit	Array: Group IP	Web service communications, including VSS for vVols integration
8443	vCenter Server and ESXi hosts	Array: Group IP	SSL communication between ESXi hosts and arrays that use vVol datastores, including VASA Provider communication

Note: If you do not enable port 8443, the ESXi hosts will not be able to mount vVol datastores.

Note: This chart contains only the required ports for integration between HPE Storage and VMware. Additional ports must be opened to allow access to the array GUI, connect iSCSI storage, or use tertiary features such as replication. Additional information about required ports can be found in the *GUI Administration Guide*, which is available on the Documentation Portal on HPE InfoSight. See <https://infosight.hpe.com/resources/nimble/docs>.

- The vCenter Plugin for web clients supports any web browser that is supported by the vCenter Server.
- VMware integration requires accurate internal name resolution to function correctly. It is a good practice to use an internal Domain Name Service (DNS). These tasks include registering the vCenter Plugin, synchronizing snapshots, or any other tasks that require a name resolution service. If any IP addresses are changed, their corresponding forward and reverse records must be updated for the VMware integration functionality to work as expected. HPE does not recommend using external DNS servers with the array.

Permissions

- You must have administrator-level privileges on the VMware vCenter.
- You must have administrator-level permissions on the storage array.
- To gather VM statistics, you must designate the System.Read privilege. You can use HPE InfoSight to gather VMware analytics.

- If you are not using the vCenter Plugin and the array is in a VMware environment, you must have ACLs (initiator group and/or CHAP username or WWPN) on all volumes.

Key Points About Volumes, Blocks, and iSCSI

- Recommended: When you use an iSCSI adapter, you should enable flow control on vNICS. This applies to both hardware-based and software-based iSCSI configurations.
- In iSCSI environments, storage arrays support jumbo frames if the network switches and other components support them.

Block Size Considerations for VMware Datastores (LUNs)

You must use the appropriate performance policies for volumes that are presented to VMware ESXi hosts.

Note: You need to apply the performance policies when you create the volume. You can change the performance policies later as long as the block size does not change.

For disks that are set up as VMFS datastores, you should use the VMware performance policy (VMware ESX 5). Datastores created using the vCenter Plugin are automatically set to VMware ESX 5 performance policy.

For disks that are set up as RDM or In-Guest iSCSI, you should use the correct application performance policy (for example, Exchange 2010 Data Store or SQL Server/SQL Server 2012).

You can also use VMware Virtual Volumes (vVols), which the vCenter Plugin supports when you run release 5.0.1.0 or later. vVols automatically leverage storage-policy based management.



Planning Your Installation and Setup

Before you set up VMware to work with your storage environment, it is a good practice to plan your installation. The following checklist provides high-level information about some of the issues you should consider.

Note: You can use the vCenter Plugin provided by HPE to set up your storage environment to work with VMware. Using this plugin can simplify many actions that you would have to perform manually otherwise.

Considerations	Details
Does your environment meet the system requirements for integrating VMware with an HPE Storage array?	<p>The online Validated Configuration Matrix contains the most current information about system requirements.</p> <p>In addition, this guide provides information about the following:</p> <ul style="list-style-type: none">• The VMware and HPE Storage features you can use and where they are located.• Tips that you should know before you set up your integrated environment. <p>More information:</p> <ul style="list-style-type: none">• The Validated Configuration Matrix, which is online on HPE InfoSight at https://infosight.hpe.com/resources/nimble/validated-configuration-matrix.• VMware Integration Features on page 11• Key Points About Setting Up an Integrated Environment on page 13•
Are your ports set up correctly?	<p>All firewall ports that deal with communication between the Windows guest operating system and the storage array must be open.</p> <p>More information:</p> <ul style="list-style-type: none">• Key Points About Setting Up an Integrated Environment on page 13



Considerations	Details
Which protocol will you be using?	<p>HPE Storage supports both iSCSI and Fibre Channel (FC) environments:</p> <ul style="list-style-type: none"> iSCSI environments. <p>Recommended: Use the HPE Storage Connection Manager to assist with setting up iSCSI.</p> <p>You need to manually configure:</p> <ul style="list-style-type: none"> ESXi networking with one or more vSwitches ESXi software adapter and bind the VMK ports to it (single iSCSI subnet) Ensure there is a 1-to-1 mapping of the VMK port to VMNIC Fibre Channel. <p>Most FC setup tasks are handled automatically by the HPE Storage Setup Manager, which is included in the HPE Storage Windows Toolkit.</p> <p>There are some best practice suggestions you should follow when using FC.</p> <p>More information:</p> <ul style="list-style-type: none"> HPE Storage Connection Manager for VMware on page 19 VMware iSCSI Configuration on page 50 VMware Fibre Channel Configuration on page 63 iSCSI Best Practices in VMware Environments on page 136 Helpful Information on page 142 Windows Integration Guide (available on HPE InfoSight)
Will you be using role-based access control?	<p>You can create role-based permissions to control access to datastores.</p> <p>More Information:</p> <ul style="list-style-type: none"> Create RBAC Roles on page 75 GUI and CLI Administration Guides (available on HPE InfoSight)
Which Web Client will you be using?	<p>The vCenter Plugin supports an HTML5 vSphere Web Client and a Flex vSphere Web Client. Starting with release 5.1.1.0, the HTML5 client enables you to use the vCenter Plugin to perform more tasks, such as create a vVols datastore and then increase its size.</p> <p>More Information:</p> <ul style="list-style-type: none"> Clients Supported in the vCenter Plugin on page 64



Considerations	Details
Have you registered the HPE vCenter plugin with the vCenter Server?	<p>You must register the vCenter Plugin with a vCenter Server before you can use it.</p> <p>More information:</p> <ul style="list-style-type: none"> • Set Up the vCenter Plugin on page 65 • Registration Requirements for Using vCenter Plugin on page 65 • Register or Add a vCenter Plugin Using the Array OS GUI on page 65 • Display a List of Registered Plugins Using the Array OS CLI on page 66
Will you be using VMware virtual volumes (vVols)?	<p>By using vVols, you can map VMware virtual disks to storage volumes. The array OS provides a VASA Provider that allows you to use vVols.</p> <p>In addition to features supported by traditional volumes, vVols also support storage policies. You can use a storage policy to enable a replication partner virtual machine (VM) that can take over if a problem occurs with your primary VM.</p> <p>More Information:</p> <ul style="list-style-type: none"> • Working with VMware Virtual Volumes on page 84 • How VASA Provider Works with vVols on page 84
Will you need to take snapshots?	<p>The array OS supports both HPE Storage VSS for vVols snapshots and VMware synchronized snapshots.</p> <ul style="list-style-type: none"> • VSS snapshots. This feature works with vVols only and requires that you install the HPE Storage Windows Toolkit 5.x or later on the Windows guest operating system. You must configure the VM iSCSI initiator for the guest operating system. • VMware synchronized snapshots. The array OS automatically supports this feature when you have the latest VMware Tools installed. <p>More information:</p> <ul style="list-style-type: none"> • VSS and vVols on page 92 • VMware Synchronized Snapshots and VMFS Datastores on page 108 • <i>Windows Integration Guide</i> (available on HPE InfoSight)
Do you want to use the vStorage APIs for Array Integration (VAAI) features?	<p>The array OS allows you to use the VAAI feature set, which enables the ESXi host to offload VM and storage management operations from the ESXi host to the storage array.</p> <p>More Information:</p> <ul style="list-style-type: none"> • VAAI Integration on page 127



Considerations

Will you be setting up disaster recovery?

Details

The HPE storage array supports the following forms of disaster recovery:

- Array-based. The HPE Storage Replication Adapter (SRA) enables an array to support the VMware vCenter Site Recovery Manager (SRM).
- VM-based. When you use vVols, you can set up a replication partner VM to provide a backup to an upstream or source VM.
- vVol-based. When you use VASA 3.0 or later, you can set up the VASA Provider disaster recovery feature, which allows you to restore information at the vVols level.

More Information:

- [SRM and Storage Integration](#) on page 112
 - [Using Replication Partner VMs with the vCenter Plugin Dashboard](#) on page 99
 - [Using VASA Provider to Provide Disaster Recovery for vVols](#) on page 87
-



HPE Storage Connection Manager for VMware

The HPE Storage Connection Manager for VMware manages connections from the host to volumes on HPE Storage systems. To simplify configuring multiple connections and Multipath Input/Output (MPIO), the array OS requires that only one IP address (the iSCSI discovery IP address) be advertised at the time of discovery, not the full set of iSCSI network interfaces.

This means that you do not need to configure static iSCSI connections to the appropriate interfaces or to worry about how many connections there are to a volume. HPE Storage Connection Manager ensures that, as connections are made to the same address (group target portal), the connections are redirected to the appropriate distribution of actual iSCSI network interfaces.

HPE Storage Connection Manager also provides support for high-performance storage configurations by allowing the use of HPE Storage striped pools in both FC and iSCSI protocol modes.

To achieve optional I/O performance on HPE Storage devices, HPE recommends using HPE Storage Connection Manager for automatic iSCSI session management and for optimal path selection in both FC and iSCSI environments. HPE Storage Connection Manager is required in environments that use striped pools as well as environments that use the volume migration feature to ensure traffic is directed to the optimal array.

Understand How HPE Storage Connection Manager Supports Groups and Pools

The HPE Storage Connection Manager for VMware can help you work with storage groups and pools. Within a storage pool, a volume can span multiple arrays.

If an I/O request is sent to an array that does not have the block or blocks requested, it must be forwarded to the correct array, which can result in a decrease in I/O performance. When you use groups and storage pools, you avoid this problem because you link arrays to create a single, logical storage entity. HPE Storage Connection Manager then helps you maintain the optimum number of iSCSI sessions and directs both Fibre Channel and iSCSI I/O requests to the efficient route.

How a Group Works


A group or cluster is a set of member arrays that are physically connected and logically represent a single storage entity for the purposes of aggregating performance, capacity, and simplifying management. For most administrative tasks, a group looks and acts like a single array.

A group is a collection of one to four storage arrays. A group contains one or more pools. The storage pools contain arrays, which contain folders, and volumes.

You interact with the group by connecting to its group management IP address, which is hosted by one of the arrays that is known as the group leader. Data is striped across the arrays in the same pool.

For more information about storage groups, see the "Array Groups" section of the *GUI Administration Guide*. This guide is available on HPE InfoSight (<https://infosight.hpe.com>).



Figure 1: Relationships of Groups, Pools, Arrays, Folders, and Volumes in Multi-array Groups

- | | |
|-----------------------|-----------------|
| 1 Group | 4 Folder |
| 2 Storage Pool | 5 Volume |
| 3 Array | |

How a Pool Works

A pool confines data to a subset of the arrays within a group. An array can be part of only one pool; it cannot span multiple pools. However, a pool can contain multiple arrays.

The data on the volumes within the pool is striped and automatically rebalanced over the members of the group, which provides the benefit of aggregated performance.

Pools are important because they dictate the physical locality and striping characteristics. You can think of a pool as a logical container that holds one or more member arrays in which volumes reside.

Volumes have the following relationships with pools:

- Volumes, their snapshots, and their clones are tied to a specific pool.
- You can migrate volumes between different pools.
- Volume collections are not tied to pools and can contain volumes that reside in different pools.

For more information on storage pools, see the "Storage Pools" section of the *GUI Administration Guide*. This guide is available on HPE InfoSight (<https://infosight.hpe.com>).

How HPE Storage Connection Manager Supports Groups and Pools

HPE Storage Connection Manager consists of two components:

Component	Function
HPE Storage Connection Service	HPE Storage Connection Service automatically calculates, balances, and maintains the optimal number of iSCSI sessions to the storage device across the ESXi host NICs.
HPE Storage Path Selection Plugin (PSP)	The PSP for VMware Pluggable Storage Architecture automatically directs the I/O request for an HPE Storage device to the most favorable route.

Within the storage group, through the group leader array, PSP determines on which arrays the volume resides and redirects communication to the appropriate paths.

If you add an array to a storage pool, the volume may be adjusted and balanced to reside partially on the newly added array. In this case, HPE Storage Connection Manager automatically creates the optimal number of iSCSI connections to the new pool member.

HPE Storage Connection Manager uses the following formula to determine the optimal number of paths to use with your Storage group:

$M * A = \text{optimal path count}$

Where:

M = number of vmknics on the ESXi host

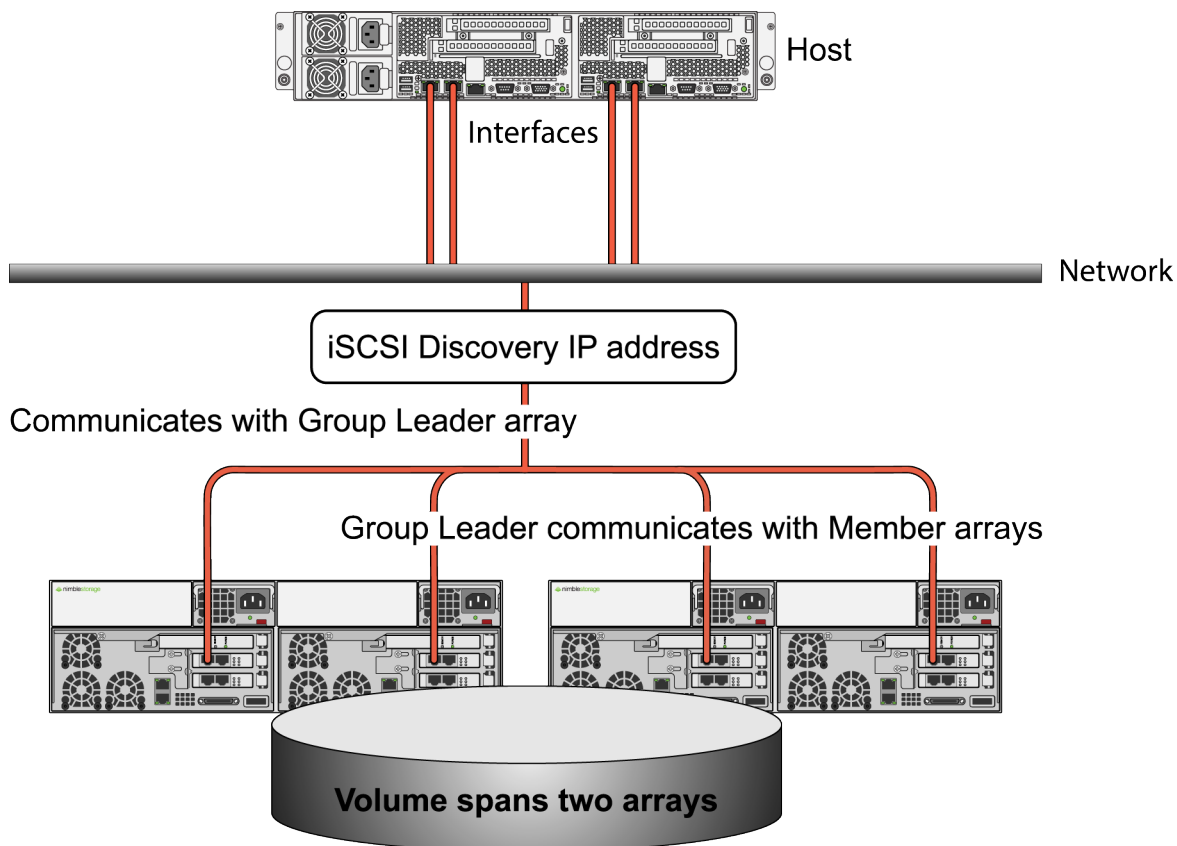
A = number of arrays in the volume's pool

For example, assume you have four vmknics port-bound to your software iSCSI initiator on your ESXi host. Additionally, you have two arrays in your Storage group that are managed as separate pools. Since the volume is located on a single array only, the calculation HPE Storage Connection Manager uses to determine the proper number of paths to form for each volume would be $4 * 1$, which creates a total of 4 paths.

In networks that have multiple subnets and where port-binding is not recommended, if you enter a single discovery address on the array, that will cause the discovery addresses on all allowed subnets to be reported back to the ESXi host. ESXi, by default, will attempt to initiate one session per subnet. The array will automatically redirect all incoming connections to the array that owns the pool for the associated volume.

HPE Storage Connection Manager improves this behavior by ensuring that ESXi continues to hold the minimum and maximum number of iSCSI sessions as defined by the HPE Storage Connection Manager formula above.

Figure 2: HPE Storage Connection Manager Diagram



All of this ensures that even with arrays being added or removed from a storage group or pool, the correct number of connections are made. You do not need to manually configure the connections.

Best Practice

To achieve optional I/O performance on HPE Storage devices, HPE Storage recommends using HPE Storage Connection Manager for automatic iSCSI session management and for optimal path selection in both FC and iSCSI environments. HPE



Storage Connection Manager is required in environments that use striped pools as well as environments that use the volume migration feature to ensure traffic is directed to the optimal array.

HPE Storage Connection Manager for VMware Installation Options

HPE Storage Connection Manager for VMware is available for download from the HPE InfoSight Software Downloads site. You install connection manager on each ESXi host connected to a Storage array.

You can download the HPE Storage Connection Manager software package for your version ESXi using the vSphere Update Manager or the vSphere Lifecycle Manager. The vSphere tool will download HPE Storage Connection Manager either directly to the ESXi host or to a Windows or Linux host. If you download HPE Storage Connection Manager to a Windows or Linux host, you must copy it to the ESXi host. You always install HPE Storage Connection Manager on the ESXi host.

Depending on whether the vCenter Server or ESXi host is connected to the internet, you can install HPE Storage Connection Manager as either an online bundle or an offline bundle.

Note: HPE Storage recommends you use the vSphere Update Manager, or, if you are using vSphere 7.0, the vSphere Lifecycle Manager because they allow you to keep HPE Storage Connection Manager up-to-date when performing your regular ESXi patching. You can also use the vCLI utility from a Windows or Linux host. See your VMware documentation for more information.

Tool	Internet Connection	Procedure
vSphere Update Manager	Online	vSphere Update Manager and vSphere Lifecycle Manager automatically download and install the HPE Storage Connection Manager bundle. The one you use depends on which version of vSphere you have. See HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager When Connected to the Internet .
	Offline	Manually download the HPE Storage Connection Manager bundle to a Windows or Linux host. Then install it using vSphere Update Manager. See HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager With No Internet Connection .
ESXCLI	Online	Use a single ESXCLI command to download and install the HPE Storage Connection Manager bundle. The ESXi host must have Internet access to the Support site. See ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Online Bundle on page 34.
	Offline	Download the HPE Storage Connection Manager bundle to a Windows or Linux host and then copy it to the ESXi host. Now use the ESXCLI utility to install the bundle. See ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Offline Bundle on page 35.

Requirements for Installing HPE Storage Connection Manager

Before you install HPE Storage Connection Manager for VMware, make sure that your system meets the requirements for installing and using it.





Important: The most current requirements for using HPE Storage Connection Manager are provided by the Validated Configuration Matrix, which is available on HPE InfoSight (<https://infosight.hpe.com/>).

System requirements include:

- A supported version of ESXi that has the VMware vSphere Standard, Enterprise, or Enterprise Plus licensing running on the host

Supported versions include 6.5, 6.7, and 7.0.

Note: The Standard license applies to VMware vSphere 6 and later.

- A supported vCenter Server

Supported versions include 6.5, 6.7, and 7.0.

- The HPE Storage Connection Manager software package for your ESXi host

For example, if you are running ESXi 6.0, download the HPE Storage Connection Manager "Software for ESXi6."

In addition, you should also have the following:

- Internet connection to the Windows or Linux host
- SFTP client such as WinSCP on the Windows or Linux host
- SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host



Important: HPE Storage Connection Manager cannot be installed on an ESXi host that is in lockdown mode. You must disable lockdown mode before installing HPE Storage Connection Manager. For more information see [VMware KB article 1008077](#).

Important Information About Updating the ESXi Host and HPE Storage Connection Manager 7.x

Not all versions HPE Storage Connection Manager for VMware work well with all versions of ESXi.

HPE Storage Connection Manager 7.0.0 provides two builds in order to support both ESXi 7.0 and ESXi 6.5 and 6.7. Starting with ESXi 7.0, ESXi deprecated support for 32-bit VIBs and uses 64-bit VIBs. ESXi 6.5 and 6.7 only support 32-bit VIBs. You must make sure you use the correct HPE Storage Connection Manager 7.0 software package for the version of ESXi you are using:

- Build 7.0.0-700012 supports ESXi 7.0 and provides 64-bit VIBs.
- Build 7.0.0-650012 supports ESXi 6.5 and 6.7 and provides 32-bit VIBs.



Important: Before you upgrade your host to ESXi 7.x, make sure you uninstall any version of HPE Storage Connection Manager prior to 7.x. Doing this ensures that all 32-bit VIBs have been removed before the ESXi upgrade. After the upgrade completes, you can install the correct HPE Storage Connection Manager 7.0.0 software package.

You must reboot the ESXi host after you install the new version of HPE Storage Connection Manager.

Manually Download the HPE Storage Connection Manager Software Package

If you are planning to install HPE Storage Connection Manager for VMware as an offline bundle, you must first download the HPE Storage Connection Manager software package for your version of ESXi.

Note: You do not need to download the HPE Storage Connection Manager software if you have the correct internet connections that allow you to use a tool such as vSphere Update Manager or ESXCLI to automatically download and install HPE Storage Connection Manager.



Before you begin

You must have:

- Internet connection to the Windows or Linux host where you want to put the downloaded version of HPE Storage Connection Manager
- Login information for the HPE InfoSight portal

You can obtain a user name and password at the portal by clicking **Create Account** and supplying the requested information.

Procedure

1. Go to HPE InfoSight (<https://infosight.hpe.com/>) and log in.
2. Choose **Resources** > **Software Downloads**.
3. From the **Integration Kits** list in the left column, choose **HPE Storage Connection Manager for VMware**.
4. In the Current Version list, select the **Downloads** software package for your version of ESXi.
5. Save the HPE Storage Connection Manager installation package to a convenient place.

The installation package name uses the format

`hpe-storage-connection-manager-esx6-7.0.2-650008.zip` where N.N is the supported VMware ESXi release family, x.x.x is the version of HPE Storage Connection Manager and yyyyyy is the HPE Storage Connection Manager build number.

Note: Do not unzip the installation package.

6. (Optional) Download the latest version of the *HPE Storage Connection Manager for VMware Release Notes*.

What to do next

Install HPE Storage Connection Manager using one of the following methods:

- [HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager With No Internet Connection](#)
- [ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Offline Bundle](#) on page 35

Manually Copy HPE Storage Connection Manager to the ESXi Host

If you downloaded HPE Storage Connection Manager for VMware and you plan to install it as an offline bundle using ESXCLI, you must move the software package to the ESXi hosts.

Note: You do not need to copy the HPE Storage Connection Manager software to the ESXi hosts if you have the correct internet connections that allow you to use a tool such as vSphere Update Manager or ESXCLI to automatically download and install HPE Storage Connection Manager.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager installation package
See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23
- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host

Procedure

1. Launch your SFTP client on the Windows or Linux host.
2. Log into the ESXi server.
 - a) Enter the following login information:

- The IP address of the ESXi server in the **Host name** field
- The **root** user in the **User name** field
- The root password into the **Password** field

b) Click **Login**.

3. Copy the HPE Storage Connection Manager installation package (filename `hpe-storage-connection-manager-esx6-7.0.2-650008.zip`) from the Windows or Linux host to the `/tmp` directory on the ESXi host.

Note: It is a good practice to make sure you have a enough space in the `/tmp` directory before you copy the package.

4. Close the SFTP client.

Requirements for Installing HPE Storage Connection Manager

Before you install HPE Storage Connection Manager for VMware, make sure that your system meets the requirements for installing and using it.



Important: The most current requirements for using HPE Storage Connection Manager are provided by the Validated Configuration Matrix, which is available on HPE InfoSight (<https://infosight.hpe.com/>).

System requirements include:

- A supported version of ESXi that has the VMware vSphere Standard, Enterprise, or Enterprise Plus licensing running on the host

Supported versions include 6.5, 6.7, and 7.0.

Note: The Standard license applies to VMware vSphere 6 and later.

- A supported vCenter Server
- The HPE Storage Connection Manager software package for your ESXi host

For example, if you are running ESXi 6.0, download the HPE Storage Connection Manager "Software for ESXi6."

In addition, you should also have the following:

- Internet connection to the Windows or Linux host
- SFTP client such as WinSCP on the Windows or Linux host
- SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host



Important: HPE Storage Connection Manager cannot be installed on an ESXi host that is in lockdown mode. You must disable lockdown mode before installing HPE Storage Connection Manager. For more information see [VMware KB article 1008077](#).

HPE Storage Connection Manager Installation Using vSphere 7.0 Lifecycle Manager When Connected to the Internet

When the vCenter Server is connected to the internet, you can use the vSphere Lifecycle Manager to install HPE Storage Connection Manager for VMware. This process allows you to install HPE Storage Connection Manager on all the required ESXi 7.x hosts at the same time.



Before you begin

You must have:

- The vSphere Lifecycle Manager.

Note: For more information about the vSphere Lifecycle Manager, see your VMware documentation.

- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.
- TLS v1.2 or later support on the vCenter Server host for `update.nimblestorage.com`.
- You must be using baselines, as you can not configure a single image. For information on creating baselines, see *Attach Baselines and Baseline Groups to Objects* and *vSphere Lifecycle Manager Baselines and Images* in the VMware documentation.



Important:

ESXi 7.0 requires 64-bit VIBs. Previous versions of ESXi worked with 32-bit libraries. You must use HPE Storage Connection Manager 7.0.0 and vSphere 7.x with ESXi 7.0. As a result, if you are upgrading your host to ESXi 7.0, you must do the following **before** you install HPE Storage Connection Manager 7.x:

- Uninstall any version of HPE Storage Connection Manager prior to HPE Storage Connection Manager 7.0.0. Doing this removes the 32-bit VIBs. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47.
- Upgrade the ESXi host to ESXi 7.x.

Procedure

1. On the VMware vSphere Web Client > Shortcuts screen, click **Lifecycle Manager**.
2. On the **Settings** tab, under **Administration Settings**, click **Patch Setup**.
3. Click the **New** button.
4. Enter the correct Source URL for the HPE Storage Connection Manager software package for your version of ESXi.

```
https://update.nimblestorage.com/esx<version>/ncm/index.xml
```

where <version> is the ESX version; for example, esx7.0.

5. Click the **Save** button.
6. Next to Lifecycle Manager, click **Actions** menu, under **Updates**, click **Sync Updates** to download the package.
7. Under **Baselines** tab, click **New** > **Baseline**.
8. Type NCM Baseline in the **Name** field.
9. Select **Extension** option for the Content, and then click **Next**.
10. Under the extensions list, check the checkbox next to the **Nimble Connection Service**.
11. Click **Next**, then click **Finish**.
12. On the VMware vSphere Web Client > Shortcuts screen, click **Hosts and Clusters**.
13. Click and select either the vCenter object, datacenter object, cluster object, or ESXi host that you would like to install HPE Storage Connection Manager on.
14. Go to **Updates tab** > **Hosts** > **Baselines**, under **Attached Baselines**, click **Attach** > **Attach Baseline or Baseline Group**.
15. In the **Attach Baseline or Baseline Group** dialog box, under the baselines list, check the checkbox next to the NCM Baseline, and click the **Attach** button.
16. In the Hosts' Compliance box, click **Check Compliance**. The scan shows that the baseline is non-compliant, meaning that the baseline has not yet been applied to the ESXi host(s). Note that you may need to refresh the vCenter UI for the Compliance Status to update
17. Under the baselines list of **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Stage**.

18. In the **Stage Patches** dialog box, select the host(s) you would like to upload HPE Storage Connection Manager to and click the **Stage** button.
19. Wait until **Stage patches to entity task** is completed under **Recent Tasks** panel.
20. On **Updates** tab, under **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Remediate**.
21. In the **Remediate** dialog box, select the host(s) you would like HPE Storage Connection Manager to be remediated, and then click the **Remediate** button.

HPE Storage Connection Manager Installation Using vSphere 6.5 and 6.7 Update Manager When Connected to the Internet

You can use the vSphere 6.5 and later Update Manager to install HPE Storage Connection Manager for VMware when the vCenter Server is connected to the internet. This process allows you to install HPE Storage Connection Manager on all the required ESXi hosts at the same time.



Important:

To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Update Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Before you begin

If you are using an earlier version of vSphere, see [HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager When Connected to the Internet](#) for information about installing HPE Storage Connection Manager.

For this installation, you must have:

- The vSphere Update Manager installed on the vCenter Server

Note: VMware vSphere 6.5 and later installs the vSphere Update Manager by default. For more information about the vSphere Update Manager, see your VMware documentation, including *vSphere Update Manager Installation and Administration Guide*.

- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.
- TLS v1.2 or later support on the vCenter Server host for `update.nimblestorage.com`



Important:

If you are upgrading your host from one major version of ESXi (for example ESXi 6.0 to ESXi 6.7) make sure to uninstall HPE Storage Connection Manager prior to doing the ESXi update. Doing this removes the connection service and the path selection policy that are installed when HPE Storage Connection Manager is installed. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47. If you do not uninstall HPE Storage Connection Manager before updating between major ESXi versions, the HPE Storage Path Selection Plugin (PSP) will fail to load and might cause boot failures.

Procedure

1. On the VMware vSphere Web Client >Shortcuts screen, click **Update Manager**.
2. On the **Settings** tab, under **Administration Settings**, click **Patch Setup**.
3. Click the **New** button.

4. Enter the correct Source URL for the HPE Storage Connection Manager software package for your version of ESXi.

```
https://update.nimblestorage.com/esx<version>/ncm/index.xml
```

where <version> is the ESX version; for example, esx6 or esx6.5 (esx6.5 is also used for ESXi version 6.7).

5. Click the **Save** button.
6. Under **Updates** tab, in **Patch Downloads**, click **Download Now** to download the package.
7. Under **Baselines** tab, click **New** > **Baseline**.
8. Type NCM Baseline in the **Name** field.
9. Select **Extension** option for the Content, then click **Next**.
10. Under the extensions list, check the checkbox next to the **Nimble Connection Service**
11. Click **Next**, then click **Finish**.
12. On the VMware vSphere Web Client > Shortcuts screen, click **Hosts and Clusters**
13. Click and select either vCenter object, datacenter object, cluster object, or ESXi host that you would like to install HPE Storage Connection Manager on.
14. Go to **Updates tab** > **Host Updates**, under **Attached Baselines**, click **Attach** > **Attach Baseline or Baseline Group**.
15. In the **Attach Baseline or Baseline Group** dialog box, under the baselines list, check the checkbox next to the NCM Baseline and click the **Attach** button.
16. In the Hosts' Compliance box, click Check Compliance. The scan shows that the baseline is non-compliant, meaning that the baseline has not yet been applied to the ESXi host(s).

Note: You may need to refresh the vCenter UI for the Compliance Status to update.

17. Under the baselines list of **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Stage**.
18. In the **Stage Patches** dialog box, select the host(s) you would like to upload HPE Storage Connection Manager to and click the **Stage** button.
19. Wait until **Stage patches to entity** task is completed under **Recent Tasks** panel.
20. On **Updates** tab, under **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Remediate**.
21. In the **Remediate** dialog box, select the host(s) you would like HPE Storage Connection Manager to be remediated, and then click **Remediate**.

Results

After you install HPE Storage Connection Manager:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each HPE Storage volume.
- HPE Storage Path Selection Plugin (PSP) automatically claims the available HPE Storage devices, including both existing devices and devices added later.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

HPE Storage Connection Manager Installation Using vSphere 7.0 Lifecycle Manager With No Internet Connection

When the vCenter Server is not connected to the internet, you can use the vSphere 7.0 and later Update Manager to perform an offline bundle installation of HPE Storage Connection Manager for VMware. This process enables you to install HPE Storage Connection Manager on all the required ESXi hosts at the same time.

Note:

To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Update Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE Storage recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager installation package

See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Copy HPE Storage Connection Manager to the ESXi Host Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.

- The vSphere 7.0 Update Manager installed on the vCenter Server

Note: For more information about the vSphere Update Manager, see your VMware documentation, including [Installing and Administering VMware vSphere Update Manager](#).

- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.

**Important:**

If you are upgrading your host from ESXi 6.x to ESXi 7.x, make sure you have done the following **before** you install HPE Storage Connection Manager 7.x:

- Uninstalled HPE Storage Connection Manager 6.x. Doing this removes the VIBs for ESXi 6.x. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47.
- Upgraded the ESXi host to ESXi 7.x.

If you do not uninstall HPE Storage Connection Manager 6.x before upgrading the host from ESXi 6.x to ESXi 7.x and installing HPE Storage Connection Manager 7.x, the HPE Storage Path Selection Plugin (PSP) will fail to load and might cause boot failures.

Procedure

1. On the VMware vSphere Web Client > Shortcuts screen, click **Lifecycle Manager**.
2. Next to Lifecycle Manager, click **Actions** menu, under **Updates**, click **Import Updates**.
3. Click **Browse** button.
4. In Open File dialog box, select the HPE Storage Connection Manager installation package, and click **Open**.
5. Wait for the uploading to be completed.
6. Under **Baselines** tab, click **New > Baseline**.
7. Type NCM Baseline in the **Name** field.
8. Select **Extension** option for the Content, then click **Next**.
9. Under the extensions list, check the checkbox next to the **Nimble Connection Service**.
10. Click **Next**, then click **Finish**.
11. On the VMware vSphere Web Client > Shortcuts screen, click **Hosts and Clusters**.
12. Click and select either the vCenter object, datacenter object, cluster object, or ESXi host on which you want to install HPE Storage Connection Manager.
13. Go to **Updates** tab > **Hosts > Baselines**, under **Attached Baselines**, click **Attach > Attach Baseline or Baseline Group**.

14. In the **Attach Baseline or Baseline Group** dialog box, under the baselines list, check the checkbox next to the NCM Baseline, click **Attach** button.
15. In the **Hosts' Compliance** box, click **Check Compliance**. The scan shows that the baseline is non-compliant, meaning that the baseline has not yet been applied to the ESXi host(s).

Note: You may need to refresh the vCenter UI for the Compliance Status to update.

16. Under the baselines list of **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Stage**.
17. In the **Stage Patches** dialog box, select the host(s) you want to upload HPE Storage Connection Manager to, and then click the **Stage** button.
18. Wait until **Stage patches to entity** task is completed under **Recent Tasks** panel
19. On **Updates** tab, under **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Remediate**.
20. In the **Remediate** dialog box, select the host(s) you would like HPE Storage Connection Manager to be remediated, and then click the **Remediate** button.

Results

After the installation:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each Storage volume.
- HPE Storage Path Selection Plugin (PSP) automatically manages the selection of paths to the volumes.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

HPE Storage Connection Manager Installation Using vSphere 6.5 and 6.7 Update Manager With No Internet Connection

When the vCenter Server is not connected to the internet, you can use the vSphere 6.5 and later Update Manager to perform an offline bundle installation of HPE Storage Connection Manager for VMware. This process enables you to install HPE Storage Connection Manager on all the required ESXi hosts at the same time.



Important:

To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Update Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE Storage recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager installation package

See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Copy HPE Storage Connection Manager to the ESXi Host](#) [Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.

- The vSphere 6.5 or 6.7 Update Manager installed on the vCenter Server

Note: For more information about the vSphere Update Manager, see your VMware documentation, including [Installing and Administering VMware vSphere Update Manager](#).

- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.



Important:

If you are upgrading your host from one major version of ESXi (for example ESXi 6.0 to ESXi 6.7) make sure to uninstall HPE Storage Connection Manager prior to doing the ESXi update. Doing this removes the connection service and the path selection policy that are installed when HPE Storage Connection Manager is installed. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47. If you do not uninstall HPE Storage Connection Manager before updating between major ESXi versions, the HPE Storage Path Selection Plugin (PSP) will fail to load and might cause boot failures.

Procedure

1. On the VMware vSphere Web Client > Shortcuts screen, click **Update Manager**.
2. On the **Updates** tab, in **Patch Downloads** box, click **Upload From File**.
3. Click the **Browse** button.
4. In Open File dialog box, select the HPE Storage Connection Manager installation package, and click **Open**.
5. Wait for the upload to complete.
6. Under **Baselines** tab, click **New** > **Baseline**.
7. Type NCM Baseline in the **Name** field.
8. Select **Extension** option for the Content, then click **Next**.
9. Under the extensions list, check the checkbox next to the **Nimble Connection Service**.
10. Click **Next**, then click **Finish**.
11. On the VMware vSphere Web Client > Shortcuts screen, click **Hosts and Clusters**.
12. Click and select the vCenter object, datacenter object, cluster object, or ESXi host on which you would want to install HPE Storage Connection Manager..
13. Go to **Updates** tab > **Host Updates**, under **Attached Baselines**, click **Attach** > **Attach Baseline or Baseline Group**.
14. In the **Attach Baseline or Baseline Group** dialog box, under the baselines list, check the checkbox next to the NCM Baseline, and then click the **Attach** button.
15. In the **Hosts' Compliance** box, click **Check Compliance**. The scan shows that the baseline is non-compliant, meaning that the baseline has not yet been applied to the ESXi host(s).

Note: You may need to refresh the vCenter UI for the Compliance Status to update.

16. Under the baselines list of **Attached Baselines**, check the checkbox next to the NCM Baseline, then click **Stage**.
17. In the **Stage Patches** dialog box, select the host(s) you want to upload HPE Storage Connection Manager to, and then click the **Stage** button.
18. Wait until **Stage patches to entity** task is completed under **Recent Tasks** panel.
19. On **Updates** tab, under **Attached Baselines**, check the checkbox next to the NCM Baseline, and then click **Remediate**.
20. In the **Remediate** dialog box, select the host(s) you would like HPE Storage Connection Manager to be remediated, and then click the **Remediate** button.

Results

After the installation:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each Storage volume.
- HPE Storage Path Selection Plugin (PSP) automatically manages the selection of paths to the volumes.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

ESXCLI Installation of HPE Storage Connection Manager on ESXi 7.x Using an Online Bundle

If you have an internet connection, you can use ESXCLI to install HPE Storage Connection Manager 7.x for VMware on an ESXi 7.x host as an online bundle. ESXCLI connects with the HPE InfoSight download portal and then installs HPE Storage Connection Manager on the ESXi host.

When you install HPE Storage Connection Manager as an online bundle, you do not need to download the HPE Storage Connection Manager software and copy it to the ESXi host.

Before you begin

You must have:

- Root access to the ESXi host
- An internet connection between the ESXi host and HPE InfoSight
- An SSH client such as PuTTY on the Windows or Linux host
- TLS v1.2 or later support on the ESXi host for update.nimblestorage.com.



Important:

If you are upgrading your host to ESXi 7.x from ESXi 6.x, make sure you have done the following **before** you install HPE Storage Connection Manager 7.x:

- Uninstalled HPE Storage Connection Manager 6.x. Doing this removes the `nimble-ncs` and `nimble-ppsp` VIBs for ESXi 6.x. ESXi 7.0 and later requires 64-bit VIBs. Earlier versions of ESXi used 32-bit VIBs. You must removed the 32-bit VIBs before installation HPE Storage Connection Manager 7.x. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47.
- Upgraded the ESXi host to ESXi 7.x.

Procedure

1. Place the ESXi host in maintenance mode using the VMware vSphere Web Client.
2. Launch the SSH client from the Windows or Linux host.
3. Enter the IP address of the ESXi server into the **Host name:** field and click **Open**.
4. Log in to the ESXi host as the root user.
5. From the root directory, run the ESXCLI command line for the HPE Storage Connection Manager 7.x software package that matches your host's version of ESXi. This command line installs the online HPE Storage Connection Manager bundle. The command line uses the format:

```
esxcli software component apply -d https://update.nimblestorage.com/esx<version>/ncm
```

For example, if you are running ESXi 7.0, your command line is:

```
esxcli software component apply -d https://update.nimblestorage.com/esx7.0/ncm
```

6. Reboot the ESXi host using either the ESXi web GUI or the VMware vSphere Web Client. Then reboot the host again. Do not use the CLI command **reboot** or **reboot -f** to restart the ESXi host.



Important:

To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Lifecycle Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Results

After you reboot the ESXi host, the following happens:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each Storage volume.
- HPE Storage Path Selection Plugin (PSP) automatically manages the selection of paths to each Storage volume.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

ESXCLI Installation of HPE Storage Connection Manager on ESXi 7.x Using an Offline Bundle

You can use the ESXCLI to install HPE Storage Connection Manager for VMware 7.x as an offline bundle. Unlike earlier versions of ESXi, ESXi 7.x uses a component packaging format for installation bundles.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager 7.x installation package on the ESXi 7.x host
See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.
- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host
- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.

**Important:**

If you are upgrading your host to ESXi 7.x from ESXi 6.x, make sure you have done the following **before** you install HPE Storage Connection Manager 7.x:

- Uninstalled HPE Storage Connection Manager 6.x. Doing this removes the VIBs for ESXi 6.x. ESXi 7.0 and later requires 64-bit VIBs. Earlier versions of ESXi used 32-bit VIBs. You must removed the 32-bit VIBs before installation HPE Storage Connection Manager 7.x. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47.
- Upgraded the ESXi host to ESXi 7.x.

Procedure

1. Place the ESXi host in maintenance mode using the VMware vSphere Web Client.
2. Launch your SSH client on the Windows or Linux host.
3. Enter the IP address of the ESXi host into the **Host name:** field and click **Open**.
4. Log in to the ESXi host as the root user.
5. From the root directory, run the **ESXCLI software component apply** command line to install HPE Storage Connection Manager. This command line takes the form:

```
esxcli software component apply -d /tmp/nimble-ncm-for-esxX-X.x.x-xxxxxx.zip
```

where

- esxX-X.x.x is the ESXi version.
- -xxxxxx is the current build number of the HPE Storage Connection Manager installation package.

The command line you use **must** include the version and build information for the version of HPE Storage Connection Manager and ESXi that you are installing. For example, if the build information was **700012**, you would enter:

```
esxcli software component apply -d /tmp/nimble-ncm-for-esx7.0-7.0.0-700012.zip
```

**Important:**

- Enter the absolute path to the HPE Storage Connection Manager download. Do **not** use a relative path.
- Do not use any spaces or special characters in the path.

6. Reboot the ESXi host using either the ESXi web GUI or the VMware vSphere Web Client. Then reboot the host again. Do not use the CLI command **reboot** or **reboot -f** to restart the ESXi host.

**Important:**

To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices.

HPE recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Results

After you reboot the ESXi host, the following happens:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each Storage volume.
- HPE Storage Path Selection Plugin (PSP) automatically manages the selection of paths to each Storage volume.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Online Bundle

If you have an internet connection, you can use ESXCLI to install HPE Storage Connection Manager for VMware on the ESXi host as an online bundle. ESXCLI connects with the HPE InfoSight download portal and then installs HPE Storage Connection Manager on the ESXi host.

When you install HPE Storage Connection Manager as an online bundle, you do not need to download the HPE Storage Connection Manager software and copy it to the ESXi host.

Note: ESXi uses the same NCM package for both ESXi 6.5 and ESXi 6.7. When installing on ESXi 6.7, use the same URL that would be used to install/download NCM for ESXi 6.5.

Before you begin

You must have:

- Root access to the ESXi host
- An internet connection between the ESXi host and HPE InfoSight
- An SSH client such as PuTTY on the Windows or Linux host
- TLS v1.2 or later support on the ESXi host for `update.nimblestorage.com`

Procedure

1. Place the ESXi host in maintenance mode using the VMware vSphere Web Client.
2. Launch the SSH client from the Windows or Linux host.
3. Enter the IP address of the ESXi server into the **Host name:** field and click **Open**.

4. Log in to the ESXi host as the root user.
5. From the root directory, run the ESXCLI command line for the HPE Storage Connection Manager software package that matches your hosts version of ESXi. This command line installs the online HPE Storage Connection Manager bundle. The command line uses the format:

```
esxcli software vib install -d https://update.nimblestorage.com/esx<version>/ncm
```

where <version> is your version of ESXi. For example, if you are running ESXi 6.5, your command line is:

```
esxcli software vib install -d https://update.nimblestorage.com/esx6.5/ncm
```

but if you are running ESXi 6.0, it is:

```
esxcli software vib install -d https://update.nimblestorage.com/esx6/ncm
```

Note: NCM is the same for both ESXi 6.5 and ESXi 6.7, so the command to run either of these is:

```
esxcli software vib install -d https://update.nimblestorage.com/esx6.5/ncm
```

6. Reboot the ESXi host after installing, updating, or uninstalling HPE Storage Connection Manager.



Important:

If you are installing or updating the HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Update Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Results

After the ESXi host is rebooted twice to complete the installation:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each volume.
- HPE Storage Path Selection Plugin (PSP) automatically manages the selection of paths to each volume.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.

ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Offline Bundle

If you do not have an internet connection, you can use ESXCLI to install HPE Storage Connection Manager 6.x for VMware on the ESXi host as an offline bundle.

Note: ESXi uses the same NCM package for both ESXi 6.5 and ESXi 6.7. When installing on ESXi 6.7, use the same URL that would be used to install/download NCM for ESXi 6.5.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager installation package on the ESXi host

See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host

Procedure

1. Place the ESXi host in maintenance mode using the VMware vSphere Web Client.
2. Launch your SSH client on the Windows or Linux host.
3. Enter the IP address of the ESXi host into the **Host name:** field and click **Open**.
4. Log in to the ESXi host as the root user.
5. From the root directory, run the command line to install the HPE Storage Connection Manager bundle. This command line takes the form:

```
esxcli software vib install -d /tmp/nimble-ncm-for-esxX-X.x.x-xxxxxx.zip
```

where

- esxX-X.x.x is the ESXi version.
- -xxxxxx is the current build number of the HPE Storage Connection Manager installation package.

The command line you use **must** include the version and build information for the version of HPE Storage Connection Manager and ESXi that you are installing. For example, if you downloaded HPE Storage Connection Manager for ESXi 6.5 in March 2019, the command line at that time would be:

```
esxcli software vib install -d /tmp/nimble-ncm-for-esx6.5-5.1.0-650006.zip
```



Important:

- Enter the absolute path to the HPE Storage Connection Manager download. Do **not** use a relative path.
- Do not use any spaces or special characters in the path.

6. Reboot the ESXi host after installing, updating, or uninstalling HPE Storage Connection Manager.



Important:

If you are installing or updating HPE Storage Connection Manager, you must reboot ESX twice. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices. The Update Manager automatically reboots the ESXi host once. You must reboot it a second time.

HPE recommends that you use the ESXi web GUI or the vSphere Web Client to reboot the host instead of entering the CLI reboot command from the ESXi host console. The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection.

Results

After the ESXi host is rebooted to complete the installation:

- HPE Storage Connection Service automatically creates the optimal number of iSCSI sessions for each volume.
- HPE Storage Path Selection Plugin (PSP) automatically claims the available devices, including both existing devices and devices added later.

What to do next

Go to [Verify the HPE Storage Connection Manager Installation](#) on page 38.



Updating HPE Storage Connection Manager for ESXi 7.x

When you update HPE Storage Connection Manager for VMware 7.x on the ESXi 7.x host, you install a full version of HPE Storage Connection Manager on top of an existing installation. You can install the update without uninstalling the current version of HPE Storage Connection Manager.

Before you begin

You must have:

- The downloaded HPE Storage Connection Manager 7.x installation package on the ESXi 7.x host
See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.
- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host
- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.



Important:

ESXi 7.x does not support the 32-bit VIBs used with earlier versions of ESXi and HPE Storage Connection Manager. If you are upgrading your host to ESXi 7.x from ESXi 6.x, remove the previous version of HPE Storage Connection Manager before updating. Once you complete the ESXi update, install HPE Storage Connection Manager 7.x for ESXi 7.x.

Procedure

1. Install the update. Make sure you use the correct HPE Storage Connection Manager software package for your ESXi host. You can use one of the following methods to install the HPE Storage Connection Manager update:
 - *HPE Storage Connection Manager Installation Using vSphere 7 Lifecycle Manager When Connected to the Internet*
 - *ESXCLI Installation of HPE Storage Connection Manager 7.x with ESXi 7.x UsHPE Storage Connection Managering an Online Bundle*
 - *ESXCLI Installation of HPE Storage Connection Manager 7.x with ESXi 7.x Using an Offline Bundle*
2. Reboot the ESXi host using either the ESXi web GUI or the VMware vSphere Web Client. Then reboot the host again. Do not use the CLI command **reboot** or **reboot -f** to restart the ESXi host.

Note: To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices.

The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection to verify the installation.

You must reboot the ESXi host whenever you update or uninstall HPE Storage Connection Manager or whenever you perform a fresh install of ESXi 6.x.

Update HPE Storage Connection Manager for ESXi 6.x

When you update HPE Storage Connection Manager 6.x for VMware on the ESXi host, you install a full version of HPE Storage Connection Manager on top of an existing installation. You can install the update without uninstalling the current version of HPE Storage Connection Manager.



**Important:**

If you are upgrading your host from one major version of ESXi (for example ESXi 6.0 to ESXi 6.7) make sure to uninstall HPE Storage Connection Manager prior to doing the ESXi update. Doing this removes the connection service and the path selection policy that are installed when HPE Storage Connection Manager is installed. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47. If you do not uninstall HPE Storage Connection Manager before updating between major ESXi versions, the HPE Storage Path Selection Plugin (PSP) will fail to load and might cause boot failures.

Before you begin

You must have:

- Downloaded the HPE Storage Connection Manager installation package and moved it to the ESXi host if you are performing an offline installation.

See [Manually Download the HPE Storage Connection Manager Software Package](#) on page 23 and [Manually Copy HPE Storage Connection Manager to the ESXi Host](#) on page 24.

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host if you are using the ESXCLI commands
- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete the installation.

Procedure

1. Install the update. Make sure you use the correct HPE Storage Connection Manager software package for your ESXi host.

You can use one of the following methods to install the HPE Storage Connection Manager update:

- [ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Online Bundle](#) on page 34
- [ESXCLI Installation of HPE Storage Connection Manager on ESXi 6.x Offline Bundle](#) on page 35
- [HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager When Connected to the Internet](#)
- [HPE Storage Connection Manager Installation Using vSphere 6.0 Update Manager With No Internet Connection](#)

2. Reboot the ESXi host using either the ESXi web GUI or the VMware vSphere Web Client. Then reboot the host again.

Do not use the CLI command **reboot** or **reboot -f** to restart the ESXi host.

Note: To properly install HPE Storage Connection Manager, two reboots of ESXi are required. The two reboots are necessary to enable the CompareLUNNumber advanced setting of HPE Storage Connection Manager to take effect on all existing storage devices.

The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection to verify the installation.

You must reboot the ESXi host whenever you update or uninstall HPE Storage Connection Manager or whenever you perform a fresh install of ESXi 6.0 or later. You do not need to reboot if you are performing a fresh install of HPE Storage Connection Manager on ESXi 5.x.

Verify the HPE Storage Connection Manager Installation

It is a good practice to verify your new HPE Storage Connection Manager for VMware installation.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and click **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, run the command line for your version of HPE Storage Connection Manager.

If your version of HPE Storage Connection Manager supports 64-bit VIBs and runs on ESXi 7.x, use the following command line:

esxcli software component list | grep -iE 'nimble|hpe'

If your version of HPE Storage Connection Manager supports 32-bit VIBs, which are supported by ESXi versions prior to 7.0, use the following command line:

esxcli software vib list | grep -iE 'nimble|hpe'

The components installed vary depending on the version of storage connection manager that was installed.

- If NCM 7.0.1 or earlier is installed, and you see nimble-ncs and nimble-psp, the installation was successful.
- If NCM 7.0.2 or later is installed, and you see HPE-Storage_Connection-Service and HPE-Storage-psp, the installation was successful.

Note: After you install HPE Storage Connection Manager, it runs the script `nimblepspd` each time the host boots. One of the tasks this script performs is to set two advanced system properties on the host:

- FailDiskRegistration to 1
 - CompareLUNNumber to 0
-

Verify Settings for CompareLUNNumber and FailDiskRegistration

When you reboot the host, the HPE Storage Connection Manager for VMware runs the script `nimblepspd`, which sets the following two advanced system properties on the host:

- FailDiskRegistration to 1
- CompareLUNNumber to 0

You can then verify that CompareLUNNumber and FailDiskRegistration have the correct values by performing the following steps:

Procedure

1. Confirm that CompareLUNNumber is disabled. Enter the following command line to verify that it is set to 0:

esxcli system settings advanced list | grep -A5 CompareLUNNumber | grep -e "Path:" -e "\s*Int Value:"

For example, you might enter:

```
[root@localhost:~] esxcli system settings advanced list | grep -A5 CompareLUN▶
Number | grep -e "Path:" -e "\s*Int Value:"
  Path: /Scsi/CompareLUNNumber
  Int Value: 0
```

2. Verify that FailDiskRegistration is enabled. Enter the following command line to verify that it is set to 1:

esxcli system settings advanced list | grep -A5 FailDiskRegistration | grep -e "Path:" -e "\s*Int Value:"

For example, you might enter::

```
[root@localhost:~] esxcli system settings advanced list | grep -A5 FailDiskReg▶
istration | grep -e "Path:" -e "\s*Int Value:"
```

```
Path: /Disk/FailDiskRegistration
Int Value: 1
```

Configure HPE Storage Connection Manager on the ESXi Host

You can view and change the configuration for your HPE Storage Connection Manager for VMware. Keep in mind that the default HPE Storage Connection Manager configuration has been tested as suitable for most users. HPE recommends that you do not change the HPE Storage Connection Manager configuration unless you have specific reasons for doing so.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- A text editor such as vi

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and click **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, change the directory.
`cd /etc/nimble/`
5. From the `/etc/nimble/` directory, type `cat ncm.conf`.

The HPE Storage Connection Manager configuration file displays information about the configuration. The following table contains information about the standard HPE configuration:

Configuration Items and Default Values	Descriptions and Options
interval: 120	Interval refers to time in seconds. Interval controls how often in seconds HPE Storage Connection Manager will poll for iSCSI connection updates: <ul style="list-style-type: none"> • Minimum: 60 • Default: 120
min_vol_sessions: 2	The minimum number of iSCSI sessions connected to each storage volume (target) that HPE Storage Connection Manager should try to maintain: <ul style="list-style-type: none"> • Default: 2
max_volsessions: 8	The maximum number of iSCSI sessions connected to each storage volume (target) that HPE Storage Connection Manager should allow before disconnecting a session: <ul style="list-style-type: none"> • Default: 8
vol_sessions: 0	Controls the target number of sessions per Storage volume (target), regardless of HPE Storage Connection Manager's automatic calculation. See the section <i>How NCM Supports Groups and Pools</i> for information about this calculation. <ul style="list-style-type: none"> • Default: 0 (disabled)

Configuration Items and Default Values	Descriptions and Options
log_level: 1	<p>Controls the logging level for HPE Storage Connection Manager.</p> <ul style="list-style-type: none"> • ERROR+WARNING only: 0 • ERROR+WARNING+INFO only: 1 • ERROR+WARNING+INFO+DEBUG: 2 • Default: 1
worker_stop: 0	<ul style="list-style-type: none"> • 0 allows the HPE Connection Service to monitor your sessions. • 1 prevents the connection service from monitoring your sessions.

Note: The valid range of values for *min_vol_sessions* and *max_vol_sessions* is from 2 to 32. If a value is specified outside this range, the HPE Storage Connection Manager reverts to the default value.

6. (Optional) If you want to make changes to the configuration, open the `ncm.conf` file in a text editor. After you make your changes, save the file and close it.

Using an HPE Dual 8GB MicroSD Enterprise Midline USB

You can set up HPE Storage Connection Manager for VMware to use an HPE Dual 8GB MicroSD Enterprise Midline USB to boot your ESXi system. This is a dual-microSD card module that provides data redundancy through a mirrored RAID-1 configuration.

To use an HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system, you must make sure that, after you install HPE Storage Connection Manager, you have the following settings:

- CompareLUNNumber is disabled. It should be set to 0.
- FailDiskRegistration is enabled. It must be set to 1.
- MASK_PATH rules are set for twp Redundant Paths for the USB drive.

If you do not disable CompareLUNNumber, save the log files located at `/var/log/nimble` and contact HPE Support.

The MASK_PATH rules are normally paths `C:0 T:0 L:0` and `C:0 T:0 L:1`. The rule IDs are normally less than 50. To verify that these are correctly set, you can use the command line:

localcli storage core claimrule list | grep MASK_PATH | grep -v "vendor="

For example, you might enter:

```
[root@localhost:~] :~] localcli storage core claimrule list | grep MASK_PATH
| grep -v "vendor="
MP 39 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0 lun=1
false false 0
MP 39 file location MASK_PATH adapter=vmhba32 channel=0 target=0 lun=1
false false 0
MP 40 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0 lun=2
false false 0
MP 40 file location MASK_PATH adapter=vmhba32 channel=0 target=0 lun=2
false false 0
```

View HPE Storage Connection Manager Logs

The log files produced by the HPE Storage Connection Manager for VMware can provide helpful troubleshooting information. It is a good practice to check the logs if you have a problem or if you need to contact HPE Support.

Note: By default, the log files are placed in the `/var/log/nimble/` directory on the ESXi host. If you are using HPE Storage Connection Manager 6.x.x or later, you can specify a different location for the log files.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- (Optional) A text editor such as vi. You can view the files using either a text editor or a command, such as **more** or **cat**.

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Type the IP address of the ESXi host in the **Host name:** field, then click **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, change to the logs directory.
`cd /var/log/nimble/`
5. From the `/var/log/nimble/` directory, view the list of logs by entering the **ls** command.

You should see the following list of HPE Storage Connection Manager logs:

- `ncm.log`
- `nimble_psp_installer.log`
- `nimble_psp_policy.log`

6. To view an individual log file, you can use a command such as **more** or **cat**. You can also open a log file in a text editor.

Configure Custom Log Locations

HPE Storage Connection Manager for VMware allows you to set a custom path for your log files by modifying the `nimble_logs_path` parameter in the `ncm.conf` file.

Note: The default location for the log files is `/var/log/nimble`.

Before you begin

You must have:

- HPE Storage Connection Manager 6.0.0 or later

Note: This feature is not supported with earlier versions of HPE Storage Connection Manager.

- An SSH client such as PuTTY on the Windows or Linux host
- A text editor

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Type the IP address of the ESXi host in the **Host name:** field, then click **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, change to the directory containing the HPE Storage Connection Manager configuration file.

```
cd /etc/nimble/
```

5. Open the `ncm.conf` file for editing.
6. Enter the directory where you want HPE Storage Connection Manager to place the log files as the value for the `nimble_logs_path` parameter:

```
nimble_logs_path=<directory_name>
```

7. Save the `ncm.conf` file and close it.

HPE Storage Connection Manager will place all future log files in the directory you specified.

Uninstall HPE Storage Connection Manager 7.x from an ESXi 7.x Host

If you need to, you can uninstall your HPE Storage Connection Manager 7.x for VMware. The steps that you use to uninstall HPE Storage Connection Manager 7.x when it is installed on an ESXi 7.x host differ from the steps you use when it is installed on an earlier version of ESXi. This is because ESXi 7.x uses 64-bit VIBs while earlier versions use 32-bit VIBs.

Note: You do not need to uninstall HPE Storage Connection Manager 7.x if you are upgrading to a new version of HPE Storage Connection Manager 7.x.

Before you begin

You must have:

- Root access to the ESXi host
- An internet connection between the ESXi host and HPE InfoSight
- An SSH client such as PuTTY on the Windows or Linux host
- Uninstalled HPE Storage Connection Manager versions that use 32-bit VIBs



Important:

If you are upgrading your host to ESXi 7.x from ESXi 6.x, make sure you have done the following **before** you install HPE Storage Connection Manager 7.x:

- Uninstalled HPE Storage Connection Manager 6.x. Doing this removes the `nimble-ncs` and `nimble-psp` VIBs for ESXi 6.x. ESXi 7.0 and later requires 64-bit VIBs. Earlier versions of ESXi used 32-bit VIBs. You must removed the 32-bit VIBs before installation HPE Storage Connection Manager 7.x. See [Uninstall HPE Legacy NCM Versions \(7.0.1 and lower\) from an ESXi 6.x Host](#) on page 47.
 - Upgraded the ESXi host to ESXi 7.x.
-

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and select **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, stop the PSPD server by entering the command line:

```
/etc/init.d/nimblepspd stop
```

5. Reset the CompareLUNNumber to 1 for Enabled (this is the default) by entering the command line:

```
esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber
```

For example, you might enter:

```
[root@localhost:~] esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber
```

To verify that the setting is correct, you can enter:

```
[root@localhost:~] esxcli system settings advanced list | grep -A2 CompareLUN▶
Number
  Path: /Scsi/CompareLUNNumber
  Type: integer
  Int Value: 1
```

6. Reset the FailDiskRegistration to 0 for Disabled (this is the default) by entering the command line:

esxcli system settings advanced set --int-value=0 --option=/Disk/FailDiskRegistration

For example, you might enter:

[root@localhost:/tmp/scratch/log/nimble] esxcli system settings advanced set --int-value=0 --option=/Disk/FailDiskRegistration

To verify that the setting is correct, you can enter:

```
[root@localhost:/tmp/scratch/log/nimble] esxcli system settings advanced list
| grep -A2 FailDiskRegistration
  Path: Disk/FailDiskRegistration
  Type: integer
  Int Value: 0
```

7. **(USB users only)** You must remove the MASK_PATH rules that you set up for an HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system. These rules set the mask out to "Size 0".



Important: Do not remove the default Dell Universal Xport vendor specific rules.

Normally the paths for the HPE USB-booted systems are **C:0 T:0 L:0** and **C:0 T:0 L:1** and the rule IDs are less than 50. To locate the rules, use the command line:

localcli storage core claimrule list | grep MASK_PATH | grep -v "vendor="

When you have the list of rules that apply to the HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system, use the following command to remove them. You need enter this command line for each rule.

localcli storage core claimrule remove --rule=<rule_number>

<rule_number> is the number of the rule.

In the following example, you need to remove rules 39 and 40.

```
[root@localhost:~] localcli storage core claimrule list | grep MASK_PATH |
grep -v "vendor="
  MP 39 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 39 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 40 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0
  MP 40 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0

[root@localhost:~] localcli storage core claimrule remove --rule=39
[root@localhost:~] localcli storage core claimrule remove --rule=40
[root@localhost:~] localcli storage core claimrule load
```

8. **(USB users only)** Confirm that all the rules have been removed by entering the command line:

localcli storage core claimrule list

9. Remove the HPE Storage Connection Manager software by running the command line:

esxcli software component remove -n HPE_NimbleConnectionManagement

10. Reboot the ESXi host through the ESXi web GUI or from the vSphere Web Client.

Note: Do not use the CLI command **reboot** or **reboot -f** to restart the ESXi host.

The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection to verify the uninstall.

Uninstall HPE Storage Connection Manager 7.x from an ESXi 6.x Host

You might want to uninstall HPE Storage Connection Manager for VMware. For example, if you upgrade your ESXi host to version 6.x, you must uninstall your current version of HPE Storage Connection Manager.

These steps explain how to uninstall versions of HPE Storage Connection Manager that use 32-bit VIBs. These steps apply to all versions of HPE Storage Connection Manager running on versions of ESXi prior to 7.0

Note: You do not need to uninstall HPE Storage Connection Manager if you are updating to a new version of HPE Storage Connection Manager as long as the current version of HPE Storage Connection Manager and the new version of HPE Storage Connection Manager support the same size VIBs. ESXi 7.x requires 64-bit VIBs. Prior to ESX 7.0, ESXi hosts supported 32-bit VIBs.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host
- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete this operation.

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and select **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, stop the HPE Storage PSPD server by entering the command line:
/etc/init.d/nimblepspd stop
5. Reset the CompareLUNNumber to 1 for Enabled (this is the default) by entering the command line:
esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber

For example, you might enter:

```
[root@localhost:~] esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber
```

To verify that the setting is correct, you can enter:

```
[root@localhost:~] esxcli system settings advanced list | grep -A2 CompareLUN▶
Number
  Path: /Scsi/CompareLUNNumber
  Type: integer
  Int Value: 1
```

6. Reset the FailDiskRegistration to 0 for Disabled (this is the default) by entering the command line:
esxcli system settings advanced set --int-value=0 --option=/Disk/FailDiskRegistration

For example, you might enter:

```
[root@localhost:/tmp/scratch/log/nimble] esxcli system settings advanced set --int-value=0
--option=/Disk/FailDiskRegistration
```

To verify that the setting is correct, you can enter:

```
[root@localhost:~/tmp/scratch/log/nimble] esxcli system settings advanced list
| grep -A2 FailDiskRegistration
  Path: Disk/FailDiskRegistration
  Type: integer
  Int Value: 0
```

- 7. (USB users only)** You must remove the MASK_PATH rules that you set up for an HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system. These rules set the mask out to "Size 0".



Important: Do not remove the default Dell Universal Xport vendor specific rules.

Normally the paths for the HPE USB-booted systems are **C:0 T:0 L:0** and **C:0 T:0 L:1** and the rule IDs are less than 50. To locate the rules, use the command line:

localcli storage core claimrule list | grep MASK_PATH | grep -v "vendor="

When you have the list of rules that apply to the HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system, use the following command to remove them. You need enter this command line for each rule.

localcli storage core claimrule remove --rule=<rule_number>

<rule_number> is the number of the rule.

In the following example, you need to remove rules 39 and 40.

```
[root@localhost:~] localcli storage core claimrule list | grep MASK_PATH |
grep -v "vendor="
  MP 39 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 39 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 40 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0
  MP 40 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0

[root@localhost:~] localcli storage core claimrule remove --rule=39
[root@localhost:~] localcli storage core claimrule remove --rule=40
[root@localhost:~] localcli storage core claimrule load
```

- 8. (USB users only)** Confirm that all the rules have been removed by entering the command line:

localcli storage core claimrule list

- 9.** Remove the `HPE-Storage-Connection-Service` bundle by running the command line:

esxcli software vib remove --vibName=HPE-Storage-Connection-Service --maintenance-mode

- 10.** Remove the `HPE-Storage-psp` bundle by running the command line:

esxcli software vib remove --vibName=HPE-Storage-psp --maintenance-mode

- 11.** Reboot the ESXi host through the ESXi web GUI or from the vSphere Web Client.

Note: HPE recommends that you use the ESXi web GUI to reboot the host instead of entering the CLI **reboot** command from the ESXi host console.

The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection to verify the uninstall.

Uninstall HPE Legacy NCM Versions (7.0.1 and lower) from an ESXi 6.x Host

You might want to uninstall HPE Storage Connection Manager for VMware. For example, if you upgrade your ESXi host to version 6.x, you must uninstall your current version of HPE Storage Connection Manager.

These steps explain how to uninstall versions of HPE Storage Connection Manager that use 32-bit VIBs. These steps apply to all versions of HPE Storage Connection Manager running on versions of ESXi prior to 7.0

Note: You do not need to uninstall HPE Storage Connection Manager if you are updating to a new version of HPE Storage Connection Manager as long as the current version of HPE Storage Connection Manager and the new version of HPE Storage Connection Manager support the same size VIBs. ESXi 7.x requires 64-bit VIBs. Prior to ESX 7.0, ESXi hosts supported 32-bit VIBs.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host
- A 15-minute interval when your ESXi host can be offline. You must then reboot the ESXi host to complete this operation.

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and select **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, stop the HPE Storage PSPD server by entering the command line:
/etc/init.d/nimblepspd stop
5. Reset the CompareLUNNumber to 1 for Enabled (this is the default) by entering the command line:
esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber

For example, you might enter:

```
[root@localhost:~] esxcli system settings advanced set --int-value=1 --option=/Scsi/CompareLUNNumber
```

To verify that the setting is correct, you can enter:

```
[root@localhost:~] esxcli system settings advanced list | grep -A2 CompareLUN▶
Number
Path: /Scsi/CompareLUNNumber
Type: integer
Int Value: 1
```

6. Reset the FailDiskRegistration to 0 for Disabled (this is the default) by entering the command line:
esxcli system settings advanced set --int-value=0 --option=/Disk/FailDiskRegistration

For example, you might enter:

```
[root@localhost:/tmp/scratch/log/nimble] esxcli system settings advanced set --int-value=0
--option=/Disk/FailDiskRegistration
```

To verify that the setting is correct, you can enter:

```
[root@localhost:/tmp/scratch/log/nimble] esxcli system settings advanced list
| grep -A2 FailDiskRegistration
Path: Disk/FailDiskRegistration
Type: integer
Int Value: 0
```

- 7. (USB users only)** You must remove the MASK_PATH rules that you set up for an HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system. These rules set the mask out to "Size 0".



Important: Do not remove the default Dell Universal Xport vendor specific rules.

Normally the paths for the HPE USB-booted systems are **C:0 T:0 L:0** and **C:0 T:0 L:1** and the rule IDs are less than 50. To locate the rules, use the command line:

localcli storage core claimrule list | grep MASK_PATH | grep -v "vendor="

When you have the list of rules that apply to the HPE Dual 8GB MicroSD Enterprise Midline USB-booted ESXi system, use the following command to remove them. You need enter this command line for each rule.

localcli storage core claimrule remove --rule=<rule_number>

<rule_number> is the number of the rule.

In the following example, you need to remove rules 39 and 40.

```
[root@localhost:~] localcli storage core claimrule list | grep MASK_PATH |
grep -v "vendor="
  MP 39 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 39 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=1 false false 0
  MP 40 runtime location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0
  MP 40 file location MASK_PATH adapter=vmhba32 channel=0 target=0
lun=2 false false 0

[root@localhost:~] localcli storage core claimrule remove --rule=39
[root@localhost:~] localcli storage core claimrule remove --rule=40
[root@localhost:~] localcli storage core claimrule load
```

- 8. (USB users only)** Confirm that all the rules have been removed by entering the command line:

localcli storage core claimrule list

- 9.** Remove the `nimble-ncs` bundle by running the command line:

esxcli software vib remove --vibName=nimble-ncs --maintenance-mode

- 10.** Remove the `nimble-psp` bundle by running the command line:

esxcli software vib remove --vibName=nimble-psp --maintenance-mode

- 11.** Reboot the ESXi host through the ESXi web GUI or from the vSphere Web Client.

Note: HPE recommends that you use the ESXi web GUI to reboot the host instead of entering the CLI **reboot** command from the ESXi host console.

The reboot operation takes 10 to 15 minutes. When the reboot is complete, you must reestablish your SSH client connection to verify the uninstall.

Verify the HPE Storage Connection Manager Uninstall

It is a good practice to verify that the operation to uninstall HPE Storage Connection Manager for VMware installation succeeded.

Before you begin

You must have:

- An SSH client such as PuTTY on the Windows or Linux host
- Root access to the ESXi host

Procedure

1. Launch your SSH client on the Windows or Linux host.
2. Enter the IP address of the ESXi host in the **Host name:** field and click **Open**.
3. Log in to the ESXi host as the root user.
4. From the root directory, run the command line for your version of HPE Storage Connection Manager.

If your version of HPE Storage Connection Manager supports 64-bit VIBs and runs on ESXi 7.x, use the following command line:

esxcli software component list | grep -iE nimblehpe

If your version of HPE Storage Connection Manager supports 32-bit VIBs, which are supported by ESXi versions prior to 7.0, use the following command line:

esxcli software vib list | grep -iE 'nimblehpe'

If `nimble-ncs` and `nimble-psp` no longer appear in the list, the uninstall operation was successful.

Note: When you uninstall HPE Storage Connection Manager, the two advanced system properties on the host that were changed by the script `nimblepspd` do not automatically return to their default values. The HPE Storage Connection Manager uninstall procedure describes how to use the `esxcli` commands to set these properties back to their default values. The default value for `FailDiskRegistration` is 0. The default value for `CompareLUNNumber` is 1.



VMware iSCSI Configuration

You can use the iSCSI protocol with a VMware ESXi host; however, you need to perform some basic setup tasks.

1 Set up your network configuration.

This involves setting up ESXi networking to work with one or more switches and configuring the ESXi software iSCSI adapter. The sections that follow explain how to perform this configuration.

2 Complete the integration of your storage array with VMware by setting the number of iSCSI sessions and the path selection policy for the volumes.

You can manually set this information; however, the best practice is to use HPE Storage Connection Manager for VMware. HPE Storage Connection Manager simplifies these tasks by automatically

- Creating the optimal number of iSCSI sessions for each storage volume using its Connection Service component.
- Managing the path selection for the volumes using its HPE Storage Path Selection Policy (PSP) component.

If you do not use HPE Storage Connection Manager, you must manually set the Path Selection Policy on each storage LUN to **Round Robin (VMware)**.

If you do not use the vCenter Plugin, you must set ACLs (initiator group and/or CHAP username or WWPN) on all volumes.

3 Decide whether to enable iSCSI Digest.

The sections that follow provide information about setting up your configuration, the advantage of using HPE Storage Connection Manager, and enabling and disabling iSCSI Digest.

High-Level Steps to Set Up the ESXi iSCSI Network Configuration

You must set up your ESXi iSCSI network configuration to enable your storage volumes to be recognized as ESXi datastores. The following is a high-level overview of the required steps. You **must** perform the steps in the specified order.

Procedure

1. Configure ESXi iSCSI networking using the appropriate procedure for your system configuration:

- [Configure ESXi iSCSI Networking with Multiple vSwitches](#) on page 51
- [Configure ESXi iSCSI Networking with a Single vSwitch](#) on page 52

2. [Configure the ESXi iSCSI Software Adapter](#) on page 53

3. [Bind VMK Ports to ESXi iSCSI Software Adapter](#) on page 53

Note: You must do this when you have a single IP subnet and the same broadcast domain.

What to do next

If you have HPE Storage Connection Manager for VMware installed, it will now set up the optimal number of iSCSI sessions for each storage volume and manage the path selection for the volumes.

If you do not have HPE Storage Connection Manager installed, you must manually set the Path Selection Policy (PSP) on each LUN to **Round Robin (VMware)**.



Configure ESXi iSCSI Networking with Multiple vSwitches

This task takes you through the process of setting up an ESXi iSCSI network configuration for multiple switches. The examples in these steps use the following configuration:

- Two vmnic ports
- Two vmk ports
- Two vSwitches. Each switch has one vmnic port and one vmk port each

Before you begin

The following requirements apply:

- Maintain a one-to-one relationship between vmnic ports and vmk ports.
For example, if you have four vmnic ports, you must have four vmk ports.
- Disable NIC Teaming.

Each vmk port can have only one Active vmnic port and no Standby vmnic ports. HPE Storage best practice is to **not** use NIC Teaming in this configuration.

Note: Your system must have a consistent, end-to-end maximum transmission unit (MTU) that flows from the host to the vSwitch to the switch to the array. You can either use a standard MTU on all the devices or you can use Jumbo Frames. For information about Jumbo Frames, see [Configure Jumbo Frames](#) on page 147. For information about confirming or modifying your MTU values, see *Validate the MTU Settings*.

Procedure

1. At the Configure Physical Adapters screen on the ESXi host, identify the vmnics that you want to use for iSCSI networking.
2. Go to the Configure Virtual Switches screen on the ESXi host and click **Add Networking**.
3. Select **VMkernel Network Adapter** and click **Next**.
4. Select **Create a virtual switch** and click **Next**.
5. Click **(+)** and select the desired vmnic. Click **OK** and then **Next**.
6. Assign a Port Group name and click **Next**.
7. Assign an IP address and subnet information for the first vmk port and click **Next**.
8. Review the proposed configuration. If everything is correct, click **Finish**.
The Configure Virtual Switches screen displays the new vSwitch and vmk port.
9. To configure the second switch, click **Add Networking**.
10. Select **VMkernel Network Adapter** and click **Next**.
11. Select **New standard switch** and click **Next**.
12. Click **(+)** and select the desired vmnic. Click **OK** and **Next**.
13. Assign a unique Port Group name and click **Next**.
14. Assign an IP address and subnet information. Click **Next**.
15. Review the proposed configuration. If everything is correct, click **Finish**.
The Configure Virtual Switches screen displays the new vSwitch and vmk port.

What to do next

Next, go to [Configure the ESXi iSCSI Software Adapter](#) on page 53.



Configure ESXi iSCSI Networking with a Single vSwitch

This task takes you through the process of setting up an ESXi iSCSI network configuration for one switch. The examples in these steps use the following configuration:

- Two vmnic ports
- Two vmk ports
- One vSwitch that contains both vmnic ports and vmk ports

Before you begin

The following requirements apply:

- Maintain a one-to-one relationship between vmnic ports and vmk ports.
For example, if you have four vmnic ports, you must have four vmk ports.
- Disable NIC Teaming.

When a vSwitch used for iSCSI has more than one vmnic, NIC Teaming must be disabled. Each vmk port can have only one Active vmnic port and no Standby vmnic ports. This procedure contains steps to disable NIC Teaming.

Note: Your system must have a consistent, end-to-end maximum transmission unit (MTU) that flows from the host to the vSwitch to the switch to the array. You can use a standard MTU on all the devices or you can use Jumbo frames. For information about Jumbo Frames, see [Configure Jumbo Frames](#) on page 147. For information about confirming or modifying your MTU values, see *Validate the MTU Settings*.

Procedure

1. At the Configure Physical Adapters screen on the ESXi host, identify the vmnics you want to use for iSCSI networking.
2. Go to the Configure Virtual Switches screen on the ESXi host and click **Add Networking**.
3. Select **VMkernel Network Adapter** and click **Next**.
4. Select **New standard switch** and the two vmnic ports you want to use. Click **Next**.
5. Assign a portgroup name. This label must be unique on the ESXi host being configured. Click **Next**.
6. Assign an IP address and subnet information for the first vmk port and click **Next**.
7. Review the proposed configuration. If everything is correct, click **Finish**.
The Configure Virtual Switches screen displays the new vSwitch and vmk port.
8. To start configure the second vmk port, click **Add Networking**.
9. Select **VMkernel** and click **Next**.
10. Select the vSwitch that you created earlier and click **Next**.
11. Assign a unique portgroup name and click **Next**.
12. Assign an IP address and subnet information, then click **Next**.
13. Review the proposed configuration. If everything is correct, click **Finish**.
The Configure Virtual Switches screen displays the new vSwitch and vmk ports. Write down which vmk port is associated with which portgroup name. In this example, the relationships are vmk1 = iSCSI1 and vmk2 = iSCSI2.
14. To disable NIC Teaming, select the vSwitch vmnic1.
15. Select the first Port Group name and click **Edit**.
16. Go to the NIC Teaming tab and check **Override vSwitch failover order**.
Click **OK** to finish.
17. Use the **Move Down** button to move the vmnic that you do **not** want to use to the **Unused Adapters** section. Then click **OK**.
18. Select the second port group/network label and click **Edit**.



19. Go to the NIC Teaming tab and check **Override vSwitch failover order**.
20. Use the **Move Down** button to move the vmnic that you do **not** want to use to the **Unused Adapters** section. Then click **OK**.

What to do next

Next, go to [Configure the ESXi iSCSI Software Adapter](#) on page 53.

Configure the ESXi iSCSI Software Adapter

After you have set up an ESXi iSCSI network configuration for all the switches you are using, you need to configure the ESXi iSCSI Software Adapter.

Note: For more details about working with the storage array and the discovery IP address, see the GUI or CLI *Administration Guide*.

Procedure

1. Navigate to the Storage Adapter Configuration menu on the ESXi host.
2. Go to **Hosts and Clusters** > **Configure** > **Storage Adapters**.
3. Select a host.
4. Scroll down and select the iSCSI Software Adapter.
5. Under Adapter Details, click **Properties**.
6. Click **Enabled**.
7. Verify that the iSCSI initiator now has an IQN in the "iSCSI Name" field.
8. Under Adapter Details, click the **Targets** tab.
9. Click **Dynamic Discovery**.
10. Click **Add**.
11. Type each discovery IP address of each data subnet and click **OK**.
12. If the array has more than one discovery target, repeat steps 9 and 10 until all the targets have been added.

What to do next

Next, go to [Bind VMK Ports to ESXi iSCSI Software Adapter](#) on page 53.

Bind VMK Ports to ESXi iSCSI Software Adapter

Binding the VMK ports to the ESXi iSCSI software adapter ensures the specified vmk and vmnic ports are used for I/O.

You use port binding when all VMK ports for iSCSI reside in the same broadcast domain and IP subnet.

If the VMK ports that are used for iSCSI are in a different broadcast domain and IP subnet, do not use port binding. For more information, see the KB article [Considerations for using software iSCSI port binding in ESX/ESXi \(2038869\)](#).

Procedure

1. Go to **Hosts and Clusters** > **Configure** > **Storage Adapters**.
2. Scroll down and select the iSCSI Software Adapter.
3. Under Adapter Details, click the **Network Port Binding** tab.
4. Click **Add** (the green plus sign).
5. Click the check box next to each vmk that needs to be bound and click **OK**.



Note: If the vmk that you want to bind is not in the list, the failover might not be set properly. See [Configure ESXi iSCSI Networking with a Single vSwitch](#) on page 52.

6. Verify that the iSCSI initiator now has the vmk adapters listed in the "Network Port Binding" section.

Use HPE Storage Connection Manager to Set iSCSI Path Selection Policy

After you set up your network configuration for iSCSI, you can use HPE Storage Connection Manager for VMware to create the optimal number of iSCSI sessions for each storage volume and to manage the path selection for the volumes.

HPE Storage Connection Manager handles these tasks automatically.

Note: HPE Storage Connection Manager installs a Path Selection Policy (PSP) named **NIMBLE_PSP_DIRECTED**. The first time you reboot the ESXi host, HPE Storage Connection Manager sets this PSP on each LUN on the storage array. It then sets a rule to set this PSP each time a new device is discovered. If you do not install HPE Storage Connection Manager, you need to set the PSP on each LUN to **Round Robin (VMware)**.

iSCSI Initiator Groups and the Array

The storage array uses iSCSI initiators and iSCSI targets as a method of communication. iSCSI initiator groups limit access to iSCSI target volumes and snapshots (data providers) on the array.

An iSCSI initiator group has three primary functions:

- To allow initiator access to volumes and snapshots
- To use IP-based ACLs to limit initiator access to volumes
- To deny initiator access to volumes

An iSCSI initiator group consists of one or more client iSCSI initiators on the host system and either all configured subnets or selected target subnets.

You can perform the steps needed to work with iSCSI initiator groups from either the VMware vSphere Web Client GUI or the array OS GUI or CLI.

Note: For detailed information about working with iSCSI initiator groups from the array OS, refer to the administration guide for your version of array OS. HPE provides both a *GUI Administration Guide* and a *CLI Administration Guide* for each major version array OS. These guides are available from the Documentation Portal on HPE InfoSight: <https://infosight.hpe.com/resources/documentation>

Create an iSCSI Initiator Group Using the Array OS GUI

You can use the array OS GUI to create an iSCSI initiator group. In a VMware environment, an iSCSI group enables access for specific clients to specific volumes over specific subnets.

Before you begin

You must know the iSCSI Qualified Name (IQN) or IP address of each client.

The client can be a Windows host, an ESXi host, or any another computer that needs to access a volume on your storage array. Configure your client initiator according to the vendor's recommendations.

Procedure

1. From the GUI, select **Manage** > **Data Access**.
2. On the Initiator Groups page, click **(+)**.
3. Type an iSCSI initiator group name in the Name field.



4. Under **Subnets**, choose one of the following:

- Use all configured subnets.
- Use selected subnets.

If you choose **Use selected subnet**, a list of available subnets is displayed.

- a) (Optional) Highlight the subnets you want to associate with this iSCSI initiator group and click Add.
- b) Type a name for the initiator in the Name field.
- c) Copy and paste the IQN of the client system into the IQN field.

Note: If you cannot copy and paste the IQN, type it very carefully. IQN names are case-sensitive.

- d) Add the host iSCSI IP address.

5. (Optional) If you want to add another initiator, click **Add**

6. If you have added all initiators, click **Create**.

Create an iSCSI Initiator Group Using the Array OS CLI

You can use the array OS CLI to create an iSCSI initiator group. In a VMware environment, an iSCSI group enables access for specific clients to specific volumes over specific subnets.

Before you begin

You must know the iSCSI Qualified Name (IQN) or IP address of each client.

The client can be a Windows host, an ESXi host, or any another computer that needs to access a volume on your storage array. Configure your client initiator according to the vendor's recommendations.

Procedure

1. At the command prompt, type:

```
initiatorgrp --create initiator group name
```

Example: **initiatorgrp -- create Dataman**

2. Add an initiator:

```
initiatorgrp --add_initiators initiator group name --label initiator label [--initiator_name IQN | --ipaddr client system IP address]
```

Example: **initiatorgrp --add_initiators Dataman --label client1 --initiator_name iqn.1991-05.com.microsoft:techops.storage.com --ipaddr 192.0.2.88**

Note: If you cannot copy-and-paste the IQN, type it very carefully. IQN names are case-sensitive

3. Add a subnet:

```
initiatorgrp --add_subnets initiator group name --label subnet label
```

Example: **initiatorgrp --add_subnets Dataman --label Subnet-198.51.100.0**

4. Verify your iSCSI initiator group:

```
initiatorgrp --info initiator groupname
```

Example: **initiatorgrp --info Dataman**

You should see a result like this:



```
Name: Dataman
Description:
Access Protocol: iscsi
Created: Feb 19 2014 16:16:32
Last configuration change: Feb 19 2014 16:16:32
Number of Subnets: 1
  Subnet Label: Subnet-198.51.100.0
Number of Initiators: 1
  Initiator Label: Basic
  Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
  Initiator IP Address: 198.0.2.88
```

Example

```
initiatorgrp -- create Dataman
initiatorgrp --add_initiators Dataman --label client1 --initiator_name iqn.1991-05.com.microsoft:techops.storage.com --ipaddr 192.0.2.88
initiatorgrp --add_subnets Dataman --label Subnet-198.51.100.0
initiatorgrp --info Dataman
```

Assign Volumes to an iSCSI Initiator Group Using the Array OS GUI

You can assign an existing volume to an iSCSI initiator group using the array OS GUI.

Note: You can also assign a volume to an iSCSI initiator group when you create the volume.

Before you begin

You must create your iSCSI initiator groups before you can assign volumes to them.

Procedure

1. From the GUI, select **Manage** > **Data Storage**.
2. On the Volumes page, check the volume you want to assign to the iSCSI initiator group.
3. Click **Edit** (the pencil icon).
4. Click **Access** on the progress bar.
5. Click **Add**.
6. In the **Add ACL** dialog box, choose whether you want to limit access to:
 - The volume and its snapshots
 - The volume only
 - The snapshots only
7. From the drop-down menu, choose the iSCSI initiator group you want this volume to use.
8. (Optional) From the drop-down menu, choose a CHAP account.

Note: You do not need to set up a CHAP account. In addition, CHAP is not supported with VMware Virtual Volumes (vVols).

9. Click **Add**.
10. When you have finished assigning your volumes, click **Save**.

Assign Volumes to an iSCSI Initiator Group Using the Array OS CLI

You can assign an existing volume to an iSCSI initiator group using the array OS CLI.

Note: You can also assign a volume to an iSCSI initiator group when you create the volume.

Before you begin

You must create your iSCSI initiator groups before you can assign volumes to them.



Procedure

1. At the command prompt, enter **vol --addacl <volume name> --apply_acl_to <volume | snapshot | both> --initiatorgrp <initiatorgrp-name>**

Example: **vol --addacl volume1 --apply_acl_to both --initiatorgrp Dataman**

2. Verify that the volume was assigned by entering **vol --info <volume name>**

Example: **vol --info volume1**

In the list of volume information, look for *Access Control List*. You should see a result similar to the following:

```
Access Control List:
  Apply to: volume & snapshot
  Initiator Group: Dataman
  CHAP user: *
```

Unassign Volumes from an iSCSI Initiator Group Using the Array OS GUI

You can use array OS GUI to unassign a volume from an iSCSI initiator group.

Before you begin

Unmount the volume from on the host. Do not attempt to remove a volume from an initiator group if it is still mounted on the host.

Procedure

1. From the GUI, select **Manage > Data Storage**.
2. On the Volumes page, check the volume you want to unassign from the iSCSI initiator group.
3. Click **Edit** (the pencil icon).
4. Click Access on the progress bar.
5. In the *iSCSI Initiator Group* list, find the iSCSI initiator group you want to unassign, and click the X icon.
6. Click **OK** to confirm.
7. When you have finished unassigning your volumes, click **Save**.

Unassign Volumes from an iSCSI Initiator Group Using the Array OS CLI**Procedure**

You can unassign a volume from an iSCSI initiator group using the array OS CLI.

1. At the command prompt, enter **vol --removeacl <volume name> --apply_acl_to <volume | snapshot | both> --initiatorgrp <initiatorgrp-name>**.

Example: **vol --removeacl volume1 --apply_acl_to both --initiatorgrp Dataman**

2. Verify that the volume was assigned by entering: **vol --info <volume name>**

Example: **vol --info volume1**

In the list of volume information, look for *Access Control List*. You should see a result similar to the following:

```
Access Control List:
  Apply to: *
  Initiator Group: *
  CHAP user: *
```

Edit an iSCSI Initiator Group Using the Array OS GUI

You can use the array OS GUI to edit an iSCSI initiator group. This feature allows you to:

- Add client iSCSI initiators
- Remove client iSCSI initiators
- Add subnets
- Remove subnets



Note: If you have VLANs configured, it is recommended that you configure initiator groups with the correct subnet association.

Procedure

1. From the GUI, select **Manage** > **Data Access**.
2. Select the iSCSI group you want to modify and click **Edit** (the pencil icon).
The Edit Initiator Group dialog box opens.
3. (Optional) Under Target subnets, choose one of the following options:
 - Use all configured subnets
 - Select target subnets
 If you chose **Select target subnets**, a list of available subnets appears.
4. (Optional) Under Available subnets, highlight the subnets you want to associate with this iSCSI initiator group and click **Add**.
5. (Optional) Under Available Subnets, highlight the subnets you want to disassociate with this iSCSI initiator group and click **Remove**.
6. (Optional) To add an iSCSI initiator, under Initiators:
 - a) Click **Add**.
 - b) Type a name for the initiator in the Name field.
 - c) Copy and paste the IQN of the client system into the IQN field.

Note: If you cannot copy and paste the IQN, type it very carefully. IQN names are case-sensitive.

- d) Enter the IP address of the client system in the IP Address field.
7. (Optional) To remove an iSCSI initiator:
 - a) Find the initiator you want to remove.
 - b) Click the X beside the initiator.
8. If you edit or remove an initiator (change of IQN or IP address), manually break the existing connections.
9. If you add a subnet, rescan the ESXi host or refresh Windows to connect to the new target subnets.
10. If you remove a subnet, manually break the existing connections.
11. Click **Save**.

Edit an iSCSI Initiator Group Using the Array OS CLI

You can use the array OS CLI to edit an iSCSI initiator group.

Note: If you have VLANs configured, it is recommended that you configure initiator groups with the correct subnet association.

Procedure

1. Enter the CLI command line for the task you want to perform:
 - Add an initiator to an initiator group:


```
initiatorgrp --add_initiators initiatorgrp_name --label label --initiator_nameiqn --ipaddr client_ipaddr
```

Note: If you cannot copy and paste the IQN, type it very carefully. IQNs are case-sensitive.

- Remove an initiator from an initiator group:


```
initiatorgrp --remove_initiator initiatorgrp_name --label label
```



- Add a subnet to an initiator group:

```
initiatorgrp --add_subnets initiatorgrp_name --label subnet_label
```

- Remove a subnet from an initiator group:

```
initiatorgrp --remove_subnet initiatorgrp_name --label subnet_label
```

2. If you edit or remove an initiator (change of IQN or IP address), manually break the existing connections.
3. If you add a subnet, rescan the ESXi host or refresh Windows to connect to the new target subnets.
4. If you remove a subnet, manually break the existing connections.
5. (Optional) Verify your changes.

```
initiatorgrp --info initiatorgrp_name
```

You should see a result like this:

```
Name: Dataman
Description:
Access Protocol: iscsi
Created: Feb 19 2014 16:16:32
Last configuration change: Feb 19 2014 16:16:32
Number of Subnets: 1
    Subnet Label: Subnet-198.51.100.0
Number of Initiators: 1
    Initiator Label: Basic
        Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
        Initiator IP Address: 198.0.2.88
```

Example

Adding an initiator.

```
$ initiatorgrp --add_initiators Dataman --label client1 --initiator_name
iqn.1991-05.com.microsoft:techops.storage.com --ipaddr 192.0.2.88
```

Removing an initiator.

```
$ initiatorgrp --remove_initiator Dataman --label client1
```

Adding a subnet.

```
$ initiatorgrp --add_subnets Dataman --label Subnet-198.51.100.0
```

Removing a subnet.

```
$ initiatorgrp --remove_subnets Dataman --label Subnet-198.51.100.0
```

Verifying the results.

```
$ initiatorgrp --info Dataman
```

Delete an iSCSI Initiator Group Using the Array OS GUI

You can use the array OS GUI to delete an iSCSI initiator group.

Procedure

1. Choose **Manage** > **Data Access**.
2. On the Initiator Groups page, check the box to the left of the iSCSI group you want and click **Delete**.
3. Click **OK** to confirm.
4. Manually break the existing connections with the client.



Delete an iSCSI Initiator Group Using the Array OS CLI

Procedure

You can use the array OS CLI to delete an iSCSI initiator group:

1. At the command prompt, enter **initiatorgrp --delete <initiator group name>**.

Example: **initiatorgrp --delete Dataman**

2. Verify that the initiator group was deleted by entering **initiatorgrp --list**.
3. Manually break the existing connections with the client.

Use a CHAP Account

You have the option of using a CHAP (Challenge-Handshake Authentication Protocol) account when you use the iSCSI protocol.

Note: CHAP is not supported with VMware Virtual Volumes (vVols).

Before you begin

You must set the Authentication Method to Use Unidirection CHAP.

If you want to use CHAP with the Group Scoped Target (GST) feature, which is supported in array OS release 5.1.x and later, keep the following in mind:

- CHAP authentication for GST is not used for established iSCSI connections. It is only used for new iSCSI connections or iSCSI connections that are renewed later.
- You must make sure that the CHAP settings are correct on both the storage array and the host. If the credentials do not match, then the current connections stay logged in, but they will fail when sessions are re-established at a later time, for example, after a host reboot, controller reboot, network issues, and so on.
- Provisioning new volumes after enabling CHAP does not cause new iSCSI connections to be created using GST. Instead, the new volumes will re-use the existing connections.

Procedure

1. Set up CHAP on the storage array. For instructions on doing this, see the GUI or CLI *Administration Guide*.
2. Select **Hosts and Clusters** > **Hosts and Clusters** > **Configure** > **Storage Adapters**
3. Select a host.
4. Scroll down and select **iSCSI Software Adapter**.
5. In the section **Adapter Details**, select the **Targets** tab.
6. Select the **Dynamic Discovery** target for the array.
7. Click the **Authentication** button.
8. In the pop-up box, uncheck the box **Inherit settings from parent**.
9. In the **Authentication Method** dropdown list, select **Use Unidirectional CHAP**.
10. Configure the CHAP credentials and select **OK**.
11. If you need to set up additional storage targets, repeat steps 5 through 8.

iSCSI Digest Considerations

You can choose whether to enable iSCSI digest when you work with datastores in VMware environments. The iSCSI digest detects errors that occur at the iSCSI layer. Doing this helps ensure the reliability of data transfers and avoid undetected data corruption.



Note: If you use iSCSI digest, you might want to evaluate what effect it has on your system performance. You need to consider how important the data is, the amount of data involved, and how often the data is accessed.

You can enable or disable iSCSI digest using the vCenter plugin provided by HPE or the ESXi host.

If you are running ESXi 6.5 or later on the host and you use the vCenter Plugin to enable or disable iSCSI digest, the change will appear at the plugin level, but not at the iSCSI adapter level.

For example, you can use the vCenter Plugin to enable iSCSI digest. However, if you then go to the ESXi 6.5 host GUI and select **Advanced Settings** at the iSCSI adapter level, iSCSI digest still appears as disabled.

The reason iSCSI digest is listed as **Prohibited** in the host settings is because the vCenter Plugin change happens at the send_target/discovery IP level only. In other words, the change has been applied to the specific discovered target at the vCenter Plugin level even though the adapter level says the change did not happen. This way, if there are iSCSI sessions going on outside of the vCenter Plugin, they are not affected by this change.

Enable iSCSI Digest

When you work with the vCenter Plugin provided by HPE, you can enable iSCSI digest on a host. Doing this helps protect the integrity of data transferred between the host and a storage array.

Note: If you want to add iSCSI digest parameters to volumes on VMware initiators, see [KB-000296 Enabling iSCSI Digest on VMware initiators](#).

Procedure

1. Log into the vSphere Web Client.
2. From the Home screen, click **Storage**.
3. Right-click datacenter, and select **Nimble Storage Actions** > **Enable iSCSI Digest** to launch the wizard.
4. Select the storage group that contains the host or hosts for which you want to enable iSCSI digest.
Click **Next**. By default, all hosts in the group are selected.
5. Uncheck the boxes until only the hosts for which you want to enable iSCSI digest are selected.
6. Click **Finish**.

The Recent Tasks area displays the progress and status of an activity titled "NIMBLESTORAGE.EnableIscsiDigest.label."

Disable iSCSI Digest

You can use the vCenter Plugin to disable iSCSI digest on a host if you do not need the additional protection that iSCSI digest provides.

Procedure

1. Log in to the vSphere Web Client.
2. From the Home screen, click **Storage**.
3. Right-click the datacenter, and select **Nimble Storage Actions** > **Disable iSCSI Digest** to launch the wizard.
4. Select the storage group that contains the hosts for which you want to disable iSCSI digest.
5. Click **Next**. By default, all hosts in the group are selected. If you do **not** want to disable iSCSI digest for certain hosts, you must de-select them.
6. Click **Finish**.

The Recent Tasks area displays the progress and status of an activity titled "Update Internet SCSI digest properties."



Additional Information for Working with iSCSI

The *Common Tasks and Best Practices* and the *Helpful Information* sections of this guide contain additional information about iSCSI and procedures that might be helpful in your system setup. These include:

- VMware iSCSI best practices, which provides information about working with VM guest machines and managing target subnets. See [iSCSI Best Practices in VMware Environments](#) on page 136.
- Configure iSCSI Discovery. See [Configure iSCSI Discovery](#) on page 143.
- iSCSI Host Connection Methods. See [iSCSI Host Connection Methods](#) on page 144.
- Set the iSCSI Host Connection Method to Manual. See [Set the iSCSI Host Connection Method to Manual](#) on page 147.
- Configure Jumbo Frames. See [Configure Jumbo Frames](#) on page 147.
- Set the host timeout values. See [Host Disk Timeout Values](#) on page 137.



VMware Fibre Channel Configuration

You can use Fibre Channel (FC) adapters to allow the host to display FC storage devices.

If you have FC cards installed, FC is automatically enabled when you run the HPE Storage Setup Manager and set up your array. The Setup Manager is included in the HPE Storage Windows Toolkit.

Note: When you use FC in a VMware environment, it is a good practice to install the HPE Storage Connection Manager for VMware. HPE Storage Connection Manager automatically performs some advanced configurations, selects the optimal path and evenly balances each I/O request. If your array uses multi-array pools, you must install HPE Storage Connection Manager.

You should keep the following points in mind when you work with your FC environment:

- Use single initiator zoning for each WWPN on each VMware host to avoid disruption and improve rescan time.
- You can include multiple WWPNs from a single storage group in a zone and improve rescan time.
- Do not include multiple storage groups or any other third-party storage devices in the same zone.
- Use HPE Storage Connection Manager, which includes the HPE Storage Path Selection Policy (PSP). This automatically sets values for round-robin and adjusts IOPS for optimal performance.
- Each volume that is shared on a FC array must have an individual LUN ID in the ACL mappings. For example, if volume "test1" uses LUN 1 as its ID, it must use that same LUN ID when presented to every host in a cluster. If a second volume is created, it must use a different LUN ID. For example, if volume "test2" is created, it cannot use LUN 1 as its LUN ID, as volume "test1" is already using that ID.

See the *HPE Storage Windows Integration Guide* for information about installing and using the Windows Toolkit.

Fibre Channel Target Limits

If you are running array OS release 4.x or later, the maximum number of Fibre Channel logins to a single port on a storage array is 256.

If you are running an earlier version of array OS, the maximum number of logins is 128.

For more information about array OS system limits, see the *GUI Administration Guide*.



The HPE Nimble Storage vCenter Plugin

The vCenter Plugin provides a web-based client plugin. You can use the plugin to work with vCenter Server to manage datastores residing on HPE storage arrays.

Note: Starting with array OS release 5.1.1.0, the vCenter Plugin no longer supports the Thick Client (also known as the Desktop or C# Client).

The plugin works with both traditional VMFS datastores and VMware Virtual Volumes (vVols). You can perform the following tasks:

- Create, clone, grow, and edit datastores
- Take, clone, and delete snapshots
- Perform vVols tasks, including:
 - Set virtual machine (VM) storage policies
 - Create protection schedules
 - Use replication partners
 - Delete and restore VMs
 - Use Volume Shadow Copy Service (VSS)
- Use role-based access control (RBAC)

Clients Supported in the vCenter Plugin

Array OS release 5.1.x and later supports two types of Web Clients with the vCenter Plugin.

- The vSphere Client (the HTML5 client)
- The vSphere Web Client (the Flash client)



Important: vSphere 7 removed support for the Flash client. It only supports the HTML5 client. You must be running array OS release 5.1.4.200 or later to use vSphere 7.

Both clients support traditional VMFS datastores and VMware Virtual Volumes (vVols).

You can use multiple vCenter Plugins simultaneously.

Note: For the most up-to-date information about support for VMware, refer to the Validated Configuration Matrix. This online tool is available on InfoSight at <https://infosight.hpe.com/resources/nimble/validated-configuration-matrix>.

There are some minor differences in the user interface depending on which Web Client you use (HTML5 vSphere Client or vSphere Web Client). As a result, some of the instructions in this document might not match your GUI.



Set Up the vCenter Plugin

The HPE Nimble Storage vCenter Plugin is part of the array OS software and allows you to use VMware features and view statistics about the virtual machines (VMs) that are hosted on the storage array.

To use the plugin, you must register it with a vCenter Server.

You can use either the array GUI or CLI to register the plugin with the vCenter Server. In addition you can register multiple plugins on the array.

After you register the plugin, it is deployed each time you log into the storage array.

The plugin enables InfoSight to collect VMware configuration data and monitoring statistics for each VM on the array. These analytics provide insights into performance and usage that you can see by logging onto HPE InfoSight (<https://infosight.hpe.com>).

Registration Requirements for Using vCenter Plugin

Registering the vCenter Web Plugin requires the following:

- A supported configuration, including ESXi. See the Validated Configuration Matrix tool on HPE InfoSight for the most current configuration information:
<https://infosight.hpe.com/resources/nimble/validated-configuration-matrix>
- Correct privileges and permissions on both the storage array and the vCenter server. You must have sufficient privileges to install and register plugins, including the vCenter Plugin and the VASA Provider plugin, and to use VMware Synchronized Snapshots. The section *RBAC and the vCenter Plugin* provides information about required privileges.

Note: The vCenter Server supports role-based access control (RBAC). You can set up roles that address specific tasks, such as performing VMware Virtual Volume (vVol) operations in the vCenter. As an administrator, you must make sure that your users have the privileges they need to perform the operations associated with their jobs.

- The vCenter Server hostname or IP address.

LUN Restrictions When Using the vCenter Plugin

The vCenter Plugin is not supported in environments that have the follow LUN configurations:

- Multiple datastores located on one LUN
- One datastore that spans multiple LUNs
- LUNs located on a storage device that is not made by HPE

Register or Add a vCenter Plugin Using the Array OS GUI

You can register the Web Client vCenter Plugin using the array GUI. You can register multiple vCenter plugins on the storage array.

Before you begin

You must meet the requirements listed in [Registration Requirements for Using vCenter Plugin](#) on page 65.

You must be logged in with administrator privileges.



Procedure

1. From the array GUI, go to **Administration > VMware Integration**.
2. If you already have a vCenter Server registered, click **Add Another vCenter** at the bottom of the page to add registration information for another vCenter Server.
3. Enter the registration information:
 - **VCENTER NAME.** This is the string that HPE uses to reference the vCenter Server; for example, `vCenter6`. It is a good practice to make the HPE name the same as the actual vCenter Server name.
 - **SUBNET.** You have the option of specifying a subnet.
 - **VCENTER HOST.** The vCenter host name is the IP address or FQDN of the system (usually a virtual machine) where the vCenter resides. For example, you might enter `esx6vs.testsystem.xyzstorage.com`.
 - **PORT.** Enter port number that is used to communicate with the vCenter. Normally, this is 443.
 - **DESCRIPTION.** This is optional, but it is a good practice to provide a description that lets you quickly identify the vCenter.
 - **CREDENTIALS.** Provide the login information for the vCenter Server.
4. Check the **Web Client** box.
5. Check the **VASA Provider** box if you want to use VMware Virtual Volumes (vVols).
6. Click **Save**.

Register or Add a vCenter Plugin Using the Array OS CLI

You can register the Web Client vCenter Plugin using the array CLI. You can register multiple vCenter plugins on the storage array. You must add the vCenter Server to the array if you want to use VMware Virtual Volumes (vVols).

Based on the CLI command you enter, you can register or add a vCenter Server.

Before you begin

You must meet the requirements listed in [Registration Requirements for Using vCenter Plugin](#) on page 65.

You must be logged in with administrator privileges.

Procedure

1. Log into your array using the array CLI.
2. At the command prompt, enter the following command to register the vCenter:

```
vcenter --register vcenter_name--extension {web | vasa}
```

If you only want to add a vCenter Server to an array, enter this command:

```
vcenter --add [--name] [--hostname {host_name|ip_addr}] [--port_number port_number] [--username user_name]
[--password password] [--description description] [--subnet_label subnet_label]
```

3. Restart the client.

Display a List of Registered Plugins Using the Array OS CLI

You can use the **vcenter** command from the array CLI to display a list of all registered plugins.

Procedure

From the CLI enter

```
vcenter --list
```



Edit the Registered vCenter Using the Array GUI

If you have registered a vCenter Server with the storage array, you can edit the vCenter Server name, description and credentials.

Procedure

1. From the array GUI, go to **Administration** > **VMware Integration**.
2. Click the **Edit** button.
3. Update the information.

Note: If you want to change the username/password for the vCenter Server, click **UPDATE CREDENTIALS**. In the dialog box that appears, you can enter the credentials.

4. Click **Save**.

Unregister the vCenter Plugin

If you no longer want to manage volumes on a storage array through the vSphere client, you can unregister the vCenter Plugin. You must unregister the storage array from all associated vCenters.

If you are using the VASA Provider, you must unregister it before you unregister the plugin.

You can choose which method you want to use to unregister plugin:

- GUI: [Use the GUI to Unregister the vCenter Plugin](#) on page 68
- CLI: [Use the CLI to Unregister the vCenter Plugin](#) on page 68
- vCenter server: [Unregister the vCenter Plugin from the vCenter Server](#) on page 67

Unregister the vCenter Plugin from the vCenter Server

You can unregister the vCenter Plugin from the vCenter Server.

Procedure

1. In the address field of your browser, enter:
https://<vCenter server name or IP>/mob
This action accesses the managed object browser (MOB) on the vCenter server.
2. Click **Content** and then select **ExtensionManager**.
3. Remove the entry (or entries) for the vCenter Plugin.
 - If you are using the vSphere Flash Client, you must remove one entry.

Note: vSphere 7 does not support the Flash client. If you are using vSphere 7, you will not see this entry.

- If you are using the vSphere HTML5 Client, you must remove two entries: one for the vSphere Flash client and one for the vSphere HTML5 Client.

The key for the Flash Client uses the format:

com.nimblestorage.hi

The key for the HTML5 Client uses the format:

com.nimblestorage.hi.h5

Removing the necessary vCenter Plugin entries unregisters all the groups that use the vCenter Plugin.





Important: You cannot remove just one group from the plugin. If you re-register the vCenter Plugin, you must register all the groups again.

4. Choose **UnregisterExtension**.
5. Paste in the name of the plugin and click **Invoke Method**.
6. Close the pop-up window.
7. To verify that the vCenter Plugin is no longer in the list, go to **Managed Object Type > ManagedObjectReference > ExtensionManager** and refresh it.

Use the GUI to Unregister the vCenter Plugin

You can use the array GUI to unregister the vCenter Plugin from the vCenter server.

Procedure

1. From the array GUI main menu, select **Administration > VMware Integration**.
You see a completed registration form for each registered vCenter.
2. Select **Edit**.
3. Uncheck all the plugins.
4. Click **Save**.

Use the CLI to Unregister the vCenter Plugin

You can run the **vcenter** command from the array CLI to unregister the vCenter Plugin from the vCenter server.

Procedure

1. Log into your storage array using the CLI.
2. At the command prompt, enter:
vcenter --unregister vcenter_name --extention {web | vasa }
3. Restart the vSphere client.

Removing Stale Entries Following Update to vSphere 7

Updating the vCenter to vSphere 7 and the storage array to array OS release 5.1.4.200 or later removes the plugin for the vSphere Web Client (the Flash client) and upgrades the plugin for the vSphere Client (the HTML5 client) to one that is compatible with vSphere 7.

In some cases, the update process leaves stale entries for the Flash Client in the vCenter MOB. This might happen if you upgrade the storage array to vSphere 7 before updating the vCenter. You can use one of the following methods to remove the stale entries:

- Unregister the array from the vCenter and then re-register it.
- Perform the procedure listed below to manually remove the Flash Client entry.

Before you begin

You must be running array OS release 5.1.4.200 or later and vCenter 7.

Procedure

1. In the address field of your browser, enter:
https://<vCenter server name or IP>/mob
This action accesses the managed object browser (MOB) on the vCenter server.



2. Click **Content** and then select **ExtensionManager**.
3. Remove the entry for the Flash client of the vCenter Plugin:

com.nimblestorage.hi

Note: Do not remove any other entries.

Removing this entry unregisters all the groups that use the vCenter Plugin.

4. Choose **UnregisterExtension**.
5. Paste in the name of the plugin and click **Invoke Method**.
6. Close the pop-up window.
7. To verify that the vCenter Plugin is no longer in the list, go to **Managed Object Type** > **ManagedObjectReference** > **ExtensionManager** and refresh it.



RBAC and the vCenter Plugin

Role-based access control (RBAC) allows you to create roles that include specific privileges. You can use these roles to protect access to the vCenter Plugin and restrict the actions taken using the plugin.

You can set up multiple roles that address the different tasks users need to perform when working with a storage array using a VMware environment. The privileges you include in a role determine the extent to which a user can access the vCenter Plugin or perform other tasks.

A task requires a certain set of base privileges. When you create a role, you must ensure that the role has sufficient privileges to perform the required tasks. For example, if a user must access the synchronized snapshot feature, the role must include snapshot privileges.

In addition, all roles must have the minimum privileges required by the array.

Note: No default roles are provided.

Privileges Required for Registering the vCenter Plugin and VASA Provider

Registering the vCenter Plugin with the vCenter Server requires specific privileges. If you also want to register the VASA Provider to a vCenter, you must include two additional privileges. The vCenter Server administrator role **administrator@vsphere.local** provides you with all the required privileges.

You can create your own role. In general, a role that allows you to register the plugin with the vCenter also contains the privileges necessary to perform tasks using the plugin, such as working with datastores. However, you should always verify that a role has all the necessary privileges to perform a task.

To create a role or view privileges, select **Administration > Roles** from the VMware vSphere Web Client **Home** page. The privileges are listed under the Roles provider called **Nimble Storage, Inc.** You can view the privileges by clicking the **+** icon at the top of the column.

Note: When you add a custom role and do not assign any privileges to it, the role is created as a Read-Only role and contains three system-defined privileges: **System > Anonymous** (Anonymous), **System > View** (System.View Task.Create), and **System > Read** (System.Read).

The following table shows you where to find the privileges you need when you are working in the VMware vSphere Web Client and want to be able to register the plugin with a vCenter.

Storage Array Privilege	VMware Privilege List
Datstore.AllocateSpace	Datstore > Allocate Space
Datstore.Config	Datstore Cluster > Configure a Datstore Cluster
Datstore.Browse	Datstore > Browse datstore
Datstore.Delete	Datstore > Remove
Datstore.FileManagement	Privilege is on the datstore.
Datstore.Move	Datstore > Move Datstore
Datstore.Rename	Datstore > Rename Datstore
Extension.Register	Extension > Register
Extension.Unregister	Extension > Unregister

Storage Array Privilege	VMware Privilege List
Extension.Update	Extension > Update
Global.CancelTask	Global > Cancel Task
Host.Config.AdvancedConfig	Host > Configuration > Advanced Settings
Host.Config.NetService	Host > Configuration > Security Profile and Firewall
Host.Config.Settings	Host > Configuration > Change Settings
Host.Config.Storage	Host > Configuration > Partition Configuration
ProfileDrivenStorage.update	Profile-driven storage > Profile-driven storage update
ProfileDrivenStorage.view	Profile-driven storage > Profile-driven storage view
StoragePod.Config	Datastore Cluster > Configuration a Datastore Cluster
Task.Create	Task > Create Tasks
Task.Update	Task > Update Tasks
VirtualMachine.Config.AddNewDisk	Virtual machine > Change Configuration > Add new disk
VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
VirtualMachine.GuestOperations.Execute	Virtual machine > Guest operation > Guest operation program execution
VirtualMachine.GuestOperations.Query	Virtual machine > Guest operation > Guest operation queries
VirtualMachine.Inventory.Create	Virtual machine > Edit inventory > Create new
VirtualMachine.Inventory.Delete	Virtual machine > Edit inventory > Remove
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
VirtualMachine.Provisioning.Clone	Virtual machine > Provisioning > Clone virtual machine

The following table shows you where to find the privileges you need when you are working in the VMware vSphere Web Client and want to be able to register the VASA Provider with a vCenter.

Storage Array Privilege	VMware Privilege List
StorageViews.ConfigureService	Storage Views > Configure Service
StorageViews.View	Storage Views > View

Privileges Required for Working with Datastores

If you create custom roles for working with the storage arrays and datastores, you must ensure the roles contain the necessary privileges.

The following table lists the required, base privileges. To create a role or view privileges, select **Administration > Roles** from the VMware vSphere Web Client **Home** page. The privileges are listed under the **Nimble Storage, Inc.** Roles provider. You can view them by clicking the **+** icon at the top of the column.

Note: When you add a custom role and do not assign any privileges to it, the role is created as a Read-Only role and contains three system-defined privileges: **System > Anonymous** (Anonymous), **System > View** (System.View Task.Create), and **System > Read** (System.Read).

Task	Storage Array Privilege	VMware Privilege List
Delete a datastore	Host.Config.Storage	Host > Configuration > Partition Configuration
	Datstore.remove	Datastore > Remove
Create a datastore	Host.Config.Storage	Host > Configuration > Partition Configuration
	Host.Config.NetService	Host > Configuration > Security Profile and Firewall
	StorageViews.ConfigureService	Storage Views > Configure Service
	Datstore.Move	Datastore > Move Datastore
	StorageViews.View	Storage Views > View
	Certificate.Manage	Certificates > Manage Certificates
Clone a datastore	Host.Config.Storage	Host > Configuration > Partition Configuration
	Host.Config.NetService	Host > Configuration > Security Profile and Firewall
	Host.Config.AdvancedConfig	Host > Configuration > Advanced settings
	HostUnresolvedVmfsResignatureSpec.se- tExtentDevicePath	
	Datstore.Rename	Datastore > Rename Datastore
	Datstore.Move	Datastore > Move Datastore
Grow a datastore	Host.Config.Storage	Host > Configuration > Partition Configuration
	System.Read	System > Read

Privileges Required for Working with VMs

You can use the vCenter Plugin to perform tasks specific to virtual machines (VMs). If you create custom roles that use these tasks, you must include the correct privileges.

Note: No storage array privileges are required for the purge operation.

The following table lists the required, base privileges. To create a role or view privileges, select **Administration > Roles** from the VMware vSphere Web Client **Home** page. You can use the vCenter Plugin to perform tasks specific to virtual page. The privileges are listed under the **Nimble Storage, Inc.** Roles provider. You can view them by clicking the **+** icon at the top of the column.

Note: When you add a custom role and do not assign any privileges to it, the role is created as a Read-Only role and contains three system-defined privileges: **System > Anonymous** (Anonymous), **System > View** (System.View Task.Create), and **System > Read** (System.Read).

Task	Storage Array Privilege	VMware Privilege List
Restore from a snapshot	VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
	VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
Clone to a new VM	Host.Config.Storage	Host > Configuration > Partition Configuration
	VirtualMachine.Inventory.Create	Virtual machine > Edit inventory > Create new
	VirtualMachine.Inventory.Delete	Virtual machine > Edit inventory > Remove
	Datastore.Browse	Datastore > Browse datastore
	VirtualMachine.Provisioning.Clone	Virtual machine > Provisioning > Clone virtual machine
	VirtualMachine.Config.AddNewDisk	Virtual machine > Change Configuration > Add new disk
	VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
	ProfileDrivenStorage.update	Profile-driven storage > Profile-driven storage update
	ProfileDrivenStorage.view	Profile-driven storage > Profile-driven storage view
Attach	VirtualMachine.Config.AddNewDisk	Virtual machine > Change Configuration > Add new disk
	VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
	ProfileDrivenStorage.update	Profile-driven storage > Profile-driven storage update
	ProfileDrivenStorage.view	Profile-driven storage > Profile-driven storage view
Delete	VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
	VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
	VirtualMachine.Inventory.Delete	Virtual machine > Edit Inventory > Remove



Task	Storage Array Privilege	VMware Privilege List
Claim	VirtualMachine.Inventory.Create	Virtual machine > Edit Inventory > Create new
	Datastore.Browse	Datastore > Browse datastore
	VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
	Datastore.FileManagement	Privilege is on the datastore.
	Host.Config.Storage	Host > Configuration > Storage partition configuration
	VirtualMachine.Config.AddNewDisk	Virtual machine > Change Configuration > Add new disk
	VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
	ProfileDrivenStorage.update	Profile-driven storage > Profile-driven storage update
ProfileDrivenStorage.view	Profile-driven storage > Profile-driven storage view	
Undelete	Host.Config.Storage	Host > Configuration > Storage partition configuration
	VirtualMachine.Inventory.Create	Virtual machine > Edit inventory > Create new
	VirtualMachine.Inventory.Delete	Virtual machine > Edit inventory > Remove
	Datastore.Browse	Datastore > Browse datastore
	VirtualMachine.Config.AddNewDisk	Virtual machine > Change Configuration > Add new disk
	VirtualMachine.Config.RemoveDisk	Virtual machine > Change Configuration > Remove disk
	ProfileDrivenStorage.update	Profile-driven storage > Profile-driven storage update
	ProfileDrivenStorage.view	Profile-driven storage > Profile-driven storage update
Configure VSS	VirtualMachine.GuestOperations.Execute	Virtual machine > Guest operation > Guest operation program execution
	VirtualMachine.GuestOperations.Query	Virtual machine > Guest operation > Guest operation queries

Privileges Required for Creating VMware Synchronized Snapshots

There are two privileges that you must include in any roles where the users need to work with VMware Synchronized Snapshots. These privileges are in addition to the standard privileges required for the role.

Note: The vCenter Server administrator role **administrator@vsphere.local** contains the privileges you need to create synchronized Snapshots.

The privileges that are specific to creating VMware Synchronized Snapshots are the following:

Storage Array Privilege	VMware Privilege List
VirtualMachine.State.CreateSnapshot	Virtual Machine > Snapshot Management > Create Snapshot
VirtualMachine.State.RemoveSnapshot	Virtual Machine > Snapshot Management > Remove Snapshot

Storage Array Roles and vVols

The storage array role-based access control (RBAC) roles have the following access with respect to VMware Virtual Volumes (vVols):


- Administrator and Power User
 - Can perform all operations - create, read, update, and delete
 - No buttons or menus are hidden
- Operator
 - Can perform create, read, and update operations
 - Can not delete a vCenter instance, so the **Remove** button is hidden on the vCenter page.
- Guest
 - Can view only the vCenter page
 - Cannot perform create, update, or delete operations
 - The **Add new vCenter**, **Edit**, and **Remove** buttons are hidden on the vCenter page.

Note: A **vcenter** command allows you to manage VMware vCenter extensions and add, edit, and delete a vCenter appserver. For more information, refer to the *Command Reference*.

Create RBAC Roles

You can use role-based access control (RBAC) to set up multiple roles that work with the vCenter Server. You create roles from the VMware Web Client GUI.

Procedure

1. Log on to the VMware vSphere Web Client.
2. From **Home**, select **Administration** in the left navigation pane.
Alternatively, in the Home tab, under the Administration section, click **Roles**.
3. From the Roles view, click the  icon to add a user role.
4. Enter a name for the new role.
Select the privileges to be associated with the role.
5. Click **OK**.
The new role appears in the list of roles. Click Privileges to view the role's associated privileges.



Using the vCenter Plugin with VMFS Datastores

You can use the HPE vCenter Plugin to perform create, mount, clone, edit, and delete operations on VMFS datastores. From the plugin, you can also expand a VMFS datastore and take a snapshot of a VMFS datastore.

Note: VMFS datastores are also referred to as traditional datastores.

When you work with a VMFS datastore, you should keep the following points in mind:

- The datastore is a formatted disk (LUN) using VMFS.
- The datastore is protocol aware.
- As administrator, you must be aware of the available storage when you grow a datastore.

Note: It is a good practice to have an access control list (ACL), such as an iSCSI initiator group and/or a CHAP username, on all the volumes when you use an array with traditional datastores.

Create a VMFS Datastore from the vCenter Plugin

You can use the HPE vCenter Plugin to create VMFS datastores that are mapped to volumes on an array. The vCenter Plugin always uses the latest VMFS version available to provision a datastore.

Note: You might need to use the scroll bar on the right side of the wizard dialog to see the entire dialog box and to get to the **Back**, **Next**, **Create**, or **Cancel** buttons.

Procedure

1. From the Home screen of the VMware vSphere Web Client, click **Hosts and Clusters**.
2. Right-click the datacenter where you want to create the datastore and select **HPE Alletra 6000 and HPE Nimble Actions > Create VMFS Datastore**.

The datastore wizard launches.

3. Select the Storage group where you want to create the datastore.

The datastore must be online. You must have the group-specific privilege enabled to be able to view the datastores that belong to that group.

Note: A single instance of the vCenter Plugin supports multiple groups.

4. Enter the general information about the new datastore:

- Provide a name for the datastore
- (Optional) Provide a description. It is a good practice to enter a description that will help you quickly identify the datastore.
- If Data Encryption is enabled on the array, you can select it here to enable it on the new datastore. Important points about using Data Encryption:
 - When Data Encryption is enabled on the group or volume, only the administrator can disable it from the array or encrypt the volume.
 - When Data Encryption is enabled on the volume, users with permissions can choose to enable or disable the option during datastore creation.
- Select the protocol from the drop-down list. If there is only one protocol enabled on the array, then the drop-down list is disabled. If you are using CHAP, you must have it enabled on both the EXSi host and the array.



- Specify the hosts:
 - Hosts can be inside or outside of a cluster, but they must be in the same datacenter.

Note: HPE recommends using hosts that are clustered. Having shared storage on devices that are clustered together is a best practice.

- If a host is within a cluster, you must select the cluster. This means that all hosts in the cluster will be granted access.
- You cannot select offline hosts.
- If the group supports Fibre Channel, only Fibre Channel hosts are listed.

5. Click **Next**.

6. Enter the size and space reserve information:

- Use the drop-down list to specify whether the size is MiB, Gib, or TiB. The size can range from 1.5 GiB to 64 TiB.
- Use the drop-down list to specify a pool or folder where the new datastore will be located.
- If Deduplication is enabled on the pool or folder you selected and if you select Thinly Provisioned, then you can check this box. Otherwise, this checkbox is disabled.
- Specify the percentage of the datastore that is reserved. This is the maximum amount of space that the datastore can consume before it is taken offline or marked unwritable. You can choose one of the following :
 - Thinly Provisioned. You can specify 0-99. The system will only use the amount of reserve you specify, even if that is less than the datastore size you specified above. The system can allocate the difference between what you specify for the reserve and the datastore size elsewhere.
 - Thickly Provisioned. You must specify 100. The system allocates the entire amount of the reserve to the datastore..

Note: You cannot specify deduplication when you select Thick Provisioned.

7. Click **Next**.

8. Select the protection options and synchronization options that you want to use with this datastore:

- No volume collection. There is no protection schedule set up for the datastore.
- Join volume collection. Use the Search option or the drop-down list to select an existing volume collection. When you select a volume collection, the wizard displays its protection schedule.
- Create new volume collection. The dialog expands to allow you to create a volume collection and a schedule for it. You must perform the following tasks. You might need to use the scroll bar to see all the options.
 - Enter a name for the volume collection. You can then use that volume collection with another datastore, if you choose.
 - Start from Protection Template. Decide whether to use one of the default protection templates as the basis for the volume collection you are creating. Otherwise, select No protection template.
 - Synchronization Service. From the drop-down list, select the application that provides the synchronization. If you select VMware vCenter, you must provide:
 - A host IP address or a hostname
 - The port number to be used for communication. The default is 443.
 - A user name and password for the host
 - Schedule name. Provide a name you can use to identify the schedule. It is a good practice to include in the name a summary of the schedule; for example, Retain-30Daily indicates that backups are made daily and are retained for 30 days.
 - Local information. Specify backup schedule information, including how often the backup is taken, during which time period it is taken, on which days it is taken, and how long snapshots are retained. For example, you might specify that a snapshot be taken every hour between 8 a.m. and 10 p.m. daily and then retained for 48 hours.
 - Remote information. Specify the Replication Type and Replication partner.

If you specify the Replication Type as **Periodic Snapshot**, you must provide a replication partner. When you use replication (downstream) partners, the vCenter Plugin takes snapshots of the upstream datastore and updates the downstream partner with that information based on the schedule you specify.

You can associate two downstream replication partners with a volume collection when you specify **Periodic Snapshot** as the replication type. You can only specify one partner for one schedule, but you can switch between two downstream partners as you set up the schedules for the volume collection. These partners can be either two on-site partners or one on-site partner and one HPE Cloud Volume partner.

You can set up multiple schedules by selecting Add Schedule. A volume collection can have a maximum of 10 schedules.

- Protect as standalone volume. The dialog expands to allow you to create a volume collection and a schedule for that is specific to that datastore. You do not need to provide a name for this schedule; however, you need to supply the other information that you would supply if you had selected Create a new volume collection.

See the GUI or CLI Administration Guide for more information about protection definitions. The Administration Guides are on the HPE InfoSight Documentation Portal (<https://infosight.hpe.com/resources/nimble/docs>).

9. Click **Next**.

10. Set **IOPS Limit** and **MB/s Limit**. You can select either No Limit or Set Limit, which allows you to enter a value for that option.

11. Click **Next**.

12. View the settings summary and click **Finish**.



CAUTION: Do not begin using the new datastore until the task finishes.

A new task appears in the Recent Tasks list to show you that the datastore is being created. After completion, the new datastore appears in the datastore list for the datacenter.

Mount an Existing VMFS Datastore Using the vCenter Plugin

You can use the HPE vCenter Plugin to mount existing VMFS datastores on new ESXi hosts. There is no disruption of I/O on the datastores being mounted during this operation. The mount operation works the same way on both iSCSI and Fibre Channel arrays.

Note: You cannot mount a datastore that is located on a different datacenter. To mount a datastore to an ESXi host, it must be mounted on at least one other ESXi host in the current datacenter.

Procedure

1. Log on to the vSphere Web Client.
2. From the Home screen, click **Hosts and Clusters**.
3. Select either the datacenter, host cluster, or ESXi host under which you want to mount the datastore.

The hosts on which you can mount the datastore are determined based on how the Datastore Wizard is launched in Step 4.

- If you right-click on a datacenter to launch the wizard, the datastores you select will be mounted on all hosts in that datacenter.
 - If you right-click on a host cluster to launch the wizard, the datastores you select will be mounted on all ESXi hosts in that cluster.
 - If you right-click on a ESXi host to launch the wizard, the datastores you select will be mounted on only that ESXi host.
4. Right-click on the datacenter, host cluster, or ESXi host and select **Nimble Storage Actions > Mount Nimble Datastore(s)** to launch the datastore wizard.

Note: You might need to use the scroll bar to see all the information in the wizard dialog.

5. Select the Storage group on which you want to mount the datastore and click **Next**.

You can select only one group at a time. The group must support mounting, or it cannot be selected.

Note: You must have the group-specific privilege enabled to be able to view the datastores that belong to that group.

6. Select one or more datastores to be mounted.

The plugin only displays datastores that are visible to the vCenter in the list. It does not include offline volumes and datastores mounted in other datacenters.

7. View the settings summary and click **Save**.

After the wizard finishes, the following takes place:

- Independent datastore tasks are created for each ESXi host.
- Each task mounts the selected datastores on the hosts.
- Each task creates an initiator group for each host and adds it to the volume. Existing initiator groups are not edited.
- Mounting issues are reported on a per-host basis.

Note: Because a task is created for each ESXi host involved, you cannot start another mount task on that host until the current task completes.

Clone a VMFS Datastore from the vCenter Plugin

You can use the HPE vCenter Plugin to clone VMFS VMware datastores that reside on an array. Clones are created from snapshots.

Procedure

1. Start the vSphere Web Client and navigate to the Datastores list view.
2. Right-click the datastore you want to clone and select **Nimble Storage Actions > Clone Datastore**.
3. Verify that the correct snapshot is selected for the datastore.
4. Enter a name for the clone.
5. Select the number of clones that you want to create.
6. (Optional) Enter a description.
7. For multiple clones, enter the **Starting number** for clones.

The Starting Number selector only appears when the number of clones is greater than one. When creating multiple clones, all clones created at this time have the same name with a number appended. That number is incremented based on the starting number. For example, if you create three clones and set the starting number to 7, the clones will be named CloneA7, CloneA8, and CloneA9.

8. Choose whether you want to use an existing snapshot or take a new one.

If you choose to use an existing snapshot, the wizard displays a list of existing snapshots.

If you choose to take a new one, enter the name for the new snapshot.

9. Click **Clone**.

The Active Tasks list shows an activity titled NimbleStorage.CloneDatastore.label and all of the subtasks.

Results

After the datastore is successfully cloned, it is added to all hosts where the original datastore is mounted.



Grow a VMFS Datastore from the vCenter Plugin

You can use the HPE vCenter Plugin to grow or resize a VMFS datastore. You must have the correct permissions to perform this task.



Important: It is a best practice to use the plugin to grow an HPE datastore. This way you do not select the wrong device during a grow operation.

Procedure

1. Start the vSphere Web client and navigate to the datastores list view.
2. Right-click the datastore you want to grow and select **Nimble Storage Actions > Grow Datastore**.
The wizard displays the current datastore size. Type in the new size and select the unit type.
3. Increase the size of the datastore to activate the **Grow Datastore** button.
The minimum growth size is 1GB. You cannot use this procedure to shrink a datastore.
4. Click **Grow Datastore**.
The Active/Recent Tasks area shows the in-progress growth of the task and any subtasks.

Take a Snapshot of an HPE Nimble Storage-Based Datastore Using the vCenter Plugin

You can use the HPE Nimble Storage vCenter Plugin to take snapshots of VMware datastores that are mapped to volumes on an HPE Nimble Storage array.

Procedure

1. Open the vCenter Web Client and bring up Storage View.
2. From the list of datastores, select the datastore to take a snapshot and click the **Configure** tab. Then either select Snapshots and click the snapshots icon, or right-click the HPE Nimble Storage datastore for which you want to create a snapshot and select **Nimble Storage Actions > Snapshot Datastore**.
3. Provide a name and optional description for the snapshot and click **Take Snapshot**.
Each new snapshot appears in the list under Snapshots within the Configure tab.

Name	Time	Origin	Schedule	New Data	Compression	Writable	Online
SnapshotA	03/31 02:43 PM	group-c8-array3		0 B		No	No
SnapshotB	03/31 02:42 PM	group-c8-array3		0 B		No	No
sample-schedule-	03/31 02:00 PM	group-c8-array3	schedule-1	0 B		No	No
sample-schedule-	03/31 02:00 PM	group-c8-array3	schedule-2	0 B		No	No

Edit a Datastore Mapped to an Array Using the vCenter Plugin

You can use the HPE vCenter Plugin to edit VMFS VMware datastores that are mapped to volumes on an array.



Procedure

1. From the Home screen of the VMware vSphere Web Client, navigate to the datastores and datastore clusters.
2. Expand the datacenter object.
3. Click the HPE datastore that you want to edit.
4. Right-click the datastore, and select **HPE Alletra 6000 and HPE Nimble Storage Actions > Edit Datastore**. The Edit Datastore wizard opens to the Size page.
5. Click in the text box under DATASTORE or SNAPSHOT to change reserve space or usage warning threshold for the datastore itself or snapshots taken from the datastore. Then click **Next**.
6. The Protection page displays information about the current schedules, Replication Type, and Replication Partner for the associated volume collection. You cannot edit this information, select whether to create (and name) or join (and choose) an existing volume collection, or change protection to use as a Standalone volume. Then click **Next**.

You have the following options on the Protection page:

- No volume collection. There is no protection schedule set up for the datastore.
- Join volume collection. Use the Search option or the drop-down list to select an existing volume collection. When you select a volume collection, the wizard displays its protection schedule.
- Create new volume collection and set up schedules for it. When you select this option, the dialog expands to allow you to perform the necessary tasks. You might need to use the scroll bar to see all the options.
 - Enter a name for the volume collection. You can then use that volume collection with another datastore, if you choose.
 - Start from Protection Template. Decide whether to use one of the default protection templates as the basis for the volume collection you are creating. Otherwise, select No protection template.
 - Synchronization Service. From the drop-down list, select the application that provides the synchronization. If you select VMware vCenter, you must provide:
 - A host IP address or a hostname
 - The port number to be used for communication. The default is 443.
 - A user name and password for the host
 - Schedule name. Provide a name you can use to identify the schedule. It is a good practice to include in the name a summary of the schedule; for example, Retain-30Daily indicates that backups are made daily and are retained for 30 days.
 - Local information. Specify backup schedule information, including how often the backup is taken, during which time period it is taken, on which days it is taken, and how long snapshots are retained. For example, you might specify that a snapshot be taken every hour between 8 a.m. and 10 p.m. daily and then retained for 48 hours.
 - Remote information. Specify the Replication Type and Replication partner.

If you specify the Replication Type as **Periodic Snapshot**, you must provide a replication partner. When you use replication (downstream) partners, the vCenter Plugin takes snapshots of the upstream datastore and updates the downstream partner with that information based on the schedule you specify.

You can associate two downstream replication partners with a volume collection when you specify **Periodic Snapshot** as the replication type. You can only specify one partner for one schedule, but you can switch between two downstream partners as you set up the schedules for the volume collection. These partners can be either two on-site partners or one on-site partner and one HPE Cloud Volume partner.

You can set up multiple schedules by selecting Add Schedule. A volume collection can have a maximum of 10 schedules.

- Protect as standalone volume. The dialog expands to allow you to create a volume collection and a schedule for that is specific to that datastore. You do not need to provide a name for this schedule; however, you need to supply the other information that you would supply if you had selected Create a new volume collection.
7. View, edit, or add schedules in the Scheduler. Then click **Next**.



Click the **Calendar view** or **List view** to change how you view existing schedules. Optionally, use the drop-down to select a Schedule Template. To edit (or delete) a schedule template, hover over the schedule details box to activate the icons.

8. If desired, set limits for IOPS or MB/s. This option is available starting with NimbleOS 5.0.x.
9. View a summary of your changes.
Click **Back** to make any further changes, or click **Finish**.

View VMFS Datastore Details from the vCenter Plugin

The View Details page allows you to quickly and easily view details about an HPE VMFS datastore.

Note: You must have the View Details privilege to perform this task. This privilege must be explicitly granted. By default, the View Details privilege is not set on a datastore.

Procedure

1. Start the vSphere Web Client.
2. Click the datastore whose details you want to view.

By default, the details page opens to the Summary tab.

3. View the volume information.

There are four HPE Storage-related sections within the Summary tab:

- Volume Information - lists the names of the Storage group, volume(s), pool, and gives a performance profile, identifier(s) and data encryption status.
- Volume Space Information - provides size, usage, reserve space of the underlying volume (which is the size and the overhead), quota, and compression data for both volumes and snapshots.
- Protection Policy - lists the volume collection name, synchronization data (plus access credentials), and any snapshot schedules created for the datastore.
- Volume Access Control - refers to access control of volumes and their snapshots:
 - The Volume tab refers to the HPE volume, and displays the CHAP User Name for iSCSI arrays and LUNs for Fibre Channel arrays:

4. Click tabs to view more information.

Tabs include Summary, Monitor, Manage, Related Objects. Within the Monitor tab, click on an action to view Issues, Performance, Tasks, Events, Connections, and Replication. (Connections action shown):

For Fibre Channel, the Connections tab displays the Initiator WWPN/Alias and Target WWPN/Alias for each connection.

To perform other actions on the datastore, click **Actions** at the top of the tab bar and select an Action.

Delete a Datastore Mapped to a Volume Using the vCenter Plugin

You can use the HPE vCenter Plugin to delete traditional VMware datastores that are mapped to volumes on an array.



CAUTION: Before you delete a datastore, make sure that there is no I/O traffic or active connection to the datastore.

Procedure

1. From the Home screen of the VMware vSphere Web Client, navigate to the datastores and datastore clusters.
2. Expand the datacenter object.
3. From the list of datastores available in the datacenter, click the datastore that you want to delete.



4. Right-click the datastore, and select **Nimble Storage Actions > Delete Datastore**.
5. Confirm that you want to delete the datastore. Then click **Delete Datastore**.



Working with VMware Virtual Volumes

VMware Virtual Volumes (vVols) simplify storage management. You do not need to know the implementation details of the underlying storage. Instead, you manage storage at the virtual machine (VM) level. vVols are essentially VMware virtual disks that you can map to a volume.

You must have a VASA Provider registered with a vCenter Server to use vVols. The vVol volumes reside in a vVol storage container, also known as a vVol datastore. You can use the VMware vSphere Web Client to manage the storage.

vVols support all VM workflows, including clone, restore, and delete operations, setting up replication partners, creating HA/DRS environments, taking synchronized snapshots, and using e Volume Shadow Copy Service (VSS) integration for VMware.

Note: When a storage vmotion is performed between vVol datastores, a volume copy is performed as opposed to a volume move from source to destination folders. This operation also deletes the previous volumes, which removes all of the snapshots for the volumes.

You can use the VMware Virtualization section of InfoSight to view analytics about vVols.

When you work with vVols, some points to keep in mind include:

- A vVol is not tied to the underlying storage, but is treated as a virtual storage system.
- vVols can grow, but do not shrink.
- To use vVols, you must have
 - vCenter Server 6.0 or later
 - ESXi 6.0 or later
 - VASA Provider registered with the vCenter Server

Note: The array OS provides a VASA Provider. The VASA Provider creates an automated initiator group for each host in the environment that the vVol datastore is mounted to.

- An HBA that supports a secondary LUN ID
- TCP port 8443 must be enabled. Certification communications use this port.
- The CA certification must be imported to the vCenter for arrays that use custom certifications

How VASA Provider Works with vVols

VASA Provider implements APIs to manage VMware Virtual Volumes (vVols) and to apply storage profiles defined in the vCenter Server. The array OS contains a VASA Provider that supports vVols with the storage arrays.

Before you can use the VASA Provider, you must register it with VMware server.

Note: After upgrading the VASA Provider, the VASA version displayed in vCenter will not update automatically. Upgrading the VASA version requires the administrator to unregister and reregister the VASA provider for vCenter's SMS service to update the field. This is a VMware limitation.

Note: VASA relies on several fundamental services in your environment including the following services:

- vCenter Server
- Directory services (for example, Microsoft Active Directory, LDAP)
- DNS



If these services are stored on a vVol, it is possible to enter a scenario where dependencies on one or more of these services will prevent management of virtual machines on a vVol datastore. As such, these services should not be stored on a vVol.

Some of the advantages of VASA Provider include the ability to:

- Use Storage Policy Based Management (SPBM) to specify the storage requirements for an application. Using SPBM ensures that the vVols associated with an application have the requirements that the application needs.
- Initiator groups are automatically created for each host in the environment where the vVols datastore is mounted.
- Perform periodic checks to ensure that the vVols remain compliant with the application needs.

Note: If a failover occurs when you are running VASA 3.0 and your failover site has multiple storage profiles in a VMware Replication Group, some VMs are incorrectly marked as being out of compliance with the current storage profile. This is only seen when there are multiple storage profiles. It occurs because VMware checks for storage profile compliance as each VM is brought back online during the failure instead of waiting until all vVols are online and have been assigned their storage profiles.

- (VASA 3.0 and later) Provide support for disaster recovery workflows.

The disaster recovery feature included in VASA 3.0 supports vVol-based disaster recovery, not datastore-based recovery. Site Recovery Manager (SRM) and the Storage Replication Adapter (SRA) continue to support array-based recovery. See [SRM and Storage Integration](#) on page 112.

The disaster recovery feature requires that the vVols on both the source site and the target site use the same protocol. An administrator can set up a Replication Group containing specific vVols. The Replication Group, which is equivalent to a volume collection, represents the minimum unit of failover.

Each Replication Group resides within a Fault Domain. A Fault Domain specifies the set of resources that fail together. Setting Fault Domains can limit a failure because not all resources must fail together.

Note: To determine which release of the array OS support VASA Provider 3.0, see the Validated Configuration Matrix.

Protocol Endpoints

VMware Virtual Volumes (vVols) are created within a storage container and are bound using a data access construct called a protocol endpoint (PE). All data access control is handled by the PE. The array OS automatically manages PEs.

PEs work in both Fibre Channel and iSCSI environments. There can be only one PE per Fibre Channel or iSCSI group. In addition, there can be only one PE per pool. This PE is created by default when the first vVol container (folder) is created in the pool. There can be two PEs in the same pool if it is a multi-protocol environment where both a Fibre Channel (FC) folder and an iSCSI folder are created.

There will also be two PEs of the same protocol type in a pool when a pool merge occurs. One PE is marked as deprecated and is deleted when the last vVol using it is unbound.

The PE of an HPE Storage pool is shared by all vCenters and ESXi hosts that need access to the containers in that pool.

A PE LUN is protected by an ACL that makes it visible by iSCSI discovery and login, and FC LUN masking to the ESXi hosts associated with the vCenter where you registered the VASA Provider. The VASA Provider creates the PE ACL entry when an ESXi host initiates communication with it. The vCenter administrator must manually add the iSCSI discovery IP to each ESXi host.

A PE is deleted when the last vVol container in the physical pool is deleted.

Support for vVols Space Reclamation Using UNMAP

VMware Virtual Volume (vVol) environments support using the VAAI UNMAP primitive to reclaim space.



VMware vSphere does not automatically use UNMAP. When a Guest Operating System issues the UNMAP command, VMware vSphere sends it directly to the storage array. The array then recognizes the command.

VMware vSphere detects support for the command for the vVol based on the Protocol Endpoint (PE) LUN to which the vVol is bound. VMware vSphere does not enforce any alignment criteria when it executes an UNMAP command.

VMware vSphere handles the UNMAP command the same way it handles the VAAI Extended Copy (Xcopy) and Atomic Test and Set Locking (ATS) primitives. It detects these commands automatically and sends them to the array, thus bypassing the host.

For more information about using UNMAP with vVols, see the VMware KB article "[How Virtual Volumes and UNMAP primitive interact? \(2112333\)](#)."

vVols and Stale Bindings

The HPE vCenter Plugin enforces certain time-outs when stale bindings occur between an ESXi host and VMware Virtual Volume (vVol).

If an ESXi host is unable to connect to a vVol for 10 minutes, the binding is marked as stale.

After seven days, all stale bindings that do not have data connections are removed.

Supported Features

vVols and HPE Storage Connection Manager

HPE Storage Connection Manager for vSphere 6.0 or later is required for VMware Virtual Volumes (vVols) on iSCSI arrays. The connection manager supports both vVols and regular storage volumes.

Path Selection Plugin (PSP) is a component of HPE Storage Connection Manager. PSP uses the round-robin policy based on IOPS to switch between paths to the PE for all I/O directed to vVols. This is the same behavior used for regular storage volumes in a single array group.

HPE Storage Connection Manager manages connections to iSCSI Protocol Endpoint (PE) LUNs in the following way:

- If a pool consists of multiple arrays in a group, the HPE Storage Connection Manager PSP chooses the optimal array paths for I/O to each vVol.
- If any paths to member arrays in the pool are not discovered, PSP does an automatic SCSI path rescan.
- For iSCSI, connections to all data interfaces are created automatically by the host itself.
- If an array data interface IP is changed, PSP automatically removes the iSCSI sessions that are connected to old data IP and creates connections with the new data IP.

Note: There are more connections to PE LUNs than to regular volumes. This is because connections to all data IPs are created by the host. NCS does **not** support dynamically adding or removing sessions based on `ncm.conf` settings.

Encryption

Encryption At Rest (EAR) is enabled at the group or volume level.

Windows Toolkit

vVols allow the installation of a Windows Toolkit on a guest OS, using guest-initiated iSCSI.

Storage Policy Based Management

Storage policy based management (SPBM) allows you to create virtual machine (VM) storage policies to define the capabilities assigned to storage. These capabilities can include disk type, encryption, compression, deduplication, snapshot management, and replication. Snapshots that you create by applying SPBM are crash-consistent snapshots.



The features that you can specify with SPBM include:

- QoS
- Performance policies
- Deduplication
- Replication
- Encryption
- Snapshots

Group and Pool Merges

VVols supports the following types of group and pool merges:

- Group merge when neither group has VVols enabled.
- Group merge when one group has VVols enabled and in use. If the group with VVols enabled is the group merged into, the operation is non-disruptive to VMs.
- Group merge when both groups have VVols enabled and in use. This operation is disruptive to VMs, as the VVols move from one VP to another.
- Adding an uninitialized array to a pool that has the VVol feature in use.
- Pool merge when neither pool has VVols in use.
- Pool merge when one of the pools has VVols in use. This operation should be non-disruptive to VMs.
- Pool merge when both pools have VVols in use. This operation should be non-disruptive and will result in two PEs in the pool.

Using VASA Provider to Provide Disaster Recovery for vVols

VASA 3.0 includes a vVol-based disaster recovery feature. It requires an upstream (protected) site and a downstream (recovery) site. You can use the VMware vSphere Web Client to enable this feature.

Before you begin

You must have set up your upstream and downstream sites. Each site must have the following:

- A vCenter Server

Note: Both sites (upstream and downstream) must be running the same version of vCenter Server.

- A VASA Provider that is running VASA 3.0
- A vVol datastore that resides within the vCenter Server

Note: To determine which versions of the array OS support VASA Provider 3.0, see the Validated Configuration Matrix.

Procedure

1. Mount the destination folder as a vVol datastore within the downstream site vCenter Server.



Tip: It is a good practice to use the same name for both the upstream folder and the downstream folder.

2. Click the checkbox to make the association between the upstream and downstream folders.

After you do this, the defined partner will be listed in the storage profile configuration wizard within the VMware vSphere Web Client.

3. Use the Web Client to define your storage profile and configure the VMware replication line of service.

Managing vVols

You use the VMware vSphere Web Client to manage VMware Virtual Volumes (vVols). All virtual machine (VM) workflows are supported. You can perform create, clone, delete, snapshot, migrate, and disaster recovery operations.

You can run certain vVol management tasks using the array OS. You can also run some of these tasks using the vCenter Plugin. For information about performing these tasks, refer to information this guide, in the *Administration Guide* and in the *CLI Administration Guide*.

Note: VMware also provides standard documentation about supported workflow tasks.

These tasks include the following:

- Create
 - Creating storage volumes
- Snapshots
 - Creating storage snapshots
 - Restoring managed snapshots using vCenter Snapshot Manager
 - Restoring unmanaged snapshots after offlining the VM (you can also restore an unmanaged snapshot using the array OS)
 - Creating Volume Shadow Copy Service (VSS) snapshots
- Clone
 - Same container - storage clone
 - Different container - new storage volume followed by ESXi host copying data
- Migrate
 - New storage volume followed by ESXi host copying data
- Delete
 - Independent vVol - volume taken offline and deleted
 - Clone parent - volume taken offline and deleted when the last clone child is deleted
- Swap File Location
 - Change the location of the swap files at a host level
 - Change the location of the swap files at a cluster level

Configuring vVols Using the vCenter Plugin

You can use the HPE vCenter Plugin to create, edit, grow, and delete VMware Virtual Volumes (vVols) datastores. These features are available when you use the HTML5 vSphere Client that is provided by the vCenter Plugin.

You must have a vCenter Server registered to the storage array.

Note: You can also perform these tasks from the array GUI.

As you work with vVols, keep in mind that the VM hosting the vCenter Server must be on a VMFS datastore, not a vVol datastore. Refer to the VMware vCenter installation documentation for more information about working with vCenter Servers.

Create a vVols Datastore Using the vCenter Plugin

You can use the HPE vCenter Plugin to create VMware Virtual Volume (vVol) datastores that are hosted on a storage array. This feature is available when you use the HTML5 vSphere Client that is provided with the vCenter Plugin.



Note: You might need to use the scroll bar on the right side of the wizard dialog to see the entire dialog box and to get to the **Back**, **Next**, **Create**, or **Cancel** buttons.

Procedure

1. From the Home screen of the VMware vSphere Web Client, click **Hosts and Clusters**.
2. Right-click the datacenter where you want to create the datastore and select **Nimble Storage Actions > Create VVol Datastore**.

The datastore wizard launches.

3. Select the Storage group where you want to create the datastore.

The group must be online. You must have the group-specific privilege enabled to be able to view the datastores that belong to that group.

Note: A single instance of the vCenter Plugin supports multiple groups.

4. Enter the general information about the new datastore:

- Enter a name for the datastore.
- (Optional) Provide a description. It is a good practice to enter a description that will help you quickly identify the datastore.
- (Optional) Specify the pool using the drop-down list. If you do not select a pool, then **default** is used.
- Select the datacenter where the new datastore will reside. Use the drop-down list.
- Select the protocol. If the array has both iSCSI and Fibre Channel enabled, you can choose one of these two. If only one protocol is enabled on the array, then the drop-down list is disabled and the array protocol appears by default.
- Specify the hosts:
 - Hosts can be inside or outside of a cluster, but they must be in the same datacenter.
 - If a host is within a cluster, you must select the cluster. This means that all hosts in the cluster will be granted access.
 - The host must be online.

5. Click **Next**

6. Enter the space limit.

7. Click **Next**

8. Set **IOPS Limit** and **MB/s Limit**. If you select Set Limit, you can enter a value for that option.

9. Click **Create**.

Note: If the time on the host and the group don't match, you will get an error message. If that happens, reset the times to match and click **Create** again.

A new task appears in the Recent Tasks list to show you that the datastore is being created. After completion, the new datastore appears in the datastore list for the datacenter.

Edit a vVol Datastore Using the vCenter Plugin

You can use the HPE vCenter Plugin to edit a VMware Virtual Volume (vVol) datastore that is hosted on a storage array. This feature is available when you use the HTML5 vSphere Client that is provided with the vCenter Plugin.

The edit feature lets you modify the performance settings for the datastore. It does not let you modify anything else. If you want to increase the size of the datastore, you must use the grow feature. See [Grow a vVol Datastore Using the vCenter Plugin](#) on page 90.

Note: You might need to use the scroll bar on the right side of the wizard dialog to see the entire dialog box and to get to the **Back**, **Next**, **Save**, or **Cancel** buttons.

Procedure

1. From the Home screen of the VMware vSphere Web Client, click **Hosts and Clusters**.
2. Right-click the vVol datastore that you want to grow and select **Nimble Storage Actions > Edit VVol Datastore**.
The Edit vVol Datastore wizards starts.
3. Click **Next** until you reach the Performance dialog.
4. Make the changes to **IOPS Limit** and/or **MB/s Limit**.
5. Click **Save**.

Grow a vVol Datastore Using the vCenter Plugin

You can use the HPE vCenter Plugin to grow a VMware Virtual Volume (vVol) datastore that is hosted on a storage array. This feature is available when you use the HTML5 vSphere Client that is provided with the vCenter Plugin.

Procedure

1. From the Home screen of the VMware vSphere Web Client, click **Hosts and Clusters**.
2. Right-click the vVol datastore that you want to grow and select **Nimble Storage Actions > Grow VVol Datastore**.
3. In the Grow VVol Datastore dialog, enter the new size for the datastore.
4. Click **Grow**.

Delete a vVol Datastore Using the vCenter Plugin

The HTML5 vSphere Client provided by the HPE vCenter Plugin allows you to delete VMware Virtual Volume (vVol) datastores that are hosted on a storage array.

Procedure

1. Navigate to the Datastores list view.
2. Right-click the vVol datastore you want to delete and select **HPE Nimble Storage Actions > Delete VVol Datastore**.
3. In the **Delete VVol Datastore** dialog, click **Delete**.

Run Configuration Checks Against the vCenter Server, Array

The HTML5 vSphere Client provided by the HPE vCenter Plugin includes a Configuration Checks feature. This feature collects configuration information about the vCenter Server and the storage array. It then checks this information against a set of rules.

Configuration Checks displays a report about how the configuration complies to the rules. Next to each rule, it provides a status. Under the severity column, it displays Major, Warning, Minor, and Information. These severity levels let you know whether the configuration is in compliance with the rule or whether you need to fix your configuration. You should resolve any issue that has a severity level of Major.

By default, the Configuration Check runs once a day. You can manually run it any time you choose. Depending on the size of your vSphere environment, this operation could take a few minutes. You can monitor its progress by viewing the vCenter task list.

Procedure

1. From the VMware Home menu select **HPE Alletra 6000 and HPE Nimble Storage**.
2. Select the storage group that is associated with the configuration you want to check.
The vCenter Plugin organizes information based on storage groups.
3. Select Configuration Checks. The Configuration Checks features appears. Click:
 - **Refresh** to see the last results that were run.
 - **Rerun** to manually run the Configuration Check.



Configuring vVols and VMs from the VMware GUI

You can create VMware Virtual Volumes (vVols) datastores and virtual machines (VMs) from the VMware GUI.

To work with vVols, you must have a vCenter Server registered to the storage array. Refer to the VMware vCenter installation documentation for more information about working with vCenter Servers.



Important: The VM hosting the vCenter Server must be on a VMFS datastore, not a vVol datastore.

Overview of the vVols Workflow

The following is a high-level overview of how to configure VMware Virtual Volumes (vVols) on a storage array. After you register VASA Provider, you must use the VMware vSphere Web Client to perform the tasks.

Before you begin

You should do the following:

-
- Confirm that port 8443 is opened for SSL communication between ESXi hosts and HPE Storage arrays. This port ensures that the ESXi hosts can mount the vVols databases.

Procedure

1. Add a vCenter Server. Include the VASA Provider option.
2. Create a vVol datastore on the storage container using the VMware vSphere Web Client.

Note: (iSCSI only) Configure the discovery IP address of the storage array. If you do not configure this IP, the Protocol Endpoints (PE) will not be visible. As a result, the vVol datastore will not be created.

3. Create a virtual machine using the VMware vSphere Web Client.

Note: If you do not already have one, you will also need to create a VM storage policy using the VMware vSphere Web Client.

Create a vVols Datastore

The Create Datastore wizard lets you create a VMware Virtual Volume (vVol) datastore and map it to a folder on an array. Before you can do this, you must discover the protocol endpoint (PE).

Note: The folders mentioned in this task should already have been created.

Procedure

1. Discover the PE.

Take the following action based on the protocol you are using:

Option

An iSCSI array

Description

Add the discovery IP and rescan.

Note: You can use the `pe target group --info` to display the iSCSI IQN for the PE.

A Fibre Channel array

Rescan.



2. Start the Create Datastore wizard by right clicking on a datacenter and selecting **Storage > New Datastore**.
3. On the Type screen, choose **vVol** and click **Next**.
You see the folders created on the array.
4. Select the folder you want to map to the vVol and click **Next**.
You see a list of accessible hosts.
5. Select the host you want to map to the vVol and click **Next**
6. Review the Summary page and click **Finish** to complete the datastore creation process.

Create a VM

You use the VMware vSphere Web client to create a virtual volume (VM) as well as the storage profile you want to associate with it.

Procedure

1. If you do not already have a storage profile, create one now using the **Create New VM Storage Policy** wizard.

Note: For information about working with storage policies, see *Create a VM Storage Policy*.

- a) From the VMware vSphere Web Client **Navigation** pane, select: **Policies and Profiles > VM Storage Policies**
- b) Click the **Create** icon.
- c) Enter the information to define the storage profile.

Note: If your storage policy includes using replication, choose the replication group where you want to place your VM.

2. Create the VM using the **New Virtual Machine** wizard.
From the VMware vSphere Web Client, right-click a datacenter, cluster, or host and select **New Virtual Machine** to start the wizard.
3. Enter the information required by the wizard. Most of the dialogs in the wizard are self-explanatory. Some key points when you work with the wizard include the following:
 - Enter an alphanumeric name for the VM. Do **not** include any unicode characters.
 - Specify a VM storage profile. Using a profile to define the VMs instead of entering the information manually helps you maintain consistency when you create new VMs.
 - Select a compatible vVol datastore. Doing this ensures that volumes are created for each VMDK.
4. Click **Next**.
5. To complete the VM creation process, click **Finish**.

VSS and vVols

Starting with array OS release 5.0.1.0 and HPE Storage Windows Toolkit 5.0.0, the HPE vCenter Plugin extends Volume Shadow Copy Service (VSS) integration to include VMware Virtual Volumes (vVols) in both iSCSI and Fibre Channel environments. You can use VSS to take application-consistent snapshots for Microsoft Exchange Server and Microsoft SQL Server when the application data is hosted on a vVol. You enable VSS by creating a storage policy that has VSS selected and assigning the policy to a virtual machine (VM).

VSS allows you to back up an application without having to back up an entire VM. This way you can restore only the vVol that contains the application data, not the entire VM.

VSS for vVols requires that you do the following:

- Install the Windows Toolkit 5.0.0 or later on a guest operating system running Windows Server 2008 R2 or later and either Microsoft Exchange Server or Microsoft SQL Server.



- Install VMware Tools on the guest OS.
- Run array OS release 5.0.1.0 or later on the array.
- Register the Web Client vCenter Plugin with the vCenter Server.
- Create a storage policy that has VSS enabled and assign it to the VM with vVols.
- Ensure that you have iSCSI access.
- Select the VM in the array dashboard and synchronize it with the guest OS using the **APP SYNC** feature.

The vCenter Plugin uses VSS when the storage policy for the VM has **Application synchronization** enabled and either Microsoft Exchange Server and Microsoft SQL Server selected.

Note: It is a good practice to create one storage policy that applies to the VM's operating system disks and another storage policy that applies to the data disks associated with the guest OS. The schedules must match exactly so that the vVols are all placed in a single volume collection.

Windows Toolkit includes VSS and the tools necessary for integration with the vCenter Plugin. These include the VSS Requestor and VSS Hardware Provider. For details about setting up Windows Toolkit, refer to the HPE Storage Windows Toolkit Release Notes and *Windows Integration Guide*.

Overview of the VSS for vVols Setup Process

Setting up Volume Shadow Copy Service (VSS) for virtual volumes (vVols) involves performing steps in the Windows guest operating system and the HPE vCenter Plugin.

As you perform the steps, you use different products and tools. Detailed information about working with these products and tools is available in the following documentation:

- HPE Storage Windows Toolkit. See the *Windows Integration Guide*, which is available on InfoSight in the Documentation section.
- Storage array running array OS. See the GUI or CLI administration guide. These guides are available on InfoSight in the Documentation section.
- vCenter Server, ESXi, and VMware Tools. See the VMware documentation for these products.
- HPE Storage vCenter Plugin. Information about working with the vCenter Plugin is in this document.
- VSS for vVols. Information about working with VSS and vVols is provided in this document.
- Microsoft VSS. See the Microsoft documentation.

The steps that follow provide a high-level overview of the actions you must perform and where you must perform them so that you can plan your installation and setup. These steps are a checklist of the actions you must take and do **not** contain all the details necessary to perform the steps. See the appropriate documentation for detailed instructions.

Keep the following in mind as you perform these steps:

- The database and log files need to be on a separate VMDK.
- You cannot use the VMware synchronous replication feature for volume collections.
- The VSS option to quiesce the operating system is a Microsoft VSS function, not a VSS function.

- 1 (VMware Web Client)** Create a storage policy that has **Application synchronization** set to VSS. For details, see the section *Create a VM Storage Policy*.

Note: The minimum time period that you can select with the **Minutely** schedule option is 30 minutes.

- 2 (VMware Web Client)** Assign the storage policy with VSS enabled to a vVols virtual machine (VM). This can be a new VM or an existing VM.

- The storage policy must have **Application synchronization** set to VSS. The VM will be listed as non-compliant to the storage policy. This state will change to compliant when you synchronize the VM with the guest OS.
- The VM must have the same login credentials as the guest OS so that it can communicate with HPE Storage Windows Toolkit.

- 3 (Guest OS)** Set up your guest OS. It must be running Windows Server 2008 R2 or later and array OS release 5.0.1.0 or later.

Note: For iSCSI arrays, iSCSI network access from within the guest VM to the array is required.

- 4 (Guest OS)** Make sure the guest OS VM is running VMware Tools.
- 5 (Guest OS)** Provision an extra network interface for the guest VM to use to access the iSCSI interface, which is required for mounting snapshots. Set up the guest OS network interface to connect to the same iSCSI subnet and add the array discovery IP address to the guest OS iSCSI initiator.
- 6 (Array)** Update the ACL to provide the guest OS IQN with snapshot access.
- 7 (Guest OS)** Install the application software you want, either Microsoft Exchange or Microsoft SQL Server.
- 8 (Guest OS)** Install Windows Toolkit 5.0.0 or later.
- a** Download the Windows Toolkit from InfoSight (<https://infosight.hpe.com/>).
 - b** Follow the installation instructions in the *Windows Integration Guide*.
 - c** Make sure you select the VSS feature when you install Windows Toolkit.
- 9 (HPE Dashboard)** Enable the VM to communicate with HPE Storage Windows Toolkit by synchronizing the VM and the guest OS. Select **Nimble Storage** > **Local** > **APP SYNC**.

Note: If you have multiple groups configured with the same vCenter Server, you should use the drop-down list in the top right. Select the group to display all the VMs and then perform the application synchronization.

You must enter the following information:

- IP address of the guest OS application server.
- The guest OS user name and password. The VM uses this information to log into the guest OS and use VSS.

Note: When entering a user name in VSS for vVols setup, you must use the format `username@domain`. The account needs to be a domain administrator, not a local machine administrator. This account is only for initial configuration and is not used for the VSS snapshots.

If this format does not work, log on using the local Windows Administrator account for the guest OS and remove any permissions issues.

- The application that you want VSS to take snapshots of. From the drop-down list, select either Microsoft Exchange Server or Microsoft SQL Server.
- Whether you want to use the Verify Backups feature. This feature is only available in environments using iSCSI. If you select this option and the application type is Microsoft Exchange Server with a database availability group (DAG), the dialog asks if you want to skip the consistency check for database files.

After you set up this environment, the vCenter Plugin automatically uses VSS when it takes snapshots of any vVols associated with that storage policy.

Note: When a VSS snapshot is taken, several flags and attributes are set on the snapshot. Volumes cloned or restored from that snapshot inherit those flags and attributes. If you perform a restore operation using a VSS snapshot, the VM or disk comes back up in an offline state. You must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You must also reset the flags and attributes before you mount a VSS snapshot or a volume cloned or restored from a VSS snapshot on a Windows host or cluster. You can use the Set-NimVolume PowerShell cmdlet. The PowerShell cmdlets are included in the Windows Toolkit. See the *Windows Integration Guide* for instructions on using these cmdlets.

VSS for vVols Limitations and Troubleshooting

There are few limitations and issues that you might encounter when using Volume Shadow Copy Service (VSS) for virtual volumes (vVols).



Exchange-verified snapshots and Fibre Channel

VSS for VVols does not support Exchange-verified snapshots in a Fibre Channel-based VVol environment. You cannot mount them using the Fibre Channel access protocol from within the guest operating system.

The snapshots can only be mounted as guest-attached iSCSI LUNs. If you need to mount this type of snapshot on the host, you must have a Microsoft iSCSI Software Initiator running in the guest operating system.

You can mount the snapshots on a proxy or backup server using Fibre Channel on the physical server or a virtual machine (VM) using VMware RDMs or iSCSI using an operating system initiator or VMware RDMs.

Warning message: VSS snapshot completed successfully, but the volume collection X needs to be reviewed. Extra volumes found in the volume collection

The warning message above appears in the array GUI if the volume collection has either volumes without application components or a volume that is a system drive when the snapshots were created.

Configuring application synchronization fails

If you cannot synchronize the VM and the guest OS, make sure that you have VMware Tools installed on the guest OS.

Setting Up the Guest Operating System for VSS for vVols

To use the Volume Shadow Copy Service (VSS) for virtual volumes (vVols), the guest operating system must have the following installed:

- Windows Server 2008 R2 or later
- Windows Toolkit 5.0.0 or later
- VMware Tools in the guest OS virtual machine (VM)
- An application (Microsoft Exchange Server or Microsoft SQL Server)
- An extra network interface for the guest VM to use to access the iSCSI interface

Note: The guest OS network interface must connect to the same iSCSI subnet. You must add the array discovery IP address to the guest OS iSCSI initiator.

The array associated with the guest OS must be running release 5.0.1.0. Make sure the ACL provides the guest OS IQN with snapshot access.

Note: For details about supported versions, including the vCenter Server, refer to the Validated Configuration Matrix. This tool is available on InfoSight (<https://infosight.hpe.com/InfoSight/#validated-configuration-matrix>).

Windows Toolkit provides VSS as an optional feature. You must select this feature when you install Windows Toolkit. The VSS feature includes the VSS requestor, the VSS hardware provider, and the PowerShell script ConfigWindowsHostForVvols.ps1.

Instructions for installing and setting up Windows Toolkit are in the *Windows Integration Guide*.

Information about working with array OS is in the GUI and CLI administration guides.

These guides are available on InfoSight in the [Documentation](#) section.

Synchronize the VM with the Guest OS

To use Volume Shadow Copy Service (VSS) for virtual volumes (vVols) you must synchronize the virtual machine (VM) so that it communicate with Windows Toolkit, which is running on the guest operating system. VSS is included in Windows Toolkit as an optional feature.

Note: When you assign a VSS-enabled storage policy to a VM, the VM shows up as non-compliant until you synchronize the VM with the guest OS.

You perform the synchronization operation from the array dashboard in the VMware vSphere Web Client.



Before you begin

You must have set up the Windows guest operating system on the VM by installing:

- Windows Toolkit 5.0.0 or later. See the *Windows Integration Guide* for instructions on doing this.
- VMware Tools. See the VMware documentation.
- Either Microsoft Exchange Server or Microsoft SQL Server. See the Microsoft documentation for these products.

Procedure

1. From the VMware vSphere Client Navigator panel, select **HPE Alletra 6000 and HPE Nimble Storage**.
2. When the dashboard appears, select the group from the drop-down **Local**.

The dashboard displays a list of all the VMs associated with the array as well as all VMs registered with the vCenter Server.

Note: If you have multiple groups configured with same vCenter Server, you should use the drop-down list in the top right. Select the group to display all the VMs.

3. Check the box next to the VM that you want to synchronize with the guest OS.
4. Select the **APP SYNC** button.
5. In the **Configure VSS SYNC** dialog, provide the following information to allow the VM to communicate with the guest OS.
 - The IP address of the guest OS.
 - The login information for the guest OS. You must enter the guest OS username and password. The VM uses this information to log into the guest OS and use VSS.
 - The application that you want VSS take snapshots of. From the drop-down list, select either Microsoft Exchange Server or Microsoft SQL Server.

You also have the option of selecting **Verify backups**. If you select this option and the application type is Microsoft Exchange Server with a database availability group (DAG), the dialog asks if you want to skip the consistency check for database files.

What to do next

After you synchronize the VM with the guest OS, you can check the VM's storage policy compliance. It will now show up as compliant.

Working with VM Storage Policies

VMware Virtual Volumes (vVols) support storage policy-based management (SPBM), which allows you to create virtual machine (VM) storage policies to define the capabilities assigned to storage. These capabilities can include disk type, encryption, compression, deduplication, snapshot management, and replication.

The policies also provides a Protection section, where you can set a backup schedule and specify a downstream VM that can take over if the upstream VM fails.

You can create a storage policy that serves as the default policy each time you create a storage object. For example, you can specify a default storage policy for a datastore. Each VM created on that datastore will automatically use that policy unless you select a different policy for the VM.

Virtual Machine Backups

When you set up storage policy-based management (SPBM), you can specify a virtual machine (VM) backup schedule. The storage policy wizard prompts you to provide a backup schedule in the Protection section of the VM storage policy. You can specify weekly, daily, hourly, and minutely backups. You can also specify a downstream replication partner that will take over if the upstream VM fails.



Each backup is a snapshot of the VM. The storage policy specifies how long these snapshots are retained. If you want to restore a disk or a VM, you select either the current or an older snapshot to use for that operation. If the storage policy has Volume Shadow Copy Service (VSS) for virtual volumes (vVols) enabled, you can select one of those snapshots.

Note: Snapshots that were created with VSS for vVols have the tag "**AppSync**". For information about VSS, see [VSS and vVols](#) on page 92.

When you schedule your backup, keep the following best practices in mind:

- Choose a time for the backup when the system is less busy.
- Keep in mind that the number of VMs included in the backup operation will affect how long the backup takes.
- You should only schedule minutely backups for a small number of critical VMs to avoid slowing the system with repeated backups.

Note: The minimum time period that you can select with the VSS **Minutely** schedule option is 30 minutes.

Create a VM Storage Policy

You use the vCenter interface to create and edit virtual machine (VM) storage policies. These policies allow you to filter the available datacenters to locate the ones that match your requirements for the VMs. They also allow you to set up the following:

- Backup schedules for the VM
- A replication partner

Note: While you protect VMFS datastores by adding a datastore volume to an array OS volume collection, you configure vVol datastore protection from within the VMware vSphere Web client itself.

If you are using HPE Storage Windows Toolkit 5.0.0 or later and VMware Virtual Volumes (vVols), you also have the option of specifying that the backups use HPE Storage Volume Shadow Copy Service (VSS). VSS provides application-consistent snapshots for Microsoft Exchange Server and Microsoft SQL Server. Using VSS allows you to restore the application data without having to restore the entire VM.

Note: For information about VSS, see [VSS and vVols](#) on page 92.

Procedure

1. From the Web Client Navigator, select **Home** > **Policies and Profiles**.
2. Under the **VM Storage Policies** tab, select the icon to create a new VM storage policy.
3. In the Create New VM Storage Policy wizard, enter the requested information.
Most of the fields are self-explanatory. The following steps provide additional information if you need it.
4. Provide a clear name and description so that you can quickly identify the policy.
5. Under **2b Rules-set 1**, check **Use rule-sets in the storage policy**.
6. Specify the storage type.
The VMs will be placed on datastores that match that storage type.
7. Specify the storage requirements for the VM.. These requirements identify the datastores that you can use for the VMs.
 - Deduplication. Select this if you want to use deduplication.
 - Application policy. The default is operating system. Each disk can have a different application policy.
 - All Flash. Specify whether the storage must be all flash.
 - Performance. Enter limits for IOPS and MB/S.
 - Data encryption cipher. Specify whether to use a cipher for data encryption.

Note: You must enable encryption within the array user interface before you can successfully apply an encryption-enabled policy to a VM.



8. Specify the Protection Schedule details for the VM. If you are using Microsoft Exchange Server or Microsoft SQL Server and have set up your environment to use VSS, you can specify that here.

You have several options when creating the backup schedules:

- How frequently backups are made. You can specify weekly, daily, hourly, or minutely. You should only schedule minutely backups for a small number of critical VMs. Otherwise, the repeated backups might impact the performance of the VM or array.

Note: The minimum time period that you can select with the **Minutely** schedule option is 30 minutes.

- The time of day when backups are made. It is a good practice to pick a time when the system will be less busy.
- The day of the week when backups are made. You can select multiple days.
- Whether to require application synchronization. When you select this option, you can specify that the HPE Storage vCenter Plugin use VSS to take snapshots of vVols that host Microsoft Exchange Server or Microsoft SQL Server application data. Your guest operating system must be running Windows Toolkit 5.0.0 or later to use this feature.

If you select **None**, the vCenter Plugin uses the standard snapshot feature.

- How many snapshots to retain locally. The vCenter plugin will store the current snapshot and as many older snapshots as needed to meet this number. It will purge the oldest snapshot to maintain this number.

9. Specify the replication partner details for the VM.

Note: Before you assign a replication partner to a storage policy, you must configure it in the array. Refer to either the *GUI Administration Guide* or the *CLI Administration Guide* for more information. You can download these guides from InfoSight.

- The replication partner name. This storage should be on an array that is registered to a different vCenter Server.
- Replication frequency. Specify how often you want the replication partner VM to be updated with a snapshot of the upstream VM. The vCenter plugin monitors the upstream VM and updates the partner VM after the specified number of snapshots of the upstream VM have been taken. For example, if you specify 3, the partner VM is updated with the third snapshot of the upstream VM. The update count resets after each update.

To replicate every snapshot, set this value to 1.

- Snapshots to retain on replica. Enter how many snapshots of the upstream VM you want to save on the partner array.
- Delete replicas from partner. Selecting this option means that, if you delete a vVol on the upstream VM, it will also be deleted from the partner VM.

10. When you have added all the necessary rule sets, select **Finish**.

Change the Assigned Storage Policy for a VM

When you create a virtual machine (VM), you can assign a VM storage policy to it. If you want to change the policy for that VM, you can use the **Edit VM Storage Policies ...** option to assign another policy to it. The **Edit VM Storage Policies ...** option also allows you to update the policy before you assign it.

Before you begin

The storage policy must exist.

Procedure

1. Right-click on the VMware Virtual Volumes (vVols) VM.
2. Select **VM Policies** > **Edit VM Storage Policies**.
3. Either select the storage policy from the drop-down list and click Apply to All, or select the disk that you want associated with that policy from the drop-down list next to it.
4. Click OK.



Verifying Policy Compliance

The vCenter server routinely checks virtual machines (VMs) to determine whether they are still in compliance with their assigned storage policy. The vCenter server displays an error message when it discovers VMs that are out of compliance. It also puts the VM in the non-compliant list.

A VM might become non-compliant because of changes to the datastore or the VM features. For example, Deduplication might have been enabled when the storage policy was assigned and later disabled. As a result, the VM now appears as being non-compliant with its storage policy.

Note: When you set up Volume Shadow Copy Service (VSS) for VMware Virtual Volumes (vVols) the VM is listed as non-compliant to the storage policy until you synchronize it with the guest operating system. For information about VSS, see [VSS and vVols](#) on page 92.

You can check the non-compliant list at any time. There are several ways to do this:

- You can view this information in the VMware vSphere Web Client.
 - 1 From the VMware vSphere Client Navigator pane, select **VMs and Templates**.
 - 2 Select the VM.
 - 3 Select **Configure Policies** > **Check VM Storage Policy Compliance Button**.
 - 4 Check the list to see if the VM you are concerned about is non-compliant.
- You can set up the **Details** list for a VM to display the **VM Storage Policies Compliance** column:
 - 1 When looking a vCenter server, datacenter, host cluster, host, or datastore select the **VM** tab.
 - 2 Select **Show/Hide Columns**.
 - 3 Enable the **VM Storage Policy Compliance** column.
- You can monitor the storage policy to see which VMs attached to it are compliant:
 - 1 Go to **Policies and Profiles** > **VM Storage Policies**.
 - 2 Select your policy from the inventory tree.
 - 3 Select the **Monitor** tab.
 - 4 Select the **VMs and Virtual Disks** button.

What to do next

If the VM is non-compliant, check the storage policy against the datastores in the VM to determine why the VM is not compliant. After you determine what the problem is, you can perform one of the following actions:

- Modify the storage policy to reflect any datastore changes since the policy was first assigned.
- Assign a different storage policy to the VM.
- Update the datastore to match the current storage policy requirements.

Using Replication Partner VMs with the vCenter Plugin Dashboard

The HPE vCenter Plugin dashboard allows you to set up a replication partner virtual machine (VM) that provides a backup for the upstream or source VM. The replication partner is on a another array that is registered with a different vCenter Server. It is not located on the same array as the upstream VM. The replication partner VM is also known as the downstream or target VM.

You can use a replication partner VM to perform the following tasks:

- Claim. The claim operation means that the partner VM becomes the upstream VM and appears in the local VM list.
- Clone. If the upstream VM is available, you can clone the replication partner VM.
- Purge. The purge operation permanently deletes the replication partner VM. You can only perform a purge operation if the upstream VM has been deleted.



Note: You cannot use this replication feature with the VASA Provider Virtual Volume (vVol) disaster recovery capability. The vCenter Plugin replication feature and the VASA 3.0 failover capability are mutually exclusive.

These operations are available when you select **HPE Alletra 6000 and HPE Nimble Storage > Replicated**.

Keep in mind that replication involves regular read and write operations, which add to the load on the storage.

Note: The vCenter Plugn is not intended to be a replacement solution for a site disaster.

Specifying a Replication Partner VM

You specify a replication partner virtual machine (VM) when you create a VM storage policy. Both the upstream VM and the replication partner VM must use the same VM storage policy.

When you set up a replication partner, you should keep the following in mind:

- Replication partners must have been configured on the storage array before you configure the VMs in the dashboard.
- When you set up the replication partner, select **Use the same pool and folder as the source location if the location also exists on the local group**.
- **(Recommended)** Both the upstream and downstream folder should use the same name.
- The downstream vCenter server should manage the folder.
- The replication feature should point to the folder.

The VM storage policy governs the replication partner VM. It includes requirements, including the following:

- How frequently the partner VM is replicated to match the upstream VM.
- How many snapshots of the upstream VM are retained on the partner VM.
- Whether deleting a virtual volume (vVol) on the upstream VM will also delete that vVol from the downstream VM.

Note: You must select Replication groups either from the VM wizard or by clicking the button on the VM when you select **Configure > Policies > Edit VM Storage Policies**.

For information about creating a VM Storage Policy, see the section [Create a VM Storage Policy](#) on page 97.

Claim a VM

If your upstream virtual machine (VM) is deleted, you can claim the replication partner VM. The replication partner must be on the downstream storage array and vCenter Server. The claim operation makes the partner VM become the upstream VM and appear in the list of local VMs.

Before you begin

You must have set up a replication partner VM.

Procedure

1. From the Downstream VMware vSphere Client Navigator, select **HPE Alletra 6000 and HPE Nimble Storage**.
2. When the dashboard appears, select **Replicated** to view all the replication partner VMs.
3. Select the VM that you want to use as the new upstream VM.
4. Select the **CLAIM** button at the top of the list.
5. In the Confirmation pop-up box, select **CLAIM** again.
The VM that you selected now appears in the local VM list.

Options for Restoring VMs from the vCenter Plugin

HPE Alletra 6000 and HPE Nimble Storage vCenter Plugin provides a restore wizard with several options for restoring and cloning virtual machines (VMs). These options include:



- Replacing the current VM. You can use a single snapshot of the VM or multiple snapshots of individual disks on the VM and replace each disk individually. The snapshot information overwrites the existing information on the VM.
- Replacing one or more disks on the VM. You can replace disks individually using snapshots. The vCenter plugin displays the timestamp of each Snapshot Collection. If multiple disks use the same storage policy, then it shows the timestamp for that snapshot collection. You can specify different snapshots to use when you replace the disks. The snapshot information overwrites the existing information on the each disk that you replace.
- Attaching a disk to the VM. You can clone a disk from a snapshot and then attach it to a VM. When you do this, you are effectively creating a new disk.
- Create a new VM. You can clone the existing VM and use the clone as a new VM. In this case, no changes are made to the existing VM.

You select the VM and run the restore wizard from the dashboard in the VMware vSphere Web Client. The restore wizard allows you to perform operations such as:

- Specifying which snapshot to use for the restore or clone operation.
- Searching a snapshot for individual disks.
- Specifying whether the restore operation replaces the current VM or clones the VM to create a new VM.
- Specifying whether the restore operation replaces an existing disk on the VM or clones the disk and attaches it as a new disk on that VM or another VM.

Note: When you clone a disk and attach it to a VM, the restore wizard applies the storage policy of the parent VM to the disk.

Note: Restoring VMs that exist on multiple datastores is unsupported. This includes storing your VSWP files on a separate datastore as the VM, or having multiple virtual disks stored on different datastores.

During a restore operation, snapshots that were created with Volume Shadow Copy Service (VSS) for virtual volumes (vVols) display the tag **AppSync**. If you perform a restore operation using a VSS snapshot, the VM or disk comes back up in an offline state. You must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags. The PowerShell cmdlets are included in the HPE Storage Windows Toolkit. See the *Windows Integration Guide* for instructions about using these cmdlets. See [VSS and vVols](#) on page 92.

Restore an Entire VM from the vCenter Plugin

The restore wizard enables you to replace the current virtual machine (VM) with data from a snapshot or to create a clone that becomes a new VM. When restoring a VM, you can use the most recent snapshot or select an older one.

You perform the restore operation from the dashboard in the VMware vSphere Web Client.

Before you begin

The VM you want to restore must be powered off before you perform an in-place restore operation.

Procedure

1. From the VMware vSphere Client Navigator, select **HPE Alletra 6000 and HPE Nimble Storage**.
2. When the dashboard appears, select **Local**.

The dashboard displays a list of all the VMs associated with the array as well as all VMs registered with the vCenter Server.

3. Select the VM you want to restore.
4. Select the **Restore** button at the top of the list to start the restore wizard.
5. In the restore wizard, select the snapshot you want to use to restore the VM.

By default, the wizard displays the most recent snapshot. The **View more** option displays all the available snapshots. The **Search** option allows you to look for a specific snapshot by entering a date and time for the snapshot.

Note: Snapshots that were created with Volume Shadow Copy Service (VSS) for virtual volumes (vVols) have the tag **AppSync**. If you use a VSS snapshot, the VM comes back up in an offline state. You must select each disk and

clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags. See the *Windows Integration Guide* for instructions about using this cmdlet.

6. Select the **Replace existing VM** option.
7. Select **Restore**.

The restore wizard powers off the VM that you are restoring. When the operation successfully completes, the wizard displays a message in the menu bar stating that the VM has been restored. It then powers the VM back on.

Restore One or More Disks From a VM

The restore wizard enables you to restore individual data disks on a VM on a vVol datastore by replacing them with information from one or more snapshots. This operation performs an in-place restore.

You can perform the operation using the most recent snapshot of the disk or an older snapshot. You can restore multiple disks at one time using the same snapshot. If you want to restore multiple disks using different snapshots, you can repeat these steps until you have restored all the disks.

You perform the restore operation from the dashboard in the VMware vSphere Web Client.

Procedure

1. From the VMware vSphere Web Client Navigator, select **HPE Alletra 6000 and HPE Nimble Storage**.
2. Select **Local**.
3. Select the VM containing the disk (or disks) you want to restore.
4. Select the **Restore** button at the top of the list to start the restore wizard.
5. In the restore wizard, select the option **Show individual disks**.

The wizard displays the list of disks that are on the VM. If the VM has multiple storage policies, the wizard groups the disks by storage policy.

6. Select the disk that you want to restore.
You can restore multiple disks at the same time using one snapshot..
7. Select the snapshot you want to use when restoring the disk.

If you selected disks with different storage policies, you must select a snapshot for each storage policy.

By default, the wizard displays the most recent snapshot. The **View more** option displays all the available snapshots. The **Search** option allows you to look for a specific snapshot by entering a date and time for the snapshot.

Note: Snapshots that were created with Volume Shadow Copy Service (VSS) for virtual volumes (vVols) have the tag **AppSync**. If you use a VSS snapshot for a restore operation, the VM and disks come back up in an offline state. You must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags. See the *Windows Integration Guide* for instructions about using this cmdlet.

8. Select **Replace existing disk in the same VM**. The wizard powers off the VM and overwrites each disk you selected with information from the snapshot. Then it powers on the VM.
9. Select **Restore**.

When the operation successfully completes, the restore wizard displays a message in the menu bar to let you know the disk has been restored.

Clone a VM

The restore wizard enables you to clone an existing VM on a vVol datastore from a snapshot to create a new VM. You can use the most recent snapshot or select an older one.



You perform the restore operation from the dashboard in the VMware vSphere Web Client.

Procedure

1. From the VMware vSphere Client Navigator panel, select **HPE Alletra 6000 and HPE Nimble Storage**.

2. When the dashboard appears, select **Local**.

The dashboard displays a list of all the VMs that are associated with the array as well as all VMs registered with the vCenter Server.

3. Select the VM that you want to clone.
4. Select the **Restore** button at the top of the list to start the restore wizard.
5. In the restore wizard, select the snapshot you want to use to create the cloned VM.

By default, the wizard displays the most recent snapshot. The **View more** option displays all the available snapshots. The **Search** option allows you to look for a specific snapshot by entering a date and time for the snapshot.

Note: Snapshots that were created with Volume Shadow Copy Service (VSS) for virtual volumes (vVols) have the tag **AppSync**. If you use a VSS snapshot for a restore operation, the VM and disks come back up in an offline state. You must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags. See the *Windows Integration Guide* for instructions about using this cmdlet.

6. Select the **Clone a new VM** option.

You must enter a name for the new VM. By default, the wizard displays a name that consists of the name of the original VM and a time stamp.

7. Select **Restore**.

When the operation successfully completes, the restore wizard displays a message in the menu bar to let you know the VM has been cloned.

Clone One or More Disks

The restore wizard enables you to clone one or more data disks. The clone operation creates new disks based on the snapshots and attaches them to either the current VM or another VM.

You perform the clone operation from the dashboard in the VMware vSphere Web Client.

Procedure

1. From the VMware vSphere Web Client Navigator select **HPE Alletra 6000 and HPE Nimble Storage**.

2. Select **Local**.

3. Select the VM containing the disk (or disks) that you want to clone.

4. Select the **Restore** button at the top of the list to start the restore wizard.

5. In the restore wizard, select the option **Show individual disks**.

The wizard displays the data disks that are on the VM. If the VM has multiple storage policies, the wizard groups the disks by storage policy.

6. Select the disk that you want to clone.

You can restore multiple disks at the same time.

7. Select the snapshot you want to use when cloning the disk.

If you selected disks with different storage policies, you must select a snapshot for each storage policy.

By default, the wizard displays the most recent snapshot. The **View more** option displays all the available snapshots. The **Search** option allows you to look for a specific snapshot by entering a date and time for the snapshot.



Note: Snapshots that were created with Volume Shadow Copy Service (VSS) for virtual volumes (vVols) have the tag **AppSync**. If you use a VSS snapshot for a restore operation, the VM and disks come back up in an offline state. You must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags. See the *Windows Integration Guide* for instructions about using this cmdlet.

8. Specify whether where you want to attach the cloned disk.

If you select the option

- **Attach as new disk in same VM**, the wizard creates a new disk for each one you cloned using in the snapshot you specified and attaches the disk to the current VM.
- **Attach as new disk in different VM**, you must specify the VM from the drop-down list. The wizard creates a new disk for each one you selected using the snapshot you specified and attaches the disk to that VM.

9. Select **Restore**.

When the operation successfully completes, the restore wizard displays a message in the menu bar to let you know the disk has been cloned and attached to the VM you specified.

Options for Deleting vVol VMs from the vCenter Plugin

You can use both the dashboard and the VMware vSphere Web Client interface to delete VMware Virtual Volume (vVol) virtual machines (VMs). After you delete a vVol VM, the dashboard provides you with options for either restoring those VMs or purging them from the system.

When you select the **DELETED** tab on the dashboard, you see a list of all the VMs used by arrays that have been deleted.

The delete operation provides added protection against accidentally deleting a VM by taking a snapshot of the VM before moving it to the deleted list. After 72 hours, it removes the VM from the deleted list and purges it from the system.

While a VM is on the deleted list, you can use the undelete operation to restore that VM to the list of available VMs.

If you are certain that you no longer need the VM, you can perform a purge operation to remove the VM before the 72 hours is up. You **cannot** restore a purged VM.

You can select a VM to delete from the list provided by the dashboard on the local pane or from the list provided by the vCenter Server in the **VMs and Templates** pane.

When you use the dashboard to delete a VM, the HPE vCenter plugin powers it off and unregisters it from the vCenter Server. It moves the VM to the Recycle Bin and adds the name of the VM to the deleted list in the dashboard.

When you use the VMware vSphere Web Client interface to delete a VM, you must manually power off the VM before you select Delete from the Action menu. The VM is then added to the deleted list, where you can use the undelete and purge operations.

Note: If there is not a snapshot of the VM, the undelete operation restores the VM to a hardware configuration that is similar to the configuration when VM was initially created. You might need to manually attach any disks, interfaces, or other things that you configured on the VM after it was created.

Delete a VM

You can use the delete feature provided by the dashboard to delete a virtual machine (VM).

Note: For information about using the VMware vSphere Web Client to delete a VM, see the VMware documentation.

Procedure

- 1.** From the VMware vSphere Web Client Navigator panel, select **HPE Alletra 6000 and HPE Nimble Storage**.



- When the dashboard appears, select **Local** to view all the available VMs.

The dashboard displays a list of all the VMs associated with the array as well as all VMs registered with the vCenter Server.

- Select each VM that you want to delete.
- Select the **DELETE** button at the top of the list.
- In the Confirmation pop-up box, select **DELETE** again.

When you confirm the delete operation, the vCenter plugin performs the following tasks:

- Powers down each VM you selected.
- Unregisters the VM from the vCenter Server.
- Puts the VM in the Recycle Bin.
- Adds the VM to the dashboard deleted list.

- Select the **DELETED** tab to view the names of all the VMs that have been deleted.

The VMs remain in the deleted list for 72 hours unless you perform either an undelete operation or a purge operation.

Undelete a VM

When you delete a virtual machine (VM), it goes on the dashboard deleted list. While it is on the that list, you can perform an undelete operation to return the VM to the list of available VMs.

If there is not a snapshot of the VM, the undelete operation restores the VM to a hardware configuration that is similar to the configuration when VM was initially created. In this case, you might need to manually attach any disks, interfaces, or other things that you configured on the VM after it was created.

Note: You cannot undelete a VM that has been purged from the system. By default, VMs are purged from the deleted list every 72 hours. You can also manually purge a VM.

Procedure

- From the VMware vSphere Client Navigator panel, select **HPE Alletra 6000 and HPE Nimble Storage**.
- When the dashboard appears, select **DELETED**.

The dashboard displays a list of all the VMs that are on the deleted list.

- Select each VM that you want to restore and click **UNDELETE**.
- In the Confirmation pop-up box, click **UNDELETE** again.

The vCenter plugin takes the VM from the Recycle Bin, re-registers it with the vCenter Server, and then puts it back in the list of active VMs. You can now power on the VM.

Purge a VM Using the vCenter Plugin

The purge operation provided by the dashboard removes the selected virtual machines (VMs) from the array. You cannot reverse this operation. You must always make sure that you do not need the VM before you purge it.

Procedure

- From the VMware vSphere Client Navigator panel,select **HPE Alletra 6000 and HPE Nimble Storage**.
- When the dashboard appears, select **DELETED**.

The dashboard displays a list of all the VMs that have been deleted.

- Select each VM you want to remove and click **PURGE**.
- In the Confirmation pop-up box, click **PURGE** again.

The vCenter removes the VM from the disk.



Troubleshooting Tips

This section describes the most common issues and their resolutions/workarounds.

Registration Error - Invalid Provider Certificate

You see the message "The provider certificate is invalid. It is either empty, malformed, or expired, not yet valid, revoked, or fails host name verification."

- First, try clicking **OK** again.
- If you still see this message, it may be that the certificate is invalid. This error indicates the address used in the URL is not in the certificate, or that the certificate start time is in future (compared to vCenter's current time).

Failure to Add VP - Time Mismatch

You see a "Failure to Add VP - Time Mismatch" error.

This could be due to a time mismatch between vCenter and the array. Check the following timestamps in the order presented below.

- 1 First, look at the current time on vCenter.

```
# date
Thu Dec 18 15:10:56 UTC 2014
```

- 2 Then look at the current time on the array. In this example the time is different.

```
# date --utc
Thu Dec 18 15:15:37 UTC 2014
```

- 3 Then look at the certificate start time in UTC. In this example, the start date and time cannot be before December 18 at 16:58. The current times on vCenter and the array, in these examples, are earlier than the certificate start time.

```
# openssl x509 -in /nimble/var/private/config/current/group/certs/host.crt
-noout
-startdate notBefore=Dec 18 00:16:58 2014 GMT
```

- 4 Then look at errors in vCenter sps.log. In the example below, the timestamp check has failed because of the above conditions.

```
# grep "timestamp check failed"
/var/log/vmware/vmware-sps/sps.log 2014-11-26 06:02:44,628 [pool-14-thread-3]
WARN opId=4106403a-680d-4cd4-80d2-8e7e66d9762a
com.vmware.vim.sms.provider.vasa.VersionHandler - [isLegacyProvider]
Failed to retrieve version information from provider:
sun.security.validator.ValidatorException: PKIX path validation failed:
java.security.cert.CertPathValidatorException: timestamp check failed
```

- 5 Finally, look for errors in vvold.log on ESXi. The example below, the certificate is not yet valid.

```
2014-12-17T17:41:37.154Z error vvold[FFD2FB70] [Originator@6876
sub=HttpConnectionPool-000000] [ConnectComplete] Connect failed to <cs
p:1f26bdc8, TCP:10.18.112.46:8443>; cnx: (null), error:
N7Vmacore3Ssl18SSLVerifyExceptionE(SSL Exception: Verification parameters
```

- Synchronize the times on vCenter and on the array.



Datastore Inaccessible

If the datastore is inaccessible, there are three things to check:

- For an iSCSI array, check to see if the discovery IP is configured on the ESX hosts. A static target with a group name should show up after performing a rescan.
- Ensure that the ESX time is not behind the array time. See [Failure to Add VP - Time Mismatch](#) on page 106.
- Ensure that the latest version of HPE Storage Connection Manager is installed on every host that should mount the datastore.

VSS Snapshots Fail with the Message: "No Volume Connected to the Host"

Snapshots taken of a VMware Virtual Volume (VVol) using Volume Shadow Copy Service (VSS) fail if the VVol does not have its array serial number registered on the workstation. The error message "No volume connected to the host" is displayed.

This error occurs when you import a virtual machine (VM) containing the VVols to a new vCenter workstation. The registration information does not get added to the new workstation. You must manually add it.

Procedure

1. Power off the VM.
2. Open the VMX file and add the following line for the VVol:

disk.EnableUUID="true"

Note: For information about locating the VMX file, see the VMware KB article [Locating a hosted virtual machine's files \(1003880\)](#)

3. Power the VM back on.



VMware Synchronized Snapshots and VMFS Datastores

The array OS contains built-in VMware integration that allows snapshots to be synchronized with VMware Virtual Machine snapshots in VMFS datastore environments. Doing this ensures that they are application-consistent. To use this feature, no additional configuration is required in the guest OS; however, you must have the latest VMware Tools installed.

Because of this synchronized snapshots integration, the array OS lets you schedule and manage application-consistent snapshots and replicas. Synchronizing snapshots coordinates with the application to quiesce I/O to ensure that the snapshot does not capture in-progress writes. The array OS gives you the ability to coordinate with the VMware vCenter Server and VSS Services Support in VMware Tools by using vSphere APIs to quiesce I/O to ensure that the snapshot does not capture any in-progress file system transactions.

Note: HPE supports VM-consistent snapshots for vVols and application-consistent snapshots for SQL and Exchange applications only.

To use VMware synchronized snapshots, select *VMware synchronization* for a volume collection and include the vCenter Server, an Administrator user name and password. The system works with vCenter to take application-consistent snapshots.

A vCenter synchronized snapshot can be scheduled in a volume collection. The system takes snapshots of all volumes in the volume collection at the same time. By using the vSphere APIs and VSS Services Support in VMware Tools, snapshots and replicas are application consistent, rather than just crash consistent. As a result, when you restore from an HPE snapshot, the application can immediately start using the data without performing extra restore recovery steps.

For testing and development, a snapshot can be cloned instantly and used to verify an application's compatibility with the system or to use with a new application.

Array clones are *zero-copy*: they share data blocks with the source volume, creating an extremely capacity- and time-efficient read/write copy for test or development purposes.

How HPE Storage Synchronization Works with VMware

Virtual Machines (VMs) are in a datastore that is on the specified volume(s). When VMware synchronization is enabled in the volume collection, the software in the array communicates directly with vCenter Server.

HPE ensures that snapshots are application-consistent for these volumes by using the vSphere APIs for backup and restore operations. Doing this also ensures that clones from these snapshots are consistent. The snapshot/replication schedule configured for a volume collection is controlled without extra requirements.

The sequence of synchronization steps is as follows:

- 1 Based on a schedule, the array OS connects to vCenter to determine the VMs present in the datastores on the LUNs in the volume collection.
- 2 vCenter Server queries ESXi to determine the list of VMs and returns the information to the array.
- 3 The array OS requests that vCenter Server take the snapshot of the VMs.
- 4 vCenter then communicates with the VMware VSS Services Support in VMware Tools to quiesce application writes to LUNs that correspond to volumes associated with the volume collection. Snapshot creation is then triggered, a VM application-consistent snapshot is created, and writes are unquiesced.
- 5 The array takes the snapshot of the volumes. The snapshot may be replicated to a remote array, depending on the schedule.
- 6 The VM snapshot is deleted.

Snapshot Exclusion and Inclusion Options for VMs and Datastores

By default, the VMware synchronized snapshots provided by array OS include all the VMs and datastores in a volume collection. Starting with array OS release 4.2.0.0, you have the option of including or excluding specific VMs and datastores from the snapshot and its restore operation.



The array OS implements this option through the VMware Tags feature when you are running vCenter Server 6.0 or later. To use this option, you must first create a VMware tag category called **NimbleVMwareSyncSnaps** and then create two tags. The exclude tag is called **NimbleSnapExclude**, and the include tag is called **NimbleSnapInclude**. You can only apply these tags to VMware vSphere objects.

You have the following choices:

- **No tags.** If you do not create or assign any tags, the synchronized snapshot and its restore operations include all VMs and datastores. You do not have to take any action or apply any tags.
- **Exclude tag.** Any VM or datastore that has an Exclude tag associated with it is omitted from all synchronized snapshots and their restore operations.
- **Include tag.** Any VM or datastore that has an Include tag associated with it is always included in all synchronized snapshots and their restore operations.

You can have a mix of no tags, Exclude tags, and Include tags. You do not need to assign a tag to a VM or datastore unless you want to exclude it or explicitly include it.

For example, you might exclude a datastore, but include some of the VMs within it by associating them with the Include tag. When you associate a datastore with an Exclude tag, any VM within the datastore that does not have an Include tag is omitted from a snapshot operation or a restore operation for that snapshot. You can also associate a datastore with an Include tag. Then, any VM with an Exclude tag is omitted from a snapshot or restore operation and all other VMs, regardless of whether they have an Include tag, would be included.

Create a NimbleVMwareSyncSnaps Category

To use the VMware Tags feature to include or exclude a VM or datastore from a synchronized snapshot and its restore operation, you must set up a VMware category for the Exclude and Include tags. This category defines how you can use these tags.

The following steps explain how to create categories using the VMware vSphere Web Client.

Note: You can also use other methods to create tag categories, including VMware vSphere PowerCLI cmdlets and the VMware vSphere APIs. See the VMware documentation for information on using those tools.

Before you begin

Make sure your environment includes the following:

- Array OS 4.2.0.0 or later
- vSphere Server 6.0 or later

Procedure

1. From the **Navigator** pane in the VMware vSphere Web Client, select **Tags & Custom Attributes**.
2. Select the **Tags** tab.
3. Select **Categories**.
4. Select the **Create** icon and perform the following tasks:
 - Create a category called **NimbleVMwareSyncSnaps** that has the **One Tag per object** attribute. This attribute ensures that only one tag from this category can be associated with a vSphere object at a time.
 - Under Associable Object Types, select **Datastore** and **Virtual Machine**. This selection ensures that these tags can be applied to only these objects.

Create Include and Exclude Tags

The Include and Exclude tags enable you to explicitly designate VMs and datastores to be excluded or included when VMware takes a synchronized snapshot or performs a restore operation for that snapshot.

Using these tags is optional.



Before you begin

Make sure your environment includes the following:

- Array OS 4.2.0.0 or later
- vSphere Server 6.0 or later
- A VMware tag category called **NimbleVMwareSyncSnaps**

Procedure

1. From the **Navigator** pane in the VMware vSphere Web Client, select **Tags & Custom Attributes**.
2. Select the **Tags** tab.
3. Select **Tags**.
4. Select the **Create** icon.
5. Create a tag called **NimbleSnapExclude** and associate it with the **NimbleVMwareSyncSnaps** category.
6. Create a tag called **NimbleSnapInclude** and associate it with the **NimbleVMwareSyncSnaps** category.

Assign Include and Exclude Tags to VMs and Datastores

You can assign **NimbleSnapExclude** and **NimbleSnapInclude** tags to VMs and datastores. The VMware tags feature allows you to explicitly designate VMs and datastores to be excluded or included when VMware takes a synchronized snapshot or performs a restore operation for that snapshot.

Using these tags is optional. You can use them with only some VMs and datastores while leaving the other VMs and datastores without tags. You can omit them completely. VMware synchronized snapshots always include VMs and datastores that have no tags.

Before you begin

Make sure your environment includes the following:

- Array OS 4.2.0.0 or later
- vSphere Server 6.0 or later
- A VMware tag category called **NimbleVMwareSyncSnaps**
- The VMware tags **NimbleSnapExclude** and **NimbleSnapInclude**

Procedure

1. From the **Navigator** pane in the VMware vSphere Web Client, select **Hosts and Clusters** > **Storage**.
2. Right-click on the vSphere object (VM or datastore) that you want to tag.
3. Select **Tags and Custom Attributes**.
4. Select **Assign Tag**.
5. Select the **NimbleSnapExclude** tag to exclude that VM or datastore from synchronized snapshots and their restore operations. Select the **NimbleSnapInclude** tab to include that object in those operations.

Note: You cannot assign both tags to the same vSphere object.

6. Repeat these steps until you have assigned a tag to all the objects that you either want to exclude or to explicitly include. Any object that does not have a tag is automatically included in all synchronized snapshots and their restore operations.

Volume Collections and VMware Objects

The array OS supports at most 16 vCenter synchronized snapshot schedules to run at a time. When 16 vCenter synchronized snapshot operations are in progress, any subsequent vCenter synchronized snapshot operation may result in only a crash-consistent snapshot being taken. An alert is issued indicating the failure to take a vCenter synchronized snapshot.



Note: HPE recommends that vCenter synchronized snapshot schedules be staggered so that no more than 16 run in parallel. In addition, snapshots should be scheduled at a frequency of one hour or more.

You should schedule as few simultaneous vCenter synchronized snapshots as possible on the same ESXi cluster to reduce resource contention and obtain timely successful application-consistent snapshots.

When you select a volume collection that includes synchronization with VMware vCenter, you must define the vCenter Server. The vCenter Server manages all the virtual machines. The first time you use this option, you must provide the name of a user with Administrator access and that user's password to access the vCenter Server.

After you complete the creation, return to the details page for the volume collection and click **Validate** to ensure that the username, password, and permissions are correct.

The following list includes the permissions that are checked on all the VMs in the datastores of the VolColl:

- VirtualMachine.State.CreateSnapshot
- VirtualMachine.State.RemoveSnapshot

The system displays either a success message or any issue that was found.

Volumes can be accessed from servers on virtual machines in the same way that physical servers access volumes. Synchronization is similar to any other synchronization-based schedules within the volume collection. When you apply a snapshot schedule to a volume collection, it applies to all volumes assigned to that collection.

Because all volumes in the volume collection share the same schedules for snapshots and replication, stagger snapshot schedules so that not all the snapshot work is being done at the same time. For example, you might set up your schedules to take hourly snapshots on the hour, daily snapshots on the half-hour, and replicated snapshots every few hours on the fifteen-minute mark.

Bringing a VSS Snapshot Online After a Restore Operation

If you use a snapshot that was created with HPE Volume Shadow Copy Service (VSS) for a restore operation, the RDM or guest operating system iSCSI disks might come back up in an offline state.

When this happens, you must select each disk and clear the VSS snapshot-related flags to bring the volumes online. You can use the **Set-NimVolume** cmdlet to clear the flags.

The PowerShell cmdlets are included in the HPE Storage Toolkit for Windows. See the Windows Integration Guide for more information about the Powershell cmdlets.



SRM and Storage Integration

The VMware vCenter Site Recovery Manager (SRM) is a plugin to the vCenter Server that enables you to create and test disaster recovery plans in a VMware environment.

SRM provides the following options for working with arrays:

- SRM works with the Replication Adapter (SRA). You can use SRM and SRA with storage arrays to perform array-based recovery operations.
- SRM 8.3 and above work with VASA Provider to fully integrate with VMware Virtual Volumes (vVols). It provides protection granularity to the level of a virtual machine (VM) on an array.

Note: In addition to using SRM, you can also use VASA Providers running VASA 3.0 or later provide a disaster recovery feature that restores vVols. For more information about using only VASA Provider with vVols, see [Using VASA Provider to Provide Disaster Recovery for vVols](#) on page 87.


How SRA Works with SRM

You can use HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM and HPE Nimble Storage Replication Adapter to create and test a Disaster Recovery (DR) plan for a storage array without affecting your production environment. After you verify that your DR plan works the way you expect, you know that your DR plan is adequate to protect your datacenter.

In DR scenarios, your DR replication partner (the array) maintains your data for immediate availability.

The array uses the VMware-provided interfaces to talk to SRM using TCP/IP.

The following block diagram shows the communication path.



When you are working with HPE Storage SRA, keep the following points in mind:

- SRA is case-sensitive. If you manually create an initiator group, the SRA IQN (iSCSI qualified name) must match both the name and the case of the host IQN. For example, if the host IQN uses upper-case letters, but you create the SRA IQN using lower-case letters, you will create a new initiator group instead of accessing the host initiator group.
- Do not use wildcards to configure an initiator group IQN that is used as part of an SRM configuration.
- SRA only supports the replication type **Periodic Snapshot**.

For details about using SRM, see the VMware documentation.

Overview of SRA Setup Process

Setting up the HPE Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) involves performing steps on the SRM Server, the vCenter Server, the ESXi host, and the HPE Storage array. Detailed information about these tasks is available in the following documentation:

- SRM. See the VMware SRM documentation.
- Array. See the Administration Guides that are available on HPE InfoSight in the [Documentation](#) section.
- vCenter Server and ESXi. See the VMware documentation for these products.
- SRA. Information about working with SRA is provided in this document.



The steps that follow provide a high-level overview of the actions you must perform and where you must perform them so that you can plan your installation and setup. These steps are a checklist of the actions you must take and do **not** contain all the details necessary to perform the steps. See the appropriate documentation for detailed instructions.

- 1 (SRM Server)** Install SRM. You must install SRM before you can install SRA.
- 2 (SRM Server)** Install SRA.
- 3 (array OS)** Set up replication between the arrays. The upstream (protected site) array is the source and the downstream (recovery site) array is the destination.
 - a** Set up at least one array in a group for the upstream site and another array in a group for the downstream site.
 - b** Use the wizard to pair the two groups as replication partners in array OS.
- 4 (Array OS)** Create and configure initiator groups on the upstream site. You do not need to create one on the downstream site unless that array has production VMs that must be included in an SRM workflow or you want to limit access to a placeholder volume.

Points to keep in mind when working with initiator groups:

- In a disaster recovery scenario, any initiator groups associated with the downstream/recovery partner will use a different IQN than the upstream array used. You must be aware that the replication partner uses a different IQN because that can affect SAN booting and datastore access and might require that you remove previous static IQNs.
 - You must configure an initiator group even if you are using Unrestricted Access. Otherwise SRM will not be able to discover the volumes.
 - Do not use wildcards to configure an initiator group IQN that is used as part of an SRM configuration. Using a wildcard or leaving the IQN field blank will result in an error. An initiator group must include an IQN for each resource mapping relationship
 - SRA is case-sensitive. If you manually create an initiator group, the SRA IQN (iSCSI qualified name) must match both the name and the case of the host IQN. For example, if the host IQN uses upper-case letters, but you create the SRA IQN using lower-case letters, you will create a new initiator group instead of accessing the host initiator group.
- 5 (Array OS)** The following volumes must exist:
 - A datastore at the upstream production site to store the VMs you have set up to be protected.
 - A placeholder datastore at both the upstream and downstream site to hold the placeholder VMs prior to a recovery operation but after a protection group is created.
 - 6 (Array OS)** Set up one or more volume collections on the upstream array. If the downstream array is hosting production VMs, you must make sure all volume collections are replicated there.
 - 7 (ESXi)** Make sure the protocol you are using is set up correctly:
 - FC environment. You must have proper zoning in place so that a rescan of the HBAs will allow SRM to discover the recovered volume.
 - iSCSI environment. You should already have the discovery IP address set up and presented to the ESXi hosts from the upstream array.

To verify that the discovery IP address is set up, use the vCenter Server wizard: **VMware vSphere Web Client > Hosts and Clusters > <Your_ESXi_Host> > Configure > Storage Adapters > Software iSCSI Adapter > Targets (under Adapter Details) > Dynamic Discovery**

For more information about working with iSCSI and VMware, see [VMware iSCSI Configuration](#) on page 50.

- 8 (ESXi)** Confirm that the datastore volume is mounted on the ESXi host(s).
- 9 (SRM Server)** Configure the Array Manager to enable array-based replication on each site in the SRM using the group name/management IP address. You add arrays at the Array Group level. If you have multiple arrays in a group, you only need to add the group once.
- 10 (SRM Server)** If you did not pair the arrays using the Array Manager wizard, you can pair them now using the wizard: **VMware vSphere Web Client > Home > Site Recovery > Array based Replication > <Select an Array> > Manage > Array Pairs**

Select the array pair from the list and click **Enable Array Pair**.



- 11 (SRM Server)** Pair the vCenter Server sites. You can use the wizard to do this: **VMware vSphere Web Client** > **Home** > **Site Recovery** > **Sites** > **<local Site>** > **Pair Site**
- 12 (SRM Server)** SRM has a built-in set of mappings to control the networks, folders, hosts or clusters, and storage policies that the VMs will use as well as the datastores that are used for placeholder files. You can use SRM to create reverse relationships automatically. You perform these actions by going to **VMware vSphere Web Client** > **Home** > **Site Recovery** > **Sites** > **<Select Site>** > **Manage**

For more information about working with these settings, see the SRM documentation.

SRA for SRM Prerequisites

While different versions of HPE Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) have different requirements and prerequisites, there are some shared prerequisites:

- SRM must be installed before you install SRA.
- You must have administrator privileges on the SRM server.
- Both the upstream and downstream replication sites must be running the same version of vCenter Server.

Note: You can use either the vCenter Server Appliance (Photon OS) or the Standard vCenter Server on Windows. You must install the same version of SRA on both sites.

- Both sites must be able to communicate with both SRM and the HPE Storage arrays.
- All virtual machines (VMs) affected by this must be mounted to ESXi hosts.
- Protocol requirements:
 - FC environments: Zoning must be set up so that a rescan of the HBAs allows SRM to discover the recovered volumes.
 - iSCSI environments: The recovery ESXi hosts must be configured for discovery IP for destination array.
- SRA is case-sensitive. If you manually create an initiator group, the SRA IQN (iSCSI qualified name) must match both the name and the case of the host IQN. For example, if the host IQN uses upper-case letters, but you create the SRA IQN using lower-case letters, you will create a new initiator group instead of accessing the host initiator group.
- SRA only supports the replication type **Periodic Snapshot**.
- (Recommended) Install HPE Storage Connection Manager for VMware on each ESXi host that has a relationship with an HPE Storage volume.

For specific information about SRA requirements for different SRA versions, see the *Validated Configuration Matrix*, which is posted on [HPE InfoSight](#) (**Resources** > **Validated Configuration Matrix** > **Support for: SRA for SRM**).

For information about SRM, see the VMware documentation for Site Recovery Manager.

Download the SRA for SRM Installation Package

You can download the HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM software from either HPE InfoSight or from VMware. The steps that follow describe how to download SRA from InfoSight.

After you download the file, place it in a local directory on the host. SRA is packaged as a .zip file that contains a .exe file for installing on hosts using Windows and a .tar.gz file for hosts using Photon OS.

Procedure

1. Go to InfoSight, which is located at <https://infosight.hpe.com/>.
2. Enter your login ID and password and click **Login**.
3. Select **Resources** > **Software Downloads**.
4. In the Integration Kits section on the left side of the page, select **HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM**.

The installation package has a name similar to "HPE Nimble-SRA-x.x.x.x.zip or HPE Alletra 6000 and Nimble Storage SRA-x.x.x.x.zip"

where x.x.x.x is the SRA version number.

5. Extract the file appropriate for your installation:

- Windows: HPE Nimble-SRA-x.x.x.x.exe installer file
- Photon OS: HPE Nimble-SRA-x.x.x.x.tar.gz installer file

Install SRA for SRM for Windows

Before you begin

Before you install and set up HPE Nimble Storage Replication Adapter for VMware SRM on a Windows host, you must have:

- Met the prerequisites for SRA for SRM.
See [SRA for SRM Prerequisites](#) on page 114.
- Made sure you are running a version of SRM that is compatible with the version of SRA you plan to install.
- Confirmed that you will install the same version of SRA on both the upstream (protected) and the downstream (recovery) sites. This version of SRA must be compatible with the currently installed versions of SRM and array OS.
- Ensured that both sites allow communication between the arrays for the purpose of replication.
- Obtained the SRA for SRM installation file from the InfoSight portal and extracted the SRA installer file.
See [Download the SRA for SRM Installation Package](#) on page 114.

Procedure

- 1.** Make sure you have administrator privileges and log into the SRM server.
- 2.** Right-click the SRA installer file and select **Run as administrator**.
This starts the SRA InstallShield Wizard.
- 3.** Follow the steps in the wizard.
These steps are self-explanatory.
You must install SRA on both the upstream and downstream sites.
- 4.** When the wizard completes, log into SRM to refresh the newly installed instances of SRA.
Select the following:

vCenter Home > Site Recovery > Sites > <your_SRA_site> > Monitor > SRAs > Refresh Button

SRA alerts you if it detects a compatible SRA installation at the remote site. If the sites are paired and each one has SRA installed, you might have to go to the remote site and select **Refresh** before the status update occurs.

Install SRA for SRM for Photon OS SRM or VAs

HPE Nimble Storage Replication Adapter and HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM for VMware Photon OS or SRM virtual appliances (VAs) require that you have deployed the SRM 8.2 or above.

The main steps for installing SRA in this environment are:

- 1** Use the **Deploy OVF Template** wizard to deploy SRM on the ESXi host.
- 2** Use the **Configure Site Recovery Manager** wizard included in the **SRM Config Service** to provide details about the vCenter and the SRM server and to register the VCSA with the SRM configuration.



- 3 Use the **Storage Replication Adapter** wizard in **SRM Config Service** to add the HPE Storage SRA for VMware Photon OS or SRM VAs file.

The following steps provide details about how to perform these tasks.

Before you begin

You must have the following:

- vServer Appliance version 6.5 or later
- SRM 8.2 or above

Procedure

1. Deploy the SRM VA on the ESXi host using the **Deploy OVF Template** wizard .
 - 1 Select **ESXi host** > **Deploy OVF Template**.
 - 2 Click **Browse** to select the local file and upload the SRM VA. Then click **Next**.
 - 3 Enter the name of the SRM virtual machine (VM) and select the datacenter. Click **Next**.
 - 4 Select the ESXi host and click **Next**.
 - 5 Review the SRM VA details and click **Next**.
 - 6 Accept the End User License's Agreement (EULA) and click **Next**.
 - 7 Specify **No** for CPU for configuration and click **Next**.
 - 8 Select the datastore and click **Next**.
 - 9 Specify the network and IP allocation setting. Click **Next**.
 - 10 Check the box to enable SSHD. Enter the administrator and root passwords. Click **Next**.
 - 11 Review you settings. Click **Finish** if they are correct.
2. Register the VCSA with the SRM configuration by using the **Configure Site Recovery Manager** wizard to enter the details about the vCenter and the SRM server.
 - 1 Log in to the **SRM Config Service**. For example, you might enter
HTTPS://<SRM_IP_Address>: 5480/drconfig
 - 2 Click the **Configure Appliance** button in the SRM Summary page.
 - 3 Follow the steps in the **Configure Site Recovery Manager** dialog wizard to enter the details about the vCenter.
 - 4 At the **Connect**.
 - 5 Select the vCenter Server to register it with the SRM server.
 - 6 When the **Security Alert** pop-up appears again, click pop-up, click **Connect**.
 - 7 Follow the steps in the wizard to enter the SRM details.
 - 8 Review the information and, if it is correct, click **Finish**. This registers the vCenter with the SRM.
3. Select the **Storage Replication Adapter** menu from the **SRM Config Service** interface to upload the HPE Storage SRA .tar file.
 - 1 Select the **Storage Replication Adapter** menu.
 - 2 Click the **New Adapter** button.
 - 3 In the pop-up box, click the **Upload** button.
 - 4 Select the SRA file, which uses the format `HPE Alletra 6000 and Nimble Storage SRA-x.x.x.x.tar`, where `x.x.x.x` is the build number for that version of SRA.
 - 5 Verify that the HPE Storage SRA information has been added to the SRM Config Service page.
 - 6 Use the command **sudo docker images** command to verify that the HPE Storage SRA has been loaded into the SRM 8.2. and above docker container.



Update SRA for SRM for Windows

Because HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM does not support Windows SRM, when you update HPE Nimble Storage Replication Adapter on a Windows host, a full version of SRA is installed on top of an existing installation. You can install the update without un-installing the current version of SRA.

Before you begin

To update HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM, you must have the downloaded HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM installation package.

See [Download the SRA for SRM Installation Package](#) on page 114.

Procedure

1. Check to make sure that no SRM operation is running.
2. Right-click the HPE Storage SRA installer file and select **Run as administrator**. This starts the HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM InstallShield Wizard. .
3. From the Windows Services, highlight *VMware vCenter Site Recovery Manager Server*, and click **Restart**.
4. Verify that the update was successful.
 - a) Make sure that there are no errors displayed in the existing Array Manager.
 - b) In SRA, click **Edit Array Manager**.
 - c) Make sure that only a Local Array is specified in the connection specification.
Previously, SRA required both Local Array and Remote Array connection specifications. There is no need to re-enter the Username and Password.
 - d) Click **Cancel**.
 - e) In the vSphere Web Client, under the Array Pairs tab, click **Refresh**.
 - f) Make sure that no errors are reported.

Update SRA for SRM for Photon

Before you begin

To update HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM, you must have the downloaded HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM installation package.

See [Download the SRA for SRM Installation Package](#) on page 114.

Procedure

1. If you have SRA for SRM for Photon installed, you must first uninstall it. For instructions, see [Uninstall SRA for SRM for Photon OS SRM or VAs](#)
2. Install SRA for Photon. For instructions, see [Install SRA for SRM for Photon OS SRM or VAs](#)

Uninstall SRA for SRM for Windows

Use this process if you need to uninstall HPE Nimble Storage Replication Adapter for VMware SRM for Windows.

Procedure

1. Log into the ESXi server as the administrator.
2. Right-click the setup icon or use **Control Panel** > **Programs** to find the SRA.



3. Right-click the SRA and select **Remove** from the menu or click **Uninstall** from the Control Panel.
4. Confirm removal if requested.

Uninstall SRA for SRM for Photon OS SRM or VAs

This procedure takes you through the steps to uninstall HPE Nimble Storage Replication Adapter or HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM for VMware Photon OS or SRM virtual appliances (VAs).

Procedure

1. Log in to the **SRM Config Service**. For example, you might enter **HTTPS://<SRM_IP_Address>: 5480/drconfig**
2. From **SRM Config Service**, select **Storage Replication Adapters**.
3. From the drop-down menu for the **HPE Nimble Storage – Storage Replication Adapter**, select **Delete**.
4. In the pop-up box, check both options.
5. Click **Delete**. Information about the Storage Replication Adapter is removed and a Success message appears.
6. Use the **sudo docker images** command to verify that the HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM has been removed from the SRM docker container.

Working with Array Manager in SRM

If you decide to use VMware Site Recovery Manager (SRM), you must configure array-based replication in SRM for each site.

After you connect a protected (upstream) site with the recovery (downstream) site, you must enable the array manager to discover replicated devices, display the replication direction, and trigger storage operations.

Arrays are added at the HPE Storage Array Group level. You only need to add the group once, even if you have multiple arrays in it.

Depending on the version of SRM and your setup, the exact configuration steps might differ slightly from the ones provided in this guide. See the VMware Site Recovery Manager documentation for specific information about SRM.

Here are some best practice tips for adding arrays:

- Use the same name for the for the downstream and upstream folders
- When you set up a replication partner, use the same pool and folder as the source location if that location also exists on the local group.

Detailed information about working with SRM is in the VMware documentation. Information about adding HPE Storage arrays to your version of SRM is available here:

- [Configure SRM 6.0, 6.1, and 6.5 for HPE Storage Arrays](#) on page 118
- [Configure SRM 8.x for Storage Arrays](#) on page 119
- [Configure SRM 8.3 and above for Storage Arrays and vVols](#) on page 121

Configure SRM 6.0, 6.1, and 6.5 for HPE Storage Arrays

These steps tell you how to configure array-based replication for each site when you are using VMware Site Recovery Manager (SRM) 6.0, 6.1, and 6.5.

You must have a pair of array managers. You can either configure them at the same time, as these steps do, or you can configure one array manager and then repeat the steps for the next array manager.

The exact configuration steps might differ slightly depending on your system setup. See the VMware Site Recovery Manager documentation for specific information.

Before you begin

You must have done the following:



- Installed SRM 6.0, 6.1, or 6.5.
- Logged into SRM and refreshed each instance of HPE Storage Replication Adapter (SRA) for SRM before you begin.
- Gotten the IP address of both the upstream array and downstream array. This is the group management IP address of the HPE Storage Array Group you are adding.
- Checked the [Validated Configuration Matrix](#) on HPE InfoSight to confirm that your version of the array OS supports this version of SRM.

Procedure

1. From the **Home** page of the VMware vSphere Web Client, select **Site Recovery**.
2. Select the **Array Based Replication** tab.
3. Click the **Add New Array Manager** icon.
4. Select **Add a Pair of Array Managers**, which lets you create the upstream and downstream array managers at once.

Note: If you choose **Add a Single Array Manager**, you must repeat these steps to add the partner array.

5. When the **Add Array Manager** wizard starts, select the pair of sites for the array managers.
6. Select **Nimble Storage - Storage Replication Adapter** as the value for **SRA type**.

Note: You can only associate one version of SRA with an array.

7. In the **Nimble Storage - Storage Replication Adapter** wizard, enter the information needed for the upstream array. Then repeat the steps and enter the information for the downstream array. You also have the option of entering the administrator credentials for the vCenter. If you plan to use VMware synchronized snapshots, it is a good practice to enter the vCenter credentials.

Upstream array:

- 1 **Display Name:** Enter the upstream array name.
- 2 **IP Address of Local Array:** Enter the upstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the upstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

Downstream array:

- 1 **Display Name:** Enter the downstream array name.
- 2 **IP Address of Local Array:** Enter the downstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the downstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

8. Enable the array pair by selecting the paired arrays.
9. Review your configuration. If it is correct, select **Finish**.
You must complete these steps for each Array Group in the SRM configuration.

Configure SRM 8.x for Storage Arrays

These steps tell you how to configure array-based replication for each site when you are using VMware Site Recovery Manager (SRM) 8.1 or above.

You must have a pair of array managers. You can either configure them at the same time, as these steps do, or you can configure one array manager and then repeat the steps for the next array manager.

The exact configuration steps might differ slightly depending on your system setup. See the VMware Site Recovery Manager documentation for specific information.



Before you begin

You must have done the following:

- Installed SRM 8.1 or above.
- Logged into SRM and refreshed each instance of HPE Alletra 6000 and Nimble Storage Replication Adapter for VMware SRM before you begin.
- Gotten the IP address of both the upstream array and downstream array. This is the group management IP address of the HPE Storage Array Group you are adding.
- Checked the [Validated Configuration Matrix](#) on HPE InfoSight to confirm that your version of the array OS supports SRM 8.1 or above.

Procedure

1. From the **Home** page of the VMware vSphere Web Client, select **Site Recovery** > **Open Site Recovery**. The SRM GUI opens in a separate, HTML5 client page.
2. Select **Site Pair**.
3. Under **Array Based Replication**, select **Array Pairs**.
4. Click **ADD**.
5. Select **Storage Replication Adapter**.
6. In the **Configure Array Manager** wizard, enter the information needed for the upstream array. Then repeat the steps and enter the information for the downstream array. You also have the option of entering the administrator credentials for the vCenter. If you plan to use VMware synchronized snapshots, it is a good practice to enter the vCenter credentials.

Upstream array:

- 1 **Display Name:** Enter the upstream array name.
- 2 **IP Address of Local Array:** Enter the upstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the upstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

Downstream array:

- 1 **Display Name:** Enter the downstream array name.
- 2 **IP Address of Local Array:** Enter the downstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the downstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

7. Enable the array pair by selecting the paired arrays the **Array Pairs** wizard.

8. Review your configuration. If it is correct, select **Finish**.

You must complete these steps for each Array Group in the SRM configuration.

Using SRM to Restore vVols

VMware Site Recovery Manager (SRM) 8.3 or above integrates with VMware Virtual Volumes (vVols) and allows you to use array-based replication. It provides protection granularity to the level of a single virtual volume (VM). If multiple VMs are replicated together, the resulting replication group is maintained as a consistency group.

SRM 8.3 and above works with VASA Provider instead of Storage Replication Adapter (SRA), which helps ensure smooth vVol operations and maintenance.

Configuring SRM 8.3 and above to work with vVols involves performing the following basic steps from the SRM GUI:

- 1 Pairing both sites. See [Configure SRM 8.3 and above for Storage Arrays and vVols](#) on page 121.
- 2 Selecting the storage profile for the Storage Policy Mappings. See [Set the Storage Policy Mappings to Use SRM 8.3 or above with vVols](#) on page 122.



- 3 Adding a protection group. See [Create Protection Groups for vVols](#) on page 122.
- 4 If you use array-based protection groups or vSphere replication protection groups, add a placeholder datastore. Do not do this if you are using only storage policy protection groups. See [Add Placeholder Datastores](#) on page 123.

The exact configuration steps might differ slightly depending on your system setup. See the VMware Site Recovery Manager documentation on the VMware website for specific information.

Configure SRM 8.3 and above for Storage Arrays and vVols

Setting up VMware Site Recovery Manager (SRM) 8.3 and above requires that you pair two arrays using the SRM GUI.

Note: The exact configuration steps might differ slightly from the steps that follow depending on your system setup. See the VMware Site Recovery Manager documentation on the VMware website for specific information about configuring SRM.

Before you begin

You must have done the following:

- Installed SRM 8.3 or above.
- Gotten the IP address of both the upstream array and downstream array. This is the group management IP address of the HPE Storage Array Group you are adding.
- Checked the [Validated Configuration Matrix](#) on HPE InfoSight to confirm that your version of array OS supports SRM 8.3 or above.

Procedure

1. From the **Home** page of the VMware vSphere Web Client, select **Site Recovery** > **Configure**. The SRM GUI opens in a separate, HTML5 client page.
2. Select **Site Pair**. You must pair both sites in the SRM GUI using the sites' vCenter credentials.
3. Select **Array Based Replication** > **Array Pairs**.
4. Click **ADD**.
5. Select **Storage Replication Adapter**.
6. In the **Configure Array Manager** wizard, enter the information needed for the upstream array. Then repeat the steps and enter the information for the downstream array. You also have the option of entering the administrator credentials for the vCenter. If you plan to use VMware synchronized snapshots, it is a good practice to enter the vCenter credentials.

Upstream array:

- 1 **Display Name:** Enter the upstream array name.
- 2 **IP Address of Local Array:** Enter the upstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the upstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

Downstream array:

- 1 **Display Name:** Enter the downstream array name.
- 2 **IP Address of Local Array:** Enter the downstream array management IP address.
- 3 **Username** and password **Password:** Enter the array administrator credentials.
- 4 (Optional) **IP Address of Local vCenter:** Enter the IP address of the vCenter associated with the downstream array.
- 5 (Optional) **Username** and password **Password:** Enter the vCenter administrator credentials.

7. Enable the array pair by selecting the paired arrays from the **Array Pairs** wizard.
8. Review your configuration. If it is correct, select **Finish**.



What to do next

To complete the SRM setup for a VMware Virtual Volumes (vVols), you must map storage policies to the primary and secondary vCenters and create protection groups for vVols. See [Set the Storage Policy Mappings to Use SRM 8.3 or above with vVols](#) on page 122 and [Create Protection Groups for vVols](#) on page 122. If you are using array-based protection groups or vSphere replication protection groups, add a placeholder datastore. See [Add Placeholder Datastores](#) on page 123.

Set the Storage Policy Mappings to Use SRM 8.3 or above with vVols

To use VMware Site Recovery Manager (SRM) 8.3 or above with VMware Virtual Volumes (vVols) you must set the storage policy mapping from the SRM GUI.

Note: The exact configuration steps might differ slightly from the steps that follow depending on your system setup. See the VMware Site Recovery Manager documentation on the VMware website for specific information about configuring SRM.

Procedure

1. From the **Home** page of the VMware vSphere Web Client, select **Site Recovery** > **Open Site Recovery** > **Configure**. The SRM GUI opens in a separate, HTML5 client page.
2. Select **Storage Policy Mappings** > **New**.
3. Select the box next to the storage profile you want to use.
4. Select **Creation Mode** to specify how you want to prepare the storage policies.

Select:

- Automatically if you want to have the system prepare the mappings for storage policies with matching names under the selected policy container.
 - Manually if you want to select the policies yourself.
5. Select **Recovery storage policies** and highlight the same policy in both columns. This is the policy that will be used with both the primary and secondary vCenter, so the name must be the same.
 6. Select **Reverse mappings** and choose the policy that you want to use to automatically create the reverse mappings.

Note: This profile mapping might overwrite existing mappings.

7. Review your configuration. If it is correct, select **Finish**.
You must complete these steps for each Array Group in the SRM configuration.

What to do next

To complete your SRM setup, you must create a protection group and add it to a recovery plan. See [Create Protection Groups for vVols](#) on page 122. If you are using array-based protection groups or vSphere replication protection groups, you also need to add a placeholder datastore. See [Add Placeholder Datastores](#) on page 123.

Create Protection Groups for vVols

To use VMware Site Recovery Manager (SRM) 8.3 with VMware Virtual Volumes (vVols) you must create protection groups for the VMware Virtual Volumes. You do this from the SRM GUI.

Note: The exact configuration steps might differ slightly from the steps that follow depending on your system setup. See the VMware Site Recovery Manager documentation on the VMware website for specific information about configuring SRM.

Procedure

1. From the **Home** page of the VMware vSphere Web Client, select **Site Recovery** > **Open Site Recovery** > **Configure**. The SRM GUI opens in a separate, HTML5 client page.
2. From the menu, select **Protection Groups** > **New**.

3. Enter the information about the protection group.

On this page, you must supply the following:

- The protection group name
- The direction for the protection group:
 - Upstream: Select **site1-srm** > **sit2-srm**.
 - Downstream: Select **sit2-srm** > **site1-srm**.

4. Select **Type** and choose the following options:

- **Virtual Volumes (vVol replication)**
- The group-wide fault domain

5. Select **Replication groups** and choose the replication groups you want to use with this protection group. The replication groups contain the virtual machines (VMs) that are recovered together.

Note: Protected VMs and their Storage policies needs to have replication enabled and the vendor type to be set "NimbleStorage.Replication" before completing this step.

6. Select **Recovery plan**. You must specify whether to add the new protection group to an existing recovery plan, a new recovery plan, or to not add it to either one.**7.** Review your configuration. If it is correct, select **Finish**.**Add Placeholder Datastores**

If you use array-based protection groups or vSphere replication protection groups, you must specify a placeholder datastore when you use VMware Site Recovery Manager (SRM) 8.2 or above.

Note: If you are using only storage policy protection groups, you do not need to set up a placeholder datastore. SRM does not create placeholder VMs for storage policy protection groups.

When you select this option, SRM creates the placeholder virtual machine (VM) on the downstream site for each VM in the protection group. The placeholder VM contains small files that are not full copies of the protected VM. When you use array-based replication to protect datastore groups, you must identify the datastore on the downstream site. SRM stores the placeholder virtual machine files in this datastore.

Some points to keep in mind when you are considering using placeholder datastores:

- You must configure a placeholder datastore on both the upstream and downstream sites to establish bidirectional protection.
- If you are using clusters, make sure the placeholder datastores are visible to all hosts in the cluster.
- When choosing a placeholder datastore, do not select a datastore that is replicated using array-based replication.
- When you are using HPE replication with VMware Virtual Volumes (vVols) and want to create a placeholder datastore, you must use a local VMFS datastore for the placeholder. Do not use a vVol datastore.

Note: The exact configuration steps might differ slightly from the steps that follow depending on your system setup. See the VMware Site Recovery Manager documentation on the VMware website for specific information about configuring SRM.

Procedure

- 1.** To set up a placeholder datastore, mount VMFS datastores to the upstream and downstream hosts.
- 2.** From the **Home** page of the VMware vSphere Web Client, select **Site Recovery** > **Open Site Recovery** > **Configure**. The SRM GUI opens in a separate, HTML5 client page.
- 3.** Pair the two sites. For information on doing this, see [Configure SRM 8.3 and above for Storage Arrays and vVols](#) on page 121
- 4.** Select **Placeholder DatastoresConfigure** > **Placeholder Datastores** > **New**. The Placeholder Datastores window displays the IP addresses for the upstream array and the downstream array.

5. Select the upstream array IP address and place a check next to the shared VMFS datastore you want to designate as its placeholder datastore.
6. Select the downstream array IP address and place a check next to the shared VMFS datastore you want to designate as its placeholder datastore.

Initiate a Recovery Plan

If your upstream (protected) site goes down, you must initiate a recovery to ensure business continuity. That action requires that you have a recovery plan, preferably one created on VMware Site Recovery Manager (SRM).

If a site failure occurs, you can make data in your volume collections available to applications from the downstream array while you restore the failed array.

In SRM, click the Recovery Plans tab to begin creating your recovery plan.

For more information, refer to the VMware Site Recovery Manager documentation..

Test the Recovery Plan

It is a good practice to test the recovery plan for HPE Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM).

Refer to the VMware Site Recovery Manager documentation to:

- Configure protection groups on the protected site
- Create a recovery plan at the recovery site

Note: SRA for SRM must be installed on the downstream array against your vCenter Server.

Procedure

To test your recovery plan:

1. Connect to SRM through the vCenter Web Client.
2. Run the SRM Test Recovery Plan in test mode.
3. To end the test, click **Cleanup**.

Results

SRA for SRM performs the necessary cleanup after the test recovery has been performed.

Using SRM Test Failover Workflows and Replication

A test failover is a VMware Site Recovery Manager (SRM) workflow that "tests" the recovery plans for a configuration by simulating a real recovery operation. If you have replication enabled, it continues to work while the Test Failover is in place.

The test failover workflows allow you to understand how HPE Replication Adapter (SRA) for SRM handles a test failover when replication is used and then later if replication is not used.

Note: Test failover operations work with the test objects; that is, the objects that were created for that test. If no test objects exist, the operations use existing objects.

When you perform a test failover in an environment that uses replication, SRA creates a snapshot that has the prefix **nimble4sra2sync-**. The cleanup operation does not remove this snapshot.

As a result, if you then perform a test failover without replication, the cloned volume uses the **nimble4sra2sync-** snapshot because it is the last snapshot that SRA created. The operation does not use a volume collection snapshot from the schedule.



Tasks Performed During a Test Failover Operation

The workflow for a test failover operation is the same as the workflow for a recovery operation, but the storage operations and API calls are slightly different. Also, HPE Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRM) uses test objects when performing test operations. SRA only uses existing objects if there are no test objects..

Note: HPE recommends that you always perform a test failover operation before you perform a recovery operation. Keep in mind, though, that a recovery operation can fail even if the test failover succeeds.

During a test failover the following actions occur:

- 1 When **replicate recent changes** is selected, SRA triggers a replication update from the protected side by creating a snapshot with the prefix **nimble4sra2sync-**.
- 2 SRA attempts to match an existing initiator group on the recovery site to the access group requested by SRM. If SRA cannot locate a matching initiator group, it creates a new one with the prefix **sra-**.
- 3 SRA clones the most recent SRA-created snapshots on each of the volumes on the recovery site that are involved in the test failover workflow.
- 4 SRM mounts the cloned volumes, registers the virtual machines (VMs) in place of the placeholder VMs, modifies the configuration as needed (for example, the network), and turns the VMs on.

Tasks Performed During a Cleanup Operation

The cleanup operation runs after a test failover operation completes, even if the failover operation failed. The goal of the cleanup workflow is to restore the recovery site to its original state.

The following actions take place during a cleanup operation:

- Virtual machines (VMs) on the recovery site are powered off and the vCenter unmounts the cloned datastore.
- The cleanup operation deletes the cloned volumes on the recovery site, but it does not delete the snapshot SRA created for this test failover operation. The snapshot has the prefix **nimble4sra2sync-**. You must manually delete the snapshot.

Note: If you do not delete the snapshot that SRA creates during the test failover, you might encounter an unmanaged snapshot situation. The next test failover that you perform when you do not have replication enabled will use the **nimble4sra2sync-** snapshot that SRA created during an earlier test failover.

Implement the Recovery Plan

The VMware Site Recovery Manager (SRM) lets you respond quickly when a site goes down.

After a site failure, you can make data in the volume collections available to applications from the downstream array while you restore the failed array. Once the failed array is restored or replaced, you can return data I/O to the original partner and restore the relationship with downstream partner at the DR site.

SRM 5.0 introduced the *Reprotect* operation to configure protection in the reverse direction automatically. Reprotect helps in preparation for failback to the primary site.

Before you begin

You must have already created a recovery plan (see [Initiate a Recovery Plan](#)) in SRM.

Procedure

To implement the recovery plan:

1. Connect to SRM through the vCenter Web client.
2. Use the SRM Run Recovery Plan feature.
3. When the primary site is resolved, create and test a new recovery plan in SRM.

In this case, failing over is equivalent to failing back.



SRA and Microsoft Volume Shadow Service

You can configure Storage Replication Adapter (SRA) to work with Microsoft Volume Shadow Service (VSS).

VSS enables application-consistent snapshots and recovery of Microsoft Exchange and SQL Database servers. You can connect the virtual storage on these applications using Raw Device Mapping (RDM), VMDK or In-guest Attached iSCSI. You must set up the option you use with SRA.



VAAI Integration

The storage array lets you take advantage of VMware's vStorage APIs for Array Integration (VAAI). VAAI reduces the time that it takes to provision new VMs. This makes it easier for you to have large-scale VMware deployments.

VAAI is automatically installed.

What is VAAI?

vStorage APIs for Array Integration (VAAI) is a set of features that provide hardware acceleration. VAAI enables the ESXi host to offload VM and storage management operations to the storage array. As a result, these operations are performed faster while consuming fewer resources.

The following VAAI primitives are supported:

- Atomic Test and Set Locking (ATS) – The ATS feature is used for file locking to control access to datastores by multiple ESXi hosts.
- Zero Blocks/Write Same – This feature is used to zero-out large numbers of blocks on VMDKs (both thick and thin) at the same time.
- Block Delete/SCSI UNMAP – The UNMAP feature is designed to efficiently reclaim any deleted space to meet continuing storage needs.



Important: On servers running ESXi 5.0u1 and later, the VAAI UNMAP primitive is disabled by default. You must enable it manually. See the ESXi server documentation for more information.

- Thin Provisioning Stun – The Thin Provisioning Stun primitive provides a mechanism by which the array is able to return warnings to the hypervisor when space thresholds are exceeded.
- Extended Copy (XCOPY) – The Extended Copy primitive requests the array to perform a full copy of blocks and is used primarily in clone and migration operations. Support for this primitive began with release 3.x.

Note: XCOPY is not supported if the destination volume has synchronous replication enabled.

[VMware KB article 1021976](#) contains more information about VAAI.

VAAI Requirements

To use vStorage APIs for Array Integration (VAAI), you must have the following:

- A vSphere Standard, Enterprise, Enterprise Plus, or Remote Office Branch Office licensing for ESXi hosts
- Storage arrays that support VAAI storage-based hardware acceleration

The following settings control the basic VAAI operations:

Advanced Parameter Name	Description
HardwareAcceleratedLocking	Atomic Test & Set (ATS). Used during the creation of files on the VMFS volume
HardwareAcceleratedMove	Clone Blocks/Full Copy/XCOP. Used to copy data
HardwareAcceleratedInit	Zero Blocks/Write Sam. Used to zero-out disk regions

These options are enabled by default.



Enable the VMware VAAI Provider to Use Storage Volumes

When you use the array OS with VMware vCenter servers, you can take advantage of the vStorage APIs for Array Integration (VAAI) "write same" (also called block zeroing) primitive. Everything necessary to enable the VAAI write-same feature with storage volumes is included with the array OS. No installation is required.

If you are using ESXi 5.0 or later, VAAI is enabled by default. After triggering a write-same operation, VAAI shows status as *supported*; for example, the "Zero Status" line appears as "supported" in the following output:

```
~ # esxcli storage core device vaa1 status get --device
eui.0844019010b05c836c9ce9003e350c78
eui.0844019010b05c836c9ce9003e350c78
VAAI Plugin Name:
ATS Status: supported
Clone Status: supported
Zero Status: supported
Delete Status: supported
```



Using InfoSight Virtualization Data to Evaluate Performance

You can view detailed information about virtual machines (VMs) hosted on storage arrays from HPE InfoSight by going to the VMware Virtualization pages. They are located under **Infrastructure** > **Virtualization:VMware**. You can use the analytics found on these pages to monitor latency and performance across the storage, network, and host stack. This information allows you to address potential issues and improve performance. For example, the data can show you where there are CPU bottlenecks and which hosts have the most CPU headroom. That way you can move a VM to a host that has more resources.

Note: The InfoSight Virtualization section has been expanded to include information that was previously included in the VMVision feature. As a result VMVision is no longer a separate menu option.

The Virtualization pages include detailed analytics about datacenters, clusters, ESXi hosts, datastores, and VMs. The analytics provide data about memory, latency, IOPS, and MB/s.

You can use the analytics to identify potential issues, such as:

- Noisy neighbor VMs
- Overloaded hosts (CPU and memory)
- Latent VMs by datastore
- Inactive VMs (an inactive VM is one that has not has any I/O activity during the last seven days)
- Location of the storage volumes on VMware datastores

To collect this data, InfoSight uses the storage array and the VMware vSphere APIs. You do not need to install anything on the vCenter Server.

You can quickly set up InfoSight to gather information about the storage array and VMware. From the array OS, you must register the HPE vCenter Plugin or enable a VMware sychronized volume collection. From InfoSight, you must enable both the Streaming button for the array you want monitor and the VMware button. For detailed information about enabling this information, see [Enable InfoSight to Collect VMware Information](#) on page 129.

InfoSight can take up to 48 hours to process the registration and report data. After that, depending on the data, InfoSight refreshes the data as frequently as every five minutes.

Note: InfoSight only provides analytics about VMs and datastores that are hosted on a storage array.

Enable InfoSight to Collect VMware Information

You can enable HPE InfoSight to provide VMware data for your system by registering either the vCenter Server used with the HPE vCenter Plugin or a VMware synchronized snapshot volume collection and then enabling InfoSight to collect data from the array and VMware. You perform steps in the array OS GUI and in InfoSight. You do not need to install anything on the vCenter Server.

Before you begin

- Your array environment must be set up to work with VMware.
- At least one VM must be on a datastore.
- You must have an InfoSight account.
- Ports 80 and 443 must be open between the vCenter Server and the array.

Procedure

1. Configure the vCenter Plugin on one or more arrays.
 - a) Log into the array.



- b) From the array OS GUI, select **Administration** > **VMware Integration**.
- c) Register the vCenter Server by entering the vCenter Server name and login credentials.

Note: If there is already a registered vCenter Server, select **Add Another vCenter** to add your vCenter Server.

Optional: If you do not want to use the HPE vCenter Plugin, you can enable a VMware synchronized snapshot volume collection and register it instead of a vCenter Server.

After you register the vCenter Server, it is a good practice to verify that the array can communicate with vCenter Server.

2. Configure InfoSight to access the VMware data.

- a) Log into InfoSight.
- b) From the settings menu (the gear icon) select **Telemetry Settings**.
- c) Locate the array you want to monitor and click the button for **Steaming to On**.

This button enables data streaming from the array.

- d) In the same row, click the button for **VMware to On**.

This button allows data to be collected from VMware.

Note: If you do not set the **VMware** button to **On**, you will only receive streaming data from the array. You will not receive any VMware data.

What to do next

It takes up to 48 hours for InfoSight to process the vCenter registration and start streaming VMware and array data.

When the data streaming begins, you can view the analytics for the different objects by choosing **Infrastructure** > **Virtualization** > **VMware** > **[Datacenters | Clusters | ESXi Hosts | Datastores | VMs]**.

Using the Virtualization Dashboard to View VMware Data

The HPE InfoSight dashboard for the Virtualization section provides analytics about your virtual environment. This section appears when you log into InfoSight and choose **Infrastructure** > **Virtualization**. You can then select the objects you want information about from the **Datacenters | Clusters | ESXi Hosts | Datastores | VMs** tabs. The main display window changes to provide information about the objects you selected.

Note: It can take up to 48 hours after you enable VMware streaming for the data to start appearing. After that InfoSight updates the VMware analytics every five minutes.

The menu bar above the main data window allows you to tailor the display to specific information. Use the **View** menu to select the display you want for an object. InfoSight provides analytics for the entire stack as well as individual objects.

The Virtualization display uses tables, graphs, and pop-up boxes to provide information. In some cases, you can specify custom date and time ranges for the information being displayed.

Some of the graphics display additional information when you hover your cursor over a spot on a graph or the name of an object.

Note: For more information about using InfoSight, see the *InfoSight User Guide for HPE Nimble Storage*.

Key Data Provided by the Virtualization Section of InfoSight

The Virtualization section of HPE InfoSight provides both cross-stack analytics and drill-down analytics. You can look at performance and latency information across the storage, network, and host stack or for a specific object. You can use this data to determine how your system is working and whether there are issues within the stack that could be improved. For example, if there is resource contention on one host, you might see improvement if you move a virtual machine (VM) to a different host.



Some of the key data includes details about VMs. You can use the VM analytics to help you locate potential performance problems and devise plans to address these problems. When you select the **VMs** tab under Virtualization, you can then use the **View** menu to display the following categories of data:

- VM Capacity Trend
- VM I/O Contention Treemap
- VM I/O Contention Trend
- VM Memory Contention Trend
- VM CPU contention Trend.

You can view the analytics provided by these options for all the VMs or a specific VM. For example, you can check performance across a datastore by looking at the VM I/O Contention Treemap or the VM I/O Contention Trend. The VM Memory Contention Trend data provides you with insights into memory issues.

When you select the VM host, you can look at the CPU usage for that host. You can also review the top VMs by CPU usage and the top VMs by CPU ready.

The VM Capacity analytics are based on data retrieved from VMware. The Performance information combines data from both VMware and the array, which gives you information about host and network latency

The analytics provided for VMs can help you narrow down why a VM might be performing poorly.

It is also useful to look at a specific vCenter and review data about all the hosts associated with that vCenter. You get to this data by selecting **ESXi Hosts** in the Virtualization section and then choosing **View > Host Performance Trend**. The details provided by Host Performance Trend include analytics about host, network, and storage latency as well as read and write comparisons, IOPS data, and throughput. This information is presented across a range of time to allow you to locate changes.

You can combine the analytics provided in the **VM** and **ESXi Hosts** sections with those provided in the **Datacenters, Clusters,** and **Datastores** sections to get detailed information about the health of your environment.



Common Tasks and Best Practices

This section provides instructions for common tasks and some best practices when integrating HPE Alletra 6000 and VMware.

VMware Partition Alignment

In virtualized file systems, multiple layers of storage are organized into blocks, which makes accessing the storage more efficient. The block size and the starting offset can differ at each layer. While block size may not be an issue across these storage layers, the starting offset is important.

For optimal performance, the starting offset of a file system should align with the start of a block in the next lower layer of storage. For example, an NTFS file system that resides on a LUN should have an offset that is divisible by the block size of the storage group presenting the LUN.

Misalignment of block boundaries at any one of these layers of storage can result in performance degradation.

Misalignment requires the group to read from or write to more blocks than necessary. VMFS partitions that are misaligned and must be manually aligned include certain VMFS partitions and certain Guest OS partitions.

There are no block alignment issues when using Nimble volumes as ESXi/VMware datastores if the datastore is created using the vCenter GUI with default settings. VMware aligns the datastore correctly.

For additional information about alignment requirements, see *KB-000010 Proper Block Alignment*. At the time this document was published, the KB was located at https://infosight.hpe.com/InfoSight/media/kb/active/sup_KB-000010_Proper_Block_Alignment.pdf.

Aligning Guest OS partitions

In the case of the Guest OS alignment on Windows, Windows Server 2008 and Vista partitions are aligned by default at 1024k. VMware does not recommend Windows boot disk alignment.

- **Windows Server 2003/2000/XP data disks**

Use Diskpart to create the disk partition. Be sure to specify a starting offset of 2,048 sectors (1 megabyte), which is the recommended offset and covers most stripe unit size scenarios. For details, see the Microsoft KB article at <http://support.microsoft.com/kb/929491>.

- **Windows Server 2008/Vista data disks** that are in-place upgrades from Windows server 2003/2000/XP

Use Diskpart to create the disk partition. Be sure to specify a starting offset of 2,048 sectors (1 megabyte), which is the recommended offset and covers most stripe unit size scenarios. For details, see the Microsoft KB article at <http://support.microsoft.com/kb/929491>.

- **Linux OS**

For Linux OS, use fdisk to align a partition manually:

- 1 Enter **fdisk -u /dev/sd <x>** where <x> is the device suffix.
- 2 Create the new partition by typing n.
- 3 Create a primary partition by typing p.
- 4 Choose partition 1 by typing 1.
- 5 For the first sector, enter 2048.
- 6 Use default value for the last sector.
- 7 Write the label and partition information to disk by typing w.



Register or Add VM to Inventory

You can register or add VMs to inventory using ESXi host or vCenter Server.

vCenter Server

Procedure

1. Log into vSphere or the VMware client.
2. If connecting to vCenter Server, click on the desired host.
3. Click the **Configuration** tab, then click **Storage**.
4. Right-click the appropriate datastore and select **Browse Datastore**.
5. Navigate to the folder for the virtual machine, and locate the <virtual machine>.vmx file.
6. Right-click the .vmx file and click **Add to inventory**.
7. Complete the steps in the wizard to add the virtual machine to inventory.

ESXi Host

Procedure

1. Use an SSH client to log in as root to the ESXi host.
2. Enter the following command:

```
# vim-cmd solo/registervm /vmfs/volumes/<datastore name>/<VM directory>/<VM name>.vmx
```

Make sure that the Virtual Machine name does not contain any Unicode characters.

Restore a VM from a Datastore

You can use the vCenter plugin to restore a single virtual machine (VM) using a snapshot. When you use this procedure, you replace the current VM information with the information in an existing snapshot.

To perform this operation, you must first create a clone of the datastore containing the VM. You can then move the desired VM to the parent datastore.



Important: When a VM is restored to an earlier snapshot, all I/O that was written after that snapshot was created will be lost.

Before you begin

This procedure applies to VMFS datastores, not VMware Virtual Volumes (vVols).

Procedure

1. Verify that a snapshot containing the data you need exists for the time period that you want to restore.
2. From the VMware vSphere Web Client, select the datastore containing the VM that you want to restore.
3. Select **HPE Alletra 6000 and Nimble Storage Actions > Clone Datastore**.
4. When the Clone Datastore wizard starts, enter the name of the new datastore and select **Use Existing snapshot**. Follow the steps in the wizard.
5. Click **OK**.
6. After the wizard completes, check the **Recent Tasks** section to verify that the cloned datastore is now mounted to the same hosts as the parent datastore.
7. From the VMware vSphere Web Client Inventory Tree, select the cloned datastore and browse the contents.



8. Select the VM folder that you want to restore.
9. Select the VM_Name .vmx file and then select **Register VM**.

Keep the following important points in mind as you perform this step:

- If the parent VM is still in the inventory, you cannot use the same name.
- If the parent VM is still online, you cannot power on the cloned VM because of possible IP conflicts.

10. Power on the VM to verify that it is working.

Note: [Restore a VM from a Datastore on page 133](#)

11. After you verify that you can power on the VM, select **Cloned VM** > **Migrate**. Use Storage vMotion to move the VM back to the parent datastore. If you cannot power on the VM, you must remove the cloned VM from the VMware inventory and use the vCenter Plugin to delete datastore.
12. If you cannot power on, verify that the cloned VM is now on parent datastore.
13. Delete the cloned datastore by selecting it and choosing:
HPE Alletra 6000 and Nimble Storage Actions > **Delete Datastore**

Restore Entire Datastore to an Existing Snapshot

This procedure restores the virtual machines (VMs) on a datastore to the values stored in the snapshot you specify. You can use this procedure to restore the datastore VMs to an earlier point in time without having to use a clone.



Important: When a volume is restored to an earlier snapshot, all I/O that was written after that snapshot was created will be lost.

When you use this procedure, **all** VMs on the datastore will be restored to the snapshot.

Procedure

1. If any VMs are running on the datastore that you are recovering, shut them down and remove them from inventory.
2. For each ESXi host that is using iSCSI, remove all connections from the host to the volume.

Note: If the host is running the Fibre Channel protocol, you do not need to do anything.

ESXi 5.x, 6.x:

You must power off the VMs and remove them from the VMware Inventory Tree.

- a) From the VMware vSphere Web Client, select the VMware vSphere Web Client Inventory Tree.
- b) Select the datastore that you want to restore and choose **Unmount Datastore** > **Select all hosts**.
- c) From the Hosts and Clusters view, select a host and then choose **Configure** > **Storage Adapters**.
- d) **(iSCSI)** Select the iSCSI Software Adapter and click the **Targets** tab at the bottom of the window.
- e) **(iSCSI)** Select the **Static Discovery** tab.
- f) **(iSCSI)** Sort by the **Target Name** column and locate the entries that match the volume name. Highlight all entries for the volume, and then click **Remove**. You can press the **Ctrl +** keys to highlight multiple entries.
- g) **(iSCSI)** Answer **Yes** to confirm removal.
- h) **(iSCSI)** When the removal completes, click **Close**.



Important: When you are prompted to rescan, answer **No. DO NOT RESCAN!** Rescanning picks up the connections you just removed.

3. Move to the GUI dashboard and select **Manage** > **Data Storage**.
4. Click on the appropriate volume. Go to the **Data Protection** tab to view the list of snapshots for that volume.

5. Check the snapshot containing the point in time you need to go back to. Click **Restore**.
6. Check the box that confirms setting the volume offline. Click **OK**.
7. An automated snapshot is created with a name that has the format: *Nss-volume_name*.
This snapshot is not managed. You should manually delete it after you validate the restoration and confirm that it was successful.
8. Re-establish all connections from all ESXi hosts to the volume. For each host:
 - a) Bring the volume back online.
 - b) Select the **Configuration / Storage Adapters** page.
 - c) Select on the iSCSI Software Adapter or Fibre Channel adapter, and choose **Rescan**.
 - d) Move to the **Configure/Datastores** page and confirm that the datastore is visible and inaccessible (that is, has not been mounted).
 - e) Right-click the datastore and choose **Mount** to remount it.
9. For each VM in the datastore, navigate to its folder, and right-click the `.VMX` configuration file associated with it. Then select **Register VM**.
Follow the onscreen prompts to add the VM back to inventory.
10. Power on each VM.
If the VM doesn't power on, look for a yellow exclamation point beside the VM name, select the VM, and choose **Answer Question**. If asked whether you copied or moved the VM, answer **MOVED**.

Results

After the VMs have powered on and are running, the restore operation is complete.

Recover a Virtual Machine from a Cloned Snapshot Using the Array OS GUI

Restoring a VMFS datastore from a cloned snapshot prevents the original volume from being overwritten with data from an earlier snapshot. A clone is actually a new volume. Use this method of restoration if you have more than one virtual machine (VM) per datastore.

The process of recovering data from a VMware snapshot is slightly different than restoring a datastore snapshot. You perform some of these steps from the array OS GUI. Then you perform the rest of the steps from the VMware vCenter.

Note: Presenting multiple cloned volumes that have the same VMFS UUID prevents VMware ESXi from offering the option for datastore re-signature.

Procedure

1. **(array OS GUI)** From the main array OS menu, select **Manage > Data Storage** to access the volume details screen and then select the appropriate volume.
2. **(array OS GUI)** Go to the **Data Protection** tab to view the list of snapshots for that volume.
3. **(array OS GUI)** Check the snapshot containing the point in time that you need want to recover.
4. **(array OS GUI)** Click Clone and insert the volume name; for example, *Volume_Name-Clone*.
5. **(array OS GUI)** Click OK.
The cloned volume is automatically set to online and should contain the initiator group assigned to the original volume.
6. **(vCenter GUI)** In VMware vCenter Inventory Tree, select the host and go to the **Configure/Storage Adapters** page. Perform a rescan.
You should see the new cloned volume in the **Static Discovery** section.
7. **(vCenter GUI)** Select the **Datastores** section and choose **Create New Datastore > VMFS > Next**.
Your newly cloned volume should appear in the list of storage available to add.

8. (vCenter GUI) Select the new clone and click **Next**.

You are prompted to keep or assign a new signature.

Note: HPE recommends that you always re-signature at this point. Your cloned volume now appears as a datastore. The format for name of the re-signed volume is `snap-xxxoriginal-name`.

9. (vCenter GUI) Click **Next** > **Finish**.**10. (vCenter GUI)** Remove your original VM from inventory.**11. (vCenter GUI)** Browse to the new cloned volume datastore for your VM and add the recovered VM to inventory.**12. (vCenter GUI)** For VMware synchronized snapshots, select the VM and choose **Snapshot** > **Manage Snapshots**.**13. (vCenter GUI)** Select the snapshot and choose **All Actions** > **Revert to**.

You only see a VMware snapshot if the volume was in a volume collection that used vCenter Synchronized snapshots. Manual snapshots or non-synchronized snapshots will not have a VMware snapshot.

14. (vCenter GUI) To move the recovered VM to its original volume, select the VM, and select the **Migrate** option.

Change Access Information Using the Array OS GUI

You can use the array OS GUI to change the VMware vCenter access information on any volume collection that uses the selected synchronization method. You need to do this if the IP address or host name of your VMware vCenter Server changes.

Procedure

1. From the array OS main menu, select **Manage** > **Data Protection**.
2. Select the volume collection that uses the synchronization option you want to change so that you can see its details.
3. Click **Pencil** icon to open the editing wizard.
4. Make any necessary changes under the Synchronization Service.
5. Click **Save**.
6. Repeat this procedure for each volume collection using synchronization.

Change Access Information Using the CLI

You can use the CLI to change the VMware vCenter access information on any volume collection that uses the selected synchronization method. You need to do this if the IP address or host name of your VMware vCenter Server changes.

Procedure

Edit the volume collection information.

```
volcoll --edit collection_name [--new new_collection_name] [--description description] --app_sync
[none|vss|vmware|generic] --app_server host-name
```

Example

```
volcoll --edit collection33 --app_sync vmware --app_server 10.15.133.70
```

iSCSI Best Practices in VMware Environments

Here are some iSCSI-specific best practices when using VMware and a storage array. Adhering to these practices will help ensure that iSCSI connectivity is maximized and that communication between the network components is optimized for best performance.

- There must be a one-to-one mapping between vmkernel (vmk) ports and physical (vmnic) NIC ports.

- If you are using multiple VLANs for iSCSI traffic, use separate vSwitches for each VLAN.
- If you are using a single vSwitch, disable NIC Teaming.
- Use the custom round robin policy and IOPS adjustment for optimal performance.
- Ensure all vmkernel (vmk) ports are only bound to the software iSCSI initiator when on a single iSCSI network, and not when multiple iSCSI networks are used.
- Ensure all VM guest machines using direct attached or Raw Device Mapped (RDM) volumes are connecting to the array over additional interfaces, not those with bound kernel ports to the iSCSI adapter.
- Use Jumbo Frames (MTU 9000 or larger) for both the physical switches and the array.

For more information, see [VMware iSCSI Configuration](#) on page 50 in this document and also see the [VMware iSCSI SAN Configuration Guide](#). Also see the switch manufacturer's documentation.



Important:

iSCSI Target Limits

The bound maximum number of iSCSI static (manually assigned IP addresses) or dynamic (IP addresses assigned to discovered targets) allowed per adapter port is between 62 and 128, depending on the initiator model.

The sum of all iSCSI software targets, either manually assigned or dynamically discovered, cannot exceed 256.

For more information, refer to the *Configuration Maximums – VSphere 6.0* and *Configuration Maximums – VSphere 6.5* documents, available at <https://www.vmware.com/pdf/vshpere6/60/vsphere-60-configuration-maximums.pdf> and <https://www.vmware.com/pdf/vsphere6/65/vsphere-65-configuration-maximums.pdf>, respectively.

Set Up 1:1 Mapping for vSphere Switches

When you have multiple NICs, you can set up a 1:1 mapping by designating a separate vSphere switch for each virtual-to-physical adapter pair.

If you are using separate vSphere switches, you must connect them to different IP subnets. If you do not do this, the VMkernel adapters might experience connectivity problems. When this happens, the host cannot discover the iSCSI LUNs.

Another way to approach this issue is to add all NICs and VMkernel adapters to one vSphere standard switch. To do this, you must override the default network setup and make sure that each VMkernel adapter maps to only one corresponding, active physical adapter.

When the VMkernel adapters are on the same subnet, you must use the single vSwitch configuration.

For more information, refer to the VMware KB article *Considerations for using software iSCSI port binding in ESX/ESXi (2038869)* (<https://kb.vmware.com/s/article/2038869>).

Host Disk Timeout Values

You can configure VMware host operating systems to allow an array failover to be transparent to data services.

The default disk multipath input/output (MPIO) configuration of many host operating systems does not allow for timeouts that permit the array to perform a non-disruptive failover, for example, in the event of a software update.

If incorrect timeout values are set on host operating systems, there is a risk that access to target volumes may time out or only be mounted if a failover occurs on the array. Host applications running on the target volumes may be seriously affected.

After a failover, if the host operating systems report connection losses to target volumes and do not recover, it is good practice to check the timeout values that are set.

Timeout configuration applies only to VMware hosts containing VMDK-based VM guests that reside on datastore volumes, RDM disks and Virtual Volumes (vVols). It does not apply to Fibre Channel environments, because HBA PCI Passthrough is not currently supported.

For guest VMs that also contain directly-attached secondary disks connected through the guest iSCSI initiator, the relevant disk timeout values must be put in place separately for that guest operating system. Refer to the articles below for guidance on this.

The following timeout values should be set for all vSphere environments attached to an array where iSCSI is used:

- LoginTimeout – 30 seconds (default is 5 seconds)
- NoopTimeout – 30 seconds (default is 10 seconds)
- NoopInterval – 30 seconds (default is 15 seconds)

Note: Version 6.0.0 or later of HPE Storage Connection Manager for VMware automatically sets each of these timeout values to 30 seconds.

Registering the HPE vCenter Plugin and performing a datastore related operation (such as a new provision, mount or clone) using the plugin sets the values above in the Advanced Settings of the iSCSI target specific to the datastore acted upon.

VMware includes support for setting iSCSI timeout values the following releases:

- VMware ESXi 6.x
- VMware ESXi 5.5 Update 1

Note:

Refer to the following articles for information about recommended timeout values for Windows and Linux/Citrix Xen Server:

[KB-000052 Windows: Host Disk Timeout Values](#)

[KB-000304 Linux: Host Disk Timeout Values](#)

The table below lists the VMware definitions of the three timeout values.

Timeout Value	Description
LoginTimeout	The default iSCSI login timeout is set to 5 seconds. This means that after 5 seconds, the ESXi host stops the iSCSI session if there is no response, and tries to log in again immediately after.
NoopTimeout	The amount of time, in seconds, that can lapse before your host receives a NOP-In message. The message is sent by the iSCSI target in response to the NOP-Out request. When the NoopTimeout limit is exceeded, the initiator terminates the current session and starts a new one.
NoopInterval	Time interval, in seconds, between NOP-Out requests sent from an iSCSI initiator to an iSCSI target. The NOP-Out requests serve as the ping mechanism to verify that a connection between the iSCSI initiator and the iSCSI target is active.

Note: In ESXCLI output, the NoopTimeout and NoopInterval may display as NoopOutTimeout and NoopOutInterval, respectively.

Managing Target Subnets

Each NIC on a VMware host attempts to connect to each of the NICs on the array over all available subnets. Many of those paths are invalid. This results in numerous dead paths that are eventually removed. But before they are removed, the dead paths count toward the maximum of 1024 paths supported by the VMware host. Having a large number of dead paths can increase the time required for boot, scan, and rescan operations.

To shorten the time required for these operations, use iSCSI initiator groups to limit the number of connections between the VMware host and the array to a specific subnet. When you use iSCSI initiator groups only the NICs on associated subnets are visible to the host. As a result, the host does not attempt to make all theoretically possible connections.

An iSCSI initiator group consists of one or more subnets and one or more host initiators. You can create an iSCSI initiator group using the GUI or CLI. After you create your iSCSI initiator groups, assign your volumes to them.

See [iSCSI Initiator Groups and the Array](#) on page 54.

Guest Disk Timeout Values

HPE recommends the following timeout values for guest operating systems:

- Linux: 180 seconds
- Windows: 60 seconds

When VMware Tools is installed, it sets these values by default using UDEV rules. You should not need to modify these settings.

The increased timeout values for the guest operating systems provide time in the event that ESXi needs to recover from an error. This way the guest OS is not performing retry and stop operations while ESXi is recovering from the error.

The Linux timeout is applied by the file `/etc/udev/rules.d/99-vmware-scsi-udev.rules` from the package **open-vm-tools**.

The Raw Disk Maps (RDMs) the timeout is applied by appending:

```
ACTION=="add", SUBSYSTEM=="block", ATTRS{vendor}=="Nimble", RUN+="/bin/sh
-c 'echo 180 >/sys/$DEVPATH/device/timeout' "
```

to the same file: `/etc/udev/rules.d/99-vmware-scsi-udev.rules`.

Using VMware RDM Disks

You can use VMware Raw Disk Maps (RDMs) with virtual machines (VMs).

The following are some important tips to keep in mind as you work with RDMs:

- To share an RDM with multiple VMs, the VMs must be clustered. For more information, see the VMware Knowledge Base article [Sharing an RDM virtual disk between multiple virtual machines \(1002782\)](#).
- An I/O error can occur on VMs with RDM devices when you are running vSphere versions 5.1, 6.0, or 6.5. This behavior might be caused by cached SCSI INQUIRY data that interferes with specific guest operating systems and applications.

If you encounter this problem, configure the virtual machine with the RDM to ignore the SCSI INQUIRY data cached by ESXi. Add the following parameter to the **.vmx** file:

scsi:x:y.ignoreDeviceInquiryCache = "true"

Where *x* is the SCSI controller number and *y* is the SCSI target number of the RDM.

See your VMware documentation for more information.

- If you are running vSphere versions 6.7 and 7.0, use the following command to ignore the SCSI INQUIRY cache at the host level. This does not require a VM reboot.

esxcli storage core device inquirycache set --device [device id] --ignore true

See your VMware documentation for more information.

- HPE Storage Connection Manager for VMware does not currently support SCSI-3 persistent reservations.

The steps in the procedure below explain how to add an RDM disk to a VM.

Procedure

1. Attach the volume to be used as the RDM to the ESXi host. Do not add a datastore.
2. Add the ESXi initiator or WWPN to the initiator group for the RDM volume.
3. Edit the settings on the VM that will use the RDM volume, adding the disk (RDM) and select the volume.

Results

After adding the RDM disk to the VM, it appears as a regular disk in the guest OS.

Denylist RDMS Using SCSI 3 Persistent Reservations

When you use a disk for a Microsoft cluster, you must denylist the VMware Raw Disk Maps (RDMs) devices associated with that cluster. The HPE Storage Connection Manager does not support SCSI 3 Persistent Reservations (SCSI 3 PRs), which are required for Microsoft clustering configurations.

Note: VMware added support for clustered virtual disks in vSphere 7.0. See the VMware document [Enable or Disable Support for Clustered Virtual Disks on the VMFS6 Datastore](#) for more information.

When you denylist a device, you prevent it from being claimed by the Path Selection Plugin (PSP). As a result, you can set SCSI 3 PRs on the RMD or clustered virtual disks without having validation failures.

To denylist an RDM or datastore, you must perform the following steps on each ESXi host where the virtual machine (VM) resides.

Before you begin

You only need to perform this procedure if you need to use RDMs in a Microsoft cluster.

Procedure

1. Reset PSP policy for all devices. Enter the following command on the ESX host:
`/etc/nimble/nimble-policy.sh resetpolicy`
2. Add each device serial number to the `/etc/nimble/nimble-bsp-exclude-devices.txt` file.

Note: The device serial numbers start with `eui.xxx`.

To add this information, open the file using a text editor and enter the changes. You can add one device to each line. Make sure you save your changes.

Note: The text editor Vi comes with ESXi.

If you need help determining which devices are used as RDMs, refer to the VMware article [Identifying virtual disks pointing to Raw Device Mappings \(RDMs\) \(1005937\)](#).

3. Set the PSP policy. This automatically excludes the listed devices. Enter the command:
`/etc/nimble/nimble-policy.sh setpolicy`
4. Repeat these steps on each host. You must also repeat them whenever you add a new device that requires SCSI 3 PR support.

Configure ESXi iSCSI Networking with Distributed Virtual Switches

HPE supports VMware Virtual Distributed Switches (vDS) with the array OS. The following steps provide an overview of the tasks required for setting up vDS. For detailed instructions, refer to the VMware documentation.

Procedure

1. Set up a vDS by following the instructions in the VMware documentation for your version of vSphere.
For example, at the time this document was created, the instructions for configuring vDS with vSphere 7.x were located here:
<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-375B45C7-684C-4C51-BA3C-70E48DFABF04.html>
2. Add the vmkernel adapter by following the instructions in the VMware documentation for your version of vSphere.



For example, at the time this document was created, the instructions for adding the vmkernel adapter to a vSphere 7 configuration were located here:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-1CAD9F77-71C5-41DE-9357-E5FAFCD20D36.html>

3. Edit the MTU so that it matches the switch and vmkernel port by following the instructions in the VMware documentation for your version of vSphere.

For example, at the time this document was created, the instructions for adding the vmkernel adapter to a vSphere 7 configuration were located here: :

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-E0FED4AB-823D-4B61-B668-9400746D52E5.html>

4. It is a good practice to set up a one to one mapping of the vmkernel adapter to the physical port (vmnic).

Storage LUNs in the Device List

Even though the vCenter iSCSI discovery was successful, you might find that none of the volumes appear in the list of datastores. You can perform the following steps to verify this on an iSCSI adapter.

Procedure

1. Ensure that the dynamic discovery IP address has been configured.
2. Select **Storage Adapters** and then the iSCSI adapter. Click **rescan**.
3. After the rescan completes, move to the **Datastores** link and click **Add Storage...**

Results

If you still don't see any LUNs in the **Select Disk/LUN** list, it is possible that the ESXi host is not in the volume access control list. Verify that the volume ACL is correctly assigned to the ESXi host. If the volume is expected to be accessed by multiple ESXi hosts, go to the Volume Edit wizard on the array and set the volume option to **Allow multiple initiator access**.

Locate vCenter Log Files

Your log files are kept on both on the Windows server and on the appliance.

- In Windows, vCenter logs are usually located in `C:\ProgramData\VMware\vCenterServer\logs`. Logs are named `vcpd-nnnn.log`, where `nnnn` is an incrementing number.
- On the appliance, the vCenter logs are located in `6.x: /var/log/vmware/`.



Helpful Information

The tasks that follow explain how to perform certain functions from the storage array.

Most of these functions are handled automatically by HPE Storage Connection Manager. However, if you do not have HPE Storage Connection Manager installed, and you have not registered the HPE vCenter Plugin with the vCenter Server, or if you prefer to make these settings manually, you can use these procedures.

Recommendations for Environments That Do Not Use HPE Storage Connection Manager

HPE Storage Connection Manager for VMware performs certain tasks and enables you to use features such as striped pools and volume migrations. If you do not install HPE Storage Connection Manager, then you should manually do some setup, including setting:

- The PSP to Round Robin for ESXi hosts on software 5.x and above (VMware)
- The iops parameter to 1. Doing this improves the I/O balance across paths when you use the built-in VMware PSP. The NCM policy NIMBLE_PSP_DIRECTED automatically applies this optimization for you.

If you are running a 5.x or 6.x version of ESXi and you did not install the HPE Storage Connection Manager for VMware, you must manually set the Path Selection policy to Round Robin. After you run this procedure, all volumes created will use Round Robin. It is the default.

Note: If you have installed the HPE Storage Connection Manager for VMware, you do not need to perform this procedure. The connection manager automatically creates the optimal number of iSCSI sessions for each volume and manages the selection of paths to the volumes.

This procedure must be done from a console session to the ESXi host, such as an SSH session to the host.

HPE recommends that when you have an ESXi host that has a volume mounted, but does not have HPE Storage Connection Manager installed, you should perform the following steps.

Note: Modifying the PSP is normally non-disruptive. However, HPE always recommends a risk-adverse approach. Before you perform the following steps, consider scheduling a maintenance window and placing ESXi hosts that need these modification into maintenance mode.

Procedure

1. Log into the CLI of the VMware ESXi host using SSH.
2. To obtain a list of the currently mounted volumes and their PSPs and device configurations, enter the following:
for i in `esxcli storage nmp device list | grep "^eui\\.\\w{16}(6c9ce9|6C9CE9)\\w{10}"`; do esxcli storage nmp device list --device \$i; done
3. Remove any existing SATP rule (in most cases, this should return an error that no existing rule is found):
esxcli storage nmp satp rule remove --psp=VMW_PSP_RR --vendor=Nimble --satp=VMW_SATP_ALUA
4. Create a new SATP rule to apply the "Round Robin (VMware)" PSP and "iops=1" to newly mounted volumes:
esxcli storage nmp satp rule add --psp=VMW_PSP_RR --satp=VMW_SATP_ALUA --psp-option='policy=iops;iops=1' --vendor=Nimble
 All newly presented volumes default to Round Robin.
 All existing volumes retain their old path selection policies.
5. Set the PSP to "Round Robin (VMware)" on already mounted HPE Nimble Storage volumes:

```
for i in `esxcli storage nmp device list | grep "^eui\.\\w{16}(6c9ce9|6C9CE9)\\w{10}"`; do esxcli storage nmp device set --device $i --psp=VMW_PSP_RR; done
```

Note: Changing the default policy to change paths after each IO is only applicable to newly added devices or after a host reboot. So for existing devices a host reboot is required for this setting to take effect.



CAUTION: This script must be typed on a single line EXACTLY as shown. If possible, copy each line of the script from this page and paste it directly into the command line.

6. Repeat the command in Step 2 to validate the settings have been applied to volume:

```
for i in `esxcli storage nmp device list | grep "^eui\.\\w{16}(6c9ce9|6C9CE9)\\w{10}"`; do esxcli storage nmp device list --device $i; done
```

Configure iSCSI Discovery

Before you begin

You must have created volumes on an array and set up any required initiator groups for access control.

To configure iSCSI Digest, see the following:

- [Enable iSCSI Digest](#) on page 61
- [KB-000296 Enabling iSCSI Digest on VMware initiators](#)



Important: Validated configurations involving iSCSI hardware adapters are limited. See the Validated Configuration Matrix, which is online at <https://infosight.hpe.com/resources/nimble/validated-configuration-matrix> for information about configurations that have been tested.

Procedure

The following steps tell you how to configure the ESXi iSCSI software adapter to discover iSCSI targets on the array.

1. From the VMware vSphere Web Client, select the **Host and Cluster** view.
2. Select **Configure** > **Select Storage Adapters** > **Select iSCSI Software Adapter**.
3. Select **Dynamic Discovery tab** > **Add** > **Input Discovery IP in the iSCSI server field** > **OK**.
4. Verify that the Discovery IP was added correctly.
5. Rescan the adapter.
6. Select **Static Discovery** to see which devices have been logged into it.
The list that appears contains information about all the volumes returned during discovery, even though some might not be accessible. This list is built from information returned by the array, not from actual connections.
7. From the Storage section, select **Storage Devices** and select an HPE Nimble Storage device.
8. From the Properties tab, you can see the Path Selection Policy (PSP). Select **Edit Multipathing**.
9. From the drop-down list, select **Round Robin (VMware)** > **OK**.
10. Select **Paths**.

The Paths view displays the actual paths that ESXi has to each volume. The Runtime Name contains four items for each path:

- **vmhba##** – Stays the same across all the iSCSI software adapter entries.
- **C#** – For each volume, the channel number starts at C0 and increments for each path found to that volume.
- **T#** – Each volume has its own Target number. It identifies a known volume between this Paths view and the Device view.
- **L#** – The LUN number is always zero for storage volumes.

The default for newly discovered devices is to use only one path for I/O.

Active (I/O) indicates that Round Robin is in effect, and all paths can be used for I/O.

11. As needed, repeat the two previous steps to change the path selection policy on the ESXi host for the other storage devices/volumes.

iSCSI Host Connection Methods

The iSCSI initiators on the host connect with targets on the array through the data ports on each controller. Each port is identified by its IP address.

Beginning with the array OS release 2.0 release, the default iSCSI host connection management method is *Automatic*. When you install the Connection Service on your Windows or VMware hosts, the process of adding target portals and connecting volumes to iSCSI targets is also simplified.

Each subnet, management or data, has a unique discovery IP address.

Figure 3: Discovery IP in the GUI



Subnet Label	Network	Netmask	Traffic Type	Traffic Assignment	Discovery IP	IP Address Zone	MTU	Bytes	VLAN ...
data1	198.51.100.0	255.255.255.0	Data only	iSCSI + Group	198.51.100.55	Single	Standard	1500	
mgmt-data	192.0.2.0	255.255.255.0	Mgmt only		192.0.2.51		Standard	1500	



Important: The discovery IP address must be accessible by at least one host initiator port.

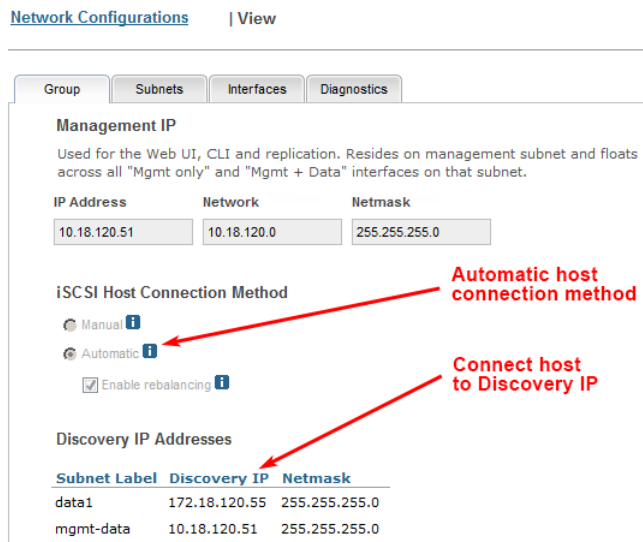
Automatic iSCSI Host Connections

The Automatic iSCSI connection method uses data subnet discovery IPs for host connection. The array then automatically redirects the connection to an appropriate iSCSI data IP. The Automatic method is the better choice for most applications.



Go to **Administration** > **Network Configuration** and click **Active Settings**.

On the array, when the iSCSI host connection method is set to Automatic:



[Network Configurations](#) | [View](#)

Group: Subnets | Interfaces | Diagnostics

Management IP
Used for the Web UI, CLI and replication. Resides on management subnet and floats across all "Mgmt only" and "Mgmt + Data" interfaces on that subnet.

IP Address	Network	Netmask
10.18.120.51	10.18.120.0	255.255.255.0

iSCSI Host Connection Method

Manual ⓘ
 Automatic ⓘ

Enable rebalancing ⓘ

Discovery IP Addresses

Subnet Label	Discovery IP	Netmask
data1	172.18.120.55	255.255.255.0
mgmt-data	10.18.120.51	255.255.255.0



At the command prompt, type **netconfig --info active**.


```

Nimble OS $ netconfig --info active
Group Management IP: 10.18.120.51/255.255.255.0
Group leader array: c20-array2
Member array(s): c20-array2
ISCSI Automatic connection method: Yes
ISCSI Connection rebalancing      : Yes

```

Automatic host connection method

```

Routes:
-----+-----+-----
Destination  Netmask      Gateway
-----+-----+-----
0.0.0.0      0.0.0.0      10.18.120.1

```

Connect Host to Discovery IP

```

Subnets:
-----+-----+-----+-----+-----+-----
Label        Network      Type      Discovery IP  VLAN  MTU
-----+-----+-----+-----+-----+-----
data1        172.18.120.0/24  Data      172.18.120.55  0     1500
mgmt-data    10.18.120.0/24  Mgmt      10.18.120.51   0     1500

```

```

Array Network Configuration: c20-array2
Controller A IP: 10.18.120.54
Controller B IP: 10.18.120.55

```

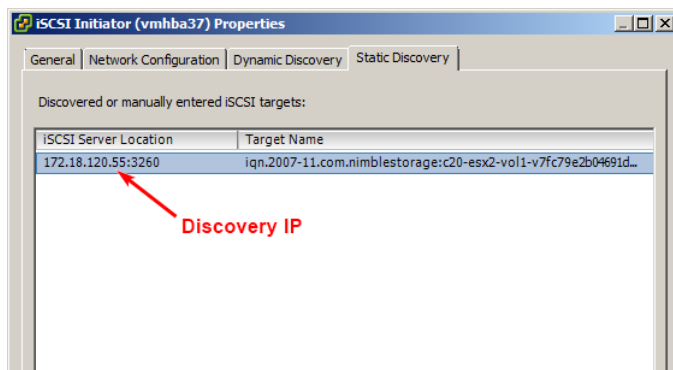
```

NIC          Subnet Label      Data IP Address Tagged
-----+-----+-----+-----+-----
eth1         mgmt-data          N/A             No
eth2         mgmt-data          N/A             No
eth3         data1              172.18.120.56  No
eth4         data1              172.18.120.57  No
eth5         data1              172.18.120.58  No
eth6         data1              172.18.120.59  No

```

Nimble OS \$ █

Use a data subnet discovery IP to connect from the host to the array. In the GUI, discovery IPs are listed on the Group and Subnets tabs.



Note: The iSCSI host to array connection process is faster and simpler when you install and use HPE Storage Connection Manager on your Windows or VMware host.

See the *Windows Integration Guide*.

Manual iSCSI Host Connections

The array OS release v1.4 and earlier releases, hosts connect manually to iSCSI data or discovery IPs. The same is true for the array OS release v2.0 when the iSCSI connection method is set to *Manual*. The Manual method is intended primarily for legacy applications.



Go to **Administration** > **Network Configuration** and click **Modify**.

On the array, when the iSCSI host connection method is set to Manual:



Network Configurations | View

Group Subnets Interfaces Diagnostics

Management IP
 Used for the Web UI, CLI and replication. Resides on management subnet and floats across all "Mgmt only" and "Mgmt + Data" interfaces on that subnet.

IP Address: 10.18.120.51 | Network: 10.18.120.0 | Netmask: 255.255.255.0

iSCSI Host Connection Method

Manual **Manual host connection method**
 Automatic
 Enable rebalancing

Discovery IP Addresses

Subnet Label	Discovery IP	Netmask
data1	172.18.120.55	255.255.255.0
mgmt-data	10.18.120.51	255.255.255.0



At the command prompt, type **netconfig --info active**.

```
Nimble OS $ netconfig --info active
Group Management IP: 10.18.120.51/255.255.255.0
Group leader array: c20-array2
Member array(s): c20-array2
iSCSI Automatic connection method: No
iSCSI Connection rebalancing : No

Routes:
-----
Destination      Netmask      Gateway
-----
0.0.0.0          0.0.0.0      10.18.120.1

Subnets:
-----
Label            Network      Type      Discovery IP  VLAN  MTU
-----
data1            172.18.120.0/24  Data     172.18.120.55  0     1500
mgmt-data        10.18.120.0/24  Mgmt     10.18.120.51  0     1500

Array Network Configuration: c20-array2
Controller A IP: 10.18.120.54
Controller B IP: 10.18.120.55

NIC      Subnet Label      Data IP Address Tagged
-----
eth1     mgmt-data          N/A             No
eth2     mgmt-data          N/A             No
eth3     data1              172.18.120.56  No
eth4     data1              172.18.120.57  No
eth5     data1              172.18.120.58  No
eth6     data1              172.18.120.59  No

Nimble OS $
```

Use a data subnet discovery IP to connect from the VMware host to the array. In the GUI, discovery IPs are listed on the Group and Subnets tabs.

ISCSI Initiator (vmhba37) Properties

General Network Configuration Dynamic Discovery Static Discovery

Discovered or manually entered iSCSI targets:

iSCSI Server Location	Target Name
172.18.120.55:3260	iqn.2007-11.com.nimblestorage:c20-esx2-vol1-v7fc79e2b04691d...

Discovery IP



Set the iSCSI Host Connection Method to Manual

The manual iSCSI host connection method is intended primarily so that you can configure legacy applications to use the automatic connection method after upgrading to release 2.0 or later. (Neither option applies to Fibre Channel arrays.)

Procedure

1. Disable automatic rebalancing of the iSCSI host connections.

```
netconfig --edit active --iscsi_connection_rebalancing no
```

2. Disable automatic iSCSI host connections.

```
netconfig --edit active --iscsi_automatic_connection_method no
```

Example

Example:

```
$ netconfig --edit active --iscsi_connection_rebalancing no
$ netconfig --edit active --iscsi_automatic_connection_method no
```

Configure Jumbo Frames

You have the option of configuring Jumbo Frames. If you configure your environment for Jumbo Frames, every device in the network path must be set up for Jumbo Frames, including:

- Physical network switches
- Arrays
- ESXi vSwitches
- ESXi vmk ports

Depending on your network infrastructure, you might need to set up additional devices.

To set Jumbo Frames on the array, go to [Change NIC Frame Size](#). To set Jumbo Frames on your physical network switches, refer to the vendor documentation.



CAUTION: When setting Jumbo Frames in ESXi, change vSwitches first. Then immediately change the vmk ports. If you do not change the vmk ports immediately after the vSwitches, unexpected path behavior can occur, such as the paths going up and down.

Procedure

You must perform this task from a console session to the ESX host, such as an SSH session (Remote Tech Support session) to the host.

This procedure enables Jumbo Frames for ESXi vSwitches and vmk ports.

1. Log into vSphere Web Client.
2. Select the ESXi host.
3. Select **Configure** > **Virtual Switches** > **Select vSwitch**.
4. Select the **Edit** button.
5. Set MTU (Bytes) to 9000 and select **OK**.
6. Select **Virtual Adapters** and choose the vmk port that is under the previous vSwitch.
7. Select **Edit**.
8. Select **NIC setting** and configure MTU to 9000. Select **OK**.
9. Configure each vmk port under that vSwitch to match that MTU size.



Validate the MTU Settings

When you use iSCSI, your system must have a consistent, end-to-end maximum transmission unit (MTU) that flows from the host to the vSwitch to the switch to the array. It is a good practice to confirm that the values are correct.

From the vSwitch, select **Configure** > **Virtual Switches** > **<select the switch>** > **Pencil icon (Edit)** > **<Modify MTU size if needed>**.

From the VMK adapter, select **Configure** > **VMKernel Adapters** > **<select VMK device>** > **Pencil icon (Edit)** > **NIC Settings** > **<Modify MTU size if needed>**.

Change NIC Frame Size

You can set the frame size for each NIC on your array. Do not select Jumbo Frames unless all of your network components support them.

Procedure

1. From the dashboard, choose **Administration** > **Network**.
2. Click **Configure Active Settings**.
3. Click the **Subnets** tab.
4. Click **Edit**.
5. Check the checkbox next to the subnet that you want to modify and click **Edit**.
The MTU settings are in the top section of the page.
6. Choose **Jumbo** from the MTU dropdown menus.
If you select **Custom**, you must also include the number of bytes to use for the frame size.
7. Click **Done**.

Configure an ESX Datastore

Procedure

To create a datastore from a newly presented volume:

1. Go to the Configuration / Datastores screen on the ESXi host and click **Create New Datastore**.
2. Under Type, choose the VMFS option and click **Next**.
Disks that can be formatted as datastores and existing datastores are listed.
Expand the Path ID column until you can see the entire volume name.
3. Select the appropriate disk, provide a name for the datastore, and click **Next**.
4. Select the VMFS version you want to use and click **Next**.
By default, the latest version is selected.
5. Select partition configuration you want to use and click **Next**.
By default, the entire disk is selected.
6. Review the proposed configuration and click **Next**.
7. The datastore is created and appears in the Configuration / Datastores screen.
8. Repeat this procedure to make additional volumes into datastores.



Enable Application-Consistent Quiescing on Windows Server 2008 VM

This procedure applies to Windows Server 2008 virtual machines (VMs) that were created on an ESX/ESXi 4.0 host and later migrated to a newer host, such as an ESXi 5.x or 6.x host.

Windows Server 2008 VMs created on an ESX/ESXi 4.0 host and migrated to a newer host are not enabled for application-consistent quiescing. Those VMs require the configuration parameter `disk.EnableUUID = TRUE`. Without this parameter, vCenter-synchronized backups complete successfully and give the impression that they are application consistent. However, the backups do not have application-consistent snapshots. Neither Microsoft Exchange or SQL support these snapshots in their restore scenarios.

Procedure

To enable the `disk.EnableUUID = TRUE` parameter, perform the following steps:

1. Log on to the vSphere Web Client.
2. Select **Virtual Machines and Templates** and click the **Virtual Machines** tab.
3. Power off the Windows 2008 VM for which you want to enable the disk UUID parameter by right-clicking it and selecting **Power > Power Off**.
The VM powers off.
4. Right-click the VM and click **Edit Settings**.
5. Click the **Options** tab, and expand the **Advanced** entry in the settings column.
6. Click **Edit Configuration**.
The Configuration Parameters window appears.
7. In the Name column, type:
disk.EnableUUID
8. In the Value column, type:
TRUE
9. Click **Add**.
10. Click **OK**.
11. Power on the VM.

Results

When the UUID parameter is added, application-consistent quiescing is enabled for the Windows Server 2008 VM that was created on an ESX/ESXi 4.0 host and later migrated to a newer host.

Mount and Unmount a Datastore Outside the Plugin

For information on how to mount a datastore outside of the plugin, refer to [Mount Datastores](#) on the VMware documentation site.

For information on how to unmount a datastore outside of the plugin, refer to [Unmount Datastores](#) on the VMware documentation site.

