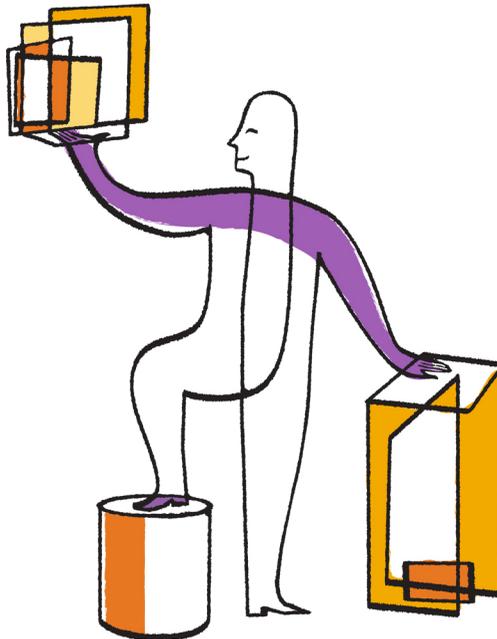




**NetApp®**

## Virtual Storage Console 4.2 for VMware® vSphere™

### Installation and Administration Guide



NetApp, Inc.  
495 East Java Drive  
Sunnyvale, CA 94089  
U.S.

Telephone: +1(408) 822-6000  
Fax: +1(408) 822-4501  
Support telephone: +1 (888) 463-8277  
Web: [www.netapp.com](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-07468\_A0  
May 2013



# Contents

<b>Changes to this document: May 2013 .....</b>	<b>9</b>
<b>VSC for VMware vSphere Overview .....</b>	<b>11</b>
Interaction between the capabilities .....	12
Architecture of Virtual Storage Console for VMware vSphere .....	13
Methods for accessing the VSC for VMware vSphere .....	14
VSC for VMware vSphere provides lock management for system resources .....	14
Online Help .....	15
<b>Installation overview .....</b>	<b>16</b>
Additional setup requirements when using RBAC .....	17
VSC for VMware vSphere supported configurations .....	18
Installing VSC for VMware vSphere using the installation wizard .....	19
Installing VSC for VMware vSphere using silent mode .....	20
Installing or uninstalling Backup and Recovery .....	21
Registering VSC for VMware vSphere with vCenter Server .....	23
VSC for VMware vSphere port requirements .....	24
Regenerating an SSL certificate for VSC for VMware vSphere .....	24
Preparation required before upgrading VSC for VMware vSphere .....	26
Upgrading VSC for VMware vSphere .....	28
Uninstalling VSC for VMware vSphere using Add/Remove Programs .....	29
Uninstalling VSC for VMware vSphere using silent mode .....	30
<b>Authentication and user management with vCenter RBAC and Data</b>	
<b>ONTAP RBAC .....</b>	<b>31</b>
vCenter Server role-based access control features in VSC for VMware vSphere ....	32
Components that make up vCenter Server permissions .....	33
Key points about assigning and modifying permissions .....	34
Advanced example of using vCenter Server permissions .....	35
Standard roles packaged with VSC for VMware vSphere .....	37
Guidelines for using VSC standard roles .....	38
Privileges required for VSC tasks .....	39
Product-level privilege required by VSC for VMware vSphere .....	40
Example of how the View privilege affects tasks in VSC for VMware	
vSphere .....	40

Data ONTAP role-based access control features in VSC for VMware vSphere .....	41
Recommended Data ONTAP roles when using VSC for VMware vSphere .....	42
How to configure Data ONTAP role-based access control for VSC for VMware vSphere .....	43
<b>Monitoring and configuring hosts and storage controllers .....</b>	<b>45</b>
VSC for VMware vSphere configuration .....	46
Storage system discovery and credentials overview .....	46
Enabling discovery and management of vFiler units .....	49
Enabling discovery and management of vFiler units on private networks ...	49
Tunneled vFiler units and Vservers discovered automatically .....	50
Differences between direct connections to Vservers and to cluster- management LIFs .....	50
Discovering and adding storage resources .....	51
Correcting controller names displayed as "unknown" .....	52
Removing controllers from Monitoring and Host Configuration .....	52
Administering the Monitoring and Host Configuration capability .....	53
Inventory panel selection limits what VSC for VMware vSphere displays .....	53
Configuring ESX server multipathing and timeout settings .....	54
Direct path access and NFS datastores .....	56
Flash Accel and VSC for VMware vSphere .....	58
NFS Plug-in for VMware VAAI requires additional installation steps .....	58
Using Web-based tools to manage storage .....	59
Monitoring and Host Configuration displays configuration details .....	60
MultiStore vFiler units are displayed differently .....	60
Collecting diagnostic information .....	61
Downloading and running tools .....	62
<b>Provisioning and cloning datastores and virtual machines .....</b>	<b>68</b>
Tips for working with Provisioning and Cloning .....	68
Cloning and managing virtual machines .....	69
Cloning virtual machines .....	69
Managing connection brokers .....	75
Redeploying clones (locally) .....	76
Reclaiming space on virtual machines .....	78
Importing virtual machines into XenDesktop .....	79

Managing storage controllers .....	80
(Data ONTAP operating in 7-Mode) Viewing storage controller details .....	80
(Data ONTAP operating in 7-Mode) Removing or adding network interfaces, volumes, and aggregates .....	80
Managing volume settings .....	81
Managing datastores .....	82
(Data ONTAP operating in 7-Mode) Replicating datastores to remote sites .....	82
Provisioning datastores .....	87
Mounting datastores .....	89
Managing deduplication .....	89
Resizing datastores .....	90
Destroying datastores .....	90
Provisioning and Cloning support files .....	91
Preferences File .....	91
Logs .....	97
Export Files .....	98
<b>Optimizing and migrating datastores and virtual machines .....</b>	<b>100</b>
Types of alignments .....	101
Important notes about using Optimization and Migration .....	102
Considerations when working with NFS datastores .....	103
The Optimization and Migration workflow .....	104
Scanning the datastores .....	105
Scheduling a scan of datastores .....	106
Performing an online alignment .....	107
Migrating virtual machines .....	109
<b>Backing up and restoring data .....</b>	<b>112</b>
Backup and Recovery requirements .....	112
Backup and Recovery requirements for optional SnapMirror protection .....	113
Configuring Backup and Recovery .....	113
Authentication methods in the Backup and Recovery capability .....	113
How Backup and Recovery discovers vFiler units .....	115
Managing backups .....	116
Considerations for adding a backup job .....	116
Backing up a virtual machine .....	117
Backing up a datastore or datacenter .....	118

Starting a one-time backup job .....	120
Editing a backup job .....	120
Deleting a scheduled backup job .....	121
Suspending an active backup job .....	121
Resuming a suspended backup job .....	122
Restoring data from backups .....	122
Where to restore a backup .....	123
Restore operations using data that was backed up with failed VMware consistency snapshots .....	123
Restoring data from backup copies .....	123
Single file restore .....	127
How Virtual Storage Console detects network connectivity .....	127
The difference between limited and direct connectivity .....	128
Types of file restore sessions .....	128
Manually creating a .sfr file for the Restore Agent .....	129
General configuration parameters for single file restore .....	129
Self-service example workflow .....	130
Limited self-service example workflow .....	134
VSC CLI commands .....	136
Launching the VSC CLI .....	137
smvi backup create .....	137
smvi backup delete .....	139
smvi backup list .....	141
smvi backup mount .....	142
smvi backup rename .....	143
smvi backup restore .....	144
smvi backup unmount .....	146
smvi discover datastores .....	147
smvi filerestore add-portgroup .....	148
smvi filerestore delete-portgroup .....	149
smvi notification list .....	149
smvi notification set .....	150
smvi notification test .....	151
smvi restoreagent set .....	151
smvi servercredential delete .....	152
smvi servercredential list .....	153

smvi servercredential set .....	153
smvi storagesystem list .....	154
smvi version .....	155
<b>Programmable APIs .....</b>	<b>156</b>
NetApp Communities provide support for API questions .....	156
What you can do with the APIs for VMware vCloud .....	156
Provisioning and Cloning programmable API .....	156
The virtual machine clone engine .....	157
The datastore management engine .....	157
The file copy/clone offload engine .....	157
Provisioning and Cloning methods .....	158
Provisioning and Cloning specifications and messages .....	180
Provisioning and Cloning sample code .....	195
Provisioning and Cloning client-side programming .....	196
<b>Troubleshooting .....</b>	<b>198</b>
Issues that apply to multiple capabilities .....	198
Information available at NetApp Communities Forum .....	198
Check the Release Notes .....	198
Uninstall does not remove standard VSC roles .....	198
VMware only supports selecting one object when using right-click actions .....	199
Issues that apply to the Monitoring and Host Configuration capability .....	199
Getting information about storage controllers with an Alert status .....	199
Getting information about an ESX and ESXi host with an Alert status .....	199
Collecting the VSC for VMware vSphere log files .....	200
Troubleshooting error message "The client cannot communicate with the Virtual Storage Console Server" .....	201
Updating vCenter credentials for background discovery .....	201
Issues that apply to the Provisioning and Cloning capability .....	202
Cloning operations fail when permission is assigned to wrong object .....	202
Resolution of Backup and Recovery issues .....	202
Backup and Recovery values that you can override .....	202
Backup and Recovery event and error logs .....	203
Email notification for scheduled backup contains a broken link .....	203
You may have reached the maximum number of NFS volumes configured in the vCenter .....	204

Error writing backup metadata to repository\backups.xml: move failed ....	204
Backup and Recovery unable to discover datastores on a Vserver without a management LIF .....	204
VMware vSphere does not remove all snapshot delta disks during a restore operation .....	205
File already exists .....	205
<b>Using ESX hosts with NetApp storage .....</b>	<b>206</b>
LUN type guidelines .....	206
Manually provisioning storage .....	206
How to set up VMware ESX .....	207
Configuring the VMware ESX host .....	208
(Data ONTAP operating in 7-Mode) Manually setting the path selection policy for Microsoft cluster configurations .....	208
Timeout values for guest operating systems .....	209
Running the GOS timeout scripts for Linux .....	210
Running the GOS timeout scripts for Solaris .....	211
Running the GOS timeout script for Windows .....	211
How to identify and fix VMDK partition alignment issues .....	212
Checking VMDK partition alignment with mbralign .....	212
VMDK partition alignment with mbralign overview .....	213
Fixing VMDK partition alignment using mbralign .....	215
Reinstalling GRUB for Linux guests after running mbralign .....	217
<b>Copyright information .....</b>	<b>219</b>
<b>Trademark information .....</b>	<b>220</b>
<b>How to send your comments .....</b>	<b>221</b>
<b>Index .....</b>	<b>222</b>

## Changes to this document: May 2013

---

This section contains information about the changes made to this guide for the 4.2 version of Virtual Storage Console for VMware vSphere. Previously, this guide was released with the 4.1 version of VSC.

The guide documents the tasks you can perform with VSC. If new information or corrections that affect this guide become available during the VSC release, then this guide is updated with the new information, and this section lists what has changed and when it was changed. Any time this guide is updated, a note is added to the Release Notes. It is a good practice to check the online Release Notes on a regular basis to determine whether there is new information about using VSC or changes to this guide. The most current versions of the Release Notes and this guide are posted on the NetApp Support Site, which is available online at [support.netapp.com](http://support.netapp.com).

### May 2013 update

In May 2013, this document was updated to add the following information in support of the 4.2 release of VSC:

- VSC has expanded its support for authentication and user management.  
VSC supports both vCenter Server role-based access control (RBAC) and Data ONTAP RBAC. For more information, see the [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31 and the information that follows it.
- The VSC-specified View privilege, which is read-only, is required for a user to view the VSC GUI.  
Without this privilege, menus, tabs, and other elements of the GUI are not visible. If you do not have this privilege and click the NetApp icon from the Home View, you get an error message. For more information, see [Product-level privilege required by VSC for VMware vSphere](#) on page 40.
- VSC provides a set of standard roles for working with VSC tasks.  
These roles provide the necessary vCenter native privileges and VSC-specific privileges to allow users to perform standard VSC tasks. For more information, see [Standard roles packaged with VSC for VMware vSphere](#) on page 37.
- Monitoring and Host Configuration manages storage discovery and credentials for the entire VSC product.  
All VSC capabilities now use the credentials entered during the initial setup and discovery process in Monitoring and Host Configuration, eliminating the need for you to manage separate sets of storage entities in Backup and Recovery and unifying the experience within VSC. For more information, see [Storage system discovery and credentials overview](#) on page 46.
- VSC provides centralized management and reporting for the Flash Accel components on the ESX host and its virtual machines.  
For more information, see [Flash Accel and VSC for VMware vSphere](#) on page 58.

- If you want to use the connection broker feature of Provisioning and Cloning, you must have .Net 3.5 available on the system where VSC is installed.

For more information, see [Managing connection brokers](#) on page 75.

- Optimization and Migration supports online alignment and migration of NFS datastores on systems running Data ONTAP 8.1.3 or higher.

There are several considerations you need to keep in mind when working with NFS datastores.

For more information, see [Considerations when working with NFS datastores](#) on page 103,

[Performing an online alignment](#) on page 107, and [Migrating virtual machines](#) on page 109.

- By default, in the Backup wizard, the option to create a VMware snapshot is not selected for each backup copy.

You must enable this option to trigger a VMware snapshot for each virtual machine during a backup.

For more information, see [Backing up a virtual machine](#) on page 117.

# VSC for VMware vSphere Overview

---

Virtual Storage Console for VMware vSphere software is a single vCenter Server plug-in that provides end-to-end lifecycle management for virtual machines in VMware environments using NetApp storage.

VSC provides different capabilities to perform the following functions:

- **Storage configuration and monitoring using Monitoring and Host Configuration**  
Monitoring and Host Configuration handles adding and removing storage controllers and assigning storage controller credentials for all capabilities. It also enables you to manage ESX and ESXi servers connected to NetApp storage. You can set values for host timeouts, NAS, and multipathing as well as view storage details and collect diagnostic information.
- **Datastore provisioning and virtual machine cloning using Provisioning and Cloning**  
Provisioning and Cloning uses FlexClone technology to let you efficiently create, deploy, and manage the lifecycle of virtual machines from an interface that has been integrated into the VMware environment.
- **Online alignments and single and group migrations of virtual machines into new or existing datastores using Optimization and Migration**  
Optimization and Migration enables you to quickly check the alignment status of virtual machines. If there are alignment issues with virtual machines in VMFS datastores, you can in most cases resolve those issues without having to power down the virtual machines.
- **Backup and recovery of virtual machines and datastores using Backup and Recovery**  
Backup and Recovery allows you to rapidly back up and recover multi-host configurations on NetApp storage.

In addition, VSC provides access control at two levels:

- vSphere objects, such as virtual machines and datastores.  
These objects are managed using vCenter role-based access control (RBAC).
- Data ONTAP storage  
Storage systems are managed using Data ONTAP RBAC.

If access-control is not an issue, you can log in as administrator and have access to all the features that VSC provides.

VSC also provides APIs for VMware vCloud. These APIs enable access to VSC in the context of vCloud Director objects and semantics. These APIs enable you to manage credentials for multiple vCenter Servers, discover vCloud Director objects for vCloud tenants, and provision and clone vApps.

VSC also supports the Provisioning and Cloning API, which is designed to be leveraged with the VI SDK.

As a vCenter Server plug-in, VSC is available to all vSphere Clients that connect to the vCenter Server. Unlike a client-side plug-in that must be installed on every vSphere Client, you install the VSC software on a Windows server in your data center.

**Note:** Do not install this software on a client computer.

The software adds a NetApp icon to the Solutions and Applications panel of the vSphere Client home page. You must have the read-only View privilege, or VSC displays an error message when you click this icon.

**Tip:** The View privilege is required for all users who do not have administrator privileges. Without the View privilege, these users cannot see the VSC GUI.

When you select the About panel, VSC displays its version information as well as version information for each of the installed capabilities.

If you have questions about VSC or want to comment about it, you can go to the NetApp Communities Forum, which is online at [https://communities.netapp.com/community/products\\_and\\_solutions/virtualization/vsc](https://communities.netapp.com/community/products_and_solutions/virtualization/vsc).

### Related concepts

*Monitoring and configuring hosts and storage controllers* on page 45

*Provisioning and cloning datastores and virtual machines* on page 68

*Optimizing and migrating datastores and virtual machines* on page 100

*Backing up and restoring data* on page 112

## Interaction between the capabilities

All of the Virtual Storage Console for VMware vSphere capabilities provide functions you can use to manage your VMware environment. In most cases, the VSC capabilities operate separately from each other. There are some areas where the capabilities interact with each other.

These areas include the following:

- There is a single installation program that you use to install the capabilities. The installation program automatically installs Monitoring and Host Configuration, Provisioning and Cloning, and Optimization and Migration. If you want to install Backup and Recovery, you must select it when the installation program starts.
- Monitoring and Host Configuration manages discovery and removal of storage systems and the default storage controller credentials for all of VSC.
- VSC provides lock management for the capabilities, which prevents two capabilities from acting on the same virtual machine or datastore at the same time. As a result, some alignment, migration, provisioning, cloning, and recovery features become unavailable when multiple capabilities attempt to use the same target virtual machine or datastore at the same time.

## Architecture of Virtual Storage Console for VMware vSphere

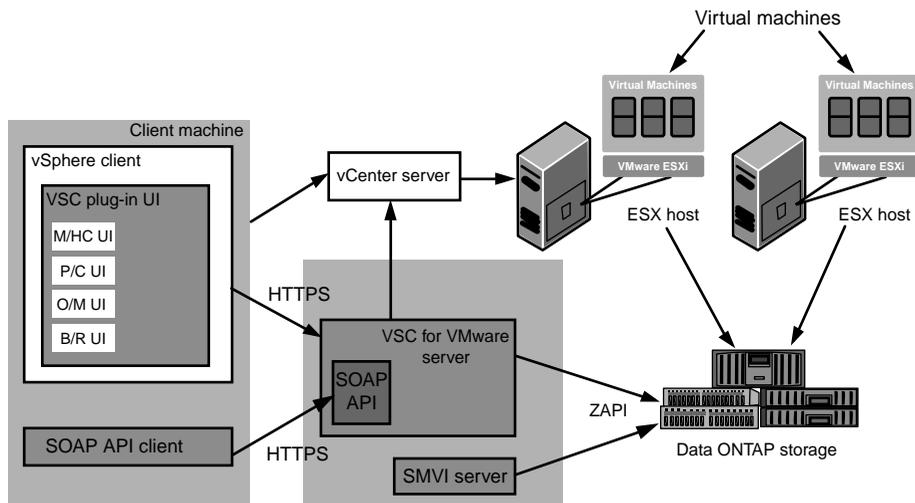
The VSC architecture includes the storage system running Data ONTAP, the vCenter Server, the vCenter client, the ESX or ESXi host, and the VSC capabilities.

VSC uses VMware-recommended web-based architecture. It consists of two major components:

- A graphical user interface (GUI) web application that displays as a plug-in within the vSphere client to provide a single management console for virtualized environments.
- A server component that is controlled by the VSC service and hosts Java Servlets to handle the GUI and API calls to and from the storage systems and the ESX/ESXi hosts.

The following diagram provides a high-level overview of the VSC architecture.

**Note:** In a production environment, you might have multiple vSphere clients on different machines. You might also put the SOAP API client on different machine than the one that the vSphere client is on.



When you run VSC, you use the VMware vSphere client and the VMware vCenter server. VSC provides the following:

- Plug-ins for the four capabilities: Monitoring and Host Configuration, Provisioning and Cloning, Optimization and Migration, and Backup and Recovery. Each capability has its own user interface and online help.
- The VSC server
- The SMVI server

In addition, you can write applications that communicate with the VSC server. For example, you can use the SOAP API to create a client. VSC supports the VSC API, which works in a VMware vCloud Director deployment and the Provisioning and Cloning API.

The vSphere client and any applications you create use the HTTPS protocol to communicate. The VSC server and the SMVI server use ZAPI to communicate with the storage systems that are running Data ONTAP. Communication from the VSC server to the vCenter is done using SOAP.

The vCenter server communicates with the physical servers where ESX or ESXi hosts are running. You can have multiple virtual machines running on the ESX or ESXi hosts. Each virtual machine can run an operating system and applications. The ESX and ESXi hosts then communicate with the storage systems.

## Methods for accessing the VSC for VMware vSphere

Virtual Storage Console for VMware vSphere provides a graphical user interface (GUI) that enables you to move between tasks. You access the GUI by clicking the NetApp icon in the vCenter Server.

**Note:** If you do not have the View privilege, you get an error message when you click the NetApp icon.

You can also right-click an object in the Inventory panel and then select the NetApp tab. From that tab, you can select the action you want to perform.

## VSC for VMware vSphere provides lock management for system resources

Virtual Storage Console for VMware vSphere uses lock management to keep multiple capabilities from performing simultaneous tasks on the same target datastores or virtual machines. As a result, certain alignment, migration, provisioning or cloning, and backup and recovery features become temporarily unavailable if the target datastore or virtual machine is being used by another capability.

You can see the lock management process when you are working with Provisioning and Cloning, Optimization and Migration, or Backup and Recovery . For example, if you are migrating virtual machines using Optimization and Migration, you cannot use Provisioning and Cloning to clone one of the virtual machines until the migration is complete.

You cannot perform the following Backup and Recovery tasks if Provisioning and Cloning or Optimization and Migration is using the target datastore or virtual machine:

- Create on-demand backup copies of individual virtual machines, datastores, or a datacenter.
- Schedule automated backup copies of individual virtual machines, datastores, or a datacenter.
- Recover a datastore, virtual machine, or virtual disk file.
- Mount a backup for a file restore session.
- Unmount a backup that was previously mounted for a file restore session.

**Note:** When a lock occurs on a mount or unmount operation for a file restore session, the lock is held from when the backup is mounted to the virtual machine until the backup is unmounted.

Conversely, if Backup and Recovery has already started one of the listed operations, the datastore or virtual machine is unavailable for Provisioning and Cloning or Optimization and Migration operations.

### Related concepts

[Provisioning and cloning datastores and virtual machines](#) on page 68

[Backing up and restoring data](#) on page 112

## Online Help

The Virtual Storage Console for VMware vSphere GUI has a separate online Help module for each capability. Each online Help module describes the fields and commands for its capability.

You can access online Help for all the VSC capabilities from the vSphere Client **Help** menu:

- **Help > NetApp > Monitoring and Host Configuration > Monitoring and Host Configuration Help**
- **Help > NetApp > Provisioning and Cloning > Provisioning and Cloning Help**
- **Help > NetApp > Optimization and Migration > Optimization and Migration Help**
- **Help > NetApp > Backup and Recovery > Backup and Recovery Help**

The Help information is displayed in a web browser.

## Installation overview

---

You can install Virtual Storage Console for VMware vSphere on a 32-bit or 64-bit Windows server.

For information about which versions of Windows are supported, see the Interoperability Matrix, which is available online at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability).

**Note:** Do not install this software on a client computer.

### Installation guidelines

To install the Virtual Storage Console (VSC) software, follow these guidelines:

- VSC must be installed on a local disk of the Windows server. Do not attempt to install VSC on a network share.
- The network must be connected between the Windows server running VSC and the management ports of the storage controllers, the ESX/ESXi hosts, and the vCenter Server.
- A reboot is not required to complete the installation. However, vSphere clients must be closed and restarted to be able to display the VSC plugin.
- At a minimum, the display must be set to 1,280 by 1,024 pixels to view VSC pages correctly.

### Installing the capabilities

By default, VSC installs the following three capabilities:

- Monitoring and Host Configuration
- Provisioning and Cloning
- Optimization and Migration

You have the option of also installing Backup and Recovery . If you choose Backup and Recovery, you must purchase a license for SnapManager for Virtual Infrastructure.

### Software licenses

The following software licenses might be required for VSC depending on which capabilities you use:

- The required protocol license (NFS, FCP, iSCSI)
- SnapManager for Virtual Infrastructure (if installing Backup and Recovery )
- SnapMirror (if using Provisioning and Cloning or if using the SnapMirror update option in Backup and Recovery )
- SnapRestore (if installing Backup and Recovery )
- A\_SIS (if using Provisioning and Cloning when configuring deduplication settings)
- MultiStore (if using Provisioning and Cloning and working with vFiler units)
- FlexClone

The FlexClone license is required in the following situations:

- You are using Provisioning and Cloning to clone virtual machines.

- You are using Backup and Recovery in NFS environments and running a version of Data ONTAP prior to 8.1.

You do not need a FlexClone license if you are running Backup and Recovery in NFS environments with one of the following versions of Data ONTAP:

- Data ONTAP 8.1 operating in 7-Mode
- Clustered Data ONTAP 8.1.1 or later

## Additional setup requirements when using RBAC

New features in version 4.2 of Virtual Storage Console for VMware vSphere mean that administrators must perform additional tasks to ensure that VSC is set up correctly.

VSC administrators must make an installation plan that includes setup for the following features:

- vCenter Server and Data ONTAP role-based access control (RBAC)
- Centralized storage management

These features require that administrators provide users with appropriate RBAC permissions, and that they verify the credentials.

### Set up access for users

Administrators must set up vCenter Server permissions for users so that they either have or do not have access to certain VSC tasks for certain vSphere objects. In addition, they must set up the Data ONTAP privileges, which provide the credentials used by the storage systems.

To simplify the process of creating vCenter Server user roles, VSC provides several standard VSC roles for key tasks. These roles contain all the VSC-specific and native vCenter Server privileges required for the tasks. As administrator, you can assign these roles to users.

**Note:** You should not edit the standard roles that VSC provides. These roles return to their default settings each time you restart the VSC service. This means that any changes you made to these roles will be lost if you modify or upgrade your installation. If you need a privilege that these roles do not provide, you can create a role containing that privilege and then use the Groups feature to combine that role with the appropriate standard VSC role. Or you can clone the standard VSC role and then modify the cloned role to meet your needs.

Unless your company's security policies require more restrictive permissions, it is a good practice to assign permissions on the root object (also referred to as the root folder). Then, if you need to, you can restrict those entities that you do not want to have permissions.

### Ensure that all VSC users have the View privilege correctly assigned

The VSC-specific View privilege is read-only and enables users to see the menus, tabs, and other elements of the VSC interface. If a user who does not have this privilege clicks on the a NetApp icon from the Home View, VSC displays an error message.

In addition, the level at which you assign this privilege to a user affects which part of the VSC GUI that user can see and access.

It is a good practice to assign the permission containing this privilege to the root object.

### **Confirm that storage systems and their credentials are available through Monitoring and Host Configuration**

Monitoring and Host Configuration manages storage discovery and credentials for the entire VSC product. All VSC capabilities use the credentials entered during the initial setup and discovery process in Monitoring and Host Configuration.

Having Monitoring and Host Configuration manage credentials for all of VSC eliminates the need for you to manage separate sets of storage entities in Backup and Recovery.

After you install Backup and Recovery, you should check the Backup and Recovery setup screen to see if there are any issues with the controller credentials. If this screen reports that a storage controller does not have the necessary permissions to perform Backup and Recovery operations, you need to go to Monitoring and Host Configuration to correct the issues.

## **VSC for VMware vSphere supported configurations**

Virtual Storage Console for VMware vSphere is supported on specific releases of ESX/ESXi and Data ONTAP software.

### **Server configuration**

Your Windows system must meet minimum hardware requirements before installing the VSC software.

The memory requirements depend on whether you install VSC on the same machine as the vCenter Server or on a different machine. When this document was created, the memory requirements for 32-bit environments and 64-bit environments where VSC was installed on a separate machine were the following:

- Minimum memory requirement: 4 GB RAM
- Recommended memory requirement: 4 GB RAM

Hardware requirements are higher if you are running VSC on the same machine as the vCenter Server.

**Note:** Please refer to VMware documentation for the current list of hardware requirements.

You should be aware of the following requirements before you install the VSC software:

- Supported Microsoft Windows software
- vCenter Server requirements
- ESX host software requirements
- Data ONTAP requirements

See the Interoperability Matrix, which is online at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability), for details.

**Note:** IPv6 is not supported on VSC. If the server on which you are installing VSC has IPv6 enabled, you should disable IPv6 before installing VSC. IPv6 should not be re-enabled after VSC is installed.

If you are using the single file restore (SFR) feature, you might need to change a configuration setting to prevent the SFR restore session email from sending an IPv6 address. See the Known issues: Backup and Recovery section of the *Virtual Storage Console for VMware vSphere Release Notes* for the configuration workaround. The Release Notes are posted on the NetApp Support Site, which is available online at [support.netapp.com](http://support.netapp.com).

### Client configuration

The client computer that runs the VMware vSphere Client software must have Microsoft Internet Explorer 8 or later installed.

## Installing VSC for VMware vSphere using the installation wizard

You can use the installation wizard to install Virtual Storage Console for VMware vSphere. By default, the VSC software installs Monitoring and Host Configuration, Optimization and Migration, and Provisioning and Cloning. It gives you the option of installing Backup and Recovery.

### Before you begin

You must be logged on as a user with administrator privileges to the machine on which you install VSC.

In addition, you must use the correct installer package for your host machine. You cannot run the 32-bit installer on a 64-bit host machine or the 64-bit installer on a 32-bit host machine.

### Steps

1. Download the VSC installer.
2. Double-click the installer icon, and click **Run** to start the installation wizard.
3. Follow the instructions in the installation wizard to install the software.

**Note:** If you want to install Backup and Recovery, you must select that option. Otherwise, the VSC installer does not install it.

4. Click **Finish** to complete the installation.
5. At the web page that appears when the installation is complete, register VSC with the vCenter Server.

You must provide the vCenter Server host name or IP address and the administrative credentials.

**Note:** To register VSC with the vCenter Server, you must have administrator privileges for your Windows login.

### After you finish

When you install Backup and Recovery, you should check the Backup and Recovery setup screen after the installation completes to see if there are any issues with the controller credentials. If this screen reports that a storage controller does not have the necessary permissions to perform Backup and Recovery operations, you need to go to Monitoring and Host Configuration to correct the issues.

If you uninstalled SMVI before you installed VSC, you must migrate the backup metadata after you complete the installation. However, before you migrate the backup metadata, you must first save it by copying the contents that were just installed in the `C:\Program Files\NetApp\SMVI\server\repository` folder into the `C:\Program Files\NetApp\Virtual Storage Console\smvi\server\repository` folder.

After you finish copying the backup metadata, you must copy the contents of the credential file from the `C:\Program Files\NetApp\SMVI\server\etc\cred` folder into the `C:\Program Files\NetApp\Virtual Storage Console\smvi\server\etc\cred` folder.

Then you must restart the VSC service.

### Related concepts

[Storage system discovery and credentials overview](#) on page 46

## Installing VSC for VMware vSphere using silent mode

You can install Virtual Storage Console for VMware vSphere using silent mode instead of the installation wizard. When you use silent mode, you can enter a command line that lets you automatically install all the capabilities at once.

### Before you begin

You must be logged on with administrator privileges to the machine where you are installing VSC.

In addition, you should make sure you are using the correct installer package for your host machine. You cannot run the 32-bit installer on a 64-bit host machine or the 64-bit installer on a 32-bit host machine.

### Steps

1. Use the following command format to install VSC: `installer.exe /s /v"/qn /Li logfileADDLOCAL=ALL INSTALLDIR=\"installation path\""`

This command installs all the VSC capabilities.

### Example

The following is a sample command line for a 64-bit host machine:

```
VSC-4.2-win64.exe /s /v"/qn /Li install.log ADDLOCAL=ALL INSTALLDIR=\"C:\Program Files\NetApp\Virtual Storage Console\" "
```

2. At the web page that appears when the installation is complete, register VSC with the vCenter Server.

You must provide the vCenter Server host name or IP address and the administrative credentials.

**Note:** To register VSC with the vCenter Server, you must have administrator privileges for your Windows login.

### After you finish

When you install Backup and Recovery, you should check the Backup and Recovery setup screen after the installation completes to see if there are any issues with the controller credentials. If this screen reports that a storage controller does not have the necessary permissions to perform Backup and Recovery operations, you need to go to Monitoring and Host Configuration to correct the issues.

If you uninstalled SMVI before you installed VSC, you must migrate the backup metadata after you complete the installation. However, before you migrate the backup metadata, you must first save it by copying the contents that were just installed in the C:\Program Files\NetApp\SMVI\server\repository folder into the C:\Program Files\NetApp\Virtual Storage Console\smvi\server\repository folder.

After you finish copying the backup metadata, you must then copy the contents of the credential file from the C:\Program Files\NetApp\SMVI\server\etc\cred folder into the C:\Program Files\NetApp\Virtual Storage Console\smvi\server\etc\cred folder.

Then you must restart the VSC service.

### Related concepts

[Storage system discovery and credentials overview](#) on page 46

## Installing or uninstalling Backup and Recovery

The installation wizard lets you install or uninstall Backup and Recovery separately from the other capabilities. The other three capabilities can only be installed or uninstalled as a group. Backup and

Recovery can either be installed or uninstalled at the same time you install or uninstall the other capabilities or at a separate time.

### About this task

The steps that follow describe how to install Backup and Recovery after you have already installed the other capabilities. If you want to install all the capabilities at the same time, you must select the Backup and Recovery check box when you first run the installation wizard.

You can also use this check box to uninstall only Backup and Recovery and leave the other Virtual Storage Console for VMware vSphere capabilities installed.

### Steps

1. On the Windows server where you installed the VSC software, select **Control Panel > Add/Remove Programs** (Windows Server 2003) or **Control Panel > Programs and Features** (Windows Server 2008).
2. Select **VSC for VMware vSphere** and click **Change** to start the installation wizard.
3. In the installation wizard, select the **Modify** option and click **Next**.
4. Select the check box for Backup and Recovery if you want to add it, or clear the check box if you want to remove it, and click **Next**.
5. If you are installing Backup and Recovery, click **Install** to start the installation.
6. Click **Finish** to complete the installation or removal of Backup and Recovery.

### After you finish

You must close the vSphere Client and restart it to either display any newly-installed capability or remove it from the GUI.

If you install Backup and Recovery separately from the other capabilities, you must restart the VSC service after the installation is complete.

You should check the Backup and Recovery setup screen after the installation completes to see if there are any issues with the controller credentials. If this screen reports that a storage controller does not have the necessary permissions to perform Backup and Recovery operations, you need to go to Monitoring and Host Configuration to correct the issues.

### Related concepts

*[Backing up and restoring data](#)* on page 112

## Registering VSC for VMware vSphere with vCenter Server

After installing the Virtual Storage Console for VMware vSphere software, you must register it with vCenter Server. By default, the registration web page opens when the VSC for VMware vSphere installation is complete.

### Before you begin

You must be logged on as a user with administrator privileges to the machine on which you install VSC. If you attempt to register VSC without having administrator privileges, the task does not complete correctly.

### About this task

IPv6 addresses are not currently supported.

### Steps

1. If the registration web page does not open automatically, type the following URL in a web browser:

```
https://localhost:8143/Register.html
```

`localhost` must be the computer where you installed VSC. If you are not performing this step from the computer where you installed VSC, you must replace `localhost` with the hostname or IP address of that computer.

2. If a security certificate warning appears, choose the option to ignore it or to continue to the web site.
3. In the Plugin service information section, select the IP address that the vCenter Server uses to access VSC.  
  
This IP address must be accessible from the vCenter Server. If you installed VSC on the vCenter Server computer, this might be the same address as the one you use to access the vCenter Server.
4. Type the host name or IP address of the vCenter Server and the administrative credentials for the vCenter Server.
5. Click **Register** to complete the registration.

If you did not enter the correct user credentials for the vCenter Server, a registration failed error message appears. If that happens, you must repeat the previous steps.

6. Close the registration page after you complete the registration process because the web page is not automatically refreshed.

## VSC for VMware vSphere port requirements

By default, VSC for VMware vSphere uses designated ports to enable communication between its components, which include storage systems and the VMware vCenter server. If you have firewalls enabled, you must ensure that the firewalls are set to allow exceptions.

For firewalls other than Windows, you must manually grant access to specific ports that VSC for VMware vSphere uses. If you do not grant access to these ports, an error message such as `Unable to communicate with the server` appears.

VSC for VMware vSphere uses the following default ports:

Default port number	Description
443	The VMware vCenter Server and the storage systems listen for secure communications using secure HTTP (SSL) on this port.
80	The VMware vCenter server and the storage systems listen for standard, unencrypted communication via standard HTTP on this port.
8143	VSC for VMware vSphere listens for secure communication on this port.
8043	The Backup and Recovery capability listens for secure communication on this port. The Backup and Recovery Restore Agent and CLI also use this port.

## Regenerating an SSL certificate for VSC for VMware vSphere

The SSL certificate is generated when you install VSC for VMware vSphere. The distinguished name (DN) generated for the SSL certificate might not be a common name (CN) ("NetApp") that the client machines recognize. By changing the keystore and private key passwords, you can regenerate the certificate and create a site-specific certificate.

### Steps

1. Before you generate a new certificate, you should stop the Virtual Storage Console for VMware vSphere (vsc) service.

There are several ways to do this. One way to stop the service is to use the Windows Services control panel.

2. Connect to the Windows console session or the Windows PowerShell console.

3. Go to the VSC for VMware vSphere installation directory and enter the following command:

```
bin\vscssl setup -cn <HOST>
```

For *<HOST>*, enter the host name of the system running VSC for VMware vSphere or a fully-qualified domain name of the system running VSC for VMware vSphere.

### Example

The following example executes the command from the installation directory and uses a host called ESXiTester:

```
C:\Program Files\NetApp\Virtual Storage Console>bin\nvpf ssl setup -cn
ESXiTester
```

4. At the prompt, enter the default keystore password:

```
changeit
```

You will also be prompted to enter a password for the private key (this can be any string you choose).

**Note:** If you do not have Java JRE version 1.6.0\_21 installed in Program Files, your command will look different.

The following files are generated:

- keystore file (default: `etc\nvpf.keystore`)  
This is the JKS keystore file.
  - keystore properties (default: `etc\keystore.properties`)  
This file contains the keystore file path and the keystore and key passwords. The administrator should secure this file and specify `http.ssl.keystore.properties` in `etc\nvpf.override` if the keystore properties file needs to be moved.
5. If you are using the Provisioning and Cloning capability or the Optimization and Migration capability, perform the following two steps:

- a) Change to the VSC for VMware vSphere installation directory and enter the following command:

```
keytool -export -alias nvpf -keystore nvpf.keystore -file nvpf.cer
```

### Example

The following example executes the command from the `etc` directory in the installation directory:

```
C:\Program Files\NetApp\Virtual Storage Console\etc>keytool -export -
alias nvpf -keystore nvpf.keystore -file nvpf.cer
```

The command creates a new file called `nvpf.cer`.

- b) Import the certificate to the local a Java keystore by entering the command:

```
c:\Program Files\NetApp\Virtual Storage Console\etc>keytool -import -
alias nvpf -file nvpf.cer -keystore "c:\Program Files\Java
\jdk1.6.0_21\jre\lib\security\cacerts"
```

### After you finish

- You must secure the `etc\keystore.properties` file and then restart the `vsc` service. There are several ways to do this, including the following:
  - If the installation directory is on a network share directory, move the file to local storage.
  - Move the file to storage accessible only to the SYSTEM user, which keeps unauthorized users from being able to view or modify the file.
- You can review and accept the SSL certificate after the vSphere Client receives the certificate when you click the NetApp icon in the vSphere Client. You can then import the SSL certificate into the Trusted Root Certification Authorities store to prevent SSL security warnings from appearing every time you launch the vSphere client. For details, see the documentation for your Windows operating system.

## Preparation required before upgrading VSC for VMware vSphere

Before you upgrade Virtual Storage Console for VMware vSphere, you should record information about your storage systems, the credentials being used, and the preferences being used. Keeping a record of the storage resources enables you to restore that information after you upgrade, if necessary.

After you perform an upgrade, you must check the following storage resources and update them if needed:

- **The storage systems available to the capabilities**  
You must make sure that all the storage systems that were available before the upgrade are still available.  
After an upgrade, Monitoring and Host Configuration exports to the other VSC capabilities all the storage systems with valid credentials that either were
  - Automatically discovered by Monitoring and Host Configuration.
  - Manually added to Monitoring and Host Configuration.

In some cases, storage systems that were available to a capability before an upgrade are not there after an upgrade. If this happens, you must either get Monitoring and Host Configuration to discover the storage systems or manually add them to Monitoring and Host Configuration

**Note:** If the storage system does not have storage mapped to an ESX/ESXi host that a vCenter Server is managing, Monitoring and Host Configuration does not automatically discover it.

- **The storage credentials required for Backup and Recovery**  
In previous versions of VSC, Backup and Recovery managed its storage credentials separately from the rest of VSC. Starting with VSC 4.2, Monitoring and Host Configuration manages the

credentials for Backup and Recovery as well as the other VSC capabilities. This means that all the VSC capabilities use the same credentials for each of the storage systems. You cannot have one set of credentials for Provisioning and Cloning, one set for Optimization and Migration, and another set for Backup and Recovery . You enter the credentials at the Overview panel for Monitoring and Host Configuration.

If you used Backup and Recovery in an earlier version of VSC, you should record the storage systems and credentials you used with it. After you upgrade to VSC 4.2, you must use Monitoring and Host Configuration to check the current storage controllers and confirm the following:

- All the storage controllers that you used with Backup and Recovery in earlier versions of VSC are included in the Monitoring and Host Configuration list.

Because Monitoring and Host Configuration did not discover storage controllers used by Backup and Recovery prior to VSC 4.2, those storage controllers might not appear in the list of controllers that Monitoring and Host Configuration manages.

In that case, you must manually add those storage controllers to Monitoring and Host Configuration.

- All storage controllers used for Backup and Recovery operations and listed with Monitoring and Host Configuration have at a minimum all the credentials required to perform those operations.

Monitoring and Host Configuration manages all of the storage controllers for VSC, which means it is possible that a storage controller could have the correct privileges for one capability, but not Backup and Recovery. If that is the case, you must use Monitoring and Host Configuration to manually add the credentials required for Backup and Recovery operations.

- **The standard roles provided by VSC**

VSC provides several VSC-specific roles that enable users to complete standard tasks as well as a read-only role that allows users to view information, but not take any action. These roles are VSC Administrator, VSC Read-only, VSC Provision, VSC CLone, VSC Scan/Migrate, VSC Backup, and VSC Restore.

Each time you install VSC, the standard roles are automatically upgraded to contain the privileges necessary to work with that version of VSC. If there are no changes to the standard roles, then they are returned to their default values. In addition, each time you restart the VSC Windows service or modify your VSC installation, the standard roles return to their default values.

You should not modify the VSC standard roles. If you make any changes to these roles, you will lose those changes when you upgrade or modify your VSC installation or restart the VSC Windows service.

- **The `etc/kamino/kaminoprefs.xml` preferences file**

Each time you install VSC, it overwrites the existing preferences file with a new preferences file. If you made changes to the preferences file for any of the capabilities, you should record those changes so that you can enter them again after you install a new version of VSC. The upgrade process does not record the changes.

## Upgrading VSC for VMware vSphere

If you currently have version 2.1.2 or higher of Virtual Storage Console for VMware vSphere, installed you can upgrade VSC. The VSC installer checks the version numbers of each of the currently installed capabilities to determine whether you are upgrading to a newer version.

### Before you begin

The VSC installer automatically upgrades Monitoring and Host Configuration, Optimization and Migration, and Provisioning and Cloning to the newer versions. If you also have Backup and Recovery installed, the VSC installer upgrades it as well. If you do not have Backup and Recovery installed, the VSC installer gives you the option of installing it.

The VSC installer does not support upgrades from the following:

- VSC prior to 2.1.2
- A standalone version of Rapid Cloning Utility (RCU), which is now called the Provisioning and Cloning capability
- A standalone version of SnapManager for Virtual Infrastructure (SMVI), which is now called Backup and Recovery.

If you have any of that software installed, you must uninstall it before you can install the current version of VSC. If the VSC installer finds RCU, or SMVI on the server, it prompts you to uninstall the software, and then aborts. It also aborts if you have VSC 1.x installed.

You must be logged on with administrator privileges to the machine where you installing VSC.

### Steps

1. Download the installer for VSC.
2. Double-click the installer icon, and click **Run** to start the installation wizard.
3. Click **Yes** on the confirmation prompt.
4. In the installation wizard, select the capabilities that you want to upgrade and click **Next** to start the installation.

The installation might take several minutes.

5. Click **Finish** to complete the installation.
6. At the web page that appears when the installation is complete, register VSC with the vCenter Server.

You must provide the vCenter Server host name or IP address and the administrative credentials.

**Note:** To register VSC with the vCenter Server, you must have administrator privileges for your Windows login.

### After you finish

When you install Backup and Recovery, you should check the Backup and Recovery setup screen after the installation completes to see if there are any issues with the controller credentials. If this screen reports that a storage controller does not have the necessary permission to perform Backup and Recovery operations, you need to go to Monitoring and Host Configuration to correct the issues.

### Related concepts

[Storage system discovery and credentials overview](#) on page 46

## Uninstalling VSC for VMware vSphere using Add/Remove Programs

You can uninstall the VSC for VMware vSphere software from your system using the Windows Add or Remove Programs list.

### About this task

The uninstall program removes the entire VSC for VMware vSphere application. You cannot specify which capabilities you want to uninstall.

### Steps

1. On the Windows server where you installed the VSC for VMware vSphere software, select **Control Panel > Add/Remove Programs** (Windows Server 2003) or **Control Panel > Programs and Features** (Windows Server 2008).
2. Select Virtual Storage Console for VMware vSphere and click **Remove** to immediately remove the program or click **Change** to start the installation wizard.
3. If you select **Change**, then click **Yes** to confirm that you want to remove the program.
4. In the installation wizard, select the **Remove** option and click **Next**.
5. Click **Remove** to uninstall the VSC for VMware vSphere software.

After the process completes, a confirmation prompt is displayed.

**Note:** At the confirmation prompt, click **Yes** to remove all the metadata files from the installation directory or click **No** so that you can manually delete the files in the directory.

## Uninstalling VSC for VMware vSphere using silent mode

You can uninstall Virtual Storage Console for VMware vSphere using silent mode instead of the Windows Add/Remove Program. When you use silent mode, you can enter a command line that lets you automatically uninstall all the capabilities at once.

### Before you begin

You must be logged on with administrator privileges to the machine from which you are uninstalling VSC.

### Step

1. Use the following command line to uninstall VSC: `installer.exe /s /v"/qn /Li logfileREMOVE=ALL INSTALLDIR="installation path\""`

This command removes all the VSC capabilities.

### Example

The following is an example of the command line you might use if you were uninstalling VSC 4.2 from a 64-bit host machine:

```
VSC-4.2-win64.exe /s /v"/qn /Li uninstall.log REMOVE=ALL
```

# Authentication and user management with vCenter RBAC and Data ONTAP RBAC

---

Role-based access control (RBAC) is a process that enables administrators to control access to and user actions on vSphere objects and storage systems running Data ONTAP.

Virtual Storage Console for VMware vSphere supports both vCenter Server RBAC and Data ONTAP RBAC. Depending on your system setup, you might have different administrators handling these types of RBAC:

- **vCenter Server RBAC**

This security mechanism restricts the ability of vSphere users to perform VSC tasks on vSphere objects, such as virtual machines, datastores, and datacenters.

The vSphere administrator sets up vCenter Server RBAC by assigning permissions to specific vSphere objects, which are listed in the vSphere inventory. In many cases, a VSC task requires that more than one object have permissions. For this reason, it is a good practice to assign permissions on the root object (also referred to as the root folder). You can then restrict those entities that do not need permissions.

**Note:** At a minimum, all users must have the VSC-specific, read-only View privilege assigned to them. Without this privilege, users cannot access the VSC GUI.

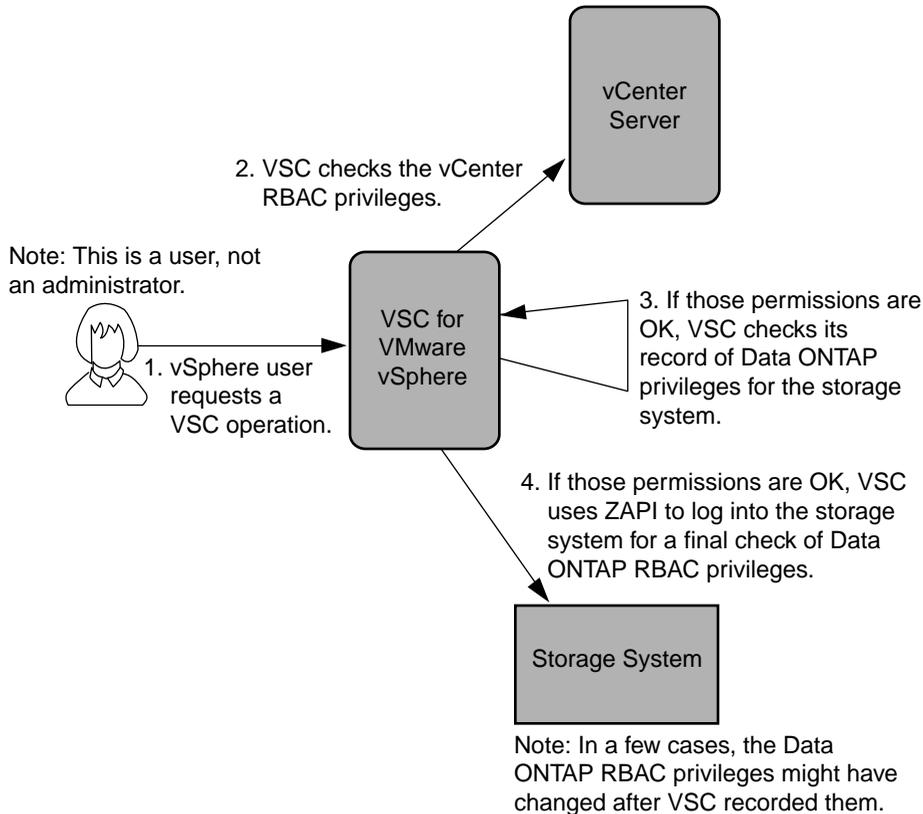
- **Data ONTAP RBAC**

This security mechanism restricts the ability of VSC to perform specific storage operations, such as creating, destroying, or backing up storage for datastores, on a specific storage system.

The storage administrator sets up Data ONTAP RBAC by defining storage credentials consisting of a user name and password in Data ONTAP. The storage credentials map to VSC storage operations. Then the administrator, usually the storage administrator, sets the storage credentials in Monitoring and Host Configuration for each storage system managed by VSC. VSC uses a single set of credentials for each storage system.

VSC checks the vCenter Server RBAC permissions when a user clicks a vSphere object and initiates an action. If a user has the correct vCenter Server RBAC permission to perform that task on that vSphere object, VSC then checks the Data ONTAP credentials for the storage system. If those credentials are also confirmed, then VSC allows the user to perform that task.

The following diagram provides an overview of the VSC validation workflow for RBAC privileges (both vCenter and Data ONTAP):



## vCenter Server role-based access control features in VSC for VMware vSphere

vCenter Server provides role-based access control (RBAC) that enables you to control access to vSphere objects. In Virtual Storage Console for VMware vSphere, vCenter Server RBAC works with Data ONTAP RBAC to determine which VSC tasks a specific user can perform on objects on a specific storage system.

If you do not have the appropriate vCenter Server RBAC permissions, you cannot perform any VSC tasks. When you initiate a task, VSC checks a user's vCenter Server permissions before checking the user's Data ONTAP privileges.

You can set the vCenter Server permissions on the root object (also known as the root folder). You can then refine the security by restricting child entities that do not need those permissions.

## Components that make up vCenter Server permissions

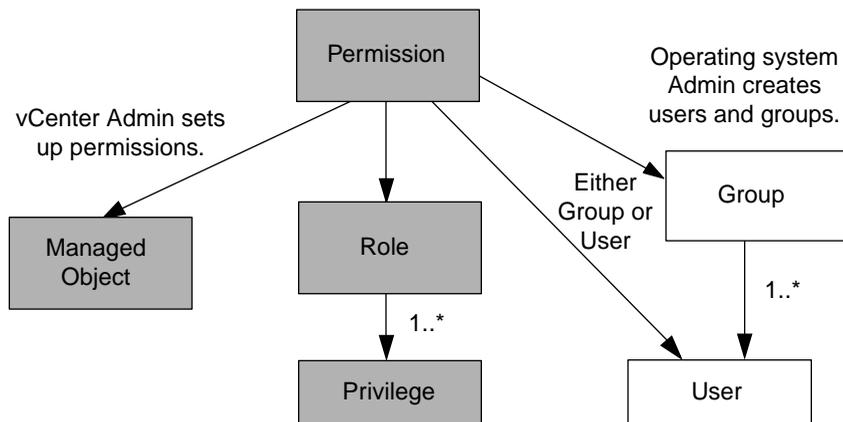
The vCenter Server recognizes permissions, not privileges. Each vCenter Server permission consists of three components.

These components are the following:

- One or more privileges (the role)  
The privileges define the tasks that a user can perform.
- a vSphere object  
The object is the target for the tasks.
- A user or group  
The user or group defines who can perform the task.

As the following diagram illustrates, you must have all three elements in order to have a permission.

**Note:** In this diagram, the gray boxes indicate components that exist in the vCenter Server and the white boxes indicate components that exist in the operating system where the vCenter Server is running.



## Privileges

From the perspective of working with Virtual Storage Console for VMware vSphere, there are two kinds of privileges:

- Native vCenter Server privileges  
These privileges come with the vCenter Server.
- VSC-specific privileges  
These privileges were defined for specific VSC tasks. They are unique to VSC.

**Note:** To make this document easier to read, it refers to the vCenter Server privileges as native privileges and the privileges defined for VSC as VSC-specific privileges. For detailed information about VSC-specific privileges, see KB article 1013941 *How to configure Storage and vCenter*

*RBAC for VSC for VMware vSphere* at [kb.netapp.com/support/index?page=content&id=1013941](http://kb.netapp.com/support/index?page=content&id=1013941). For information about vCenter Server native privileges, see VMware's *vSphere Security* guide. At the time this document was created, that guide was online at <http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-security-guide.pdf>. NetApp follows the VMware recommendations for creating and using permissions.

VSC tasks require both VSC-specific privileges and vCenter Server native privileges. These privileges make up the "role" for user. A permission can have multiple privileges.

**Note:** To simplify working with vCenter Server RBAC, VSC provides several standard roles that contain all the required VSC-specific and native privileges to perform VSC tasks.

If you change the privileges within a permission, the user associated with that permission should **log out and then log back in** to enable the updated permission.

## vSphere objects

Permissions are associated with vSphere objects, such as the vCenter Server, ESXi hosts, virtual machines, datastores, datacenters, and folders. You can assign permissions to any vSphere object. Based on the permission assigned to a vSphere object, the vCenter Server can lock down who can perform which tasks on that object.

## Users and groups

You can use Active Directory (or the local vCenter Server machine) to set up users and groups of users. You can then use vCenter Server permissions to grant access to these users or groups to enable them to perform specific VSC tasks.

**Note:** These vCenter Server permissions apply VSC vCenter users, not VSC administrators. By default, VSC administrators have full access to the product and do not need to have permissions assigned to them.

Users and groups do not have roles assigned to them. They gain access to a role by being part of a vCenter Server permission.

You can assign only one permission to a vCenter user or group. You can, however, set up high-level groups and assign a single user to multiple groups. Doing that allows the user to have all the permissions provided by the different groups. In addition, using groups simplifies the management of permissions by eliminating the need to set up the same permission multiple times for individual uses.

## Key points about assigning and modifying permissions

There are several key points to keep in mind when you are working with vCenter Server permissions. How well a Virtual Storage Console for VMware vSphere task succeeds can depend on where you assigned a permission or what actions a user took after a permission was modified.

You only need to set up vCenter Server permissions if you want to limit access to vSphere objects and tasks. Otherwise, you can log in as an administrator. This login automatically allows you to access all vSphere objects.

**Note:** For detailed information on working with vCenter Server permissions, see the VMware *vSphere Security* guide. At the time this document was created, that guide was online at <http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-security-guide.pdf>. NetApp follows the VMware recommendations for creating and using permissions.

## Assigning permissions

Where you assign a permission determines the VSC tasks that a user can perform.

Sometimes, to ensure that a task completes, you must assign the permission at a higher level, such as the root object. This is the case when a task requires a privilege that does not apply to a specific vSphere object (for example, tracking the task) or when a required privilege applies to a non-vSphere object (for example, a storage system).

In these cases, you can set up a permission so that it is inherited by the child entities. You can also assign other permissions to the child entities. The permission on a child entity always overrides the permission inherited from the parent entity. This means that you can assign child entity permissions as a way to restrict the scope of a permission that was assigned to a root object and inherited by the child entity.

**Tip:** Unless your company's security policies require more restrictive permissions, it is a good practice to assign permissions on the root object (also referred to as the root folder). Then, if you need to, you can restrict those entities that you do not want to have the permission so that you have more fine-grained security.

## Permissions and non-vSphere objects

In some cases, a permission applies to a non-vSphere object. For example, a storage system is not a vSphere object. If a privilege applies to a storage system, you must assign the permission containing that privilege to the VSC root object because there is no vSphere object to which you can assign it.

For example, any permission that includes a privilege such as the Monitoring and Host Configuration privilege "Add/Modify/Skip storage systems" must be assigned at the root object level.

## Modifying permissions

You can modify a permission at any time.

If you change the privileges within a permission, the user associated with that permission should **log out and then log back in** to enable the updated permission.

## Advanced example of using vCenter Server permissions

The privileges you include in a vCenter Server permission determine the role that the Virtual Storage Console for VMware vSphere user has and the tasks associated with that role. In addition, the

vSphere object to which you assign the permission also affects the tasks the user can perform. This example illustrates those points.

**Note:** Depending on your system setup and company's security requirements, you might either log in as administrator and not use vCenter Server RBAC, or you might set permissions on the root object (folder). Setting permissions on the root object normally allows all the child objects to inherit those permissions, unless you place a restriction on a child object to exclude it.

If you plan to use a more restrictive implementation of RBAC, the following example illustrates how assigning permissions can affect which tasks you can complete. This example uses two datacenters (Datacenter A and Datacenter B) that are managed by a single vCenter Server. Users of one datacenter are not allowed to perform tasks in the other datacenter.

For Datacenter A, you create the following vCenter Server permission:

- A user named "Pat"
- The vCenter Server privileges, both VSC-specific and vCenter Server native, required to allow the user to perform rapid cloning
- A managed object within the vSphere inventory (Datacenter A), which has the "Propagate to Child Objects" check box selected

Because the permissions for rapid cloning are assigned to Datacenter A, VSC displays an error message if you attempt to clone a virtual machine from Datacenter B.

If you attempt to clone a virtual machine in Datacenter A, VSC allows you to complete the Rapid Clone Wizard steps. However, when you click **Apply**, the Rapid Clone Wizard disappears without completing the cloning task. The log file includes a message stating that you do not have permission to create a task.

This clone task failed because the permission was applied to Datacenter A instead of the root object (folder). The task required the native vCenter Server privilege "Create Task." Permissions containing the "Create Task" privilege must always be applied at the root object level. During the cloning task, Provisioning and Cloning confirmed that "Pat" had permission to work with a virtual machine in Datacenter A. But, as the clone task continued, Provisioning and Cloning asked the vCenter Server to create a task to track its progress. The vCenter Server rejected this request because user "Pat" did not have permission to create tasks on the root object. This caused the task to fail.

If you had assigned the permission at the root object, the privileges would have propagated down to the child vSphere objects, in this case Datacenter A and Datacenter B. You could have prevented "Pat" from working with a virtual machine in Datastore B by applying a permission restricting "Pat" from accessing Datastore B.

For detailed information on working with vCenter Server permissions, see VMware's *vSphere Security* guide. At the time this document was created, that guide was online at <http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcserver-51-security-guide.pdf>. NetApp follows the VMware recommendations for creating and using permissions.

## Standard roles packaged with VSC for VMware vSphere

To simplify working with vCenter Server privileges and role-based access control (RBAC), Virtual Storage Console for VMware vSphere provides a set of standard VSC roles that enable users to perform key VSC tasks. There is also a read-only role that allows users to view VSC information, but not perform any tasks.

The standard VSC roles have both the required VSC-specific privileges and the native vCenter Server privileges to ensure that tasks complete correctly. In addition, the roles are set up so that they have the necessary privileges across all supported versions of the vCenter Server.

As an administrator, you can assign these roles to the appropriate users.

**Note:** VSC returns these roles to their default values (initial set of privileges) each time you restart the VSC Windows service or modify your installation. If you upgrade VSC, the standard roles are automatically upgraded to work with that version of VSC.

You can see the VSC standard roles when you click Roles. VSC provides the following standard roles:

Role	Description
VSC Administrator	Provides all native vCenter Server and VSC-specific privileges necessary to perform all VSC tasks.
VSC Read-only	Provides read-only access to all of VSC. These users cannot perform any VSC actions that are access controlled.
VSC Provision	Provides all native vCenter Server and VSC-specific privileges necessary to provision storage. The user can perform the following tasks: <ul style="list-style-type: none"> <li>• Create new datastores</li> <li>• Destroy datastores</li> </ul> The user cannot perform the following tasks: <ul style="list-style-type: none"> <li>• Create clones</li> <li>• Reclaim space</li> <li>• Distribute templates</li> </ul>

Role	Description
VSC Clone	<p>Provides all native vCenter Server and VSC-specific privileges necessary to clone storage.</p> <p>The user cannot perform the following tasks:</p> <ul style="list-style-type: none"> <li>• Provision storage</li> <li>• Reclaim space</li> <li>• Distribute templates</li> </ul>
VSC Scan/Migrate	<p>Provides all native vCenter Server and VSC-specific privileges necessary to scan databases and migrate virtual machines.</p> <p>With this role, the user can perform all of the Optimization and Migration tasks.</p> <p>The user also has access to the configure privilege.</p>
VSC Backup	<p>Provides all native vCenter Server and VSC-specific privileges necessary to back up vSphere objects (virtual machines and datastores).</p> <p>The user also has access to the configure privilege.</p> <p>The user cannot perform the following task:</p> <ul style="list-style-type: none"> <li>• Restore storage</li> </ul>
VSC Restore	<p>Provides all native vCenter Server and VSC-specific privileges necessary to restore vSphere objects that have been backed up.</p> <p>The user also has access to the configure privilege.</p> <p>The user cannot perform the following task:</p> <ul style="list-style-type: none"> <li>• Back up vSphere objects.</li> </ul>

## Guidelines for using VSC standard roles

When you work with standard Virtual Storage Console for VMware vSphere roles, there are certain guidelines you should follow.

You should not directly modify the standard roles. If you do, VSC will overwrite your changes each time one of the following events occurs:

- You upgrade VSC.  
The installer updates the standard role definitions each time you upgrade VSC. Doing this ensures that the roles are current for your version of VSC as well as for all supported versions of the vCenter Server.
- You restart the VSC Windows service  
When you restart the VSC Windows service, VSC returns the standard roles to their default values. Doing this overwrites any changes that you have made to the roles.

You can, however, use the standard roles to create roles that are tailored to your environment. To do this, you should clone the VSC standard role and then edit the cloned role. By creating a new role, you can maintain this role even when you restart or upgrade the VSC Windows service.

Some of the ways that you might use the VSC standard roles include the following:

- Use the standard VSC roles for all VSC tasks.  
In this scenario, the standard roles provide all the privileges a user needs to perform the VSC tasks.
- Combine roles to expand the tasks a user can perform.  
If the standard VSC roles provide too much granularity for your environment, you can expand the roles by creating higher-level groups that contain multiple roles.  
For example, if the same person performs provisioning tasks and cloning tasks, you could set up a group for each of these tasks and then create two permissions.

Permission 1:

- The "provision" group
- The VSC provision role
- The root object

Permission 2:

- The "clone" group
- The VSC clone role
- The root object

You would then need to add the user to the "provision" and "clone" groups.

If a user needs to perform other, non-VSC tasks that require additional native vCenter Server privileges, you can create a role that provides those privileges and add it to the group also.

- Create more fine-grained roles.  
If your company requires that you implement roles that are more restrictive than the standard VSC roles, you can use the VSC roles to create new roles.

In this case, you would clone the necessary VSC roles and then edit the cloned role so that it has only the privileges your user requires.

## Privileges required for VSC tasks

Different Virtual Storage Console for VMware vSphere tasks require different combinations of VSC-specific privileges and native vCenter Server privileges.

Information about the privileges required for VSC tasks is available in Knowledgebase article 1013941 *How to configure Storage and vCenter RBAC for VSC for VMware vSphere*, which is online at [kb.netapp.com/support/index?page=content&id=1013941](http://kb.netapp.com/support/index?page=content&id=1013941).

## Product-level privilege required by VSC for VMware vSphere

To access the Virtual Storage Console for VMware vSphere GUI, you must have the product-level, VSC-specific View privilege assigned at the correct vSphere object level. If you log in without this privilege, VSC displays an error message when you click the NetApp icon and prevents you from accessing VSC.

The following information describes the VSC product-level View privilege:

Privilege	Description	Assignment level
View	<p>You can access the VSC GUI.</p> <p>This privilege does not enable you to perform tasks within VSC. To perform any VSC tasks, you must have the correct VSC-specific and native vCenter Server privileges for those tasks.</p>	<p>The assignment level determines which portions of the UI you can see.</p> <p>Assigning the View privilege at the root object (folder) enables you to enter VSC by clicking the NetApp icon.</p> <p>You can assign the View privilege to another vSphere object level; however, doing that limits the VSC menus that you can see and use.</p> <p>The root object is the recommended place to assign any permission containing the View privilege.</p>

### Example of how the View privilege affects tasks in VSC for VMware vSphere

The View privilege allows you to see the Virtual Storage Console for VMware vSphere GUI. Where you assign a permission containing the View privilege determines which parts of the GUI are visible.

**Note:** Depending on your company's security requirements, you might either log in as administrator and not use RBAC, or you might set permissions on the root object (folder). Setting permissions on the root object normally allows all the child objects to inherit those permissions, unless you place a restriction on a child object.

If you plan to implement RBAC in a more restrictive way, the following example illustrates how assigning a permission containing the View privilege can affect which parts of the GUI you can access. In this example, the following vCenter Server permission is assigned to Datacenter A:

- A user named "Pat"
- Privileges that include the View privilege

- A managed object within the vSphere inventory (Datacenter A), which has the "Propagate to Child Objects" check box selected

If you log in to VSC as user "Pat", the NetApp icon appears. However, if you click the icon, VSC displays an error message. This is because "Pat" only has View permission on Datacenter A. To see the VSC menus and toolbars, you must navigate to Datacenter A and right-click it.

To access the main VSC GUI and avoid an error message when you click on the icon, you must assign any permission containing the View privilege to the root object.

## Data ONTAP role-based access control features in VSC for VMware vSphere

Data ONTAP role-based access control (RBAC) enables you to control access to specific storage systems and the actions a user can perform on those storage systems. In Virtual Storage Console for VMware vSphere, Data ONTAP RBAC works with vCenter Server RBAC to determine which VSC tasks a specific user can perform on objects on a specific storage system.

VSC uses credentials (user name and password) that you set up in Monitoring and Host Configuration to authenticate each storage system and determine which storage operations can be performed on that storage system. VSC uses one set of credentials for each storage system. These credentials determine all VSC tasks that can be performed on that storage system; in other words, the credentials are for VSC, not an individual VSC user.

Data ONTAP RBAC applies only to accessing storage systems and performing VSC tasks related to storage, such as cloning virtual machines. If you do not have the appropriate Data ONTAP RBAC privileges, you cannot perform any tasks involving storage. You can use Data ONTAP RBAC in conjunction with the VSC-specific privileges to control which VSC tasks a user can perform:

- Monitoring and configuring storage or vCenter Server objects residing on storage
- Provisioning and cloning storage
- Scanning, optimizing, and migrating storage
- Backing up and recovering storage

Using the Data ONTAP RBAC with the VSC-specific privileges provides a storage-oriented layer of security that the storage administrator can manage. As a result, you have more fine-grained access control than either Data ONTAP or vCenter Server supports alone. For example, with vCenter Server RBAC, you can allow vCenterUserB, but not vCenterUserA, to provision a datastore on NetApp storage. However, if the storage system credentials for a specific storage system do not support creating storage, then neither vCenterUserB nor vCenterUserA can provision a datastore on that storage system.

When you initiate a VSC task, VSC first confirms that you have the correct vCenter Server permission for that task. If the vCenter Server permission is not sufficient to allow you to perform the task, VSC does not need to check the Data ONTAP privileges for that storage system because you did not pass the initial, vCenter Server security check. As a result, you cannot access the storage system.

If the vCenter Server permission is sufficient, VSC then checks the Data ONTAP RBAC privileges (your Data ONTAP role) associated with the storage system's credentials (your user name/password) to determine whether you have sufficient privileges to perform the storage operations required by that VSC task on that storage system. If you have the correct Data ONTAP privileges, you can access the storage system and perform the VSC task. The Data ONTAP roles determine the VSC tasks you can perform on the storage system.

Each storage system has one set of Data ONTAP privileges associated with it. To determine which Data ONTAP privileges are associated with a storage system, right-click a storage system name in Monitoring and Host Configuration.

Using both Data ONTAP RBAC and vCenter Server RBAC provides the following benefits:

- Security  
The administrator can control which users can perform which tasks on both a fine-grained vCenter Server object level and a storage system level.
- Audit information  
In many cases, VSC provides an audit trail on the storage system that lets you track events back to the vCenter user who performed the storage modifications.
- Usability  
You can maintain controller credentials in one place.

## Recommended Data ONTAP roles when using VSC for VMware vSphere

There are several recommended Data ONTAP roles that you can set up for working with Virtual Storage Console for VMware vSphere and role-based access control (RBAC). These roles contain the Data ONTAP privileges required to perform the necessary storage operations executed by the VSC tasks.

To create these roles, you can use a tool such as "RBAC User Creator for Data ONTAP," which is posted on the NetApp Communities Forum at:

<https://communities.netapp.com/docs/DOC-19074>

Each role has a user name/password pair associated with it. These are the role's credentials. If you do not log in using these credentials, you cannot access the storage operations associated with the role.

As a security measure, the VSC-specific Data ONTAP roles are ordered hierarchically. This means that the first role is the most restrictive role and has only the privileges associated with the most basic set of VSC storage operations. The next role includes both its own privileges and all of the privileges associated with the previous role. Each additional role is less restrictive with regard to the supported storage operations.

The following are the recommended Data ONTAP RBAC roles when using VSC. After you create these roles, you can assign them to users who need to perform tasks related to storage, such as provisioning and cloning storage and optimizing and migrating virtual machines:

1. Discover  
The Discover role enables you to add storage systems.
2. Create Clones

This role enables you to clone virtual machines. It also includes all of the privileges associated with the Discover roles.

### 3. Create Storage

This role enables you to create storage. It also includes all of the privileges associated with the previous two roles.

### 4. Modify Storage

This role enables you to modify storage. It also includes all of the privileges associated with the previous three roles.

### 5. Delete Storage

This role enables you to destroy storage. It also includes all of the privileges associated with all of the above roles.

If you use VSC only to perform backups, then the following Data ONTAP roles are recommended:

#### 1. Discover

The Discover role enables you to add storage systems.

#### 2. Backup and Recover

This role enables you to back up information on storage systems that you can recover later. It also includes all of the privileges associated with the Discover role.

Each Data ONTAP role that you create can have one user name associated with it. You must log in to the storage system using the appropriate user name/password pair if you want to perform those role-based tasks on the storage system.

To create new users, you must log in as an administrator on storage systems running clustered Data ONTAP or root on storage systems running Data ONTAP operating in 7-Mode.

## How to configure Data ONTAP role-based access control for VSC for VMware vSphere

You must configure Data ONTAP role-based access control (RBAC) on the storage system in order to use it with Virtual Storage Console for VMware vSphere.

From within Data ONTAP, you must perform the following tasks:

- Create the roles.

**Note:** You can use the "RBAC User Creator for Data ONTAP" tool to create the Data ONTAP RBAC roles for VSC. This tool is posted on the NetApp Communities Forum at:

<https://communities.netapp.com/docs/DOC-19074>

- Create a user name/password login (the storage system credentials) in Data ONTAP for each role. You need these storage system credentials to configure the storage systems for VSC. You do this by entering the credentials in Monitoring and Host Configuration. Monitoring and Host Configuration manages all credentials for VSC. Each time you log in to a storage system using

these credentials, you are presented with the set of VSC functions. Those are the functions that you set up in Data ONTAP when you created the credentials.

Monitoring and Host Configuration performs an upfront privilege validation for Data ONTAP RBAC when you log into VSC. Monitoring and Host Configuration does not perform the upfront validation if the storage system is directly connected to a Vserver or a vFiler unit. Instead, VSC checks and enforces the privileges later in the task workflow.

You can use the administrator or root login to access all the VSC tasks; however, it is a good practice to use the RBAC feature provided by Data ONTAP to create one or more custom accounts with limited access privileges.

### **Related concepts**

*[Authentication methods in the Backup and Recovery capability](#) on page 113*

# Monitoring and configuring hosts and storage controllers

---

Monitoring and Host Configuration enables you to work with storage controllers and ESX and ESXi hosts. It also manages the task of adding storage controllers and setting the credentials for all the Virtual Storage Console for VMware vSphere capabilities.

The tasks that you can perform with Monitoring and Host Configuration include the following:

- Add, manage, and remove storage controllers.
- Set credentials to access storage controllers.
- View status for the following:

- Storage controllers

You can quickly see the status of storage controllers from the following perspectives:

- SAN (FC, FCoE, and iSCSI)
- NAS (NFS)

**Note:** Monitoring and Host Configuration displays the status of storage controllers running clustered Data ONTAP and Data ONTAP operating 7-Mode as well as other versions of Data ONTAP.

- SAN and NAS datastore capacity utilization
- VMware vStorage APIs for Array Integration (VAAI) support in the storage controller
- ESX and ESXi hosts, including the operating system version and overall status
- Flash Accel components on the ESX host and its virtual machines
- Configure the correct values for the following:
  - Storage adapter timeouts
  - Multipathing settings
  - Flash Accel settings
  - NFS settings
- Determine whether the paths to NFS nodes provide direct and indirect data access to NFS datastores
- Collect diagnostic information from the ESX and ESXi hosts, storage controllers, and Fibre Channel switches
- Access tools to set guest operating system timeouts and to identify and correct misaligned disk partitions

**Note:** The tools for aligning disk partitions that are provided by Monitoring and Host Configuration require that you power down the VM. You can use the online alignment tool provided by Optimization and Migration to align disk partitions without having to power down the virtual machines

- Install and enable the NFS Plug-in for VMware VAAI software library
- Serve as the management console for Flash Accel

You perform these tasks from the Monitoring and Host Configuration GUI. To access the GUI, you click the NetApp icon in the vCenter Server and then click Monitoring and Host Configuration in the navigation pane.

## VSC for VMware vSphere configuration

You can configure and manage your ESX and ESXi hosts and virtual machines (VMs) by first specifying the physical storage systems on which the active images of the datastores and VMs that are managed by the vCenter Server reside.

### Storage system discovery and credentials overview

Virtual Storage Console for VMware vSphere uses Monitoring and Host Configuration as the single mechanism to discover storage systems for all capabilities and set storage credentials. VSC tasks require certain Data ONTAP privileges or they will fail.

Before VSC can display and manage storage resources, it must discover the storage systems that provide the storage. As part of the discovery process, you must supply Data ONTAP storage system credentials. These are the privileges (or role) associated with the user name/password pair assigned to each storage system. They use Data ONTAP role-based access control (RBAC). Because the user name/password pairs must be set up from within Data ONTAP, you cannot change their credentials at the Modify Credentials dialog box.

**Note:** If you log in as an administrator, you automatically have all privileges for that storage system.

When Monitoring and Host Configuration adds a storage system, it displays a pop-up box that lists the credentials associated with the username/ password pair you entered when you logged in. You can either set up default credentials that Monitoring and Host Configuration will use during its storage system discovery process or manually enter credentials when the storage system is discovered.

**Note:** If you have vFiler units on storage systems running Data ONTAP 8.x software, you must set the `httpd.admin.enable` for the vFiler unit in order to enable discovery.

Discovery happens in one of the following ways. In each case, you must supply credentials for any newly discovered storage system.

- When the VSC Windows service starts, Monitoring and Host Configuration begins its automatic background discovery process.
- You click **Update** on the Monitoring and Host Configuration Overview panel. This runs another automatic discovery.

**Note:** IPv6 addresses are not supported.

All of the capabilities require specific permissions to perform VSC tasks. You can limit what users can do based on the credentials associated with their vSphere Client account. All users with the same storage system user name/password pair share the same set of storage system credentials, and all can perform the same operations.

### Related tasks

[Discovering and adding storage resources](#) on page 51

[Enabling discovery and management of vFiler units](#) on page 49

[Enabling discovery and management of vFiler units on private networks](#) on page 49

## Default controller credentials simplify administration of capabilities

You can set up default Virtual Storage Console for VMware vSphere storage controller credentials in Monitoring and Host Configuration. You do not have to manually specify credentials for any storage controller for which the default credentials are valid. These credentials apply to all the VSC capabilities.

When Monitoring and Host Configuration discovers a new storage controller, it attempts to log in using the default credentials. If the login fails, the controller status is set to Authentication Failure, and you must enter credentials manually from the Overview panel by right-clicking the controller name and choosing **Modify Credentials** from the pop-up menu.

You can set the default credentials by clicking **Set Default Controller Credentials** on the Monitoring and Host Configuration Discovery Status panel.

Any time you change the default credentials and run **Update**, Monitoring and Host Configuration uses the new credentials and attempts to log in to any controller that has a status of either Authentication Failure or SSL is not configured.

## Specifying credentials with Monitoring and Host Configuration

You must use Monitoring and Host Configuration to set up credentials for a storage controller. These credentials then apply to all Virtual Storage Console for VMware vSphere capabilities.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

You either set up default credentials that Monitoring and Host Configuration will use each time it discovers a storage controller, or manually assign the credentials as they are discovered.

### Steps

1. Select the **Overview** panel of Monitoring and Host Configuration.

2. Right-click a storage controller that needs credentials.
3. From the pop-up dialog menu that appears, select **Modify Credentials**.
4. Fill in the following information:
  - Management IP address  
VSC uses the management IP address to communicate with the controller. VSC lists the available addresses.
  - Management port number  
The default management port number is 443 if the SSL box is checked and 80 if it is not checked. These are the Data ONTAP defaults. If you toggle the SSL check box, the port number switches between 443 and 80. You can specify a different port number. If you do that, then toggling the SSL check box only changes the SSL state in the dialog box.
  - Whether SSL is enabled
  - User name/password  
Storage controller credentials are assigned based on the user name/password pair. This can be the root account or a custom account that uses role-based access control (RBAC). You cannot change the roles associated with that user name/password pair at the Modify Credentials dialog box.
  - Whether the controller is skipped  
If you chose to **not** provide credentials for this controller, you must select the Skipped check box.  
If a controller is skipped, Monitoring and Host Configuration does not export it to the other VSC capabilities. If it exists for these capabilities, it is deleted from their controller lists and no longer available for their work flows.  
However, if you uncheck the Skipped check box, Monitoring and Host Configuration adds the controller back to the export list and the capabilities can add it to their work flows again.
5. When you have filled out the information in the dialog box, click **OK**.  
VSC displays the list of allowed and disallowed roles as indicated by the credentials.
6. If the allowed roles support your needs, then click **OK**.  
If the roles are not sufficient, click **Cancel**. Doing this returns you to the **Modify Credentials**. You can either enter a different user name/password pair that will provide different credentials or talk with your system administrator about expanding the privileges for this user name/password pair.  
Once a storage controller has credentials, you can view the privileges by right-clicking the controller name on the Overview panel.

## Enabling discovery and management of vFiler units

If you are using Data ONTAP 8, you must set the `httpd.admin.enable` option for vFiler units in order to enable discovery and management with the Virtual Storage Console for VMware vSphere.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

This task is not required for vFiler units created with Data ONTAP 7.x.

### Steps

1. From the storage system, enter the following command to switch to a particular vFiler context:  

```
vfiler context vfiler_name
```
2. Enter the following command in the vFiler context to set the required option that enables discovery in VSC:  

```
options httpd.admin.enable on
```
3. Repeat for each vFiler unit you want to manage using VSC.

## Enabling discovery and management of vFiler units on private networks

If vFiler units are isolated in private networks to which the Virtual Storage Console for VMware vSphere has no network connectivity, you must manually add the pFiler to Monitoring and Host Configuration.

### Before you begin

The VSC server must have network connectivity to the parent of the vFiler.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. On the Monitoring and Host Configuration **Overview** panel, right-click within the **Storage Controllers** section and select **Add Controller**.
2. Enter the management IP address and credentials for the pFiler and then click **OK**.

## Result

Any vFiler units belonging to the parent of the vFiler that provide storage to ESX hosts are discovered by Monitoring and Host Configuration.

## Tunneled vFiler units and Vservers discovered automatically

The Monitoring and Host Configuration capability automatically supports vFiler and Vserver tunneling for the storage systems it manages. You do not need to manually add these vFiler units and Vservers to the Monitoring and Host Configuration capability.

When you enter information for a cluster administrative LIF or vfiler0, the Monitoring and Host Configuration capability discovers all the subordinate vFiler units and Vservers..

## Differences between direct connections to Vservers and to cluster-management LIFs

Virtual Storage Console for VMware vSphere supports connecting a storage controller directly to either a Vserver or a cluster-management LIFs. When the storage connects directly to a Vserver, not all of the VSC features are supported. To use all the features, you must connect the storage to a cluster-management LIF.

VSC does not provide the following features when the storage controller connects directly to a Vserver:

- Upfront validation of Role-Based Access Control (RBAC)  
While RBAC is fully supported, VSC does not perform the initial privilege validation on storage that is directly connected to a Vserver.
- NFS path checking  
When you are running clustered Data ONTAP and using a cluster-management LIF, Monitoring and Host Configuration can query the storage controller to determine whether the storage controller is using a direct or indirect path. Monitoring and Host Configuration then reports this information, and supplies information you can use to set up a direct path. Better performance is normally seen when direct paths are used. If a storage controller connects directly to a Vserver, Monitoring and Host Configuration cannot query the storage controller to determine the path.
- Reports on space that is shared by volumes using data deduplication  
VSC is not able to check the space shared by volumes that have data deduplication enabled when the storage controller is directly attached to a Vserver.
- EMS logging  
VSC cannot perform EMS logging when the storage controller is directly attached to a Vserver.
- Storage-side log collections for the nSANity Diagnostic and Configuration Data Collector program  
The nSANity program collects data that can be used to resolve problems. You should only use the nSANity program if you are directed to do so by Technical Support. When the storage connects directly to a Vserver, the nSANity program cannot collect information on that storage.

These features are supported when the storage controller connects to a cluster-management LIF.

## Discovering and adding storage resources

When you first run Virtual Storage Console for VMware vSphere in a vSphere Client, Monitoring and Host Configuration discovers ESX and ESXi hosts, their LUNs and NFS exports, and the NetApp storage systems that own those LUNs and exports. You must provide the storage system credentials.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

You can discover new resources and get the latest capacity and configuration information at any time by clicking **Update** on the Overview panel of Monitoring and Host Configuration in the vSphere Client.

The discovery process collects information from the ESX and ESXi hosts managed by the vCenter Server. Make sure all ESX and ESXi hosts are shown as powered on and connected.

This discovery process is for all VSC capabilities.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the **NetApp** tab.
3. If the discovery process does not start automatically, or if you want to discover new resources and update information, click **Update** on the **Overview** panel of Monitoring and Host Configuration.
4. Right-click any discovered storage controllers with the status Authentication Failure and select **Modify Credentials**.
5. Fill in the information in the **Modify Credentials** dialog box.

### After you finish

After discovery is complete, use Monitoring and Host Configuration to configure ESX or ESXi host settings for any hosts displaying an Alert icon in the Adapter Settings, MPIO Settings, or NFS Settings columns.

### Related concepts

[Storage system discovery and credentials overview](#) on page 46

### Related tasks

[Getting information about an ESX and ESXi host with an Alert status](#) on page 199

## Correcting controller names displayed as "unknown"

If Monitoring and Host Configuration displays a controller name as "unknown" on the Overview panel, you can modify the credentials and add the management IP address of the controller in the Modify storage system -unknown- pop-up box.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

This issue can occur if an NFS datastore is mounted over a private network.

If you are running clustered Data ONTAP and working with NFS datastores that are mounted using an NFS data LIF, this issue can occur with either a private network or a public network.

### Steps

1. Right-click the controller and select **Modify Credentials**.
2. Enter the management IP address of the storage controller and the storage controller credentials in the **Modify storage system -unknown-** pop-up box.

VSC must have network connectivity to the management port you specify.

## Removing controllers from Monitoring and Host Configuration

You can remove a skipped or unmanaged storage controller that is not attached to a host. When you remove a storage controller, it no longer appears in the Virtual Storage Console for VMware vSphere display.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

If a storage controller has storage mapped to an ESX or ESXi host managed by Monitoring and Host Configuration and you attempt to remove that storage controller, Monitoring and Host Configuration displays an error message and does not remove the storage controller. You can only remove storage controllers that are not attached to hosts.

**Steps**

1. Open the vSphere Client and log in to your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. Select Monitoring and Host Configuration and then select the **Overview** panel.
4. Right-click the storage controller and select **Remove Controller**.

If the storage controller is ...	Monitoring and Host Configuration
Not attached to host	Removes the storage controller
Attached to a host	Displays an error message and does not change the controller

## Administering the Monitoring and Host Configuration capability

The Monitoring and Host Configuration capability lets you work with hosts and controllers and configure values for them.

The sections that follow apply to using the Monitoring and Host Configuration capability.

### Inventory panel selection limits what VSC for VMware vSphere displays

The Monitoring and Host Configuration capability displays only the resources associated with what is selected in the vSphere Client Inventory panel. If you click the **Update** button, the Monitoring and Host Capability updates only those selected resources.

You can access the Monitoring and Host Configuration panels in two ways. If you select the **NetApp** icon in the **Solutions and Applications** section of the vSphere Client **Home** page, the Inventory panel is not displayed, and the Monitoring and Host Configuration panels display all discovered resources.

If instead you access the Monitoring and Host Configuration panels using the NetApp tab in the **Inventory** section of the vSphere Client, the Monitoring and Host Configuration panels display only those resources associated with the selection in the Inventory panel. To display all resources, you must select the Datacenter object in the Inventory panel.

If you cannot find an expected resource in one of the Monitoring and Host Configuration panels, the first thing to check is the selection in the Inventory panel.

The effect of the **Update** button is also limited by the selection in the Inventory panel. Because actions are limited to what is selected in the Inventory panel (a subset of the entire configuration), you get faster results. To update the entire configuration, you need to select the Datacenter object in the Inventory panel before clicking **Update**.

## Configuring ESX server multipathing and timeout settings

Monitoring and Host Configuration checks and sets the ESX or ESXi host multipathing and HBA timeout settings that ensure proper behavior with NetApp storage systems.

### Before you begin

This process might take a long time, depending on your configuration and system load. The task progress is displayed in the **Recent Tasks** panel. As tasks complete, the host status Alert icons are replaced by Normal or Pending Reboot icons.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the **NetApp** tab.
3. Select the **Overview** panel.
4. Select one or more ESX hosts that have an Alert icon in the **Adapter Settings**, **MPIO Settings**, or **NFS Settings** columns.  
Use Ctrl-click or Shift-click to select multiple hosts.
5. Right-click the selected hosts and select **Set Recommended Values**.
6. Select the types of settings you want to update and then click **OK**.

## ESX host values set by Monitoring and Host Configuration capability

Monitoring and Host Configuration sets ESX or ESXi host timeouts and other values to ensure best performance and successful failover.

Monitoring and Host Configuration sets the following values on an ESX or ESXi host:

### NFS Settings

<b>Net.TcpipHeapSize</b>	If you are using vSphere 5.0 or later, set to 32. For all other NFS configurations, set to 30.
<b>Net.TcpipHeapMax</b>	If you are using vSphere 5.0 or later, set to 128. For all other NFS configurations, set to 120.
<b>NFS.MaxVolumes</b>	If you are using vSphere 5.0 or later, set to 256. For all other NFS configurations, set to 64.

<b>NFS.HeartbeatMaxFailures</b>	Set to 10 for all NFS configurations.
<b>NFS.HeartbeatFrequency</b>	Set to 12 for all NFS configurations.
<b>NFS.HeartbeatTimeout</b>	Set to 5 for all NFS configurations.

## FC/FCoE Settings

<b>Path selection policy</b>	Set to RR (round robin) for ESX 4.0 or 4.1 and ESXi 5.0, FC paths with ALUA enabled. Set to FIXED for all other configurations.
<b>Disk.QFullSampleSize</b>	Set to 32 for all configurations. This setting is available with ESXi 5.0 and ESX 4.x.  <b>Note:</b> vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0. For more information on QFull settings in vSphere 5.1, see knowledgebase article 1013944, which is online at <a href="http://kb.netapp.com/support/index?page=content&amp;id=1013944">kb.netapp.com/support/index?page=content&amp;id=1013944</a> .
<b>Disk.QFullThreshold</b>	Set to 8 for all configurations. This setting is available with ESXi 5.0 and ESX 4.x.  <b>Note:</b> vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0. For more information on QFull settings in vSphere 5.1, see knowledgebase article 1013944, which is online at <a href="http://kb.netapp.com/support/index?page=content&amp;id=1013944">kb.netapp.com/support/index?page=content&amp;id=1013944</a> .
<b>Emulex FC HBA timeouts</b>	For ESX 4.0 or 4.1 or ESXi 5.0 or later, use the default value.
<b>QLogic FC HBA timeouts</b>	For ESX 4.0 or 4.1 or ESXi 5.0 or later, use the default value.

## iSCSI Settings

<b>Path selection policy</b>	Set to RR (round robin) for all iSCSI paths.
<b>Disk.QFullSampleSize</b>	Set to 32 for all configurations. This setting is available with ESX 4.x and ESXi 5.0.  <b>Note:</b> vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0. For more information on QFull settings in vSphere 5.1, see knowledgebase article 1013944, which is online at <a href="http://kb.netapp.com/support/index?page=content&amp;id=1013944">kb.netapp.com/support/index?page=content&amp;id=1013944</a> .
<b>Disk.QFullThreshold</b>	Set to 8 for all configurations. This setting is available with ESX 4.x and ESXi 5.0.

**Note:** vSphere 5.1 handles Task Set Full (QFull) conditions differently from vSphere 4.x and 5.0. For more information on QFull settings in vSphere 5.1, see knowledgebase article 1013944, which is online at [kb.netapp.com/support/index?page=content&id=1013944](http://kb.netapp.com/support/index?page=content&id=1013944).

**QLogic iSCSI HBA  
IP\_ARP\_Redirect**

Set to ON for all configurations.

**QLogic iSCSI HBA  
timeouts**

ql4xportdownretrycount (qla4022 driver), ka\_timeout (qla4xxx driver), and KeepAliveTO timeout settings are set to 14 for iSCSI SAN booted ESX hosts, and set to 60 for non-SAN-boot configurations.

## UNMAP setting turned off in ESX 5.x

On hosts running ESX 5.0, the Monitoring and Host Configuration capability automatically turns off the UNMAP (`VMFS3.EnableBlockDelete`) parameter by setting it to 0.

On ESX 5.1 hosts, 0 is the default value. If you change this value to 1 on either ESX 5.0 or 5.1, the Monitoring and Host Configuration capability automatically resets it to 0. For this value to take effect, you must apply the HBA/CNA Adapter Settings on the host.

To avoid any potential performance impact due to UNMAP operations, VMware disabled this feature beginning in ESXi 5.0 Patch 2. VSC 4.1 for VMware vSphere ensures that this feature is disabled with all versions of ESXi 5.x.

## Direct path access and NFS datastores

When you are running clustered Data ONTAP, it is possible for a client to access a data LIF with an indirect data path to a FlexVol. Indirect data paths can affect I/O performance and should be corrected. Monitoring and Host Configuration provides tools to scan for direct and indirect NFS paths and provide you with the information you need to manually correct the paths.

This situation can occur when a data LIF is bound to a different physical node than the one that owns the exported FlexVol. The NFS virtual client does not have the path selection intelligence that is native to physical clients. In order to have a direct data path, the client must access a data LIF that is local to the node that owns the exported FlexVol.

Monitoring and Host Configuration monitors which LIFs NFS is using to access the volume. You can see whether a LIF uses a direct data path or an indirect data path by going to the Storage Details - NAS panel fields and commands window and viewing the Data Path Access column. This column displays the path setting as Direct (green check), Indirect (red exclamation point (!)), N/A, or (unknown).

If the path setting is indirect, you can right-click that row and select the **View Direct Data Path Choices** option. This option displays the **Direct Data Path Choices** pop-up, which contains a list of ports using direct paths to access data.

**Note:** Virtual Storage Console for VMware vSphere does not check these ports to ensure that they are connected to the network. You must do that manually.

Anytime the data path access changes, either to direct from indirect or to indirect from direct, Monitoring and Host Configuration writes the path information to a log file.

If a direct Vserver connection is made, Monitoring and Host Configuration cannot query the storage controller to determine the path.

An N/A (not applicable) entry indicates a path to a storage controller running Data ONTAP operating in 7-Mode, so there is no issue about whether the path is direct.

An unknown path occurs if the discovery data is incomplete.

## Changing NFS data paths to direct access

If you have a cluster node that is accessing a data LIF with an indirect data path, you can change the path to one that is direct.

### Before you begin

Only a storage administrator should change the path.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

This issue only occurs if you are running clustered Data ONTAP and you have an NFS datastore using a remote data LIF that is bound to a different physical node than the one that owns the exported FlexVol.

### Steps

1. In the **Storage Details - NAS panel fields and commands** window, right-click a row that has an indirect data path (shown as Indirect) and select the **View Direct Data Path Choices** option.

This option displays the **Direct Data Path Choices** pop-up, which contains a list of ports to data paths providing direct access. You cannot use this window to change the path, but you can use it to get information about the available ports.

2. Manually check to make sure the port you want to use is connected to the network.

Monitoring and Host Configuration displays the ports without checking their network connectivity. If you try to use a port that is not connected to the network, your datastore will go offline.

3. After you have confirmed that the path you want to use is connected to the network, collect the information displayed in the **Direct Data Path Choices** pop-up and give it to a storage administrator.

Only the storage administrator should change the path. The **Direct Data Path Choices** pop-up contains all the other information a storage administrator needs to move the LIF.

To create a data path with direct access, you must have the correct credentials.

**Note:** If multiple datastores are using that LIF, moving the LIF will cause the other datastores to have data paths with indirect data access.

4. Use either the storage controller console or a NetApp tool such as System Manager to change the path.

**Note:** Anytime the path value changes, Monitoring and Host Configuration writes the information to a log file.

## Flash Accel and VSC for VMware vSphere

You can use Virtual Storage Console for VMware vSphere to manage the Flash Accel host cache software on ESX hosts and their virtual machines. The Monitoring and Host Configuration Flash Accel Cache panel provides centralized management and reporting for the Flash Accel components on the ESX host and its virtual machines.

**Note:** You must install the Flash Accel Base Package for VSC in order for Monitoring and Host Configuration to display the new Flash Accel Cache panel. You can download the Flash Accel Base Package for VSC from the NetApp Support Site at [support.netapp.com](http://support.netapp.com). Monitoring and Host Configuration only displays the Flash Accel Cache panel if VSC determines that you have Flash Accel installed.

To determine which versions of the Flash Accel Base Package and VMware vSphere are supported with VSC, see the notes and alerts in the VSC storage solution in the Interoperability Matrix (IMT), which is online at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability).

For details about using Flash Accel with VSC, see the *Flash Accel Base Package Installation and Administration Guide* and the *Flash Accel Base Package Release Notes*, which are available for downloading from the product documentation library on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

For the most up-to-date information about Flash Accel and VSC, see the *Virtual Storage Console for VMware vSphere Release Notes*.

## NFS Plug-in for VMware VAAI requires additional installation steps

The NFS Plug-in for VMware VAAI is not shipped with Virtual Storage Console for VMware vSphere; however, you can get the plug-in from the NetApp Support Site and then use the Monitoring and Host Configuration capability to install it.

The plug-in is supported on systems running ESXi 5.0 or later with vSphere 5.0 or later and clustered Data ONTAP 8.1 or later or Data ONTAP 8.1.1 or later operating in 7-Mode.

After you install the plug-in, you must reboot the host. To remind you of this, Monitoring and Host Configuration changes the status of the host in the Overview panel to **Pending Reboot**.

After you install the plug-in, VSC automatically uses it. You do not need to perform additional tasks to enable it.

For more information on the plug-in, see *Installing the NetApp NFS Plug-in for VMware VAAI*, which is available from the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

## Using Web-based tools to manage storage

Data ONTAP provides tools that you can use to create LUNs and manage storage systems. These are Web-based tools that you can use to manage common storage system functions from a Web browser.

### Before you begin

The tool that Monitoring and Host Configuration recommends depends on which version of Data ONTAP you are running on the storage system:

- For storage systems running Data ONTAP 8.1 or later and operating in 7-Mode, Monitoring and Host Configuration supports OnCommand System Manager.
- For storage systems running clustered Data ONTAP, Monitoring and Host Configuration supports Element Manager.
- For storage controllers running Data ONTAP 8.0 and earlier and operating in 7-Mode, Monitoring and Host Configuration supports FilerView.

**Note:** FilerView is not supported with vFiler units, storage systems running clustered Data ONTAP, or Data ONTAP 8.1 and later.

Your client computer must have network connectivity to the management port of a storage system to be able to run the tool for that storage system.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the Inventory panel, and then select the **NetApp** tab.
3. Select the Monitoring and Host Configuration **Overview** panel.
4. Right-click the storage system you want to manage and Monitoring and Host Configuration allows you to select either **Open Element Manager** or **Open FilerView**.

Monitoring and Host Configuration automatically displays the appropriate tool for your version of Data ONTAP. If you select **Open FilerView** and the storage system is running Data ONTAP 8.1 or later, Monitoring and Host Configuration displays a pop-up box recommending that you download OnCommand System Manager because FilerView is not available on those releases. Monitoring and Host Configuration does not provide OnCommand System Manager.

**Note:** If you are using a version of Data ONTAP operating in 7-Mode that supports FilerView and you set up your storage to use a Secure Sockets Layer (SSL) protocol, FilerView automatically opens using the HTTPS protocol. Otherwise it opens using the HTTP protocol.

5. If the storage system requires a password, enter the user name and password when prompted.

### Result

For versions of Data ONTAP that support them, either Element Manager or the FilerView GUI opens in a new browser window. You can then use the tool to work with LUNs on that storage system. For storage systems using Data ONTAP 8.1 or later, you must download OnCommand System Manager and use that to work with LUNs.

## Monitoring and Host Configuration displays configuration details

You can display details of your storage configuration using Monitoring and Host Configuration.

Monitoring and Host Configuration displays several pages of configuration information.

**Note:** See the online Help explanations of the fields displayed and available commands.

<b>Overview</b>	Displays the status of storage controllers and ESX and ESXi hosts. You can right-click a host to get detailed configuration information.
<b>Storage Details - SAN</b>	Displays information about the VMFS datastores. If you select a datastore, you can see a list of all LUNs mapped to managed ESX and ESXi hosts. When you select a LUN, the LUN details are displayed.
<b>Storage Details - NAS</b>	Displays a list of all NFS exports mounted on managed ESX and ESXi hosts. When you select a datastore, the details are displayed.
<b>Discovery Status</b>	Displays details of the storage resources discovered by the Monitoring and Host Configuration capability.

### Related concepts

[Online Help](#) on page 15

## MultiStore vFiler units are displayed differently

Monitoring and Host Configuration displays vFiler units differently than physical storage controllers.

When you have a vFiler unit created with the optional MultiStore feature of Data ONTAP software, Monitoring and Host Configuration displays the following information:

- The hostname displays a "MultiStore" prefix to identify vFiler units.
- The **Supported Protocols** column reports the storage protocols actually in use by ESX and ESXi hosts instead of the protocols licensed for the storage controller.
- The **Alert** icon in the **Status** column means that the vFiler unit does not respond to Monitoring and Host Configuration. The **Normal** icon means that Monitoring and Host Configuration is able to communicate with the vFiler unit.

- No detailed status is returned for vFiler units. The **Status Reason** column displays `This controller is a MultiStore vFiler unit`. You can connect to the physical controller that owns the vFiler unit to get more status information.
- In the Storage Details screen for vFiler units, no aggregate information is displayed. Direct vFiler units and direct vServers do not have aggregates to display.

## Collecting diagnostic information

You can use Monitoring and Host Configuration to collect diagnostic information about storage controllers, ESX and ESXi hosts, and Fibre Channel switches.

### Before you begin

If Virtual Storage Console for VMware vSphere is running on a Windows Server 2008 or Server 2008 R2 system, the VSC for VMware vSphere service must run on an Administrator account to enable the data collection programs to run correctly. This change is not required for Windows Server 2003.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. In Monitoring and Host Configuration, select the **Data Collection** panel.
4. Select the component you want to collect data from.
5. Select or enter the hostname or IP address of the component, and enter a User name and Password with root or administrator rights.
6. Clear the **Save File Locally** check box if you do not want the file copied to your local workstation.
7. Click **Submit**.

The diagnostic data is collected in a `.tar.gz` file.

8. When the Download file dialog box is displayed, select a location to save the file and the click **Save**.

This dialog box is displayed only if the **Save File Locally** check box is selected.

### Result

The `.tar.gz` file is stored on the Windows server running VSC, and, optionally, copied to your local workstation. On the VSC server, the file is saved to the `C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web\support`.

### After you finish

Send the file to Technical Support for analysis.

## Changing the service account for data collection on Windows Server 2008

To enable running the data collection programs on Windows Server 2008 or Server 2008 R2, you must change the account used by the Virtual Storage Console for VMware vSphere service to an Administrator account.

### About this task

The Virtual Storage Console for VMware vSphere service normally runs under the Local System account. For Windows Server 2008 and Server 2008 R2, the Windows security features prevent the data collection programs from running correctly under the Local System account.

### Steps

1. Log on to the Windows system on which you installed VSC for VMware vSphere.
2. Open the Services application by selecting **Start > Administrative Tools > Services**
3. Right-click Virtual Storage Console for VMware vSphere service and select **Properties**.
4. On the **Log On** tab, select **This account**.
5. Enter the credentials for an Administrator account on the Windows system and then click **OK**.

### After you finish

Switch back to the Local System account if desired after running the data collection programs.

## Downloading and running tools

The Monitoring and Host Configuration capability includes tools for detecting and correcting misaligned disk partitions and for setting virtual machine timeouts.

**Note:** The tools provided by the Monitoring and Host Configuration capability can only be used when the virtual machine (VM) is powered off. The Optimization and Migration capability of VSC allows you to perform online alignments on VMFS-based datastores without having to take your VM down. This capability also lets you review the alignment status of VMs and migrate groups of VMs.

The MBR (master boot record) tools enable you to detect and correct misaligned disk partitions for guest operating systems. You must download these tools before you can use them. There is a set of tools for ESX hosts and one for ESXi hosts. You must download the correct tool set for your hosts.

## Enabling the ESXi secure shell

When you are using ESXi, it is a good practice to enable the Secure Shell (SSH) protocol before you download the MBR tools. That way you can use the `scp` command if you need to copy the files. ESXi does not enable this shell by default.

### Steps

1. From an ESXi host, press the key combination ALT F2 to access the Direct Console User Interface (DCUI) screen.
2. Press the F2 function key to get to the Customize System screen.
3. Go to **Troubleshooting Options**.
4. Press **Enter** at the Enable SSH prompt.
5. Press **Enter** at the Modify ESX Shell timeout prompt.
6. Disable the timeout by setting the value to zero (0) and pressing **Enter**.
7. Go to **Restart Management Agents** and press **Enter**.
8. Press F11.

## Downloading and installing MBR tools for ESXi hosts

If you have an ESXi host, you must download and install the version of the MBR (master boot record) tools for ESXi. The MBR tools enable you to detect and correct misaligned disk partitions for guest operating systems. These tools must be installed and run directly on the ESXi host. Before you can install them, you must extract them from the .tar file into the root directory on the ESXi host.

### Before you begin

You must be able to open a console connection to the ESXi host.

**Note:** The MBR tools can only be used when the virtual machine (VM) is powered off. If you want to perform online alignments on VMFS-based datastores without having to take your VM down, you can use Optimization and Migration. In that case, you do not need to download the MBR tools.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. In Monitoring and Host Configuration, select the **Tools** panel.

4. Under **MBR Tools**, click the **Download (For ESXi 4.x and ESXi 5.0)** button.

Make sure you download the MBR Tools for ESXi. If you download the wrong MBR tools file, the tools will **not** work.

5. When the File Download dialog is displayed, click **Save**.
6. (**ESXi 4.x**) If you are using ESXi 4.x, manually enable the ESXi shell and SSH so that you can use the `scp` command to copy the files to the correct directories if needed.

ESXi 4.x does not enable the ESXi shell and SSH by default. You can enable these options from the physical host or from the vCenter. The following steps enable these options from the vCenter.

**Note:** vCenter creates a configuration alert for each ESXi host that has the options enabled.

To enable the ESXi shell, perform the following steps:

- a. From vCenter, highlight the appropriate ESXi host.
- b. Go to the **Configuration Tab**.
- c. In the left pane under **Software**, select **Security Profile**.
- d. Select **Properties** from the **Services** pane.
- e. Highlight the **ESXi Shell** service and select **Options**.
- f. Select **Start and Stop with Host**.
- g. Click **Start**.

To enable the ESXi SSH, perform the following steps:

- a. From vCenter, highlight the appropriate ESXi host.
- b. Go to the **Configuration Tab**.
- c. In the left pane under **Software**, select **Security Profile**.
- d. Select **Properties** from the **Services** pane.
- e. Highlight the **SSH** service and select **Options**.
- f. Select **Start and Stop with Host**.
- g. Click **Start**.

7. Copy the MBR tools for ESXi file to the root (/) directory of the ESXi host. If you are using ESXi 4.x, use the Troubleshooting Console. If you are using ESXi 5.x, use the Technical Service Console.

You might need to open ESXi firewall ports to enable copying the tools to the host.

**Note:** The MBR tools libraries must be located in specific directories on the host. Be sure to download the file to the root directory of the ESXi host.

8. Extract the files by entering the following command:

```
tar -zxvf mbrtools_esxi.tgz
```

If you did not download the file to the root directory, you must manually move the files to that directory.

**Note:** ESXi does not support `-P` with the `tar` command.

### After you finish

Run the `mbralign` tool to check and fix the partition alignment.

## Downloading and installing MBR tools for ESX hosts

If you have an ESX host, you must download and install the version of the MBR tools for ESX. These tools must be installed and run directly on the ESX host. You cannot run these tools from the vSphere Client, vCenter Server, or Virtual Storage Console for VMware vSphere server.

### Before you begin

You must be able to open a console connection to the ESX host.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. In Monitoring and Host Configuration, select the **Tools** panel.
4. Under **MBR Tools**, click the appropriate **Download** button that states for **(For ESX 4.x)**.

If you download the wrong MBR tools file for your host, the tools will **not** work.

5. When the File Download dialog is displayed, click **Save**.
6. Copy the MBR tools file to the root (`/`) directory of the host computer.

You might need to open ESX firewall ports to enable copying the tools to the host.

**Note:** The MBR tools libraries must be located in specific directories on the host. Be sure to download the file to the root directory.

7. Extract the files on the ESX host using the by entering the following command:

```
tar -Pzxf mbrtools.tar.gz
```

The `-P` option places the files in the required directories. You can move the binary files to any location, but the library files must be located in the specific directories to which they are originally extracted.

### After you finish

Run the `mbralign` tool to check and fix the partition alignment.

### Related concepts

[How to identify and fix VMDK partition alignment issues](#) on page 212

[Optimizing and migrating datastores and virtual machines](#) on page 100

## Installing GOS scripts

The ISO images of the guest operating system (GOS) scripts are loaded on the Virtual Storage Console for VMware vSphere server. Mount and run them from the vSphere Client to set the storage timeouts for virtual machines.

### Before you begin

The virtual machine must be running.

The CD-ROM must already exist in the virtual machine or it must be added.

The script must be installed from the copy of the VSC registered to the vCenter Server that manages the VM.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Open the vSphere Client and log into your vCenter Server.
2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. In the Monitoring and Host Configuration capability, select the **Tools** panel.
4. Under **Guest OS Tools**, right-click the link to the ISO image for your guest operating system version and select **Copy to clipboard**.
5. In the vSphere Client, select the desired VM and click the **CD/DVD Connections** icon.
6. Select **CD/DVD Drive 1 > Connect to ISO image on local disk**.
7. Paste the link you copied into the **File Name** field and then click **Open**.

If you receive an authorization error, be sure you select the NetApp tab and click **Yes** to proceed if a security certificate warning is displayed.

Also, be sure that the link you are using is from the copy of the VSC for VMware vSphere running on the vCenter Server that manages the VM.

### After you finish

Log on to the VM and run the script to set the storage timeout values.

**Related tasks**

*Running the GOS timeout scripts for Linux* on page 210

*Running the GOS timeout scripts for Solaris* on page 211

*Running the GOS timeout script for Windows* on page 211

**Adding the CD-ROM to a VM**

Add the CD-ROM to a virtual machine if it does not exist to enable installing the guest operating system scripts.

**Steps**

1. In the vSphere Client, select the desired VM and power it off.
2. Right-click the virtual machine and select **Edit Settings**.
3. On the **Hardware** tab, click **Add**.
4. Select **CD/DVD Drive** and then click **Next**.
5. Click **Use physical drive**.
6. Click **Next** several times to accept the default values.
7. Click **OK** to finish adding the CD-ROM.
8. Power on the VM.

# Provisioning and cloning datastores and virtual machines

---

The Provisioning and Cloning capability of VSC for VMware vSphere enables you to provision datastores and quickly create multiple clones of virtual machines in the VMware environment.

The tasks you can perform with the Provisioning and Cloning capability include the following:

- Create clones of virtual machines and place them in new or existing datastores
- Create, resize, or delete datastores
- Apply guest customization specifications and power up new virtual machines
- Run deduplication operations
- Monitor storage savings
- Reclaim space on virtual machines stored in NFS datastores
- Redeploy virtual machines from a baseline image
- Replicate NFS datastores across sites
- Import virtual machines into virtual desktop infrastructure connection brokers and management tools

## Tips for working with Provisioning and Cloning

There are several methods for accessing the Provisioning and Cloning features.

To manage datastores and clone virtual machines, right-click an object in the Inventory panel of the vSphere Client and select **NetApp > Provisioning and Cloning**. Next you must right-click the correct object for the task you want to perform:

- To create clones, right-click a virtual machine or template.
- To provision datastores, right-click a datacenter, cluster, or host.

To manage controllers and connection brokers, replicate datastores, or redeploy clones, click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.

- To add, remove, or modify properties of storage controllers, select **Storage controllers**.
- To add or remove connection broker definitions, select **Connection brokers**.
- To clone NFS datastore templates to multiple target sites, select **DS Remote Replication**.
- To redeploy virtual machines, select **Redeploy**.

**Note:** As you work with Provisioning and Cloning, keep in mind that IPv6 is not supported. If you have IPv6 configured on a LIF, Provisioning and Cloning cannot use any Vservers from that cluster.

## Cloning and managing virtual machines

You can use Provisioning and Cloning to clone virtual machines, manage connection brokers, redeploy clones locally, and reclaim unused space on virtual machines.

### Cloning virtual machines

You can use Provisioning and Cloning to create theoretically thousands of virtual machine clones and hundreds of datastores at one time. In practice, however, multiple executions of fewer requests is recommended. The ideal size of the requests depends on the size of the vSphere deployment and the hardware configuration of the vSphere Client managing the ESX hosts.

#### Before you begin

Before you perform a cloning operation, it is a good practice to enable the NFS Plug-in for VMware VAAI. This plugin is available from the NetApp Support Site. After you get the plug-in and place it in the correct directory, you can install it using the Monitoring and Host Configuration capability.

**Note:** Details on locating and installing the plug-in are in Knowledge base article 3013414, which is online at [kb.netapp.com/support/index?page=content&id=3013414](http://kb.netapp.com/support/index?page=content&id=3013414).

The following restrictions apply to this feature:

- The cloned virtual machine always has one virtual CPU (vCPU) no matter how many vCPUs the source virtual machine has.
- If you attempt to clone a virtual machine that has been functionally aligned using Optimization and Migration, the clone will be misaligned.

Provisioning and Cloning warns you when you attempt to clone a functionally aligned virtual machine. This is because a functional alignment uses a prefix to get the virtual machine to align on the correct boundary. As a result, the virtual machine performs as though it has been aligned, but no changes have been made to the hard disk to ensure that the virtual machine is aligned to the storage system.

- You cannot use the cloning feature when the target virtual machine is being used by either Backup and Recovery or Optimization and Migration.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

#### Steps

1. In the **vSphere Client Inventory**, right-click a virtual machine or template and select **NetApp > Provisioning and Cloning > Create rapid clones**. Cloning completes faster if the virtual machine or template is powered down.

The Create Rapid Clones Wizard launches.

2. In the **Storage Controller** details window, select the target storage controller for the new clones from the drop-down list and also specify:
  - If you are using clustered Data ONTAP, specify a Vserver. Provisioning and Cloning provides a drop-down list of the available Vservers.
  - If you are using Data ONTAP operating in 7-Mode, specify a vFiler. To identify vFiler units for Provisioning and Cloning operations, select the **Set vFiler Context** check box and select a unit from the drop-down list.
3. If the source virtual machine or template has snapshot copies, the **Clone Source** window displays. Select the template as the source for the new clones.
 

**Note:** If you specify a snapshot as the source, Provisioning and Cloning clones the configuration file settings for the virtual machine, not the virtual machine itself.
4. In the **Clone destination** window, select the destination for the new clones.
 

If you want to specify the virtual machine folder for the new clones, select the check box at the bottom of the window.

**Note:** If the destination is a datacenter, cluster, vApp, or resource pool, new clones are distributed across available servers as quickly as possible. New clones might be unevenly distributed, with more clones created on faster servers.
5. If selected, identify the virtual machine folder in which to place clones in the **Virtual machine folder** window.
6. In the **Disk format** window, select the disk format for the new clones.
7. In the **Virtual machine details** window, provide details for each virtual machine.
  - To specify details manually, select the **Specify VM details** radio button.
  - To import details from a file, select the **Import VM Details** radio button.

You can import the following virtual machine details from a `.csv` file:

- Non-contiguous virtual machine names
- Guest customization specifications
- virtual machine name as computer name (if guest customization specification is provided)
- Power-on setting

The file must contain the following fields: `cloneName`, `customSpecName`, `useVmNameAsPcName`, `powerOn`.

Fields 2 (`customSpecName`) and 3 (`useVmNameAsPcName`) are optional. Blank lines and lines that start with a hash (#) are ignored. The following is an example of a virtual machine details file:

```
# This is a sample VM details file
vm1, dewey, true, false
```

```
vm2, , , true
vm3, customSpecA, false, true
```

- **Create new datastores?** Creates new datastores for the virtual machine clones. Enabled for users with role of create or higher.
- **Import into connection broker?** Automatically imports clone data into a VMware View Server, or creates a .csv file for Citrix XenDesktop manual import. The .csv file is created in the directory [VSC\_home]\etc\kamino\exports\xenDesktop\_timestamp.csv, where [VSC\_home] is your VSC for VMware vSphere installation directory.

**Note:** XenDesktop 5 does not support importing a virtual machine from vApps.

- **Connection broker:** Select the desired output type.
- **Virtual Processors:** Select number of virtual processors to apply to the new virtual machines.
- **Memory Size (MB):** Enter the amount of memory to apply to new virtual machines.
- **Upgrade hardware version?:** If clone source was created on ESX 3.5 or 4.x host and the destination is on a newer version of ESX, allows upgrade of clones to the new hardware version.
- **Clone operation details path:** (Import VM details only) Enter a file name or browse to the .csv file containing clone names.
- **Number of clones:** (Specify VM details only) For new datastores, maximum 250. For existing, maximum depends on available space. Number of clones must be evenly divisible by number of datastores being created.

**Note:** Success of 2,000 or more virtual machines depends on the size and performance of the vCenter Server.

- **Clone name:** (Specify VM details only) Prefix for each clone. By default, the clone number is placed at the end of the clone name. To force clone number to a different position, use %CLONE\_NUM% where you want the number to appear. For example, new%CLONE\_NUM%clone.
- **Starting clone number:** (Specify VM details only) -Maximum of eight digits.
- **Clone number increment:** (Specify VM details only) Increment clone numbers by 1, 2, 3, 4 or 5.
- **Power on?** (Specify VM details only) Select the check box if you want all the virtual machines to power on when the operation completes.  
If you do not select this option, Provisioning and Cloning leaves the clones powered off after the cloning process completes.
- **Stagger VM booting** (Specify VM details only) This option is available only when you select the **Power on?** option.

When you select this option, Provisioning and Cloning staggers the start-up of the clones. In the **VMs per minute** box, you must supply an integer value indicating the number of clones per minute that you want the capability to power on.

Depending on your system setup and the number of clones you created, it is a good practice to stagger starting the clones so that you do not overwhelm your system. Having a large number of virtual machines start at once can slow down your system. The value you supply for this

feature depends on your system environment at the time you perform the clone operation and how many clones you are creating.

**Note:** Provisioning and Cloning keeps track of the total number of clones that should have powered on based on the number of minutes that have elapsed since the cloning operation completed. If a problem prevents Provisioning and Cloning from starting the clones on schedule, it uses this cumulative total to specify how many clones to power on once the delay ends. Depending on the length of the delay and the number of clones per minute that you specified, any delay could result in Provisioning and Cloning powering on a large number of clones at one time. After that, though, it only starts the specified number of clones per minute.

- **Apply customization specification?** (Specify VM details only) Applies a pre-defined specification to the new virtual machines. Select from specifications used for the native cloning process.
8. If you requested a new datastore, click one of the blue links in the **Datastore creation** window: **Create NFS datastore(s)** or **Create VMFS datastore(s)**. To continue without creating a new datastore, click **Next**.

You can create both NFS and VMFS datastores for the new virtual machines. Any restrictions for your configuration appear at the bottom of the window.

**Note:** Maximum VMFS datastore size and maximum number of NFS datastores depend on your version of VMware vSphere. In mixed version environments, maximums revert to earlier limits.

9. In the pop-up window, specify details for the new datastores.

- **Protocol:** (VMFS only) FCP or iSCSI.
- **Number of datastores:** Maximum 256. Number of clones must be evenly divisible by number of datastores.
- **Datastore name:** (single datastores only) Use the default or replace with a custom name.

**Note:** For multiple datastores, the golden volume name (for NFS) or base name (for VMFS) is used here and in the Summary window as the datastore name.

- **Size:** Maximum depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
- **Create new volume container:** (VMFS only) Create a volume with the same name as the LUN. If a volume with that name already exists, the volume name is appended with a number; for example, Volname01.
- **Volume:** (VMFS only) Select available volume from drop-down list.
- **Aggregate:** Select available aggregate from drop-down list.
- **Thin provision:** Sets space reserve to none and disables space checks.

**Note:** Cloning and datastore creation can fail if the size request uses too much of the aggregate. Capacity is not reserved for individual datastores. Instead, the aggregate is treated as a shared pool with capacity used as each datastore requires it. By eliminating unused but provisioned

storage, more space is presented than is available. It is expected that the datastores will not utilize all provisioned storage at once.

- **Block size:** (VMFS only) Select block sizes. For VMFS-5, block size is fixed at 1 MB.
- **Auto-grow:** (NFS only) When space is needed, automatically expands the datastore by increment you specify, up to size limit you specify.
- **Grow increment:** (NFS only) Amount of storage added to datastore each time space is needed.
- **Maximum datastore size:** (NFS only) Limit at which Auto-grow stops.
- **Datastore cluster:** (vCenter 5 or later) Select a datastore cluster in which to add the datastore. Requires SDRS on the vCenter Server.
- **Set datastore names?** (multiple datastores only) Enables modification of default datastore names.
- **Group name:** (VMFS) or Golden volume name (NFS) Use defaults or replace with custom names.

**10.** In the **Datastore selection** window, select the datastore to house the new virtual machines.

You can clone the VMDK files that comprise the virtual machine to different datastores.

- To place all virtual machines in a single datastore, select the datastore and click **Next**.
- To distribute virtual machine files across multiple datastores, click **Advanced**. In the new window, select a virtual machine file, open the corresponding Datastore pull-down list and select the datastore to house that file. Repeat for each virtual machine file. Then click **Next**

**11.** If you specified a connection broker format, enter the details for the connection broker operation.

Option	Description
<b>VMware View Server</b>	<p>For VMware View Server, clone data is imported into View Server at the end of the clone operation. You can specify unique desktop names and select from existing desktop pools.</p> <p><b>Note:</b> Refer to the <i>View Manager Administration Guide</i> at <a href="http://www.vmware.com/pdf/view40_admin_guide.pdf">http://www.vmware.com/pdf/view40_admin_guide.pdf</a> for details on desktop types and access modes.</p>
<b>Citrix XenDesktop</b>	<p>A .csv file is created in the directory [VSC_home]\etc\kamino\exports, where [VSC_home] is your VSC for VMware vSphere installation directory, with the following format:</p> <p><b>XenDesktop 4</b></p> <pre>[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostID]</pre> <p><b>XenDesktop 5</b></p> <pre>[VirtualMachinePath],[ADComputerAccount],[AssignedUsers]</pre> <p>See <i>Importing virtual machines into XenDesktop</i> on page 79 for details on importing this file into Citrix XenDesktop.</p>

**Note:** XenDesktop 5 does not support importing from vApps.

- **View Server hostname or IP Address:** (VMware View Server only) Enter the hostname or IP address of the VMware View Server.
- **Connection name:** (Citrix XenDesktop 5.0 only) Enter the name of the VMware connection.
- **Domain name:** Enter a fully qualified domain name.
- **Domain username:** (VMware View Server only) Enter the domain user name.
- **Domain password:** (VMware View Server only) Enter the domain password.
- **New desktop pool:** Allows you to import new virtual machine clones into multiple new View Server pools.
  - **Number of pools:** Enter the number of pools to create.
  - **Name pools automatically?** Select the checkbox to use the specified pool name as prefix for each pool name.
  - **Distribute VMs evenly?** Select the checkbox to distribute new virtual machines evenly across new desktop pools.

To distribute varying quantities of virtual machines to pools, do not select the **Distribute VMs evenly** checkbox. Instead, enter the desired values directly in the table. The sum of the new distribution must equal the number of new clones.

To provide unique names to the new pools, do not select the **Name pools automatically** checkbox. Instead, enter desired names directly in the table.

- **Desktop pool type:** Select access type for new desktop pools:
  - **Dedicated:** (View Server 4.5 and higher) gives users connection to the same desktop for every session.
  - **Floating:** (View Server 4.5 and higher) allocates desktops dynamically.
- **Existing desktop pool:** (View Server only) Allows you to select from existing pools on the specified View Server.

If you have changed the View server or updated credentials, click Refresh to update the list of pool names.

- **Desktop pool name:** Select an existing pool on the specified View Server from the drop-down list.

**12.** Review the summary page and click **Apply** to proceed. To return to previous pages and modify settings, click **Back**.

## Result

The Recent Tasks pane of vSphere Client is populated as clone creation proceeds.

A list of new virtual machines is written to a .csv file in the directory [VSC\_home]\etc\kamino\exports, where [VSC\_home] is your VSC for VMware vSphere installation directory. The filename is in the following format: import\_generic\_date\_time.csv

## Managing connection brokers

You can use the Connection brokers panel to view and manage the connection brokers available for importing clone data at the end of the clone operation.

- For VMware View Server, clone data is imported into View Server at the end of the clone operation.
- For Citrix XenDesktop, a .csv file is created in the directory `c:\program files\netapp\virtual storage console\etc\kamino\exports`. See [Importing virtual machines into XenDesktop](#) on page 79 for details.

To work with connection brokers in Provisioning and Cloning, you must have .Net 3.5 available on the system where you have Virtual Storage Console for VMware vSphere installed. For some versions of Windows, such as Windows 2008, .Net 3.5 is included as part of the installation. For other versions, such as Windows 2003, it is not part of the base install, so you must manually install it.

## Adding connection brokers

You can add connection brokers to Provisioning and Cloning.

### Before you begin

To work with connection brokers in Provisioning and Cloning, you must have .Net 3.5 available on the system where you have Virtual Storage Console for VMware vSphere installed. For some versions of Windows, such as Windows 2008, .Net 3.5 is included as part of the installation. For other versions, such as Windows 2003, it is not part of the base install, so you must manually install it.

To ensure that your configuration is supported, check the Interoperability Matrix, which is available online at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability).

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Connection brokers**.
3. Click the blue **Add Connection Broker** link.
4. In the **Add Connection Broker** window:
  - a) **Connection broker** - Select the name and version of the desired connection broker.
  - b) **Domain** - Enter the domain containing the connection broker.
  - c) **Connection name** (XenDesktop 5.0 only) - Enter the name given the Citrix XenDesktop 5.0 connection.

- d) **Hostname or IP Address** (VMware View Server only) - Enter the connection broker hostname or IP address.
- e) **Username** (VMware View Server only) - Enter the domain user name.
- f) **Password** (VMware View Server only) - Enter the domain password.

## Removing connection brokers

You can remove a connection broker from the list of available brokers.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Connection brokers**.
3. Click the blue **Remove Connection** broker link.
4. Click **Yes** to confirm.

## Redeploying clones (locally)

You can reset to their original state all virtual machines that are based on a selected "gold" virtual machine or template. You can also propagate changes made in the original gold virtual machine or template to all of its clones and, optionally, reapply customization specifications as well.

### Before you begin

This feature has the following requirements:

- It is available only for virtual machines that reside entirely on NFS destinations.
- It only applies templates to local virtual machines.  
To replicate templates from a vCenter to subordinate vCenters across multiple sites, see [Replicating remote datastores](#) on page 82.
- You cannot use this feature when the target virtual machine is being used by Backup and Recovery or Optimization and Migration.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Redeploy**.

3. Select the virtual machine or template to use, then click the **Redeploy...** link on the right side of the window.

**Note:** To refresh the list, click the **Update table** link.

4. In the **Select clones** window, select the check boxes for the clones to redeploy.
5. In the **Redeploy clones** window, specify the settings to apply to the redeployed clones.

**Note:** If the selected virtual machines are individual desktops or a dedicated pool in VMware View 4.5 and higher, the redeploy may result in loss of user data.

- a) **Power on?** Powers on all the virtual machines after the operation completes.

If you do not select this option, Provisioning and Cloning leaves the virtual machines powered off after the cloning process completes.

- b) **Stagger VM booting:** This option is available only when you select the **Power on?** option.

If you select this option, Provisioning and Cloning staggers the start-up of the cloned virtual machines. In the **VMs per minute** box, you must supply an integer that indicates how many virtual machines the Provisioning and Cloning capability should start each minute.

Depending on your system setup and the number of clones you created, it is a good practice to stagger starting the virtual machines so that you do not overwhelm your system. Having a large number of virtual machines start at once can slow down your system. The values you supply for this feature depend on your system environment at the time you perform the clone operation and how many clones you are creating.

**Note:** Provisioning and Cloning keeps track of the total number of virtual machines that should have powered on based on the number of minutes that have elapsed since the cloning operation completed. If a problem prevents Provisioning and Cloning from starting the virtual machines on schedule, it uses this cumulative total to specify how many virtual machines to power on when the delay ends. Depending on the length of the delay and the number of clones per minute that you specified, any delay could result in Provisioning and Cloning powering on a large number of virtual machines at one time. After that, though, it only starts the specified number of virtual machines per minute.

- c) **Apply customization specification?** From the drop-down list of available customized specifications, select the one that you want to apply to these clones.
- d) **Use the virtual machine name as the computer name?** If you are using a customization specification with a custom `sysprep` answer file, select this check box to insert the virtual machine name in the answer file.
- e) Review the summary of choices for this operation and click **Apply** to proceed.

## Removing a baseline

The Redeploy window displays available baselines that have been used to create clones. You can remove baselines from this list.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **Redeploy**.
3. Select the baseline and click the blue **Remove** link on the right side of the window.

## Reclaiming space on virtual machines

You can use the Reclaim space feature to find free clusters on NTFS partitions and make them available to the operating system.

### Before you begin

The Reclaim space feature allows Data ONTAP to use space freed when data is deleted in guest operating systems.

This feature has the following requirements:

- VMDKs attached to the virtual machine must be on NFS-backed datastores.
  - Note:** The Reclaim space feature is not supported if the NFS datastore is backed by a qtree on a vFiler unit.
- VMDKs must have NTFS partitions.
  - Note:** If the VMDK is unpartitioned or FAT, the Provisioning and Cloning capability incorrectly lists the disk as having an NTFS partition after the task completes and displays a "Yes" in the "Has NTFS partition(s)?" column. Even though the VMDK now appears to be partitioned, it is still unpartitioned or FAT, and you cannot reclaim space on it.
- ISOs mounted to the virtual machine must be contained in an NFS datastore.
- Storage systems must be running Data ONTAP 7.3.4 or later.
- You should have the VMware guest tools installed.
- When the Reclaim space feature is running, you must not power on the virtual machine.
- You cannot use the cloning feature when the target virtual machine is being used by either Backup and Recovery or Optimization and Migration.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Right-click a datastore or virtual machine and select **NetApp > Provisioning and Cloning > Reclaim space**.
2. Click **OK**.

If the virtual machine is powered on, the Reclaim space feature powers it off. After the process completes, the Reclaim space feature returns the virtual machine to its previous state.

**Note:** If you are using this feature when the virtual machine is powered on, make sure you have the guest operating system tools installed. Without these tools, the Reclaim space feature does not work when it has to power down the virtual machine

If you do not want to install these tools, then you should power down the virtual machine before running the Reclaim space feature.

## Importing virtual machines into XenDesktop

You can manually import virtual machines into XenDesktop.

To create a Citrix XenDesktop import file, select one of the Citrix XenDesktop versions as the connection broker in the Virtual Machine Details window of the Create Rapid Clones Wizard.

The Provisioning and Cloning capability creates the following XenDesktop import file, where [VSC\_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\exports\xenDesktop_timestamp.csv
```

The file contents are in the following format:

```
[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostID]
```

**Note:** XenDesktop 5 does not support importing from vApps.

## Importing the file into XenDesktop 4

You can manually import a Citrix XenDesktop import file into XenDesktop 4.

### Steps

1. Copy the import file to the XenDesktop system.
2. Using the Citrix Access Management Console, choose to create a new desktop group. A wizard launches.
3. Follow the wizard prompts to the fifth panel (Virtual Desktops page), and select the option to import the desktops from a .csv file.

4. Browse to the import file and click **OK**.

## Importing the file into XenDesktop 5

You can manually import a Citrix XenDesktop import file into XenDesktop 5.

### Steps

1. Copy the export file to the XenDesktop system.
2. In Desktop Studio, create a new catalog or modify an existing one in Desktop Studio.  
A wizard launches.
3. Follow the wizard prompts and select the option to import existing virtual machines.
4. Browse to the import file and click **OK**.

## Managing storage controllers

You can view usage and deduplication statistics, and manage volume settings and network interfaces with the Provisioning and Cloning capability.

### (Data ONTAP operating in 7-Mode) Viewing storage controller details

You can view usage and deduplication statistics for a storage controller's aggregates, volumes, and LUNs.

#### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

#### About this task

These steps apply to storage systems that are running Data ONTAP operating in 7-Mode.

#### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage controllers**.
3. Right-click a storage controller and select **View Storage Details**.

### (Data ONTAP operating in 7-Mode) Removing or adding network interfaces, volumes, and aggregates

Volumes and aggregates created outside Provisioning and Cloning must be added in the Resources window to be available for provisioning and cloning operations. You can also use this window to

restrict or expand the available network interfaces, volumes and aggregates. This feature is available only on storage systems running Data ONTAP operating in 7-Mode.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

**Note:** If you are running clustered Data ONTAP, the recommended best practice is to create a Vserver that is suited to the needs of the users. That way you use the security and RBAC capabilities provided by Data ONTAP to control what is available to Virtual Storage Console for VMware vSphere and the ESX or ESXi hosts.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage controllers**.
3. Right-click a storage controller and select **Resources**. Or click the blue **Resources...** link.

All network interfaces, volumes, and aggregates for the physical filer (filer0) and vFiler units are presented.

  - a) To restrict or expand availability, select components and click the directional buttons to move them to the desired column.
    - Items in the left column will *not* be used.
    - Items in the right column will be used.

When the right column is empty, all network interfaces, volumes, and aggregates will be used.

**Note:** New volumes and aggregates created outside Provisioning and Cloning appear in the left column as unused.
  - b) To lock these settings and require storage system credentials for modifications, select the **Prevent further changes** check box at the bottom of the window.
4. Click **Save** when finished.

## Managing volume settings

You can establish advanced settings for new volumes on the storage controller.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

## Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the left pane under Provisioning and Cloning, click **Storage Controllers**.
3. Right-click a storage controller and select **Settings**. Or click the blue **Settings...** link.

If storage controller properties have been locked to prevent changes, you will be prompted for the username and password for that storage controller. You must enter them before you can change volume settings.

4. Set the following options for new volumes created on this storage controller.

If a thin provisioned LUN is deployed into a FlexVol with volume autogrow or snapshot autodelete disabled, it is possible to over-commit the LUN to the volume. This creates an out-of-space condition.

Use these advanced options to modify the FlexVol efficiency settings to match those on the LUN being deployed:

- **Create a new volume for a new LUN** - creates a FlexVol with the same name as the LUN. If a volume with that name already exists, the volume name is appended with a number; for example, Volname01.
  - Note:** If you want the volume container to be resized with the LUN, select this option.
- **Reserve space for volumes that contain thin provisioned LUNs** - results in a thin LUN in a thick volume when a thin LUN is chosen.
- **Thin provision volume clones** - sets space reservation policy to thin provisioning for clones created from this volume.
- **Delete a volume if the last LUN in it has been deleted** - destroys the volume when its last LUN is deleted.
- **Buffer space between volume and LUN (GB)** - Amount of additional capacity in a volume that contains a LUN based datastore.

5. Click **Save** when finished.

## Managing datastores

You can use the Provisioning and Cloning capability to replicate datastores to remote sites, provision, mount, resize, and destroy datastores, and manage deduplication on datastores.

### (Data ONTAP operating in 7-Mode) Replicating datastores to remote sites

The Datastore Remote Replication feature uses Asynchronous SnapMirror to clone NFS datastores from a source vCenter to one or more remote vCenter sites. This feature is available only on storage systems running Data ONTAP operating in 7-Mode.

Setting up a replication target establishes the SnapMirror schedule and creates a FlexVol volume at the target site that is used as the SnapMirror destination. After the initial SnapMirror transfer is

complete, datastore synchronization takes place at the target site. Subsequently, the SnapMirror schedule continues to push incremental updates to the target site so the data is available for any future datastore synchronization operations. After the initial setup, a datastore is not set up at the target site until a manual synchronization is performed from the Datastore Remote Replication panel.

**Note:** The target vCenter version must be the same as or later than the source vCenter version. Otherwise, synchronization succeeds but virtual machines registered on the target are invalid.

The Datastore Remote Replication panel displays the following fields:

- **Sources** - Datastores available to be replicated. Available datastores reside on storage controllers that have SnapMirror licenses.

**Note:** vFiler context-based NFS datastores are not supported.

- **vCenter** - vCenter containing the hosts where the datastore's virtual machines will be registered.
- **Storage Controller** - Storage controller where source datastore will be replicated. Available storage controllers meet the following requirements:
  - have a valid SnapMirror license
  - are not vFiler units
  - have the same Data ONTAP version as the storage controller where the source datastore resides (SnapMirror requirement)
- **Aggregate** - Aggregate where the source datastore will be replicated.
- **Datastore** - Name for the target datastore. The name of the FlexVol volume created as the SnapMirror destination will have `_distribution` appended to it. During synchronization, a clone of the FlexVol volume is created using the name specified in this field.
- **SnapMirror Lag** - Difference between the current time and the timestamp of the snapshot copy last transferred to the destination successfully.
- **SnapMirror Schedule** - Minute, hour, month, and day defined for Asynchronous SnapMirror update. Incremental data is sent from source datastore to a FlexVol volume at the target destination and held until the next Synchronize is performed.

**Note:** This feature replicates datastores from a vCenter to subordinate vCenters across sites. To apply templates to local virtual machines, refer to [Redeploying clones \(locally\)](#) on page 76.

## Setup Replication

You can define a new datastore replication relationship using Provisioning and Cloning.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

**Note:** After establishing a remote replication relationship, if you uninstall and re-install the destination vCenter, you will have to re-create the relationship.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Select the blue **Add** link.
4. In the **Source Datastore** window, select a source datastore.  
Qualified datastores reside on **NetApp** storage controllers with SnapMirror licenses.
5. In the **Target vCenter** window, add or select the vCenter containing the target datastore.  
You can also use this window to remove a selected vCenter from the list.
  - Click **Add vCenter** if the list is empty or to add another vCenter for use as a target. Enter the **vCenter name**, **Username**, and **Password**.
  - Click **Remove** to remove a selected vCenter from the list.
6. In the **Target Destination** window, identify the destination where the source datastore will be distributed.
  - a) First select the host, cluster, or datacenter in which to register the new virtual machines.
  - b) In the **Target Storage Controller** list box, select the storage controller where the target datastore will be created.
  - c) In the **Aggregate** list box, select the aggregate where the target datastore will be created.  
Available aggregates contain space equivalent to or greater than the source volume, and have the same bit type as the source.
  - d) In the **Target Datastore** box, enter a name for the datastore to contain the new virtual machines.
7. In the **Network Mapping** window, establish the network mappings between source and destination virtual machines. Select a destination network from the drop-down for each source network.
8. In the **SnapMirror Schedule** window, establish the schedule for Asynchronous SnapMirror update.  
Incremental data is sent from the source datastore to a FlexVol volume at the target destination and held until the next Synchronize is performed.
 

**Note:** All fields accept comma-separated sequences as well as ranges. For example, 1, 2, 3, 4-7, 8, 10. To represent all valid digits, enter an asterisk (\*).

  - a) **Minute** - Enter the minute past the hour for Asynchronous SnapMirror update to begin.  
Accepts 0 - 59.

- b) **Hour** - Enter the hour of the day for Asynchronous SnapMirror update to begin. For example, enter 3 for 3 a.m., and 15 for 3 p.m.  
Accepts 0 - 23.
  - c) **Month** - Enter the month of the year for Asynchronous SnapMirror update to occur.  
Accepts 1 - 12.
  - d) **Day of Week** - Enter the day of the week for Asynchronous SnapMirror update to occur.  
Accepts 0 - 6.
  - e) **Day of Month** - Enter the day of month for Asynchronous SnapMirror update to occur.  
Accepts 1 - 31.
9. Review the summary and click **OK** to proceed or **Cancel** to exit the wizard without setting up a replication relationship.

## Remove a datastore replication relationship

This feature places the SnapMirror relationship between source and target in Broken-off status. SnapMirror activity ends, but the volume remains at the target site. The deleted relationship remains in the output of the CLI command `snapmirror status` until you delete all snapshots created on the source volume as part of the Datastore Remote Replication process.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Select a target and then select the blue **Remove** link.

## Synchronize

You can manually initiate datastore replication to the designated targets, or write stored incremental updates to the target.

### Before you begin

This action powers off all virtual machines in the source datastore.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

**Steps**

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. To replicate templates or template changes, right-click a target and select **Synchronize**.

**Define or modify a SnapMirror schedule**

You can modify the SnapMirror schedule using Provisioning and Cloning. The SnapMirror schedule defines the schedule for Asynchronous SnapMirror updates. Incremental data is sent from the source datastore to a FlexVol volume at the target destination and held until the next synchronization is performed.

**Before you begin**

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

**About this task**

**Note:** All fields accept comma-separated sequences as well as ranges. For example, 1, 2, 3, 4-7, 8, 10. To represent all valid digits, enter an asterisk (\*).

**Steps**

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Click the blue **SnapMirror Schedule** link.
4. In the **SnapMirror Schedule** window:
  - a) **Minute** - Enter the minute past the hour for Asynchronous SnapMirror update to begin.  
Accepts 0 - 59.
  - b) **Hour** - Enter the hour of the day for Asynchronous SnapMirror update to begin.

**Example**

For example, enter 3 for 3 a.m., or 15 for 3 p.m.

Accepts 0 - 23.

- c) **Month** - Enter the month of the year for Asynchronous SnapMirror update to occur.  
Accepts 1 - 12.
- d) **Day of Week** - Enter the day of the week for Asynchronous SnapMirror update to occur.  
Accepts 0 - 6.

- e) **Day of Month** - Enter the day of month for Asynchronous SnapMirror update to occur.  
Accepts 1 - 31.

## Refreshing the display

You can refresh the Datastore Remote Replication display to show the most recent changes.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the **NetApp** icon for the vCenter Server.
2. In the navigation pane under Provisioning and Cloning, click **DS Remote Replication**.
3. Click the blue **Refresh** link.

## Provisioning datastores

You can create new datastores at the datacenter, cluster, or host level. The new datastores appear on every host in the datacenter or the cluster. For clustered Data ONTAP, you can either select the cluster and choose the Vserver through which to provision the new volume, or provision directly to the Vserver.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

You cannot provision NFS datastores to a direct-connect vFiler unit. A direct-connect vFiler unit cannot add volumes to itself. To enable NFS provisioning, you must add vfiler0 and then select the vFiler unit to which you want to provision storage.

There are also limitations with direct-connect Vservers:

- To provision datastores on direct-connect Vservers, you must create a new role and a new user with the required privileges to provision. The default vsadmin role assigned to Vservers does not contain the volume efficiency commands needed by Provisioning and Cloning.
- Server privileges are not visible to Provisioning and Cloning and cannot be checked prior to the operations. As a result, when connecting directly to a Vserver, Provisioning and Cloning operations might begin but fail later due to insufficient privileges.

- Datastore creation is not supported on striped aggregates. When you add a direct-connect Vserver, any striped aggregates associated with that vServer appear as available. Provisioning on those striped aggregates will fail.

## Steps

1. In **vSphere Client Inventory**, right-click a datacenter, cluster, or host and select **NetApp > Provisioning and Cloning > Provision datastores**.
2. In the **Storage Controller details** window:
  - a) Select the target physical storage controller, vFiler unit, cluster, or Vserver.
  - b) If you selected a cluster as the target, select the **Vserver** through which to provision datastores onto the cluster.
  - c) To identify vFiler units for provisioning and cloning operations, select the **Set vFiler Context** check box and select a unit from the **vFiler** drop-down list.
3. In the **Datastore type** window, select the datastore type.
4. In the **Datastore details** window, specify the following details for the new datastore.
  - **Protocol:** (VMFS only) **FCP** or **iSCSI**.
  - **Size:** Maximum datastore size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
  - **Datastore name:** Use the default or replace with a custom name.
  - **Create new volume container:** (VMFS only) Create a FlexVol to contain the new datastore.
 

**Note:** If you want the volume container to be resized with the LUN, select this option.
  - **Volume:** (VMFS only) Select an available volume from drop-down list.
  - **Aggregate:** Select an available aggregate from the drop-down list.
  - **Thin provision:** Sets space reserve to none, and disables space checks.
 

**Important:** Cloning and datastore creation can fail if your size request uses too much of the aggregate. Capacity is not reserved for an individual datastore. The aggregate is treated as a shared resource pool, where capacity is consumed as each datastore requires it. By eliminating unused but provisioned areas of storage, more space is presented than is actually available. It is expected that all datastores will not use all of their provisioned storage at the same time.
  - **Block size:** (VMFS only) Select an available block size from the drop-down list. For VMFS-5, block size is fixed at 1 MB.
  - **Auto-grow:** (NFS only) - If more space is required, automatically expands the datastore by the increment you specify, up to the size limit you specify.
  - **Grow increment:** (NFS only) Amount of storage added to datastore each time space is needed.
  - **Maximum datastore size:** (NFS only) Limit at which Auto-grow stops.
  - **Datastore cluster:** (vCenter 5 only) Select a datastore cluster in which to add the datastore. Requires SDRS on the vCenter Server.

5. Review the datastore configuration summary and click **Apply** to begin provisioning the new datastore.

## Mounting datastores

You can add currently mounted datastores to a new host.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

**Note:** This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

### Steps

1. In the **vSphere Client Inventory**, right-click the new host and select **NetApp > Provisioning and Cloning > Mount datastores**.
2. Select the datastores to mount on the host.  
The drop-down list contains datastores on the cluster containing the new host. Use CTRL-click to select multiple datastores.
3. Click **OK**.

## Managing deduplication

You can enable deduplication on a selected datastore to optimize space utilization.

### Before you begin

For more about deduplication, refer to your *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. In the **vSphere Client Inventory**, right-click a datastore and select **NetApp > Provisioning and Cloning > Deduplication management**.
2. Select the appropriate check boxes for your configuration.

The Deduplicated column presents previously used space saved by this feature.

- **Enable/Disable** toggles the deduplication feature on or off.
- **Start** begins deduplication from the last marker position.
- **Scan** begins deduplication at the beginning of the volume.

## Resizing datastores

You can increase or decrease NFS datastore sizes. VMFS datastore sizes can be increased but not decreased.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

If you want the volume container to be resized with the LUN, go to **Storage Controllers > Settings**. Select the **Advanced** check box, and select the option to **Create a new volume for a new LUN**. Otherwise, the container volume will not be resized with the LUN.

**Note:** This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

### Steps

1. In the **vSphere Client Inventory**, right-click a datastore and select **NetApp > Provisioning and Cloning > Resize**.
2. Enter the new datastore size and click **OK**.
3. Click **Yes** to confirm the operation.

## Destroying datastores

You can permanently destroy a datastore, including all virtual machines, datastores, Snapshot copies and FlexVolumes.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

The destroy datastore feature performs the following actions:

- destroys all virtual machines in a datastore
- unregisters and detaches the datastore from the vSphere Client environment

- frees the space on the storage controller

**Note:** This feature is unavailable when the target virtual machine is in use by the Backup and Recovery capability or the Optimization and Migration capability.

### Steps

1. In the **vSphere Client Inventory**, right-click a datastore and select **NetApp > Provisioning and Cloning > Destroy**.
2. Click **OK** to proceed.

### About multiple datastores (NFS only)

When you destroy the last datastore of a golden volume, the golden volume is destroyed only if you select the option to **Delete a volume if the last LUN in it has been deleted** for that volume.

When you create more than one new NFS datastore, the Provisioning and Cloning capability first creates a golden volume on the controller. The datastores that are then attached to the vSphere Client are FlexClones of the golden volume. Because the FlexClones share the storage with the golden volume, this space is not wasted. When the last datastore (FlexClone) of a golden volume is destroyed, the golden volume is destroyed only if you select the option to **Delete a volume if the last LUN in it has been deleted** for that volume. For details, see [Managing volume settings](#) on page 81.

## Provisioning and Cloning support files

There are several support files used and created by the Provisioning and Cloning capability.

### Preferences File

The `preferences` file contains options that control Provisioning and Cloning capability operation.

This file is stored by default in the following location, where `[VSC_home]` is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\kaminoprefs.xml
```

The preferences file contains the following options.

Type	Value	Description
log4j.config.file	log4j.properties	Filename of the log configuration file. This file contains parameters that determine the level of events that are written to the log file. The log configuration filename can be changed here in the <code>kaminoprefs</code> file, and the new configuration file must be stored in <code>[VSC_home]\etc\kamino</code> .

Type	Value	Description
default.create.copyNvram File	false	Clones the VM BIOS file (.nvram). This option is necessary when using Citrix Provisioning server because the boot device order and PXE settings need to be cloned as well.
default.create.exportToAllVmKernelNics	false	Determines if the storage controller exports file will contain all IP addresses of all VMkernel ports. For example, if a dedicated storage network with IP address 192.168.0/24 and a second VMkernel network 10.10/16 on a public network for ISO mounting have this value set to true, the newly created datastores are exported to VMkernel addresses on both networks. The default behavior of false restricts the export to only the VMkernel addresses that share the same network as the IP address of the controller used to mount the datastore.
default.create.volume.option.fractional_reserve	0	Fractional reserve allows the system administrator to reduce the amount of guaranteed available reserve space. By doing this, the administrator can tune the amount of space guaranteed to be available for LUN overwrites based on application requirements and make better use of available volume space. It is adjustable between 0 and 100% of the size of the LUN.
default.create.volume.option.no_atime_update	on	Disables updating of access time when files are read. Leave on for better performance in a heavy read traffic environment.
default.create.volume.option.nosnap	on	Disables Snapshot copies for new FlexClone volumes.
default.create.volume.option.sis	true	Enables deduplication on new volumes.
default.create.volume.snapshot.preserve	0	Sets the snapshot reserve size for new volumes to the specified percentage.

Type	Value	Description
default.create.volume.lunSpaceBuffer	5	Sets 5 GB as the amount of additional capacity in a volume that contains a LUN-based datastore.
default.create.volume.useThickForVolThatContainLun	false	<p>Space reservation override for thin LUN-based datastore.</p> <p>Default volume space reservation for a thin LUN is thin. Setting this option to true results in a thin LUN in a thick volume when a thin LUN is chosen.</p> <p><b>Note:</b> Applies only to Datastore Create. It has no effect on Datastore Clone.</p>
default.create.volume.useThinForVolClone	true	<p>Space reservation override for volumes created using FlexClone technology.</p> <p>Datastores will all be thin FlexClones unless this option is set to false.</p> <p><b>Note:</b> Applies only to Datastore Clone. It has no effect on Datastore Create.</p>
default.destroy.checkPossibleCrossMount	true	If an administrator mounts a volume for use as a datastore to an ESX host using a different IP address or volume path than what was used to mount the volume to other ESX hosts, there is a potential for the Destroy feature to incorrectly report which VMs are present on the datastore. This is because the datastore UUID is generated by the ESX host based on the IP address and volume path. checkPossibleCrossMount addresses this issue and therefore should be left set to true. The only time this should be set to false is if the Destroy window stays in the "Loading..." state for an extended period of time and there is no chance of the issue described above.
default.destroy.destroyParentIfLastClone	true	When last FlexClone volume is destroyed, the golden volume is destroyed.
default.destroy.onlyOffline	false	By default, the Destroy feature takes offline and destroys the volume that was mounted

Type	Value	Description
		as a datastore. Setting this option to <code>true</code> leaves the volume intact, just offline.
<code>default.destroy.onlyUnreg</code>	<code>false</code>	By default, the Destroy feature chooses the most efficient VM cleanup method. Setting this option to <code>true</code> forces it to unregister the VM in all cases.
<code>default.destroy.destroyVolumeIfLastLun</code>	<code>false</code>	When last LUN is removed, volume is not destroyed. Setting this option to <code>true</code> causes volume to be destroyed when last LUN is removed.
<code>default.mount.volume.mountUsingHostname</code>	Not present	If <code>true</code> , hostname is retrieved by reverse name lookup using system-configured lookup service. Hostname is then used to mount the datastore.
<code>default.restrict.nfs.mount.networks</code>	None	By default, any matching network between the controller and the ESX VMKernel is used to mount datastores. In some implementations, certain networks should not be used for mounting, even if they contain the required interfaces on both the controller and ESX hosts. This value prevents a network from being used to mount NFS datastores. The example below tells the Provisioning and Cloning capability to ignore the 192.168.2.0, 10.1.0.0, and 2001:5a0:400:200::0 networks when determining the IP address on the controller to use for the next NFS datastore:  <pre>&lt;entry key="default.restrict.nfs.mount.networks" value="192.168.2.0;10.1.0.0;2001:5a0:400:200::0"/&gt;</pre>
<code>default.restrict.iscsi.mount.networks</code>	None	By default, any matching network between the controller and the ESX VMKernel is used to mount datastores. In some implementations, certain networks should not be used for mounting, even if they contain the required interfaces on both the

Type	Value	Description
		<p>controller and ESX hosts. This value prevents a network from being used to mount iSCSI datastores. The example below tells the Provisioning and Cloning capability to ignore the 192.168.2.0, 10.1.0.0, and 2001:5a0:400:200::0 networks when determining the IP address on the controller that will be used for the next iSCSI datastore.</p> <pre data-bbox="767 534 1239 673">&lt;entry key="default.restrict.iscsi.mount.networks" value="192.168.2.0;10.1.0.0;2001:5a0:400:200::0"/&gt;</pre>
default.allow.nfs.mount.networks	all	<p>Any matching network between controller and ESX VMKernel is used to mount NFS datastores. In some implementations, it is important to fence off the networks that should be used for NFS and those that should be used for iSCSI. This option allows only specified networks to mount NFS datastores. The network specified is that of the VMKernel interface, not an interface on the storage system. The <code>all</code> keyword means that all matching networks between the VMkernels and the storage system interfaces can be used for mounting NFS datastores. You can specify different networks here to enable mounting across different subnets.</p>
default.allow.iscsi.mount.networks	all	<p>Any matching network between controller and ESX VMKernel is used to map iSCSI datastores. In some implementations, it is important to fence off the networks that should be used for NFS and those that should be used for iSCSI. This option allows only the specified networks to map iSCSI-based VMFS datastores. The network specified is that of the VMKernel interface, not an interface on the storage system.</p>

Type	Value	Description
default.create.igroup. alua	yes	Configures the ALUA setting to apply to new iGroups.
default.option. copythreads	4	Controls number of simultaneous file copies. Maximum is 16.
default.option. powerOffThreads	10	Controls number of simultaneous VM power off operations. Maximum is 16.
default.option. checkAlignment	true	Controls whether alignment is checked on the source VM.
default.option. fallBackToNdmpCopy	true	Specifies whether the clone operation should fall back to using NDMP copy if the flex-clone operation begins to copy blocks rather than clone them. This happens when a block has reached the max shared blocks limit.
default.option. fallBackThreshold	0	Specifies a threshold for copying blocks before falling back to using an NDMP copy. Blocks copied threshold is specified as a percentage. This is used if the <code>default.option.fallBackToNdmpCopy</code> preference is set to true. For example, setting this value to 20 will fall back to using NDMP copy if 20% of the blocks during the FlexClone operation have been copied. Default behavior is to fall back immediately when any blocks have been copied.
default.option. controllerValidationTimeout	120	By default, controller validation is cancelled if it takes longer than two minutes to validate a controller. In some large environments, it may be necessary to increase this timeout. The example below tells the Provisioning and Cloning capability to wait five minutes (300 seconds) before canceling the validation operation.  <pre>&lt;entry key="default.option.controllerValidationTimeout" value="300" /&gt;</pre>

## Provisioning and Cloning logs

The Provisioning and Cloning capability creates a log file that includes the capability version, build number, build date, and all event messages and errors.

The log file is created in the following location, where [VSC\_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\log\kamino.log
```

Each log file is rotated at a size of 8,192 KB. Up to nine log files are retained.

## Log configuration file

The level of events recorded in the log file is determined by parameters set in the log configuration file.

The log configuration file is stored in the following location, where [VSC\_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\log4j.properties
```

## Modifying logging levels

You can change the level of events that are logged for the Provisioning and Cloning capability at runtime by modifying the log configuration file.

Edit the log configuration file [VSC\_home]\etc\kamino\log4j.properties, where [VSC\_home] is your VSC for VMware vSphere installation directory.

The new settings take effect immediately. It is not necessary to restart the VSC for VMware vSphere service.

## Using a different log configuration file

You can define a different log configuration file by entering the new configuration file name in the `kaminoprefs.xml` file.

### Before you begin

To perform these steps, you must have access to the machine where Virtual Storage Console for VMware vSphere is installed.

### Steps

1. Open the following file for editing:

```
[VSC_home]\etc\kamino\kaminoprefs.xml, where [VSC_home] is your VSC for VMware vSphere installation directory.
```

2. Modify the following entry by replacing `log4j.properties` with the new log configuration filename:

```
<entry key="log4j.config.file" value="log4j.properties" />
```

3. Place the new log configuration file in the following directory:  
[VSC\_home]\etc\kamino
4. Restart the VSC for VMware vSphere service.

**Note:** The VSC for VMware vSphere service must be restarted only if the `kaminoprefs.xml` is modified to point to a different configuration file. Changes in the configuration file itself do not require an VSC for VMware vSphere service restart.

## Export Files

The Provisioning and Cloning capability creates a network configuration file with information about each virtual machine created. When requested, the Provisioning and Cloning capability creates Citrix XenDesktop export files.

### Network configuration file

After each cloning procedure, Provisioning and Cloning creates a `.csv` network configuration file.

The network configuration file is created in the following location, where `[VSC_home]` is your Virtual Storage Console for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\exports
```

The configuration file is written in the following format:

```
import_generic_date_time.csv
```

For example: `import_generic_2009_03_06_11_04.csv`

The file includes the following information for each virtual machine:

- VM Name
- UUID, VMX Path
- Number of CPUs
- Memory in MB
- Guest Full Name
- Guest Alternate Name
- Datastore Name

The file also includes the following information for up to ten NICs:

- NIC x
- NIC x Network Label
- NIC x Address Type
- Manual - Statically assigned MAC address.

- Generated - Automatically generated MAC address.
- Assigned - MAC address assigned by Virtual Center.
- NIC x WOL Enabled
- Wake-on-LAN enabled or disabled on this virtual network adapter
- NIC x MAC Address

### **XenDesktop export file**

When XenDesktop connection broker output is requested in the cloning wizard, the Provisioning and Cloning capability creates a .csv file in the exports directory.

The XenDesktop connection broker output file is created in the following location, where [VSC\_home] is your VSC for VMware vSphere installation directory:

```
[VSC_home]\etc\kamino\exports
```

The filename is in the following format:

```
xenDesktop_timestamp.csv
```

The file includes the following information:

```
[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostID]
```

# Optimizing and migrating datastores and virtual machines

---

Optimization and Migration provides a non-disruptive (interim) solution for the performance penalty introduced by misaligned virtual machines. Rather than align the misaligned VMDKs, which requires downtime, Optimization and Migration aligns the I/O -- without requiring downtime -- by offsetting the VMDKs within specialized (optimized) datastores.

Optimization and Migration also enables you to migrate one or more virtual machines (optimized and unoptimized) into a new or existing datastore. In addition, you can review the alignment status of the virtual machines.

The interface for Optimization and Migration is integrated with VMware vCenter Server and works in conjunction with the VMware Storage vMotion feature to enable a running virtual machine to be moved between datastores.

You can use Optimization and Migration to perform the following tasks:

- **Align running virtual machines**

The online alignment tool lets you align I/O on virtual machines without having to power them down.

- **Scan one or more datastores**

By scanning datastores, you can determine which virtual machines are misaligned. You can scan all the datastores at the same time, or only the datastores you select.

- **Schedule scans**

You can schedule scans of the datastores so that they occur automatically at specific times.

**Note:** It is a good practice to schedule the scans to occur during non-critical production times. The time required to perform a scan can increase as more virtual machines are scanned.

- **Migrate virtual machines**

The migration feature allows you to migrate virtual machines either singly or as a group.

**Note:** Migrating multiple virtual machines at one time is very I/O intensive. Your system is likely to slow down while the migration is in process. You might want to limit the number of virtual machines that you migrate at one time to avoid over-stressing your system during the migration.

- **View virtual machines organized based on their state**

The built-in sort feature places the virtual machines in folders that let you quickly determine their states and the actions you can take. Based on the folders, you know the following:

- Whether a virtual machine is functionally (I/O only) or actually aligned.  
The **Aligned > Functionally Aligned** and the **Aligned > Actually Aligned** folders contain these virtual machines.
- If a misaligned virtual machine can be aligned using the online feature of the Optimization and Migration capability.

The **Misaligned > Online Migration** folder contains these virtual machines.

- If a misaligned virtual machine must be aligned using an offline tool, such as VMware vCenter Converter.

To use an offline alignment tool, you must power off the virtual machine.

The **Misaligned > Offline Migration** folder contains these virtual machines.

- If a virtual machine is not aligned and cannot be aligned.

Optimization and Migration places the virtual machine in the **Misaligned > Other** folder if it cannot be aligned. This can happen if the virtual machine:

- Has disks that are inaccessible
- Has a disk size of 0
- Has more than one disk with different offsets
- Has multiple disks spanning multiple datastores
- Lacks partitions
- Is an independent disk
- Reports an error during read attempts
- Is a dynamic disk

**Note:** If you use the Optimization and Migration capability when you have a dynamic disk, the capability might give you a false indication of alignment.

## Types of alignments

You can use Optimization and Migration to perform online alignments and to sort virtual machines so that you can see which ones require you to use another tool and perform an offline alignment. The type of alignment you perform determines whether the virtual machine is functionally aligned or actually aligned.

### Online and offline alignments

When you use Optimization and Migration to perform an online alignment, you do not need to power down a virtual machine. If Optimization and Migration is not able to perform an online alignment, you must power down the virtual machine and use a tool such as the VMware vCenter Converter to align the virtual machines offline. Optimization and Migration sorts the virtual machines into folders based on whether you can perform an online alignment, must perform an offline alignment, or cannot align the virtual machine.

During an online alignment, Optimization and Migration moves the virtual machine you specify to a single, optimized VMFS datastore and performs a functional alignment. The datastore can be either a new or existing datastore.

If you have multiple virtual machines that have the same misalignment, you can migrate them in a batch. Optimization and Migration moves all the virtual machines to the same, optimized datastore.

**Note:** Migrating multiple virtual machines at one time is very I/O intensive. Your system is likely to slow down during the migration. You may want to limit the number of virtual machines that you migrate at one time to avoid over-stressing your system during the migration.

### Functional and actual alignments

Depending on whether you use Optimization and Migration to perform an online alignment or another tool to perform an offline alignment, your virtual machines are aligned either functionally or actually.

In a functional alignment, Optimization and Migration moves the misaligned virtual machines to an optimized datastore that uses a prefix to ensure that the virtual machines align on correct boundaries. As a result, the virtual machines perform as though they are aligned.

**Note:** If you clone a virtual machine that has been functionally aligned, the clone will be misaligned. When you use Provisioning and Cloning to clone virtual machines, it warns you if the source is misaligned. Other cloning tools do not provide a warning. As a result, you should use only Provisioning and Cloning to clone datastores.

In an actual alignment, the partitions of the virtual machine's hard disk are aligned to the storage system and start at the correct offset. For this type of alignment, you perform an offline alignment, which modifies the contents of a virtual hard disk to align the virtual machine.

Optimization and Migration does not modify the contents of a virtual hard disk.

## Important notes about using Optimization and Migration

When you use Optimization and Migration, it is important to understand some key points about how it works.

There are several things that can affect the results you see when you use Optimization and Migration, including the following:

- **Multiple partitions on a disk and different offsets for the partitions**  
When a virtual machine has multiple partitions on a disk and the partitions have different offsets, Optimization and Migration functionally aligns the virtual machines to the offset of the largest partition by migrating them to that offset.
- **Multiple disks with different offsets**  
If a virtual machine has more than one disk with different offsets, you cannot use Optimization and Migration to align that virtual machine.  
You can use Optimization and Migration to migrate that virtual machine to another datastore. You can also power down the virtual machine and use an offline alignment tool such as VMware vCenter Converter to perform an alignment.  
Optimization and Migration places these virtual machines in the **Misaligned > Offline Migration** folder in the Virtual Machine Alignment panel. Optimization and Migration can only align virtual machines listed in the **Misaligned > Online Migration** folder.
- **VMware's Storage Distributed Resource Scheduler (SDRS) feature**

If you want to use VMware's SDRS feature with datastores created by Optimization and Migration, you should not put datastores with varying offsets in a single datastore cluster. In addition, do not mix optimized and non-optimized datastores in the same datastore cluster.

- **Lock management feature**

The lock management feature of Virtual Storage Console for VMware vSphere ensures that when one capability is using a target datastore or virtual machine, other capabilities cannot use that datastore or virtual machine at the same time.

For example, if you are using Provisioning and Cloning to perform a cloning operation, you can use Optimization and Migration to scan and create a datastore. However, you cannot use Optimization and Migration to align the same virtual machine that Provisioning and Cloning is using.

- **Migration of multiple virtual machines**

You can migrate multiple virtual machines at one time; however, each migration is I/O intensive and can slow down your system while it is underway.

Before you migrate multiple virtual machines, consider what your environment can handle. Because of the I/O load, you might want to limit the number of virtual machines you include in a single migration to avoid over-stressing your system.

- **Deduplication temporarily turned off**

When you use Optimization and Migration to migrate a virtual machine, deduplication is temporarily turned off. As a result, the migration might require more space.

- **VAAI extended copy feature**

If a datastore contains virtual machines that have been aligned using Optimization and Migration, you cannot use the VAAI extended copy feature.

The extended copy feature is not available on these optimized datastores. To determine whether a datastore has been optimized, check the Scan Manager panel to see if there is a *yes* in the Optimized column.

In addition, datastores that have been provisioned by Optimization and Migration use the format `<original_datastore_name>_optimized_<group_number>` where group number is the alignment offset of the virtual machine.

- **VMware Storage VMotion**

If you use VMware Storage VMotion, you must make sure that you specify the correct optimized datastore as the destination. Performance problems can occur if you specify the wrong datastore.

Whenever you migrate misaligned virtual machines to optimized datastores using VMware Storage VMotion, you should perform the task from within Optimization and Migration.

## Considerations when working with NFS datastores

There are several considerations you need to be aware of if you want to use Optimization and Migration with NFS datastores.

- The storage systems must be running Data ONTAP 8.1.3 or later.
- You should not migrate virtual machines that are currently aligned to a volume that has just been created using Optimization and Migration.

- When new volumes are created with an optimized offset at the source, a new baseline mirror or vault must be created.  
This is normal behavior that must occur any time a new volume is created.
- Optimization and Migration is not supported for Snaplock volumes.
- NDMP or dump restores to volumes created with Optimization and Migration are up to 60 percent slower.
- Using the `ndmpcopy` command to copy a volume to an Optimization and Migration volume is significantly slower.

## The Optimization and Migration workflow

The Optimization and Migration capability lets you scan datastores and correct the alignment of certain misaligned virtual machines (VM) without having to power down the VM.

The following high-level workflow takes advantage of the Optimization and Migration capability features.

**Note:** The sections that follow provide detailed information on these steps.

1. Go to the Scan Manager panel and click **Scan now** to scan the datastores and VMs.  
This action provides information about the datastores and whether they have been optimized. You can specify which datastores you want to include in the scan. Depending on how many datastores are scanned, it may take a few minutes before the Virtual Machine Alignment panel folders are populated.
2. Go to the Virtual Machine Alignment panel to view the alignment status of the VMs.
3. Check which misaligned VMs can be corrected with the online alignment tool.
4. Use the Optimization and Migration capability to migrate those VMs to a new or existing datastore to correct the alignment issue.
5. Alternatively, select VMs that you wish to migrate.  
You can migrate the VMs as a group as long as the VMs have the same misalignment. When you migrate VMs, the Virtual machine migration wizard starts. This wizard provides a series of dialog boxes that prompt you for the information you need to provide.
6. Maintain an optimized environment by scanning the datastores on a regular basis.

## Scanning the datastores

You can use Optimization and Migration to scan datastores and see which ones are optimized.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

It is a good practice to schedule the scans to occur during non-critical production times. When scanning virtual machines in VMFS datastores, Optimization and Migration uses VMware snapshots. Taking these snapshots increases the time required for the scan as well as requiring more space. Optimization and Migration quickly deletes the VMware snapshots as soon as they are no longer needed, which frees up the space they were using.

### Steps

1. Select the Optimization and Migration capability **Scan Manager** panel.
2. Specify which datastores you want the Optimization and Migration capability to scan.

You have the following options:

- Click **Scan now** without selecting anything.
- Exclude certain datastores from being scanned.  
Check the box next to the names of those datastores and click **Exclude**.
- Select specific datastores for scanning.  
Check the box next to the names of those datastores.

**Note:** **Scan now** changes to **Scan selected** when you check the box next to a datastore.

- Include datastores that previously were excluded.  
Check the box next to their names and click **Include**.
3. Click either **Scan selected** or **Scan now**.
- When you choose **Scan selected**, the Optimization and Migration capability scans the datastores with checks next to their names. If you have not selected any datastores, choose the **Scan now** option, which scans the entire datacenter.

After the scan completes, the Optimization and Migration capability displays information about the datastores, including the following:

- Whether the datastore is optimized
- Whether it was included in the most recent scan
- When the last scan of the datastore occurred

### After you finish

After you have scanned the datastores, you can go to the Virtual Machine Alignment panel and check the alignment status of the virtual machines.

## Scheduling a scan of datastores

Optimization and Migration lets you set up automatic scans of the datastores.

### About this task

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select the Optimization and Migration **Scan Manager** panel.
2. Specify which datastores you want Optimization and Migration to scan.

You have the following options:

- Exclude certain datastores from being scanned.  
Check the box next to the names of those datastores and click **Exclude**.
- Select specific datastores for scanning.  
Check the box next to the names of those datastores.
- Include datastores that previously were excluded.  
Check the box next to their names and click **Include**.

3. Click **View/edit schedule**.
4. In the pop-up box that appears, specify the day and time when you want the scan to occur.

You can specify multiple days. All scans will occur at the same time on each day that you specify.

If you do not use the **Exclude** button, Optimization and Migration scans the all of the datastores.

**Note:** It is a good practice to schedule the scans to occur during non-critical production times. When scanning virtual machines in VMFS datastores, Optimization and Migration uses VMware snapshots. Taking these snapshots increases the time required for the scan.

5. When you have the schedule for scans set up, click **Save**.

Optimization and Migration automatically scans the datastores according to the schedule you set up.

## Performing an online alignment

In many cases, Optimization and Migration allows you to perform online alignments of virtual machines. As a result, you do not have to power the virtual machines down during the alignment procedure. Optimization and Migration aligns the virtual machines by using the Virtual Machine Migration Wizard to migrate them into appropriately offset optimized datastores.

### Before you begin

You should perform a full scan of the datastores in order to populate the database with information about which virtual machines need to be aligned. Because a scan can slow a system down, it is a good practice to schedule the scans to occur during non-critical production times

### About this task

The following are some points to keep in mind about this task:

- Optimization and Migration provides a non-disruptive (interim) solution for the performance penalty introduced by misaligned VMs.  
Instead of aligning the misaligned VMDKs, which requires downtime, Optimization and Migration aligns the I/O – without requiring downtime - by offsetting the VMDKs within specialized (optimized) datastores.
- In some cases, virtual machines cannot be aligned using the Optimization and Migration online tool.  
When a virtual machine cannot use the online alignment (migration) feature, Optimization and Migration places the virtual machine in the **Offline Migration** folder. You must power down the virtual machine and use an offline tool such as the VMware vCenter Converter to align it.
- If you are working with a Storage Distributed Resource Scheduler (DRS), you must create a storage DRS cluster for each required offset.
- The VAAI extended copy feature is not available on optimized datastores.
- You should also make sure that you do not mix optimized datastores with non-optimized datastores.
- You cannot cancel an online alignment that is in progress.
- Optimization and Migration only supports online alignment of NFS datastores if you are running Data ONTAP 8.1.3 or higher.

**Note:** If you are using Optimization and Migration to align NFS datastores, you should review the information provided in *Considerations when working with NFS datastores* on page 103.

- In some cases, if the virtual machine is being used by another capability, you cannot perform an online alignment.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see *Authentication and user management with vCenter RBAC and Data ONTAP RBAC* on page 31.

## Steps

1. Select **Virtual Machine Alignment** panel of Optimization and Migration.
2. Open **Misaligned > Online Migration**.
3. Select the virtual machines that you want to align.

You have the following options:

- To align all the listed virtual machines, select **Migrate all**.
  - To align specific virtual machines, check the box next to the names of the virtual machines you want to align and then click **Migrate**.
4. The **Virtual machine migration wizard** launches and prompts you for information about the migration.
    - a) Select the target storage controller that will contain the virtual machines.
    - b) If you are using vFiler units, select the **Set vFiler Context** check box and select a unit from the drop-down list.
  5. Specify whether you want to migrate the virtual machines to an existing datastore or create a new datastore.
  6. If you are using an existing datastore, select the datastore from the list of datastores displayed in the **Choose a datastore for the virtual machines** dialog box.

This dialog displays information about the existing datastores, including the following:

- **Datastore name**
  - **Datastore type**
  - **Capacity** of the datastore in GB
  - **Free space** available on the datastore in GB
7. If you are creating a datastore, specify its type: **VMFS** or **NFS**.  
To continue without creating a new datastore, click **Next**.
  8. In the **Specify the details for the new datastore** dialog box, provide details about the datastore you want to create.
    - **Protocol** (VMFS only) - FCP or iSCSI.
    - **Size (GB)** - Specify the maximum size for the datastore.  
The maximum size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
    - **Datastore name** - Supply a name for the datastore you are creating.
    - **Create new volume container** (VMFS only) - Create a volume with the same name as the LUN. If a volume with that name already exists, the Optimization and Migration capability appends a number to the volume name; for example, `volname01`.
    - **Volume** (VMFS only) - Select a volume from drop-down list of available volumes.
    - **Thin provision** - Select this option to set space reserve to none and disable space checks.
    - **Block size** (VMFS only) - Select the block sizes. For VMFS-5, block size is fixed at 1 MB.

- **Aggregate** (NFS only) - Select an aggregate from the drop-down list of available aggregates.
  - **Auto-grow** (NFS only) - Check this box if you want to use the auto-grow feature.  
When you check auto-grow, the datastore is automatically expanded by the increment you specify up to the size limit you specify anytime space is needed.
  - **Grow increment (GB)** (NFS only) - Specify the amount of storage added to datastore each time space is needed.
  - **Maximum datastore size (GB)** (NFS only) - Specify in GB the limit at which the auto-grow feature stops.
9. Review the summary page and click **Apply** to proceed. To return to previous pages and modify settings, click **Back**.
- The Optimization and Migration capability displays the status of the virtual machines as they are aligned.
10. After the online alignment completes, the virtual machines you just aligned no longer appear in the **Misaligned > Online Migration** folder.

## Migrating virtual machines

You can use Optimization and Migration to migrate a group of virtual machines at one time.

### About this task

Migrating multiple virtual machines at one time is very I/O intensive. Your system is likely to slow down while the migration is in progress. You might want to limit the number of virtual machines that you migrate at one time to avoid over-stressing your system during the migration.

You cannot cancel a migration that is in progress.

Migration is not restricted to optimized datastores. You can select groups of virtual machines to migrate to any datastore as long as the source and destination datastore both have the same offset or both have no offset.

In some cases, when the virtual machine is being used by another capability, you cannot use the migration features.

Optimization and Migration only supports migration of NFS datastores if you are running Data ONTAP 8.1.3 or higher.

**Note:** If you are using Optimization and Migration to migrate NFS datastores, you should review the information provided in [Considerations when working with NFS datastores](#) on page 103.

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

## Steps

1. From the **Virtual Machine Alignment** panel, select the virtual machines that you want to migrate.

Click **Migrate all** to migrate all the virtual machines or **Migrate** to migrate only the virtual machines you have checked.

2. The **Virtual machine migration wizard** launches and prompts you for information about the migration.
  - a) Select the target storage controller that will contain the virtual machines.
  - b) If you are using vFiler units, select the **Set vFiler Context** check box and select a unit from the drop-down list.
3. Specify whether you want to migrate the virtual machines to an existing datastore or create a new datastore.
4. If you are using an existing datastore, select the datastore from the list of datastores displayed in the **Choose a datastore for the virtual machines** dialog box.

This dialog displays information about the existing datastores, including the following:

- **Datastore name**
- **Datastore type**
- **Capacity** of the datastore in GB
- **Free space** available on the datastore in GB

5. If you are creating a datastore, specify its type: **VMFS** or **NFS**.

To continue without creating a new datastore, click **Next**.

6. When you create a new datastore, the **Specify the details for the new datastore** dialog box prompts you for details.
  - **Protocol** (VMFS only) - FCP or iSCSI.
  - **Size (GB)** - Specify the maximum size for the datastore.  
The maximum size depends on the controller and space available. For details, see the *Data ONTAP Storage Management Guide* for your Data ONTAP release.
  - **Datastore name** - Supply a name for the datastore you are creating.
  - **Create new volume container** (VMFS only) - Create a volume with the same name as the LUN. If a volume with that name already exists, the Optimization and Migration capability appends a number to the volume name; for example, Vol1name01.
  - **Volume** (VMFS only) - Select a volume from drop-down list of available volumes.
  - **Thin provision** - Select this option to set space reserve to none and disable space checks. Selecting thin provisioning for a datastore can let you over subscribe the storage controller. The limit for resizing the datastore is 10 times its initial size.
  - **Block size** (VMFS only) - Select block sizes. For VMFS-5, block size is fixed at 1 MB.
  - **Aggregate** (NFS only) - Select an aggregate from the drop-down list of available aggregates.
  - **Auto-grow** (NFS only) - Check this box if you want to use the auto-grow feature.

When you check auto-grow, the datastore is automatically expanded by the increment you specify up to the size limit you specify anytime space is needed.

- **Grow increment (GB)** (NFS only) - Specify the amount of storage added to datastore each time space is needed.
- **Maximum datastore size (GB)** (NFS only) - Specify in GB the limit at which the auto-grow feature stops.

7. Review the summary page and click **Apply** to proceed.

To return to previous pages and modify settings, click **Back**.

## Backing up and restoring data

---

Backup and Recovery enables you to rapidly back up and recover multihost configurations running on NetApp storage systems.

You can use Backup and Recovery to perform the following tasks:

- Create on-demand backup copies of individual virtual machines, datastores, or a datacenter
- Schedule automated backup copies of individual virtual machines, datastores, or a datacenter
- Support virtual machines and datastores that are located on either NFS directories or VMFS file systems
- Mount a backup copy to verify its content prior to restoring it
- Restore datastores to their original locations
- Restore virtual machines to their current locations
- Restore virtual machine disks (VMDKs) to their current locations or alternate locations

### Related concepts

[Configuring Backup and Recovery](#) on page 113

[Managing backups](#) on page 116

[Restoring data from backups](#) on page 122

[Single file restore](#) on page 127

[VSC CLI commands](#) on page 136

### Related references

[Resolution of Backup and Recovery issues](#) on page 202

## Backup and Recovery requirements

Your datastore and virtual machines must meet Snapshot and SnapRestore requirements before you can use the Backup and Recovery capability.

Following are some Snapshot and SnapRestore requirements:

- Snapshot protection is enabled in the volumes where those datastore and virtual machine system images reside.
- SnapRestore technology is licensed for the storage systems where those datastore and virtual machine system images reside.

## Backup and Recovery requirements for optional SnapMirror protection

Optional SnapMirror protection requires that you perform certain configuration tasks on virtual objects in your environment prior to using Backup and Recovery.

Following are some SnapMirror requirements:

- The volumes containing the active datastore and virtual machine images must be configured as SnapMirror source volumes.
- The source volumes must have a SnapMirror relationship with target volumes on a second SnapMirror-licensed storage system that is located a safe distance from the source storage system.
- The host names and IP addresses of the SnapMirror source and destination storage systems must be resolvable for the SnapManager for Virtual Infrastructure server, either through a configured DNS server or through host entries added to the host file on the SnapManager for Virtual Infrastructure server.
- Cluster or Vserver administrators must create a Vserver management LIF, which is required by Backup and Recovery to update SnapMirror relationships for storage systems using clustered Data ONTAP.

## Configuring Backup and Recovery

Before you begin using Backup and Recovery to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must use Monitoring and Host Configuration to add the storage systems that contain the datastores and virtual machines for which you are creating backup copies. You can also create a custom user account for a storage system.

Backup and Recovery uses the storage credentials that you specify in Monitoring and Host Configuration. You must confirm that all previous storage credentials from Backup and Recovery now exist in the list of VSC storage credentials.

You should check the Setup panel in Backup and Recovery after the installation is complete to see if there are any issues with the storage system credentials. A list of the storage systems on this page shows any systems that do not have the necessary permissions to perform Backup and Recovery operations and how to resolve any issues using Monitoring and Host Configuration.

## Authentication methods in the Backup and Recovery capability

You can override the default authentication method by creating a custom user account through the CLI, which enables you to log in with authentication credentials other than your Windows credentials.

## Custom user accounts for accessing a storage system

A non-root or non-administrator account might be required for the Backup and Recovery capability to access a specific storage system.

In this case, you must create a custom storage system account with a new storage system role, group, and user as described in the following table.

Item	Description
Role	The new role must allow the Backup and Recovery capability to access the storage system data through its APIs.
Group	A storage system maintains groups as a collection of roles. The group you create must contain your new role.
User	A user account that the Backup and Recovery capability uses to access a storage system must be a member of a group that contains a role. You can create this user and assign a password to it, then you are able to add a storage system to Backup and Recovery with the assigned user name and password.

For more information on how to manage users on your storage system, refer to your storage system's administrator guide.

## Creating custom users

You can use the CLI to create a custom user. The credentials of a custom user provide the same access to commands and features as an administrator who logs in using the default Windows credentials authentication method.

### Steps

1. Double-click the **VSC CLI** desktop icon or navigate to **Start > All Programs > NetApp > Virtual Storage Console > VSC CLI**.
2. Enter the `smvi servercredential set` command.
3. Enter a user name and a password.

## Creating a custom storage system user account

You can use your storage system's CLI to create a custom storage system account with a new storage system role, group, and user.

### About this task

Complete the following steps from the CLI of the storage system that the Backup and Recovery capability needs to access.

### Steps

1. Use the following command to create a role named `api-access` with the minimum configuration required for the Backup and Recovery capability to access the storage system.

```
useradmin role add api-access -a api-*,login-http-admin,cli-ifconfig
```

2. Use the following command to create a group named `api-group` which contains the `api-access` role.

```
useradmin group add api-group -r api-access
```

3. Use the following command to create a user named `smvi-user` as a member of the `api-group` group.

```
useradmin user add smvi-user -g api-group
```

4. To set the user's password, run the `passwd` command as root.

The storage system prompts you for the account name that you want to change, followed by the new password for this account.

## How Backup and Recovery discovers vFiler units

Backup and Recovery supports vFiler tunneling on the physical storage system that contains the vFiler units. In certain environments, Backup and Recovery can access only the physical storage system, not the vFiler unit, for communication on a storage network.

With vFiler tunneling enabled, you can use Backup and Recovery to create Snapshot copies. When you add a physical storage system to Backup and Recovery, it automatically discovers all the associated vFiler units. You do not need to manually add the vFiler units.

### Related tasks

[Enabling discovery and management of vFiler units](#) on page 49

[Enabling discovery and management of vFiler units on private networks](#) on page 49

## Managing backups

You can use Backup and Recovery to perform on-demand backups or schedule automated backups of individual virtual machines, datastores, or a datacenter.

After you add a new backup job using the Backup wizard, you can set a backup job schedule and specify a retention policy for the backup copy. You can also change the schedule and retention policy and delete or suspend a backup job.

### Considerations for adding a backup job

Before you use the Backup wizard to add, schedule, and automate backups, you should be aware of the information that you can specify to ensure that the backup schedule, the backup retention policy, and the alert notifications of backup activity for your job perform as expected.

When you add a new backup job, you can specify whether you want to initiate a SnapMirror update on the virtual entities that are to be backed up or create VMware snapshots for every backup. If you select virtual entities that span multiple datastores, you can specify one or more of the spanning datastores to be excluded from the backup.

If you want to run a backup script that is installed on the server with this job, you can choose the scripts that you want to use. If you create a pre-backup or post-backup script that results in an output file when the script is run, the output file is saved to the same directory where you initially installed the pre-backup or post-backup script. You can specify the hourly, daily, weekly, or monthly schedule that you want applied to your backup job, or you can add a backup job without attaching a schedule to the backup.

You can specify the maximum number of days and maximum number of backups and the email alerts for this backup job. You must first set the SMTP server and the destination email addresses to receive an alert notification when an alarm is triggered or the system status changes. You can enable or disable alarms and specify how often you receive email alerts when an error or warning occurs. The **Notify on** options in the Backup wizard include the following:

<b>Always</b>	An alert notification is always sent.
<b>Errors or Warnings</b>	A backup failure or a warning such as a SnapMirror data transfer did not complete triggers an alert notification. The email that is sent for a warning does not indicate what triggered it, even if the backup job is successful.
<b>Errors</b>	A backup failure or partial failure triggers an alert notification.
<b>Never</b>	An alert notification is never sent.

## Backing up a virtual machine

You can use the Backup wizard to add a new backup job for a virtual machine. You can schedule and automate your backup jobs, specify a retention policy, and create an automated policy for email alerts.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

The vSphere Client must be connected to a vCenter Server to create backup copies.

### About this task

Backup and Recovery uses the storage credentials that you specify in Monitoring and Host Configuration. You must confirm that all previous storage credentials from Backup and Recovery now exist in the list of VSC storage credentials. You can use the `smvi storagesystem list` command to verify the list of storage systems.

During backup or recovery of a virtual machine, Backup and Recovery does not allow other backup or recovery operations on that virtual machine to start. Backup and Recovery delays any other backup or recovery operations until the current backup or recovery is finished.

### Steps

1. Select a virtual machine in the **Inventory** panel, and then right-click the node and select **NetApp > Backup and Recovery > Schedule a Backup**.

The Backup wizard appears.

2. Type a backup job name and description and click **Next**.
3. Select one of the following options:
  - If you want to start a SnapMirror update on the selected entities concurrent with every backup, select **Initiate SnapMirror update**.  
For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.  
The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.
  - If you want to create a VMware snapshot for each backup copy, select **Perform VMware consistency snapshot**.
  - If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.
4. Select the virtual entities available for this backup job and click **Next**.
5. Select one or more backup scripts and click **Next**.

If an error message appears, indicating that at least one of the selected scripts has been deleted, you can save the backup job without any script in the selected scripts list, thereby removing the deleted script from the job. Otherwise, the backup job continues to use the deleted script.

6. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.
7. Select **One time only** and click **Delete this job after backup is created** if you do not want to retain this backup job.
8. Type the user name and password for the vCenter Server and click **Next**.
9. Select the maximum number of days or the maximum number of backup jobs and click **Next**.
10. Select the frequency for receiving email alerts and click **Next**.

You can add multiple email addresses by using semicolons to separate each email address.

11. Optional: Click **Send test email** to verify that the outgoing email server where the alert notifications are to be sent is working correctly.
12. Review the summary page and click **Finish**.
13. Optional: Select the **Run Job Now** check box to immediately run the job.

## Backing up a datastore or datacenter

You can use the Backup wizard to add a new backup job for an entire datacenter, a datastore, a particular set of datastores, or a particular set of virtual machines. You can schedule and automate your backups, specify a retention policy, and create an automated policy for email alerts.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

The vSphere Client must be connected to vCenter Server to create backup copies.

### About this task

Backup and Recovery uses the storage credentials that you specify in Monitoring and Host Configuration. You must confirm that all previous storage credentials from Backup and Recovery now exist in the list of VSC storage credentials. You can use the `smvi storagesystem list` command to verify the list of storage systems.

During backup or recovery of a virtual machine, Backup and Recovery does not allow other backup or recovery jobs or for that virtual machine to start. Backup and Recovery delays any other backup or recovery operations until a current backup or recovery operation is finished.

## Steps

1. Select a datacenter or datastore in the **Inventory** panel, and then right-click the node and select **NetApp > Backup and Recovery > Schedule a Backup**.

The Backup wizard appears.

2. Type a backup job name and description and click **Next**.
3. Select one of the following options:
  - If you want to start a SnapMirror update on the selected entities concurrent with every backup, select **Initiate SnapMirror update**.  
For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.  
The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.
  - If you want to create a VMware snapshot for each backup, select **Perform VMware consistency snapshot**.
  - If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.
4. Select the virtual entities available for this backup job and click **Next**.
5. Select one or more backup scripts and click **Next**.  
If an error message appears, indicating that at least one of the selected scripts has been deleted, you can save the backup job without any script in the selected scripts list, thereby removing the deleted script from the job. Otherwise, the backup job continues to use the deleted script.
6. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.
7. Select **One time only** and click **Delete this job after backup is created** if you do not want to retain this backup job.
8. Type the user name and password for the vCenter Server and click **Next**.
9. Select the maximum number of days or the maximum number of backup jobs to retain and click **Next**.
10. Select the frequency for receiving email alerts and click **Next**.  
You can add multiple email addresses by using semicolons to separate each email address.
11. Optional: Click **Send test email** to verify that the outgoing email server where the alert notifications are to be sent is working correctly.
12. Review the summary page and click **Finish**.
13. Optional: Select the **Run Job Now** check box to immediately run the job.

## Starting a one-time backup job

You can perform a one-time backup job on a virtual machine or datastore. This type of backup is useful if you do not want to schedule regular backups for a particular virtual machine or datastore.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then right-click the node and select **NetApp > Backup and Recovery > Backup Now**.

The Backup Now window appears.

2. Clear the **Automatically name backup** check box if you do not want the backup name to be automatically generated.
3. Type a backup name in the **Name** text box and click **Backup Now**.
4. Select one of the following options:
  - If you want to start a SnapMirror update on the selected entities concurrent with every backup job, select **Initiate SnapMirror update**.  
For this option to execute successfully, the selected entities must reside in volumes that are already completely configured as SnapMirror source volumes.  
The SnapManager for Virtual Infrastructure server should be able to resolve the host name and IP address of the source and destination storage systems in the `snapmirror.conf` file.
  - If you want to create a VMware snapshot for each backup copy, select **Perform VMware consistency snapshot**.
  - If you want to include independent disks from datastores that contain temporary data, select **Include independent disks**.

## Editing a backup job

You can use the Job Properties dialog box to modify the name and description, the datastores and virtual machines that are assigned, the backup scripts, the user credentials, the schedule, the retention policy, or the email alerts for an existing backup job.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the backup job whose properties you want to modify.
4. Click **Edit**, and then click the appropriate tab for the properties that you want to modify for this backup job.
5. Modify backup job properties as necessary and click **OK** to change the properties.

## Deleting a scheduled backup job

You can select and delete one or more backup jobs from the list of scheduled jobs, but you cannot delete any backup jobs that are running.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel and click the **NetApp** tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select one or more backup jobs that you want to delete.

In the Entities page, note the existing datastore and virtual machines currently associated with the selected backup job. When you delete the selected backup job, its backup operations are no longer performed on these entities.

4. Click **Delete**, and then click **Yes** at the confirmation prompt.

## Suspending an active backup job

You can suspend the scheduled operations of an active backup job without deleting the job. This gives you the ability to temporarily halt backup jobs in case of planned maintenance or during periods of high activity, or for other reasons.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.

2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the active backup job that you want to suspend.

In the Entities page, note the existing datastore and virtual machines currently associated with the selected backup job. Backup operations are no longer carried out on these entities when you suspend the selected backup job.

4. Right-click the selected backup job, and then select **Suspend**.
5. Click **Yes** to confirm when prompted.

## Resuming a suspended backup job

You can resume and run a suspended backup job.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Click **Backup** under Backup and Recovery in the navigation pane.
3. Select the suspended backup job that you want to resume.
4. Right-click the selected backup job, and then select **Resume**.
5. Click **Yes** to confirm when prompted.

## Restoring data from backups

You can use the Backup and Recovery capability to restore your virtual machines and datastores from backups that contain them. VSC for VMware vSphere supports restoring from local backups and from backups that contain VMware snapshots.

The Backup and Recovery capability uses FlexClone technology to restore data on systems running Data ONTAP 7.x and, in conjunction with `snap restore hidden flag`, Data ONTAP 8.1.

For VMware Virtual Machine File System (VMFS) datastores, if FlexClone is not licensed, the Backup and Recovery capability uses LUN clone technology on systems running Data ONTAP 7.x and 8.1.x. For Network File System (NFS) datastores in cluster configurations, the Backup and Recovery capability uses `snap restore hidden flag`, while VMFS datastores use SIS clone technology.

## Where to restore a backup

VSC for VMware vSphere enables you to select the original (default), the current, or an alternate location as the restore destination by using the Restore wizard.

- **Original location (default)**  
The backup copy of an entire datastore is restored to the original location. You set this location by choosing an entire datastore. This type of restore operation is referred to as an “in-place restore” operation.
- **Current location**  
The backup copy of a single virtual machine, or a copy of a virtual machine's disk files, is restored to a new, current location after it has been migrated from its original location. You set this location by choosing an entire virtual machine.
- **Alternate location**  
The backup copy of the virtual machine disk files is restored to an alternate location. This type of restore operation is referred to as an “out-of-place restore” operation. You set this location by selecting the destination location as a different datastore from the one on which the virtual disk currently resides by setting the **Particular virtual disks** option.

## Restore operations using data that was backed up with failed VMware consistency snapshots

Even if a VMware consistency snapshot for a virtual machine fails, the virtual machine is nevertheless backed up. You can view the backup copy in the Restore panel and use it for restore operations.

During backup operations, a VMware quiesced snapshot is created by default. When creating a snapshot, the virtual machine pauses all running processes on the guest operating system so that file system contents are in a known consistent state when the snapshot is taken. Despite the VMware snapshot failure, the virtual machine is still included in the Snapshot copy. You can restore a virtual machine without manually removing it from a Snapshot copy.

You can restore a potentially functional virtual machine despite failing the VMware process of creating a quiesced snapshot, and view its status during quiesced backup operations in the backed-up entities section of the Restore panel. The Quiesced column can display the following values:

- Yes, if a VMware snapshot operation is successful and the guest operating system was quiesced.
- No, if the VMware snapshot operation failed because the guest operating system could not be quiesced.
- Not Applicable, for entities that are not virtual machines.

## Restoring data from backup copies

You can restore a datastore, virtual machine, or its disk files to their original location or an alternate location. From the Restore panel, you can sort the backup listings by name, date, or other search criteria to help you find your backup copies.

From the Restore panel, you can do the following:

- Restore a datastore, virtual machine, or its disk files/VMDKs from a local backup copy to the original location
- Restore VMDKs from a local backup to a different location
- Restore from a backup copy that has a VMware snapshot

**Note:** Because of the time it takes to update the backup listings in the Restore panel, you can complete this task much faster when you right-click the name of a vSphere object, such as a virtual machine or a datastore, from the vCenter Inventory panel.

## Searching for backup copies

You can locate a specific backup copy by searching the backup table. After you locate a backup copy, you can then restore using it, view its status, or delete it.

### Steps

1. From the navigation pane, click **Restore** to display the available backup copies.
2. In the **Search** text box, type one or more search terms and press **Enter**.

Available criteria for search are the backup ID, VMware snapshot name, mount state, or resource name.

A filtered list displays in the Restore panel.

## Restoring a datastore

You can use Backup and Recovery to restore a datastore. By doing so, you overwrite the existing content with the content of the backup copy that you select.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. In the **Restore** panel, select a backup copy of the datastore.
2. Click **Restore**.

Restoring a datastore powers off all virtual machines.

3. In the **Restore dialog box**, click **Restore**.

## Restoring a virtual machine or its VMDKs

You can restore an entire virtual machine or particular virtual disks of a given virtual machine. By doing so, you overwrite the existing content with the backup copy you select.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

You must have already backed up a virtual machine or its VMDKs before you can restore it.

### About this task

If you are restoring a virtual machine to a second ESX server, the virtual machine is unregistered from the first ESX server and the restored virtual machine is placed on the second ESX server. Both ESX servers must share the same datastore.

### Steps

1. Select a virtual machine in the **Inventory** panel, and then right-click the node and select the **NetApp > Backup and Recovery > Restore** from the context menu.

The Restore wizard opens and lists all backup copies that include the virtual machine.

2. From the Backups table, select a backup copy of the virtual machine and click **Next** to display the **Virtual Machine Component Selection** page.
3. Select one of the following recovery options:

Option	Description
<b>The entire virtual machine</b>	Restores the contents of your virtual machine from a Snapshot copy. The <b>Start VM after restore</b> check box is enabled if you select this option and the virtual machine is registered.
<b>Particular virtual disks</b>	Restores the contents of the VMDKs on a virtual machine. This option is enabled when you clear the <b>The entire virtual machine</b> option. <ul style="list-style-type: none"> <li>• If a backup copy of a virtual machine exists in the Backups table, you cannot select an ESX host. You can only restore the VM to its parent ESX host.</li> <li>• If a backup copy of a virtual machine does not exist in the Backups table, then you can select an ESX host. You can restore the VM to any ESX host.</li> </ul>

4. In the **ESX Host Name field**, select the name of the ESX host.

This option is available if you want to restore VMDKs or if the virtual machine is on a VMFS datastore.

5. Click **Next** to display the **Summary** page, which displays information about the restore operation that is about to be performed.

6. From this page, verify the settings and click **Restore** to begin the restore process.

The virtual machine or its VMDKs are restored to the state that the Snapshot copy captured.

### After you finish

After you finish using the Restore wizard and start the restore operation, you can track the progress of the restore operation from the Task tab of the Status panel to monitor the job for possible errors.

## Mounting a backup copy

Backup and Recovery enables you to mount an existing backup copy onto an ESX server to verify the contents of the backup prior to completing a restore operation or to restore a virtual machine to an alternate location.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

All of the datastores and the virtual machines contained in the backup are mounted to the ESX server that you specify. After it is mounted, the backup copy appears in the vSphere Client.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel and click the **NetApp** tab.
2. Click **Restore** under **Backup and Recovery** in the navigation pane.
3. In the **Restore** panel, select the name of an unmounted backup copy that you want to mount.
4. Click **Mount**.
5. In the **Mount Backup** dialog box, select the name of the ESX server to which you want to mount the backup.

You can mount only one backup each time, and you cannot mount a backup that is already mounted.

6. Click **Mount**.

## Unmounting a backup copy

After you are done verifying the contents of a mounted backup copy, you can unmount it from the ESX server. When you unmount a backup, all of the datastores in that backup copy are unmounted and are no longer visible from the ESX server.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### Steps

1. Select a virtual machine or datastore in the **Inventory** panel and click the **NetApp** tab.
2. Click **Restore** under **Backup and Recovery** in the navigation pane.
3. In the **Restore** panel, select the name of a mounted backup that you want to unmount.
4. Click **Unmount**.
5. At the confirmation prompt, click **Yes**.

## Single file restore

The Backup and Recovery capability and Restore Agent of Virtual Storage Console provide tools that help you restore backups of your virtual machine disks at the file level.

How these applications coordinate the restore process depends on the network connection between the systems that run Virtual Storage Console and Restore Agent and the level of access that the user has to the destination virtual machine.

## How Virtual Storage Console detects network connectivity

Virtual Storage Console assumes that the port group on the VMware vSphere virtual network has a direct connection between the server and the destination virtual machine.

### When to manage port group settings

Virtual Storage Console collects information about the port group after it authenticates the user credentials of the vCenter Server. Therefore, to communicate with a Restore Agent instance that is installed on the destination virtual machine, you must check the port group settings on the Single File Restore panel whenever there is a change to the port group settings on the virtual network.

### What happens if you do not update the port group

Virtual Storage Console assumes that all port groups can communicate with the SnapManager for Virtual Infrastructure server. If the administrator does not change the virtual network configuration after registering the Restore Agent, then you do not need to manage these settings.

If there are changes to the network configuration after you create a restore session and after you install Restore Agent on the guest virtual machine, in addition to checking the port groups settings, you need to create a new restore session to send the new configuration file to the user to use it with Restore Agent.

### The difference between limited and direct connectivity

Users have a varying degree of access to a backup depending on whether there is a connection between the physical and virtual network.

In virtual environments in which the virtual machines are in a separate network from the SnapManager for Virtual Infrastructure server, Restore Agent cannot communicate with Virtual Storage Console. As a result, the user has an *offline* connection and cannot access the list of backups that were made by Virtual Storage Console. If users with limited connectivity want access to these backups, those users must have an administrator create a restore session and send a configuration file to them in an e-mail message.

When there is a direct or *online* connection between the SnapManager for Virtual Infrastructure server and the destination virtual machine, the user can search for mounted backups without administrative help.

By default, all virtual machines have direct connectivity, unless you change the port group setting.

### Types of file restore sessions

Backup and Recovery automates the process of restoring single files based on the relationship between the source virtual machine (which was backed up) and the destination virtual machine. Backup and Recovery supports three types of restore sessions.

<b>Self-service</b>	The Backup and Recovery administrator creates a restore session using Backup and Recovery capability. Users can then install Restore Agent on the destination virtual machine, browse the mounted backups on a guest virtual machine, and recover the individual disk file.
<b>Administrator-assisted</b>	This type of file recovery is basically the same as self-service, except that the Backup and Recovery administrator runs Restore Agent and copies the recovered files to a shared location that the user has access to.
<b>Limited self-service</b>	The Backup and Recovery administrator finds the backup copy within a user-specified range of backups and attaches the backed-up disks to the destination virtual machine. The user can then run Restore Agent on a destination virtual machine, browse the mounted backups, and recover the individual disk file.

## Manually creating a .sfr file for the Restore Agent

In some environments, security restrictions prevent the email delivery of the single file restore (.sfr) file to users. You can use Backup and Recovery to perform a flexible clone copy of a datastore, mount the datastore to an ESX host, and add the virtual disk to a guest machine as a new virtual disk.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

### About this task

After you create a GuestRestoreClientSession.sfr file, it is delivered to users via email. You can manually create this file by extracting a portion of the Backup and Recovery log file.

### Steps

1. After you create a new single file restore session, navigate to the Backup and Recovery installation directory that contains a server.log file such as `C:\program files\Netapp\SMVI\server\log`.
2. Copy the file to a location where you can access it, such as a shared network drive, and then open the file using a text editor.
3. In the open file, copy all of the text from the string `<standalone>` that corresponds to the date and time that you created the single file restore job session, such as `<?xml version=1.0 encoding=UTF-8 standalone=no?>`.
4. Create a new text file named `<filename.sfr>`, and paste the copied text into it.

You can use this file to load the configuration for the Restore Agent onto your guest machine.

## General configuration parameters for single file restore

Using the General tab of the Setup panel, you must set some general configuration parameters before you can use the single file restore feature.

### Setting session defaults

When you set a session default from the Setup panel in the Backup and Recovery capability, any restore session that you subsequently create uses the default settings that you specified.

### Steps

1. In **Setup** panel, click **Single File Restore**.
2. Click **Edit** and then change the following session defaults:
  - Type the location of the download directory for Restore Agent.

- Select the number of hours or days before the restore session expires.  
The **Default Session Expiration Time** field displays the time in hours.

3. Click **OK**.

## Changing the network connection for a port group

You can change the type of restore session for a port group. By doing so, you change the network connection for all the virtual machines on the same subnet.

### Steps

1. In the navigation panel, click **Set up**.
2. In the **Setup** panel, click **Single File Restore**.
3. Select **Admin Assisted** for the port group.

### Result

Backup and Recovery updates the information for the virtual machines that are reported for the port group and in the list of restore sessions in the Single File Restore panel.

## Setting the SnapManager for Virtual Infrastructure server address

The restore session configuration file includes the SnapManager for Virtual Infrastructure server IP address and fully qualified domain name, but you might want to change the address or name with a specific one when you have a multi-homed server that has multiple IP addresses.

### Step

1. Add the SnapManager for Virtual Infrastructure server IP address to the `/etc/smvi.override` file.  
Virtual Storage Console will use this value, which can be an IP address or host name, instead of the one previously configured.

## Self-service example workflow

To help you understand how a self-service restore session works, imagine that you are a VMware administrator and you need to restore from a backup copy of a virtual machine a critical data file that has become corrupted.

A Virtual Storage Console administrator responds to your ticket and creates a single file restore session. The SnapManager for Virtual Infrastructure server sends you an e-mail message that provides information you need to access his computer, a link to download Restore Agent application software, and a configuration file.

The configuration file contains data that allows your workstation to communicate with the SnapManager for Virtual Infrastructure server. Rather than the administrator doing the connection work, you can select which disks to connect to and find a backup.

## Create a self-service restore session

You must create a restore session that you or a VMware administrator can use in order to access a guest operating system and find a backup copy to use.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

You must have available the information necessary to use the Add Single File Restore Session wizard effectively:

- The name or IP address of the virtual machine: the backup source and destination
- The email message recipient
- The mount expiration time: 240 hours (or 10 days)

### Steps

1. Click the **Single File Restore** link.
2. Click **Add** to start the **Add Single File Restore Session** wizard.
3. Complete the wizard, using the following values:
  - Source VM Name: **VM-XP-EXAMPLE**
  - Destination VM Name: **VM-XP-EXAMPLE**
  - Recipient Email Address(s): **JenniferKWhite@example.com**
  - Sender Email Address: **JerolimOlszewsk@example.com**
  - Mount Expiration: **10 days**
  - File Restore Access Type: **Self-service**

**Note:** In this example, the source virtual machine and the destination virtual machine are the same.

4. Confirm the details of the restore session, and then click **Finish** to complete the wizard.

### Result

Your new restore session is listed in the Single File Restore panel.

### After you finish

You receive an email notification that contains a link to download Restore Agent, and you install the software.

## Installing Restore Agent

After you create a restore session, you receive an email message that provides a link to the Restore Agent installation file and has a restore session configuration attached as an .sfr file. Before you can restore single files on a guest operating system, you must install Restore Agent.

### Before you begin

The system upon which you are going to install Restore Agent must have the following software installed:

- Microsoft Management Console 3.0
- Microsoft .NET Framework 3.5 Service Pack 1

In addition, to enable the single file restore feature for an NFS datastore and perform a mount operation, the storage system must have an installed FlexClone license.

For the most current software requirements, see the Interoperability Matrix at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability).

### Steps

1. Click the link in the email message to download and start the installation process.
2. Follow the displayed instructions.

## Load the configuration file

You must load the configuration file that you received in the e-mail message, along with the link to download Restore Agent, onto the destination virtual machine to access the mounted backup virtual machine disk files.

### About this task

You can only load the configuration file after you have already installed the Restore Agent on the destination virtual machine.

### Steps

1. Double-click the Restore Agent shortcut icon on your desktop.
2. In the resulting **Load Configuration** window, look for the configuration (.sfr) file.
3. Click **OK**.
4. Choose the virtual machine source to restore from.

## Recovering single files from a virtual machine

After you have been given access to files on the destination virtual machine, you can use the Restore Agent window to view the backed up files created by Virtual Storage Console and find the ones that you need to restore from.

### Before you begin

When you are restoring data, Restore Agent displays the first-available disk drive letter. If your guest operating system is Windows XP, 7, or 8, you must go to the Windows Disk Management snap-in application to manually assign the disk drive letter to other partitioned space on the disk. For more information about using the disk management utility, select **Help** from Disk Management.

### Steps

1. From the **Restore Agent** window, select a file from the **Disk tab**.

Alternatively, click the **Backup** tab to find the Snapshot copy by name.

2. Right-click the name of the backup copy and select **Mount**.

The **Backup** tab shows the drive letter for each backup copy only if the source virtual machine and destination virtual machine are the same.

The contents of the backup copy are written to the new location.

### After you finish

After finishing this task, you should clear the configuration cache.

## Clear the configuration

After restoring your data from a backed up virtual machine, you should clear the configuration so you can upload another one later.

### Steps

1. From the Action pane in the **Restore Agent window**, click **Clear Configuration**.
2. In the resulting window, click **OK**.

Backup and restore metadata will be removed from Restore Agent and Virtual Storage Console.

## Limited self-service example workflow

To help you understand how a limited self-service restore session works, imagine that you are a VMware administrator and you must restore a failed disk on your virtual machine with the help of a Backup and Recovery administrator.

It is Thursday afternoon and one of the files on the virtual machine is corrupt. You must restore the disk over the weekend. You complete a ticket and request access to another workstation so that you can recover the data from a backup copy.

A Virtual Storage Console administrator responds to your ticket and makes the system available from Saturday morning for 48 hours. The SnapManager for Virtual Infrastructure server sends you an email message that provides information you need to log in as an administrator to the Virtual Storage Console computer, a link to a site from which you can download Restore Agent application software, and a configuration file.

The configuration file contains attached disks so that your operating system can see which disks are attached and assign drive letters that enable your workstation to connect to that of the Virtual Storage Console administrator.

## Create a limited self-service restore session

Before you can restore a failed disk on your virtual machine, you must create a restore session that a user can use to access a guest operating system for restoring a virtual machine disk file.

### Before you begin

Unless you log in as an administrator, you must have the appropriate RBAC privileges correctly assigned to complete this task successfully. For more information, see [Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 31.

Before creating a restore session, you must gather the information necessary to complete the Add Single File Restore Session wizard:

- The name or IP address of the virtual machine: the backup source and destination
- The email message recipient and sender
- The mount expiration time: three hours

### About this task

In this example, the source virtual machine and the destination virtual machine are the same.

### Steps

1. Select the **Single File Restore** panel.
2. From the **Single File Restore** panel, click **Add** to start the **Add Single File Restore Session** wizard.
3. Complete the wizard, using the following values:

- Source VM Name: **VM-WXP-EXAMPLE**
- Destination VM Name: **VM-WXP-EXAMPLE**
- To Email Address(s): **JordanEKanode@example.com**
- From Email Address: **LinYaoHuang@example.com**
- Mount Expiration: **3 days**
- File Restore Access Type: **Limited Self-Service**

4. Confirm the details of the restore session and then click **Finish** to complete the wizard.

### Result

Your new restore session is listed in the Single File Restore panel.

### After you finish

You receive an email notification that contains a link to download Restore Agent, and you install the software.

## Installing Restore Agent

After you create a restore session, you receive an email message that provides a link to the Restore Agent installation file and has a restore session configuration attached as an .sfr file. Before you can restore single files on a guest operating system, you must install Restore Agent.

### Before you begin

The system upon which you are going to install Restore Agent must have the following software installed:

- Microsoft Management Console 3.0
- Microsoft .NET Framework 3.5 Service Pack 1

In addition, to enable the single file restore feature for an NFS datastore and perform a mount operation, the storage system must have an installed FlexClone license.

For the most current software requirements, see the Interoperability Matrix at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability).

### Steps

1. Click the link in the email message to download and start the installation process.
2. Follow the displayed instructions.

## Recovering single files from a virtual machine

After you have been given access to files on the destination virtual machine, you can use the Restore Agent window to view the backed up files created by Virtual Storage Console and find the ones that you need to restore from.

### Before you begin

When you are restoring data, Restore Agent displays the first-available disk drive letter. If your guest operating system is Windows XP, 7, or 8, you must go to the Windows Disk Management snap-in application to manually assign the disk drive letter to other partitioned space on the disk. For more information about using the disk management utility, select **Help** from Disk Management.

### Steps

1. From the **Restore Agent** window, select a file from the **Disk tab**.

Alternatively, click the **Backup** tab to find the Snapshot copy by name.

2. Right-click the name of the backup copy and select **Mount**.

The **Backup** tab shows the drive letter for each backup copy only if the source virtual machine and destination virtual machine are the same.

The contents of the backup copy are written to the new location.

### After you finish

After finishing this task, you should clear the configuration cache.

## Clear the configuration

After restoring your data from a backed up virtual machine, you should clear the configuration so you can upload another one later.

### Steps

1. From the Action pane in the **Restore Agent window**, click **Clear Configuration**.
2. In the resulting window, click **OK**.

Backup and restore metadata will be removed from Restore Agent and Virtual Storage Console.

## VSC CLI commands

The VSC for VMware vSphere command-line interface (CLI), which is labeled "VSC CLI" on your Windows desktop, provides you the benefits of a command-based view of the user interface. You can

use this CLI to perform specific Backup and Recovery tasks, such as creating or deleting a backup of a virtual machine or datastore, as well as mounting a backup.

You should keep in mind the following information about the commands that you see in the interface:

- Virtual Storage Console commands are case-sensitive.
- There are no privilege levels; any user with a valid user name and password can run all commands.

For some commands, the following two parameters control the amount of output displayed:

**verbose** This optional parameter provides detailed output when displaying information.

**quiet** This optional parameter stops any output from displaying.

**Note:** Even with the quiet parameter specified, failed commands still display their failure messages.

## Launching the VSC CLI

You can use either of two methods to launch the VS for VMware vSphere command-line interface (CLI), which is labeled "VSC CLI" on your Windows desktop. The first time you launch the VSC CLI, the application uses your Windows user credentials to grant you server access. Subsequent launches use stored credentials, speeding your access to the server.

### About this task

When you issue your first CLI command, the CLI prompts you for your password and then runs the command. If the command succeeds, the CLI caches your user credentials and stores the information locally in an encrypted format.

### Step

1. Double-click the **VSC CLI** icon or navigate to **Start > All Programs > NetApp > VSC CLI**.

An alternative method to using your Windows user credentials is to use the `smvi servercredential set` command to create custom user credentials that allow you to log in to the server.

## smvi backup create

The `smvi backup create` command creates a backup copy of a virtual machine or datastore. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi backup create -id {name | id} [name | id ...] [-backup-name {backup name}] [-server {server name}] [-include-independent] -exclude-datastores {name | id} [name | id ...] [-scripts {script name}]
```

```
[-vmware-snapshot] [-no-vmware-snapshot] [-update-mirror] [-quiet]
[-verbose] [-user] [-help]
```

## Parameters

**[-id {*name* | *id*}] [*name* | *id* ...]**

This mandatory parameter specifies the name or identification of the datastore or virtual machine that you are backing up. You can specify names or identifications of multiple datastores or virtual machines.

**[-backup-name {*backup name*}]**

This optional parameter specifies a backup copy name. After adding the flag, add a name for the backup copy. If you specify no name with this flag, the command fails. If you specify a name that contains only spaces, VSC for VMware vSphere automatically generates a name. If you specify a name that contains both spaces and other characters, VSC for VMware vSphere removes all leading and trailing spaces from the name.

**[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-include-independent]**

This optional parameter specifies that datastores containing only independent disks for a virtual machine are included in the backup.

**[-exclude-datastores {*name* | *id*}] [*name* | *id* ...]**

This optional parameter specifies the name or identification of the datastores or virtual machines to be excluded from the backup.

**[-scripts {*script name*}]**

This optional parameter specifies the name of the scripts to run with this backup.

**[-vmware-snapshot]**

This optional parameter creates VMware snapshots of virtual machines during a backup. If you specify this parameter along with [-no-vmware-snapshot] or if you do not specify either one of the parameters, it prevents the creation of any VMware snapshots.

**[-no-vmware-snapshot]**

This optional parameter prevents the creation of VMware snapshots of virtual machines during a backup.

**[-update-mirror]**

This optional parameter initiates a SnapMirror image on the secondary storage.

**[-quiet]**

This optional parameter stops any output from displaying.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-help]**

This optional parameter displays help for this command.

**Example: Creating a backup copy from a virtual machine**

The following example creates a backup copy from a virtual machine named testInd, without specifying a backup name:

```
smvi backup create -id testInd

[18:19] Starting backup request
[18:19] Backing up datastore(s) ([DS_vee_cmode_03-04
(4fd6aad6-579ff358-676a-e41 f13ba10f6), nfs_cmode_test1 (netfs://
10.60.189.79///vol_nfs_test1)])
[18:19] Backing up the following virtual machine(s) ([testInd])
[18:20] Creating storage snapshots for all datastores/virtual
machines that are being backed up.
[18:20] Storing logs for backup_43e100cdf8182d1ee79bae431f1c608f in
file .\repository\logs\unscheduled
\backup_backup_43e100cdf8182d1ee79bae431f1c608f.xml
[18:20] Backup backup_43e100cdf8182d1ee79bae431f1c608f of
datastores/virtual machines is complete.
SMVICLI-0100: Command completed successfully
```

## smvi backup delete

The `smvi backup delete` command removes a virtual machine or datastore backup copy. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi backup delete    -backup-name {backup name}    [-server {server
name}]    [-quiet]    [-verbose]    [-noprompt]    [-user]    [-help]
```

## Description

When you delete the most recent backup associated with a backup job, then the Last Run Status value displayed for that backup job in the Schedule Backup Jobs window is that of the most recent remaining undeleted backup copy associated with the backup job.

## Parameters

**`[-backup-name {backup name}]`**

This mandatory parameter specifies the backup copy you want to delete. After adding the flag, add the name of the backup copy.

**`[-server {server name}]`**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**`[-quiet]`**

This optional parameter stops any output from displaying.

**`[-verbose]`**

This optional parameter provides detailed output when displaying information.

**`[-noprompt]`**

This optional parameter disables the default prompt that asks for confirmation when deleting a backup.

**`[-user]`**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**`[-help]`**

This optional parameter displays help for this command.

### Example: Deleting a backup copy

The following example deletes a backup copy named new-one:

```
smvi backup delete -backup-name new-one

Are you sure you want to proceed and remove backup named 'new-one'?
[yes|NO] y
[15:15] Removed backup with name "new-one"
SMVICLI-0100: Command completed successfully
```

## smvi backup list

The `smvi backup list` command displays information, such as the file path on a storage system to the Snapshot copy, about all of the created and saved backups within a virtual machine or datastore. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi backup list [-id {name | id} [name | id ...]] [-mounted] [-failed] [-recent] [-with-vmware-snapshot] [-sfr-mounted] [-server {server name}] [-user] [-help]
```

### Parameters

**[-id {*name* | *id*} [*name* | *id* ...]]**

This mandatory parameter specifies the name or identification of the datastores or virtual machines that you want to list.

**[-mounted]**

This optional parameter lists all mounted backups.

**[-failed]**

This optional parameter lists all failed backups. The default list is only successful backups.

**[-recent]**

This optional parameter lists the most recent backup.

**[-with-vmware-snapshot]**

This optional parameter lists the backups that were taken with a VMware snapshot.

**[-sfr-mounted]**

This optional parameter lists the backups that were mounted for SFR.

**[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-help]**

This optional parameter displays help for this command.

**Example: Listing backups in a datastore**

The following example lists all of the backups within a datastore named data-store1:

```
smvi backup list -id data-store01

Id Name Date Entities Mounted VMware Snapshot Snapshot Name
-----
-----
      backup_sch_1_20090122233100 Jan 22, 2009 23:31 vmfs_vml No
10.72.248.38:/vol/kas1_102_iscsi:smvi_backup_sch_1
20090122233100_36d2d99a-9ee0-4841-80c0-846698463e78_kas_sw_iscsi_dsl
```

**smvi backup mount**

The `smvi backup mount` command mounts a backup to verify its contents.

**Syntax**

```
smvi backup mount      -backup-name {backup name}      -esx-server {esx
server name}      [-server {server name}]      [-quiet]      [-verbose]      [-
user]      [-help]
```

**Privilege level**

**Note:** To mount a VMFS datastore backup, the supplied ESX server must have SAN or iSAN access to the storage system, including required FC zoning or iSCSI discovery. To mount an NFS datastore backup, the supplied ESX server must be in the NFS export list of the original datastore.

**Parameters**

**[-backup-name {*backup name*}]**

This mandatory parameter specifies the backup you want to mount. After adding the flag, add the name of the backup.

**[-esx-server {*esx server name* | *IP address*}]**

This mandatory parameter specifies the name or IP address of the ESX server. This information describes where the backup resides on an ESX server.

**Note:** The server name is the name of the ESX server as viewed through the vSphere Client. This name might differ from the ESX server's host name or IP address.

**[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-quiet]**

This optional parameter stops any output from displaying.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-help]**

This optional parameter displays help for this command.

### Example: Mounting a backup

The following example mounts a backup named vmfs2\_vm1 on an ESX server with the IP address of 123.12.1.23:

```
smvi backup mount -backup-name vmfs2_vm1 -esx-server 123.12.1.23
[12:12] Starting mount request
SMVICLI-0100: Command completed successfully
```

## smvi backup rename

The `smvi backup rename` command changes the name of a backup. Changing the name of a backup also changes the name on the corresponding storage Snapshot copy on the associated NetApp storage system. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi backup rename name -new-backup-name {new name} -backup-name {backup name} [-server {server name}] [-user {user}] [-help]
```

### Parameters

**[-backup-name *{backup name}*]**

This mandatory parameter specifies the backup you want to rename. After adding the flag, add the name of the backup.

**[-new-backup-name {*new name*}]**

This mandatory parameter specifies the new name of the backup. After adding the flag, add a new name for the backup.

**[-server-name {*server name*}]**

This optional parameter specifies the name of the server to send the command to.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-help]**

This optional parameter displays help for this command.

**Example: Renaming a backup**

The following example renames a backup named vmfs2-vm1 to volume-2:

```
smvi backup rename -backup-name vmfs2-vm1 -new-backup-name volume-2
[15:52] Backup "vmfs2-vm1" has been renamed to "volume-2"
SMVICLI-0100: Command completed successfully
```

**smvi backup restore**

The `smvi backup restore` command enables you to restore a virtual machine or datastore from a backup copy. You can also perform this operation using the VSC for VMware vSphere user interface.

**Syntax**

```
smvi backup restore -id {name | id} [-esx-server {esx server
name] [-backup-name {backup name}] [-vmdk {hard disk name}] [-
server {server name}] [-restart-vm] [-quiet] [-verbose] [-
noprompt] [-user] [-help]
```

**Parameters****[-id {*name* | *id*}]**

This mandatory parameter specifies the name or identification of the datastore or virtual machine that you are restoring.

**[-esx-server {*esx server name* | *IP address*}]**

This mandatory parameter specifies the name or IP address of the ESX server. The parameter is required when restoring a VMFS datastore, or a virtual machine that

resides on a VMFS datastore, as well as when restoring an NFS virtual machine. The server name is the name of the ESX server as viewed through the vSphere Client. This name might differ from the host name or IP address of the ESX server.

**`[-backup-name {backup name}]`**

This optional parameter specifies which backup to restore. After adding the flag, you can add the name of the backup. If not specified, the latest available backup for the specified datastore or virtual machine is restored.

**`[-vmdk {hard disk name}]`**

This optional parameter specifies which hard disks are to be restored.

**`[-server {server name}]`**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**`[-restart-vm]`**

This optional parameter restarts the virtual machine after the restore operation.

**`[-quiet]`**

This optional parameter stops any output from displaying.

**`[-verbose]`**

This optional parameter provides detailed output when displaying information.

**`[-noprompt]`**

By default, a prompt appears, asking for confirmation when restoring a backup. This optional parameter disables the prompt.

**`[-user]`**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**`[-help]`**

This optional parameter displays help for this command.

### Example: Restoring a backup

The following example restores a virtual machine named `nfs1-vm1` from a backup named `backup-411`:

```
smvi backup restore -id nfs1_vm1 -backup-name backup-411
Are you sure you want to proceed with this operation? [yes|NO] y
[11:04] Starting restore request
[11:04] [WARN] No active mounts found for datastore vmfs_dsl
```

```
(47ab69d8-e7c72da0-d6c5-001a6412251d)
[11:05] Restoring nfs virtual machine on folder 'nfs1_vm1'
[11:07] Reloading virtual machine
[11:07] Restore is complete
SMVICLI-0100: Command completed successfully
```

## smvi backup unmount

The `smvi backup unmount` command unmounts a mounted virtual machine or datastore backup.

### Syntax

```
smvi backup unmount -backup-name {backup name} [-server {server name}] [-quiet] [-verbose] [-user] [-help]
```

### Description

**Note:** You must unmount a mounted backup in order to delete the backup or any of its preceding Snapshot copies.

### Parameters

**`[-backup-name {backup name}`**

This mandatory parameter specifies which backup to unmount. After adding the flag, add the name of the backup.

**`[-server {server name}`**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**`[-quiet]`**

This optional parameter stops any output from displaying.

**`[-verbose]`**

This optional parameter provides detailed output when displaying information.

**`[-user]`**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**`[-help]`**

This optional parameter displays help for this command.

**Example: Unmounting a backup**

The following example unmounts a backup named `vmfs2_db`:

```
smvi backup unmount -backup-name vmfs-2-db

[11:55] Starting unmount request
[11:55] Unmount is complete
SMVICLI-0100: Command completed successfully
```

**smvi discover datastores**

The `smvi discover datastores` command lists the datastores that are managed by the current vCenter Server and that reside on the storage systems currently assigned to your SnapManager for Virtual Infrastructure server.

**Syntax**

```
smvi discover datastores [-help]
```

**Parameters**

`[-help]`

This optional parameter displays help for this command.

**Example: Listing the datastores**

The following example lists all the datastores managed by the current vCenter Server that reside on storage systems assigned to SnapManager for Virtual Infrastructure:

```
smvi discover datastores

Password for NETAPP\vanib: *****
Datacenter: Aladdin
Datastore: nfs_datastore6
NFS: 172.17.170.21:/vol/nfs_vol6
Datastore: nfs_datastore7
NFS: 172.17.170.21:/vol/nfs_vol7
Datastore: nfs_datastore7 (Backup test1)
NFS: 172.17.170.21:/vol/
nfs_vol7_mount_33e49878c5e74363825e84652a724aef
Datastore: nfs_datastore7 (Backup test0)
NFS: 172.17.170.21:/vol/
nfs_vol7_mount_90a6ble7d6f948beaa6735af9692b3d4
Datastore: nfs_datastore7 (Backup backup_fgfdgfdgf_20080707134801)
NFS: 172.17.170.21:/vol/
nfs_vol7_mount_e50fc0eda0674cfbbf200f87f83ba8eb
```

```

Datastore: nfs_datastore8
NFS: 172.17.170.21:/vol/nfs_vol8
Datastore: nfs_datastore8 (Backup
backup_7d8597b0dfffd5c81806728dd45aea48)
NFS: 172.17.170.21:/vol/
nfs_vol8_mount_e7df47fbde00446cb6b589c821adc4dd
Datastore: vmfs_datastore5
LUN: 172.17.170.21:/vol/vmfs_vol5/lun5 Partition: 1
LUN: 172.17.170.21:/vol/vmfs_vol6/lun6 Partition: 1
Datastore: vmfs_datastore2
LUN: 172.17.170.21:/vol/vmfs_vol2/vmfs_lun2 Partition: 1
Datastore: vmfs_datastore3
LUN: 172.17.170.21:/vol/vmfs_vol3/lun3 Partition: 1
Datastore: vmfs_datastore4
LUN: 172.17.170.21:/vol/vmfs_vol4/lun4 Partition: 1
Datastore: vmfs_datastore7
LUN: 172.17.170.21:/vol/vmfs_vol7/qtrees_vol7/lun7 Partition: 1
Datastore: snap-00000002-vmfs_datastore
LUN: 172.17.170.21:/vol/vmfs_vol1/vmfs_lun1 Partition: 1
Datastore: vmfs7_testAJ_1
LUN: 172.17.170.21:/vol/volaj1/lun1 Partition: 1
Datastore: vmfs7_testAJ-2
LUN: 172.17.170.21:/vol/volaj1/lun2 Partition: 1
Datacenter: Bellagio

```

## smvi filerestore add-portgroup

The `smvi filerestore add-portgroup` command assigns virtual machines to a port group. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```

smvi filerestore add-portgroup [-name{port group name}] [-server{server
name}] [-user] [-verbose] [-help]

```

### Parameters

**[-name {port group name}]**

This mandatory parameter specifies the name of the port group, or network, that is used to enable or disable administrator-assisted file-level restore operations.

**[-server {server name}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is `localhost`.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

## smvi filerestore delete-portgroup

The `smvi filerestore delete-portgroup` command removes the port group and disables file restore sessions for the virtual machines assigned to the port group. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi filerestore delete-portgroup [-name{port group name}] [-server
{server name}] [-user] [-verbose] [-help]
```

### Parameters

**[-name {port group name}]**

This mandatory parameter specifies the name of the port group, or network, that is used to enable or disable administrator-assisted file-level restore operations.

**[-server {server name}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

## smvi notification list

The `smvi notification list` command displays information about the alert notification.

### Syntax

```
smvi notification list [-server {server name}] [-user] [-
verbose] [-help]
```

**Parameters****[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

**smvi notification set**

The `smvi notification set` command displays information about the alert notification.

**Syntax**

```
smvi notification set    [-smtp server {dns name / ip address}]    [-from
{from email address}]    [-to {to email address}]    [-server
{server name}]    [-user] [-verbose]    [-help]
```

**Parameters****[-smtp server {*dns name* / *ip address*}]**

This mandatory parameter specifies the name or IP address of the SMTP server that handles the test notification e-mail.

**[-from {*from email address*}]**

This mandatory parameter specifies the sender e-mail address.

**[-to {*to email address*}]**

This mandatory parameter specifies the comma-separated list of recipient e-mail addresses.

**[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

## smvi notification test

The `smvi notification test` command displays information about the test notification.

### Syntax

```
smvi notification test    [-server {server name}]    [-user] [-
verbose]    [-help]
```

### Parameters

**[-server {server name}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

## smvi restoreagent set

The `smvi restoreagent set` command sets the default installation URL of the restore agent. You can also perform this operation using the VSC for VMware vSphere user interface.

### Syntax

```
smvi restoreagent set    [-url]    [-server {server name}]    [-user]
[-verbose]    [-help]
```

**Parameters****[-url]**

This mandatory parameter provides an URL that points to a customer location for the restore agent installer.

**[-server {*server name*}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-user]**

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

**[-verbose]**

This optional parameter provides detailed output when displaying information.

**[-help]**

This optional parameter displays help for this command.

**smvi servercredential delete**

The `smvi servercredential delete` command deletes a user account created by the `smvi servercredential set` command.

**Syntax**

```
smvi servercredential delete -username {user name} [-help]
```

**Description**

**Note:** You cannot run this command from a remote host.

**Parameters****[-username {*user name*}]**

This mandatory parameter specifies the internal user account that you want to delete.

**[-help]**

This optional parameter displays help for this command.

**Example: Deleting a user account**

The following example deletes the `olduser2` user account:

```
smvi servercredential delete -username olduser2
SMVICLI-0100: Command completed successfully
```

## smvi servercredential list

The `smvi servercredential list` command lists a user account created by the `smvi servercredential set` command.

### Syntax

```
smvi servercredential list [-help]
```

### Description

**Note:** You cannot run this command from a remote host.

### Parameters

`[-help]`

This optional parameter displays help for this command.

#### Example: Listing the server credentials

The following example lists the current SnapManager for VI server credentials:

```
smvi servercredential list
Username
-----
administrator
```

## smvi servercredential set

The `smvi servercredential set` command adds a user account for Backup and Recovery capability to use for authentication instead of your Windows user credentials.

### Syntax

```
smvi servercredential set [-help]
```

## Description

**Note:** You cannot run this command from a remote host.

## Parameters

`[-help]`

This optional parameter displays help for this command.

### Example: Adding a user account

The following example adds a user account named administrator and sets a seven character password:

```

smvi servercredential set

Username: administrator
Password: *****
SMVICLI-0100: Command completed successfully

```

## smvi storagesystem list

The `smvi storagesystem list` command lists the added NetApp storage systems.

## Syntax

```

name}          smvi storagesystem list          [-server {server
help]          [-user]                          [-verbose]          [-

```

## Parameters

`[-server {server name}]`

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

`[-user]`

This optional parameter enables you to log in to the SnapManager for Virtual Infrastructure server with different user credentials than you are currently logged in with.

`[-verbose]`

This optional parameter provides detailed output when displaying information.

`[-help]`

This optional parameter displays help for this command.

### Example: Listing the storage systems

The following example lists the NetApp storage systems that reside in the local SnapManager for VI server; in this case, a single storage system with an IP address of 123.17.170.21:

```
smvi storagesystem list
```

Name	IP Address
-----	-----
123.17.170.21	123.17.170.21

## smvi version

The `smvi version` command displays the version of the VSC for VMware vSphere CLI and the SnapManager for VI server.

### Syntax

```
smvi version [-server {server name}] [-help]
```

### Parameters

**[-server {server name}]**

This optional parameter specifies the name of the SnapManager for Virtual Infrastructure server to which you are sending this command. The default value is localhost.

**[-help]**

This optional parameter displays help for this command.

### Example: Displaying the version

The following example displays the VSC for VMware vSphere CLI and server version:

```
smvi version
SnapManager for Virtual Infrastructure CLI Rballys.4N_120127_0000
(Build: 120127
)
SnapManager for Virtual Infrastructure Server Rballys.
4N_120127_0000 (Build: 120
127)
```

## Programmable APIs

---

Virtual Storage Console for VMware vSphere supports programmable application interfaces (APIs) for VMware vCloud and Provisioning and Cloning.

The APIs are exposed using Simple Object Access Protocol (SOAP). They provide a layer above the NetApp Manageability SDK, the VMware VI SDK, and the VMware vCloud SDK, but do not require any of these in the customer application or script.

### NetApp Communities provide support for API questions

You can submit general questions related to the VSC API for VMware vCloud and the Provisioning and Cloning API to the NetApp Communities Forum.

The NetApp Communities Forum is online at:

[https://communities.netapp.com/community/products\\_and\\_solutions/virtualization/vsc](https://communities.netapp.com/community/products_and_solutions/virtualization/vsc)

### What you can do with the APIs for VMware vCloud

Virtual Storage Console for VMware vSphere provides synchronous APIs that enable you to manage credentials for multiple vCenter Servers, discover vCloud Director objects for vCloud tenants, and provision and clone vApps. You must have the appropriate vCloud Director privileges to perform these tasks.

These APIs are Storage Profile unaware. This means that you are responsible for placing the virtual machines into appropriate Storage Profiles after a cloning or provisioning operation is performed.

For more information on working with the VMware vCloud APIs, go to the NetApp Communities Forum, which is online at:

[https://communities.netapp.com/community/products\\_and\\_solutions/virtualization/vsc](https://communities.netapp.com/community/products_and_solutions/virtualization/vsc)

### Provisioning and Cloning programmable API

The Provisioning and Cloning Application Programming Interface (API) is designed to be leveraged with the VI SDK. It provides end-to-end automated datastore provisioning and offloads the intricacies of storage object cloning while cloning virtual machines.

**Note:** Provisioning and Cloning performs the same role-based access control checks that Virtual Storage Console for VMware vSphere performs.

The managed object reference returned by the VMware VI SDK is used to identify components in the vCenter Inventory. You can view this information using the Managed Object Browser on the vCenter Server.

This version of the Provisioning and Cloning API exposes the virtual machine clone creation engine (which includes the redeploy feature), the datastore management engine (create, destroy, resize), and the file copy/clone offload engine. There are also two general-purpose utility methods included:

- `getVmFiles` returns a list of files that make up the virtual machine. This is useful for creating the list of files required in the `cloneSpec` API.
- `getMoref` returns the managed object reference of the requested object based on name and type. The `getMoref` returns the first object that matches the name and type combination. For this reason, this method should not be used in production environments unless all object names are unique.

## The virtual machine clone engine

The virtual machine clone engine provides two clone creation and routing methods: `createClones` and `redeployVMs`.

- `createClones` can be used to create virtual machine clones on new or existing datastores. When more than one datastore is created, the FlexClone feature on the controller is leveraged to create clones of the datastore.
- `redeployVMs` provides the ability to redeploy the virtual hard drives of the source virtual machine to the virtual machines specified. This feature leverages the FlexClone feature on the controller as well.

## The datastore management engine

The datastore management engine provides three methods for managing datastores: `createDatastore`, `resizeDatastore`, and `destroyDatastore`.

- The `createDatastore` method provides the ability to provision storage on the controller, present it to the ESX hosts, and create a datastore.
- The `resizeDatastore` method provides the ability to grow and shrink NFS-based datastores and grow VMFS-based datastores.
- The `destroyDatastore` method provides the ability to delete all virtual machines associated with the datastore, unmount it from ESX hosts, destroy the storage objects on the controller, and free the space.

## The file copy/clone offload engine

The file copy/clone offload engine provides four methods. These methods provide the ability to execute and monitor file copy and clone operations.

This engine provides the ability to offload file copy and clone operations to the controller for NFS-based datastores. This functionality is unique compared to that provided by the other engines in that it does not require a Virtual Center session. An ESX host session can be used instead.

The input to the methods is a combination of complex (specification and message) and simple (string, int, long, boolean, and so on) data types. The specifications and messages are described below.

**Note:** Very little verification or validation is done in the API. For example, if there is not enough space to create the requested datastore(s), the API method will fail.

## Provisioning and Cloning methods

This section describes all the available Provisioning and Cloning methods.

### Virtual machine clone creation and redeploy engine

This section describes the APIs for interfacing with the virtual machine clone creation and redeploy engine.

#### createClones

You can use the `createClones` method to create virtual machine clones on new or existing datastores.

The source can be a virtual machine or a virtual machine template. The source can be further refined by specifying a virtual machine snapshot. The following options cause at least one native clone (built into Virtual Center) operation to occur:

- clone source is powered on
- virtual machine snapshot is specified
- hard drive transformation is specified

The virtual machine or template must not contain any RDMs, must not contain any devices that use `VMDirectPath`, and must be connected.

The mix of `VirtualIDEController` attached hard drives and `VirtualSCSIController` hard drives in the same virtual machine may result in the drives being reordered in the resulting clones, therefore this is not supported. The creation of virtual machines based on hardware version `vmx-07` will fail on ESX 3.5 hosts.

#### Status

Current (added in version 2.1)

#### Type

Asynchronous

## Parameters

Name	Type	Value	Description
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.  <b>Note:</b> See <a href="#">RequestSpec</a> on page 180.

### XML

```

complexType name="createClones">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

## Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task. This task can be monitored and altered using the VI SDK

### Return XML

```

<complexType name="createClonesResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

```

    </restriction>
  </complexContent>
</complexType>

```

## RedeployVMs

You can redeploy the virtual hard drives associated with a virtual machine to other virtual machines.

The source can be a virtual machine, a virtual machine template, or a virtual machine snapshot. The following options cause a native (built into Virtual Center) clone operation before it can use the rapid clone methodology:

- clone source is powered on
- virtual machine snapshot is specified
- hard drive transformation is specified

The virtual machine or template must not contain any RDMs or any devices that use `VMDirectPath`, and must be in a good state (connected).

### Status

Current (added in version 3.0)

### Type

Asynchronous

### Parameters

Name	Type	Value	Description
requestSpec	Object	RequestSpec	<b>Note:</b> See <a href="#">RequestSpec</a> on page 180.

### XML

```

complexType name="redeployVMs">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}controllerSpec" maxOccurs="unbounded"
minOccurs="0"/>
      </sequence>

```

```

    </restriction>
  </complexContent>
</complexType>

```

**Returns**

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task. This task can be monitored and altered using the VI SDK

**Return XML**

```

<complexType name="redeployVMsResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**Datastore management engine**

This section describes the APIs for interfacing with the datastore management engine.

**createDatastore**

You can use the `createDatastore` method to provision storage on the storage controller, attach it to one or more ESX hosts and create a datastore.

More than one ESX host can be chosen by specifying the managed object reference of a cluster or datacenter in the `DatastoreSpec`.

**Status**

Current (added in version 3.0)

**Type**

Synchronous

## Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	The specification describing the datastore to create.  <b>Note:</b> See <a href="#">DatastoreSpec</a> on page 185.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.  <b>Note:</b> See <a href="#">RequestSpec</a> on page 180.

## XML

```

complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.netapp.com/}controllerSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## Returns

Type	Value	Description
String	newDatastore	The name of the new datastore that was created.

**Return XML**

```
<complexType name="createDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

**resizeDatastore**

You can use the `resizeDatastore` method to grow or shrink NFS-based datastores (and associated storage objects on the controller), and grow VMFS-based datastores (and associated storage objects on the controller).

**Status**

Current (added in version 3.0)

**Type**

Synchronous

**Parameters**

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	Specification describing datastore resize request.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.

**XML**

```
complexType
name="resizeDatastore"
```

```

>
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}datastoreSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

### Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task.

### Return XML

```

complexType name="resizeDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

### destroyDatastore

You can use the `destroyDatastore` method to delete any virtual machine with a file on the specified datastore, delete the datastore (after detaching from each ESX host), take the storage objects offline, and destroy the datastore (to free space).

### Status

Current (added in version 3.0)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	Specification describing datastore resize request.  <b>Note:</b> See <a href="#">DatastoreSpec</a> on page 185.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.  <b>Note:</b> See <a href="#">RequestSpec</a> on page 180.

### XML

```

complexType name="destroyDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}datastoreSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

### Returns

Type	Value	Description
String	Task:task-2	A managed object reference to a vCenter task.

**Return XML**

```

<complexType name="destroyDatastoreResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**Connection Broker features**

This section describes the APIs for interfacing with connection broker features.

**performViewImport**

You can use the `performViewImport` method to import the specified virtual machines into a VMware View Server.

**Status**

Current (added in version 3.2)

**Type**

Synchronous

**Parameters**

Name	Type	Value	Description
dsSpec	Object	DatastoreSpec	The specification describing the connection broker information.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.

Name	Type	Value	Description
vmsForImport	List <String>		A list of the virtual machines (by name) that should be imported into the View server.

### XML

```

complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}connectionBrokerSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg2" type="{http://
server.kamino.netapp.com/}vmsForImport" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

### Returns

Type	Value	Description
Void	N/A	Nothing returned

### Return XML

```

complexType name="createDatastore">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}void" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

## createXenImportFile

You can use the `createXenImportFile` method to create a file for importing virtual machines into a Citrix XenDesktop server.

### Status

Current (added in version 3.2)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
dsSpec	Object	ConnectionBrokerSpec	The specification describing the connection broker information.
requestSpec	Object	RequestSpec	Request can specify a vCenter server only. This method does not support direct connections to ESX hosts.
vmsForImport	List		A list of the virtual machines (by name) that should be imported into the View server.

### XML

```

complexType name="createXenImportFile">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}connectionBrokerSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg2" type="{http://
server.kamino.netapp.com/}vmsForImport" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

```

</complexContent>
</complexType>
    
```

## Returns

Type	Value	Description
Void	N/A	Nothing returned

### Return XML

```

complexType name="createXenImportFileResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}void" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

## Copy/Clone offload engine

This section describes the APIs for interfacing with the Copy/Clone offload engine.

The `fileCopyOffload` and `fileCloneOffload` take `VmFileSpec` and `RequestSpec` as arguments. The `DatastoreSpec` should contain only the datastore-managed object reference and a reference to the controller. This is defined as *Existing Datastore* in the formulas section of the *DatastoreSpec* documentation.

### fileCopyOffload

You can use the `fileCopyOffload` method to offload the copy of an NFS datastore file to the controller. This method should be used in cases where a full copy (all unique blocks) is required. In all other cases, the `fileCloneOffload` should be used.

This process involves a start-up time, which is quickly recovered when copying large files (because the offloaded controller base copy is very efficient). This start-up time may cause the offloaded copy of small files to take longer than using a host-based copy.

This method supports copying a file within the same controller. The `VmFileSpec` for the source and destination must specify the same controller.

**Status**

Current (added in version 3.0)

**Type**

Asynchronous

**Parameters**

Name	Type	Value	Description
source	Object	VmFileSpec	Specification describing the source file (datastore and controller).
destination	Object	VmFileSpec	Specification describing the destination file (datastore and controller).
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

**XML**

```

complexType name="fileCopyOffload">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}vmFileSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}vmFileSpec" minOccurs="0"/>
        <element name="arg2" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**Returns**

Type	Value	Description
Integer	876234	The operation identifier to monitor using <code>getFileOpOffloadStatus</code> .

**Return XML**

```

complexType name="fileCopyOffloadResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}int"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**fileCloneOffload**

You can use the `fileCloneOffload` method to offload the clone of an NFS datastore file to the controller.

This process uses the file level FlexClone feature of the controller. This process automatically falls back to the controller offloaded copy as needed if the `fallBackToCopy` parameter is set to `true`.

This method supports only cloning the file within the same volume. If `fallBackToCopy` is set to `true`, this method supports copying file within the same controller. In both cases, the `vmFileSpec` for the source and destination must specify the same controller.

The most effective use of this method is to employ a strategy where the output of the first operation (the destination file) becomes the input (the source file) for the next operation. For example, to create three clones of `test-flat.vmdk`, the following process (pseudo code) is the most efficient:

```

clone(test-flat.vmdk,test1-flat.vmdk)
clone(test1-flat.vmdk,test2-flat.vmdk)
clone(test2-flat.vmdk,test3-flat.vmdk)

```

**Status**

Current (added in version 3.0)

**Type**

Asynchronous

**Parameters**

Name	Type	Value	Description
source	Object	VmFileSpec	Specification describing the source file (datastore and controller).
destination	Object	VmFileSpec	Specification describing the destination file (datastore and controller).
fallBackToCopy	Boolean		If set to <code>true</code> , engine runs in "fully automatic mode" which falls back to an offloaded copy as needed. If <code>false</code> , conditions that would normally fall back to a copy result in an error (which the caller must deal with).
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

**XML**

```

complexType name="fileCloneOffload">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://
server.kamino.netapp.com/}vmFileSpec" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}vmFileSpec" minOccurs="0"/>
        <element name="arg2" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        <element name="arg3" type="{http://www.w3.org/2001/
XMLSchema}boolean" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

```

</complexContent>
</complexType>

```

## Returns

Type	Value	Description
Integer	876234	The operation identifier to monitor using <code>getFileOpOffloadStatus</code> .

## Return XML

```

complexType name="fileCloneOffloadResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}int"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## getFileOpOffloadStatus

You can use the `getFileOpOffloadStatus` method to track the progress of a `fileCopyOffload` or `fileCloneOffload` operation.

The status will be `complete`, `failed`, or `running`. When this method returns a `StatusMessage` with a status of `complete` or `failed`, the operation information is marked for cleanup, which occurs five minutes later. After the operation information has been cleaned up, it is no longer visible using this method. The progress field displays information about the progress of the operation.

## Status

Current (added in version 3.0)

## Type

Synchronous

**Parameters**

Name	Type	Value	Description
opId	Integer	876234	The operation identifier returned from fileCopyOffload or fileCloneOffload.

**XML**

```

complexType name="getFileOpOffloadStatus">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/
XMLSchema}int"/>/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**Returns**

Type	Value	Description
String	StatusMessage	Information describing status, progress, and reason for error (if operation fails).  <b>Note:</b> See <a href="#">StatusMessage</a> on page 194.

**Return XML**

```

complexType name="getFileOpOffloadStatusResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://
server.kamino.netapp.com/}statusMessage" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## clearAllFinishedOpOffloadStatus

You can use the `clearAllFinishedOpOffloadStatus` method to start the cleanup timer described in `getFileOpOffloadStatus` for all operations that have a status of complete or failed.

### Status

Current (added in version 3.0)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
opId	Integer		The operation identifier returned from <code>fileCopyOffload</code> or <code>fileCloneOffload</code> .

### XML

```
complexType name="clearAllFinishedOpOffloadStatus">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

### Returns

Void

### Return XML

```
complexType name="clearAllFinishedOpOffloadStatusResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
      </sequence>
    </restriction>
  </complexContent>
```

```

    </restriction>
  </complexContent>
</complexType>

```

## Utility methods

This section describes utility methods that return output, such as a list of created virtual machines and the managed object reference of each virtual machine.

## getVms

You can use the `getVms` method to obtain the list of virtual machines that were created using the `createClones` method. This list can be used in the `redeployVms` method.

### Status

Current (added in version 3.0)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
vmMorRef	Object	The managed object reference of the VM.	The managed object reference of the VM.

### XML

```

complexType
  name="getVms"
>
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="arg1" type="{http://server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## Returns

Type	Value	Description
vmMorRef	String	The managed object reference of the VM.
requestSpec	RequestSpec	Request can specify a vCenter server or ESX host.

### Return XML

```

<complexType name="getVmsResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/
2001/XMLSchema}string" maxOccurs="unbounded" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

## getVmFiles

You can use the `getVmFiles` method to obtain a skeleton list of `VmFileSpec` to be completed and used in the submission to `createClones`.

### Status

Current (added in version 2.1)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
vmMorRef	String		The managed object reference of the VM.
requestSpec	Object	RequestSpec	Request can specify a vCenter server or ESX host.

**XML**

```

complexType
name="getVmFiles"
>
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="arg0" type="{http://www.w3.org/2001/
XMLSchema}string" minOccurs="0"/>
        <element name="arg1" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

**Returns**

Type	Value	Description
List	Object <VmFileSpec>	A list of VmFileSpec based on the VM specified. This information should be modified and submitted using the CloneSpec.

**Return XML**

```

<complexType name="getVmFilesResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://
server.kamino.netapp.com/}vmFileSpec" maxOccurs="unbounded"
minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## getMoref

You can use the `getMoref` method to obtain the managed object reference of the requested object based on name and type.

The `getMoref` method returns the first object that matches the name and type combination. For this reason, this method should not be used in production environments unless all object names are unique.

### Status

Current (added in version 2.1)

### Type

Synchronous

### Parameters

Name	Type	Value	Description
name	String		Name of object to look for.
type	String		Managed object type.
requestSpec	String		Request can specify a vCenter server or ESX host.

### XML

```
complexType
  name="getMoref"
  >
    <complexContent>
      <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
        <sequence>
          <element name="arg0" type="{http://www.w3.org/2001/
XMLSchema}string" minOccurs="0"/>
          <element name="arg1" type="{http://www.w3.org/2001/
XMLSchema}string" minOccurs="0"/>
          <element name="arg2" type="{http://
server.kamino.netapp.com/}requestSpec" minOccurs="0"/>
        </sequence>
      </restriction>
```

```

</complexContent>
</complexType>
    
```

**Returns**

Type	Value	Description
String		Managed object reference in string format.

**Return XML**

```

<complexType name="getMorefResponse">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="return" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

**Provisioning and Cloning specifications and messages**

The Provisioning and Cloning API provides several specifications and messages.

**RequestSpec**

The RequestSpec specification describes the URL of the VMware vCenter SDK or ESX host as well as the authentication information. The authentication information may be in the form of a user name and password combination or a VMware Session ID. An optional clone specification may also be present.

**Properties**

Type	Value	Description
serviceUrl	String	URL for the VMware vCenter SDK
vcUser	String	VMware vCenter username (null ok if using vcSession)

Type	Value	Description
vcPassword	String	VMware vCenter password (null ok if using vcSession)
vcSession	String	VMware session (null ok if using vcUser/vcPassword)
cloneSpec	Object CloneSpec	A clone specification

## Notes

- cloneSpec may be null when using this spec with anything other than redeployVMs or createClones.
- vcSession should be null if vcUser and vcPassword are used.
- vcUser and vcPassword should be null if vcSession is used.

## XML

```
complexType
  name="requestSpec"
  >
    <complexContent>
      <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
        <sequence>
          <element name="cloneSpec" type="{http://
server.netapp.com/}cloneSpec" minOccurs="0"/>
          <element name="serviceUrl" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
          <element name="vcPassword" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
          <element name="vcSession" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
          <element name="vcUser" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
        </sequence>
      </restriction>
    </complexContent>
  </complexType>
```

## CloneSpec

The CloneSpec specification describes a request to create clones of a virtual machine or template or to redeploy the virtual hard drives.

When used with the redeployVMs method, each virtual machine in the map named clones will have its virtual hard drives replaced with those of the source. The string in this map is the name of the virtual machine to be redeployed and the vmSpec describes this virtual machine.

When `CloneSpec` is used with the `createClones` method, a new virtual machine is created for each entry in the clones map. The string in this map is the name of the clone and the `vmSpec` describes the new clone configuration. The list named files describes the files that make up the source virtual machine or template. This list can be used to specify different destinations for each file as well as to create new datastores.

Type	Value	Description
templateMoref	String	Source VM or template of cloning operation. String representation of type and value of <code>ManagedObjectReference</code> from VMware VI API.
snapshotMoref	String	The managed object reference for a snapshot of the source virtual machine to base the clones on.
containerMoref	String	Destination for resulting clones. Valid destination types: <code>Datacenter</code> , <code>ResourcePool</code> , <code>ClusterComputeResource</code> , and <code>ComputeResource</code> . A string representation of type and value of <code>ManagedObjectReference</code> from the VMware VI API.
destVmFolderMoref	String	Virtual machine folder the clones should be created in. If null, clones are created at the root virtual machine folder.
vmTransform	String	Transforms all virtual hard drives to specified format. Should be specified only when there is actual work to do. Specifying a transform when one is not required causes unnecessary work. Options are <code>null</code> , <code>flat</code> , and <code>sparse</code> .

Type	Value	Description
hardwareVersion	String	Upgrade hardware version from a previous version to vmx-04 or vmx-07.  <b>Note:</b> vmx-04 is supported by ESX 3.5 and both are supported by ESX 4.0.
clones	Map <String, VmSpec>	Map of new virtual machine name to virtual machine specification (VmSpec).
files	List <VmFileSpec>	List of files that make up source virtual machine or template specified in templateMoref.
memMB	Long	Override the source virtual machine (or template) amount of memory during cloning process. Value is in MB.
numberCPU	Int	Override the source virtual machine (or template) number of CPUs during cloning process.

## XML

```

<complexType name="cloneSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="clones">
          <complexType>
            <complexContent>
              <restriction base="{http://www.w3.org/
2001/XMLSchema}anyType">
                <sequence><element name="entry"
maxOccurs="unbounded" minOccurs="0">
                  <complexType>
                    <complexContent>
                      <restriction
base="{http://www.w3.org/2001/XMLSchema}anyType">
                        <sequence>
                          <element
name="key" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>

```

```

                                <element
name="value" type="{http://server.netapp.com/}vmSpec"
minOccurs="0"/>
                                </sequence>
                                </restriction>
                                </complexContent>
                                </complexType>
                                </element>
                                </sequence>
                                </restriction>
                                </complexContent>
                                </complexType>
                                </element>
                                <element name="connBroker" type="{http://
server.netapp.com/}connectionBrokerSpec" minOccurs="0"/>
                                <element name="containerMoref" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
                                <element name="destVmFolderMoref" type="{http://
www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
                                <element name="files" type="{http://
server.kamino.netapp.com/}vmFileSpec" maxOccurs="unbounded"
minOccurs="0"/>
                                <element name="hardwareVersion" type="{http://
www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
                                <element name="snapshotMoref" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/
>
                                <element name="templateMoref" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
                                <element name="vmTransform" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
                                <element name="memMB" type="{http://www.w3.org/2001/
/XMLSchema}long" minOccurs="0"/>
                                <element name="numberCPU" type="{http://www.w3.org/2001/
/XMLSchema}int" minOccurs="0"/>
                                </sequence>
                                </restriction>
                                </complexContent>
                                </complexType>

```

## VmFileSpec

The `VmFileSpec` specification describes the source configuration file (`vmx`) or the source virtual hard disk files (`vmdk`) as well as the destination datastore specification.

### Properties

Type	Value	Description
<code>sourcePath</code>	String	Path to <code>vmx</code> or <code>vmdk</code> file. The string Configuration File can be passed in place of an actual <code>vmx</code> file.
<code>destDatastoreSpec</code>	<code>DatastoreSpec</code>	Destination datastore specification.

### XML

```
<complexType name="vmFileSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="destDatastoreSpec" type="{http://
server.netapp.com/}datastoreSpec" minOccurs="0"/>
        <element name="sourcePath" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

## DatastoreSpec

The `DatastoreSpec` specification describes the destination datastore. This can describe new datastores to be created as well as existing datastores.

See the “Required parameters” section below for valid parameter combinations. The number of clones must be evenly divisible by the number of datastores.

**Properties**

Type	Value	Description
mor	String	Destination datastore. String representation of type and value of <code>ManagedObjectReference</code> from the VMware VI API.
targetMor	String	The managed object reference of the vCenter object in which to attach the new datastore. Can be an ESX host, cluster or datacenter.
goldVolume	String	Name of volume used when creating more than one NFS-based datastore. This volume is not permanently presented to the ESX hosts. This volume becomes the parent of the FlexClones.
protocol	String	The protocol being used. Valid values are NFS, FCP, iSCSI.
containerName	String	Name of the aggregate for new NFS datastores, or name of volume for new VMFS datastores.
sizeInMB	Long	Size of the datastore in MB. An additional 256 MB is added for VMFS datastores to cover metadata overhead.
thinProvision	Boolean	If <code>true</code> , space will not be reserved for the storage object. For NFS, the volume will <code>guarantee=none</code> . For VMFS, LUN will be created with <code>'-o noreserve'</code> .
volAutoGrow	Boolean	If <code>true</code> , the <code>volAutoGrowInc</code> and <code>volAutoGrowMax</code> values are applied to the volume.

Type	Value	Description
volAutoGrowInc	Long	Increment in which to grow volume automatically as needed in MB.
volAutoGrowMax	Long	Maximum size to which to grow the volume automatically in MB.
datastoreNames	List <String>	List of datastore names. Care should be taken by the application to prevent duplicate datastore, volume or LUN names. For NFS, datastore name is used as volume name. For VMFS, datastore name is used as LUN name.
numDatastores	Int	Number of datastores. This should indicate the size of the list of names in datastoreNames.
blockSize	Int	VMFS block size in MB.
controller	ControllerSpec	The controller.
wrapperVol	Boolean	When true, new volume is created to contain the new LUN to be used for a new VMFS datastore (containerName must contain aggregate name if true). When true, the volume containing the LUN (VMFS datastore) will be resized to make room for the new size of the LUN (if required).

### Required parameters

Some actions require the use of multiple parameters.

### Specifying an existing datastore

- mor
- controller

### **Specifying new NFS datastores using createClones or createDatastore**

- targetMor - only required for createDatastore
- containerName
- sizeInMB
- thinProvision
- volAutoGrow
- volAutoGrowInc
- volAutoGrowMax
- protocol - must be NFS
- controller
- datastoreNames - only one name in the list
- numDatastores - should be 1

### **Specifying new NFS datastores using createClones**

- goldVolume
- containerName
- sizeInMB
- thinProvision
- volAutoGrow
- volAutoGrowInc
- volAutoGrowMax
- protocol - Must be NFS
- controller
- datastoreNames
- numDatastores

### **Specifying new VMFS datastores using createClones or createDatastore**

- targetMor - only required for createDatastore
- containerName
- sizeInMB
- thinProvision
- protocol - must be FCP or iSCSI
- controller
- datastoreNames - only one name in the list
- numDatastores - should be 1

### **Specifying new VMFS datastores using createClones**

- containerName
- sizeInMB

- thinProvision
- protocol - must be FCP or iSCSI
- controller
- datastoreNames
- numDatastores

## XML

```

<complexType name="datastoreSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/XMLSchema}anyType">
      <sequence>
        <element name="containerName" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="blockSize" type="{http://www.w3.org/2001/XMLSchema}int" minOccurs="0"/>
        <element name="controller" type="{http://server.kamino.netapp.com/controllerSpec}" minOccurs="0"/>
        <element name="datastoreNames" type="{http://www.w3.org/2001/XMLSchema}string" maxOccurs="unbounded" minOccurs="0"/>
        <element name="goldVolume" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="mor" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="numDatastores" type="{http://www.w3.org/2001/XMLSchema}int"/>
        <element name="protocol" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="sizeInMB" type="{http://www.w3.org/2001/XMLSchema}long" minOccurs="0"/>
        <element name="targetMor" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="temp" type="{http://www.w3.org/2001/XMLSchema}string" minOccurs="0"/>
        <element name="thinProvision" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
        <element name="volAutoGrow" type="{http://www.w3.org/2001/XMLSchema}boolean"/>
        <element name="volAutoGrowInc" type="{http://www.w3.org/2001/XMLSchema}long" minOccurs="0"/>
        <element name="volAutoGrowMax" type="{http://www.w3.org/2001/XMLSchema}long" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>

```

## ControllerSpec

The `ControllerSpec` specification describes the controller connection and authentication data. This information is used by the Provisioning and Cloning capability to connect to the controller using the ZAPI interface. No other protocol is used to connect to the controller.

### Properties

Type	Value	Description
ipAddress	String	IP or host name of the controller.
username	String	User name (does not need to be root).
password	String	Password.
ssl	Boolean	If <code>true</code> , use HTTPS. If <code>false</code> , use HTTP to connect to the controller.
passthroughContext	String (optional)	Name of the vFiler or Vserver on which to create the new storage.
actuallyOpsMgr	Boolean (optional)	If <code>true</code> , connects to Operations Manager. If <code>false</code> , connects to controller.

### XML

```
<complexType name="controllerSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="ipAddress" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
        <element name="password" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
        <element name="ssl" type="{http://www.w3.org/2001/
XMLSchema}boolean"/>
        <element name="username" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
```

## VmSpec

The `VmSpec` specification describes configuration and action options for each virtual machine created.

The MAC address must be in a range defined by VMware. Refer to VMware documentation for more information.

### Properties

Type	Value	Description
<code>macAddress</code>	<code>Map &lt;String, String&gt;</code>	Virtual network adapter to MAC address information (optional).
<code>custSpecName</code>	<code>String</code>	Name of guest customization specification to be applied (optional).
<code>vmMoref</code>	<code>String</code>	The managed object reference of the virtual machine to be redeployed.
<code>powerOn</code>	<code>Boolean</code>	If <code>true</code> , the new virtual machines are powered on after all have been created.

### XML

```
<complexType name="VmSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="custSpec" type="{http://
server.kamino.netapp.com/}guestCustomizationSpec" minOccurs="0"/>
        <element name="domain" type="{http://
server.kamino.netapp.com/}domainSpec" minOccurs="0"/>
        <element name="macAddress">
          <complexType>
            <complexContent>
              <restriction base="{http://www.w3.org/
2001/XMLSchema}anyType">
                <sequence>
                  <element name="entry"
maxOccurs="unbounded" minOccurs="0">
                    <complexType>
                      <complexContent>
                        <restriction
base="{http://www.w3.org/2001/XMLSchema}anyType">
```

```

                <sequence>
                    <element
name="key" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
                    <element
name="value" type="{http://www.w3.org/2001/XMLSchema}string"
minOccurs="0"/>
                </sequence>
            </restriction>
        </complexContent>
    </complexType>
</element>
</sequence>
</restriction>
</complexContent>
</complexType>
</element>
<element name="powerOn" type="{http://www.w3.org/
2001/XMLSchema}boolean"/>
    <element name="vmMoref" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
</sequence>
</restriction>
</complexContent>
</complexType>

```

### GuestCustomizationSpec

The GuestCustomizationSpec specification identifies the guest customization specification.

#### Properties

Type	Value	Description
name	String	Name of the guest customization specification.
useVmName	Boolean	If guest customization specification is of type CustomizationSysprepText, this option can be used to make the guest hostname match the virtual machine name.  <b>Note:</b> It is the responsibility of the implementer to ensure that the virtual machine name results in a valid host name.

**XML**

```

<complexType name="guestCustomizationSpec">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
      <sequence>
        <element name="name" type="{http://www.w3.org/2001/
XMLSchema}string" minOccurs="0"/>
        <element name="useVmName" type="{http://www.w3.org/
2001/XMLSchema}boolean"/>
      </sequence>
    </restriction>
  </complexContent>
</complexType>
    
```

**ConnectionBrokerSpec**

The `ConnectionBrokerSpec` specification describes the connection broker and authentication data. This information is used by VSC for VMware vSphere to connect to the connection broker (VMware View or Citrix XenDesktop).

**Properties**

Type	Value	Description
type	ConnectionBrokerType	The type of connection broker to import into. This can be VMWARE_VIEW_4_0, VMWARE_VIEW_4_5, or XEN_DESKTOP.
host	String	Hostname or IP address of the connection broker (used only for View import).
username	String	A user who can access the View Server (used only for View import).
password	String	Password of the specified user (used only for View import).
domain	String	FQDN where the connection broker resides.
desktopType	DesktopType	INDIVIDUAL_DESKTOP or DESKTOP_POOL (used only for View import).

Type	Value	Description
accessMode	AccessMode	PERSISTENT or NON_PERSISTENT. This corresponds to dedicated and floating in View 4.5 and higher, respectively (used only for View import).
poolType	PoolType	NEW or EXISTING. Create a new pool or use an existing one (used only for View import).
poolName	String	The name of the new pool if the PoolType is set to NEW (used only for View import).

## StatusMessage

You can use the `StatusMessage` specification to obtain the progress and status of an operation.

## Properties

Type	Value	Description
id	Int	Operation identifier.
progress	Int	Valid values are 0-100. Indicates how much of the copy or clone process has completed at the time of the query.
status	String	Valid values are <code>complete</code> (finished without error), <code>failed</code> (finished with error), or <code>running</code> (operation in progress).
reason	String	If the status is <code>failed</code> , this contains the reason for the failure.

## XML

```
<complexType name="statusMessage">
  <complexContent>
    <restriction base="{http://www.w3.org/2001/
XMLSchema}anyType">
```

```

        <sequence>
            <element name="id" type="{http://www.w3.org/2001/
XMLSchema}int"/>
            <element name="progress" type="{http://www.w3.org/
2001/XMLSchema}int"/>
            <element name="reason" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
            <element name="status" type="{http://www.w3.org/
2001/XMLSchema}string" minOccurs="0"/>
        </sequence>
    </restriction>
</complexContent>
</complexType>

```

## Provisioning and Cloning sample code

```

<ns2:createClonesxmlns:ns2="http://server.kamino.netapp.com/">
<arg0>
    <cloneSpec>
        <clones>
            <entry>
                <key>apiTestClone1</key>
                <value>
                    <powerOn>>false</powerOn>
                </value>
            </entry>
            <entry>
                <key>apiTestClone2</key>
                <value>
                    <powerOn>>false</powerOn>
                </value>
            </entry>
            <entry>
                <key>apiTestClone3</key>
                <value>
                    <powerOn>>false</powerOn>
                </value>
            </entry>
        </clones>
        <containerMoref>Datacenter:datacenter-2</containerMoref>
        <files>
            <destDatastoreSpec>
                <controller>
                    <ipAddress>10.10.10.2</ipAddress>
                    <password></password>
                    <ssl>>false</ssl>
                    <username>root</username>
                </controller>
                <mor>Datastore:datastore-17</mor>
                <numDatastores>0</numDatastores>
                <thinProvision>>false</thinProvision>
                <volAutoGrow>>false</volAutoGrow>
            </destDatastoreSpec>

```

```

        <sourcePath>[unitTestSourceNFS]demoSource/
demoSource.vmx</sourcePath>
    </files>
    <files>
        <destDatastoreSpec>
            <controller>
                <ipAddress>10.10.10.2</ipAddress>
                <password></password>
                <ssl>>false</ssl>
                <username>root</username>
            </controller>
            <mor>Datastore:datastore-17</mor>
            <numDatastores>0</numDatastores>
            <thinProvision>>false</thinProvision>
            <volAutoGrow>>false</volAutoGrow>
        </destDatastoreSpec>
        <sourcePath>[unitTestSourceNFS]demoSource/
demoSource.vmdk</sourcePath>
    </files>
    <templateMoref>VirtualMachine:vm-255</templateMoref>
</cloneSpec>
<serviceUrl>https://10.10.10.2/sdk</serviceUrl>
<vcPassword>pass123</vcPassword>
<vcUser>Administrator</vcUser>
</arg0>
</ns2:createClones>

```

## Provisioning and Cloning client-side programming

Various client-side programming environments allow you to access the SOAP service. You can use this service for your own client-side programming.

### Accessing the SOAP API through Java

You can access the SOAP API by using Java-based tools or the Java programming language.

### Generating a certificate for use with the `wsimport` tool

You must generate a certificate for use with `wsimport` so that the tool will be able to read the WSDL that is generated by the build process.

#### About this task

Complete the following steps to generate a certificate for use with `wsimport`.

#### Steps

1. Stop the NVPF service. This step is optional if you have already generated an SSL certificate for the environment.

2. Run the following command in the VSC for VMware vSphere installation directory:
 

```
c:\Program Files\NetApp\Virtual Storage Console>bin\vsc ssl setup -domain <domain>
```

For *<domain>*, enter the host name of the system running VSC for VMware vSphere or a fully qualified domain name of the system running VSC for VMware vSphere.

**Note:** This step is optional if you have already generated an SSL certificate for the environment.

3. From the VSC for VMware vSphere installation directory, change to `\etc` and run the command:
 

```
keytool -export -alias nvpf -keystore nvpf.keystore -file nvpf.cer
```

**Note:** If you moved the keystore file from the `c:\Program Files\NetApp\Virtual Storage Console\etc` directory, enter the path to the keystore file.

This command creates a new file called `nvpf.cer`.

This certificate will be imported to the local Java keystore. If you have the Java JRE version 1.6.0\_21 installed in Program Files, the command to execute will look like this:

```
c:\Program Files\NetApp\Virtual Storage Console\etc>keytool -import -alias nvpf -file nvpf.cer -keystore "c:\Program Files\Java\jdk1.6.0_21\jre\lib\security\cacerts"
```

4. Enter keystore password:
 

```
changeit
```
5. Run `wsimport` to grab the WSDL and generate the Java classes to write your own client. Enter the commands:

```
cd \dev
```

```
c:\dev>wsimport -verbose -s . -p com.netapp.kamino.api https://<domain>:8143/kamino/public/api?wsdl
```

Be sure to change *<domain>* to a valid host or domain name. If that name does not resolve through DNS, add it to your hosts file (not localhost).

## Accessing SOAP through C#

To begin using the SOAP API, you must first add the web reference to your project.

After the web service reference has been added, you can start accessing the client side objects to make API calls into the SOAP service.

## Troubleshooting

---

This section describes how to troubleshoot general VSC for VMware vSphere installation and usage issues.

### Issues that apply to multiple capabilities

There are several troubleshooting suggestions and issues that apply to all of the VSC for VMware vSphere capabilities. This section lists those issues.

### Information available at NetApp Communities Forum

You can submit general questions related to Virtual Storage Console for VMware vSphere to the NetApp Communities Forum.

The NetApp Communities Forum is online at:

[https://communities.netapp.com/community/products\\_and\\_solutions/virtualization/vsc](https://communities.netapp.com/community/products_and_solutions/virtualization/vsc)

### Check the Release Notes

The *Release Notes* contain the most up-to-date information about known problems and limitations. The *Release Notes* also contain information about how to look up information about known bugs.

The *Release Notes* are updated when there is new information about Virtual Storage Console for VMware vSphere. It is a good practice to check the *Release Notes* before you install VSC, and any time you encounter a problem with VSC.

You can access the *Release Notes* from the the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

### Uninstall does not remove standard VSC roles

When you uninstall Virtual Storage Console for VMware vSphere, the standard VSC roles remain. This is expected behavior and does not affect the performance of VSC or your ability to upgrade to a new version of VSC. You can manually delete these roles, if you choose.

While the uninstall program does not remove the roles, it does remove the localized names for the VSC-specific privileges and append the following prefix to them: "XXX missing privilege". For example, if you open the vSphere Edit Role dialog box after you install VSC, you will see the VSC-specific privileges listed as `XXX missing privilege.<privilege name>.label not found XXX`.

This behavior happens because the vCenter Server does not provide an option to remove privileges.

When you reinstall VSC or upgrade to a newer version of VSC, all the standard VSC roles and VSC-specific privileges are restored.

## VMware only supports selecting one object when using right-click actions

VMware limits how many objects can be selected using right-click actions when running a vCenter plug-in such as VSC for VMware vSphere. When you select multiple VMware objects using a right-click action, VMware performs the operation on the first object only.

For example, the Provisioning and Cloning capability uses a wizard-driven workflow to select a VMware object in the Inventory panel. If you use a right-click action to select multiple objects, VMware ignores all but the first object selected. This is why the Provisioning and Cloning capability only supports space reclamation on a single virtual machine (VM) or datastore.

This limitation applies only to VMware objects that are selected through the Inventory panel. If you are working with an object that is owned by VSC for VMware vSphere, such as an object owned by the Monitoring and Host Configuration capability, you can use a multi-select action.

## Issues that apply to the Monitoring and Host Configuration capability

This section contains information on troubleshooting tips and issues that affect the Monitoring and Host Configuration capability.

### Getting information about storage controllers with an Alert status

Display the Status Reason column to get an explanation of why a storage controller displays the Alert status on the Monitoring and Host Configuration Overview panel. By default, the Status Reason column is hidden.

#### Steps

1. On the Monitoring and Host Configuration **Overview** panel, click the arrow in the column header and select **Columns**.
2. Select **Status Reason**.

Note that vFiler units do not return a detailed status reason; This controller is a MultiStore vFiler unit is displayed for vFiler units. The displayed text is truncated if the status reason is too long, and ellipses (...) are appended to the displayed text.

### Getting information about an ESX and ESXi host with an Alert status

You can display the causes of Alert icons for ESX hosts on the Monitoring and Host Configuration capability Overview panel. There are several options for learning more about an Alert.

#### About this task

The Status Reason column contains an explanation of why an Alert is being displayed. By default, the Status Reason column is hidden.

Also, the Monitoring and Host Configuration capability displays an Alert icon in the **Status** column for a host if vCenter reports alarms for that host.

In addition, the Monitoring and Host Configuration capability displays an Alert icon the **Adapter Settings**, **MPIO Settings**, or **NFS Settings** columns if a current host setting is different than the supported value set by the capability.

### Step

1. To check the Alert status, take one of the following actions:

Option	Description
<b>To check the status of the host</b>	<p>Display the Status Reason column.</p> <ol style="list-style-type: none"> <li>a. On the Monitoring and Host Configuration capability Overview panel, click the arrow in the column header and select <b>Columns</b>.</li> <li>b. Select <b>Status Reason</b>.</li> </ol>
<b>For an Alert in the Status column</b>	See the Alarms tab in vCenter for the host.
<b>For an Alert in the other columns on the Monitoring and Host Configuration capability Overview panel,</b>	<p>Right-click the ESX or ESXi host and select <b>Show Details</b>.</p> <p>Compare the host settings with the values set by the Monitoring and Host Configuration capability.</p> <p>For NFS settings, the current value is shown in red if it does not match the recommended value, and the recommended value is displayed.</p> <p>For some settings, the default is used and no explicit value is set. These default values are shown as empty strings ("").</p>

### After you finish

Correct any alarms reported for the host. Select **Set Recommended Values** from the right-click menu to have the Monitoring and Host Configuration capability update the desired ESX or ESXi host settings.

### Related references

[ESX host values set by Monitoring and Host Configuration capability](#) on page 54

## Collecting the VSC for VMware vSphere log files

You can collect the log files from all installed VSC for VMware vSphere components and capabilities using the Data Collection panel in the Monitoring and Host Configuration capability. Technical support might ask you to collect the log files to help troubleshoot a problem.

### Steps

1. Open the vSphere Client and log into your vCenter Server.

2. Select a Datacenter in the **Inventory** panel, and then select the **NetApp** tab.
3. In the Monitoring and Host Configuration capability, select the **Data Collection** panel.
4. Select **Export VSC Logs** and click **Submit**.
5. When prompted, save the file to your local computer.

#### After you finish

Send the .zip file to technical support.

## Troubleshooting error message "The client cannot communicate with the Virtual Storage Console Server"

The vSphere Client displays the message "The client cannot communicate with the Virtual Storage Console Server. Verify that the server is running and that the client has network connectivity to the server." Correct any connectivity problems, and if necessary restart the affected servers.

#### Steps

1. If the VSC for VMware vSphere and vCenter servers run on separate systems, verify that the two systems have basic network connectivity between them.
2. Verify that any firewalls on the server allow communication on TCP port 443 (HTTPS/SSL).
3. Verify that the vCenter account credentials used for background discovery have not expired.
4. Log into the Windows server or servers verify that the VSC for VMware vSphere service is running and the VMware vCenter services are running.
5. If a service is not running, restart the service.
6. If none of the previous steps solved the problem, reboot the Windows server or servers running VSC for VMware vSphere and vCenter Server.

#### Related tasks

[Updating vCenter credentials for background discovery](#) on page 201

## Updating vCenter credentials for background discovery

If the vCenter credentials specified when Virtual Storage Console for VMware vSphere was installed expire, Monitoring and Host Configuration can no longer run background discovery tasks. Monitoring and Host Configuration then displays an error message. Re-register VSC to enter updated credentials.

#### Before you begin

The vCenter account must be an administrator-level account.

### Steps

1. Click the link in the error message about expired credentials, or point a Web browser to the registration Web page:  
`https://hostname:8143/Register.html`  
*hostname* is the host name or IP address of the server where VSC is installed.  
  
If a security certificate warning is displayed, choose the option to ignore it or to continue to the Web site.  
  
The Plugin registration Web page is displayed with the current credentials.
2. Enter the new password for the user name shown, or enter a new user name and password.
3. Restart all vCenter Clients.

## Issues that apply to the Provisioning and Cloning capability

This section contains information on troubleshooting tips and issues that affect the Provisioning and Cloning capability.

### Cloning operations fail when permission is assigned to wrong object

Provisioning and Cloning operations such as redeploying clones and creating rapid clones display error messages and fail if the permission is not assigned to the correct vSphere object.

For example, VSC displays the error message

No Template

if you start the Create Rapid Clones Wizard and the permission containing the privileges for this task is not assigned to the root object.

The error occurs because the cloning task uses vCenter privileges, such as Create Tasks. The permission containing that privilege must always be assigned to the root object; otherwise, the task will not complete successfully.

To prevent these tasks from failing, assign the permission to the root folder. If you do not want all the child objects in the hierarchy to inherit the permission, you can assign a different permission to those child objects. Permissions assigned to a child object always override inherited permissions.

## Resolution of Backup and Recovery issues

If you encounter unexpected behavior while configuring Backup and Recovery or during a backup or restore operation, you can follow specific troubleshooting procedures to identify and resolve the cause of such issues.

### Backup and Recovery values that you can override

To improve operational efficiency, you can modify the `smvi.override` configuration file to override the default values specified in the file. These values control such settings as the number of

VMware snapshots that are created or deleted during a backup or the amount of time before a backup script stops running.

The `smvi.override` configuration file is located in the installation directory at `C:\Program Files\NetApp\VSC\smvi\server\etc\smvi.override`.

You can modify the following entries if you need to override the default values and if you modify any of these entries, you must restart the server for the changes to take effect.

<code>vmware.max.concurrent.snapshots=6</code>	Specifies six as the default maximum number of VMware snapshots created or deleted per datastore during a backup.
<code>vmware.quiesce.retry.count=0</code>	Specifies zero as the maximum number of retry attempts for VMware snapshots.
<code>vmware.quiesce.retry.interval=5</code>	Specifies the amount of time, in seconds, between retry attempts for VMware snapshots.
<code>vim.client.log.verbose=true</code>	When the value is true, logs the interactions between the SMVI server and the vCenter server.
<code>smvi.script.timeout.seconds=120</code>	Specifies the SMVI timeout value for a pre-backup or post-backup script, which is when the SMVI server stops waiting for the script to finish running.

## Backup and Recovery event and error logs

Backup and Recovery logs both server messages and messages between the server and the user interface, including detailed information about event messages and errors. Reviewing these logs helps you troubleshoot any errors that occur during backup or restore operations.

The log files are stored under the installation directory at the following locations:

- The server log messages are at `C:\Program Files\NetApp\Virtual Storage Console\smvi\server\log\server.log`.
- The log messages between the user interface and the server are at `C:\Program Files\NetApp\Virtual Storage Console\log\smvi.log`.

## Email notification for scheduled backup contains a broken link

<b>Issue</b>	When you click the link to view the log files in the email notification for a backup job, an error occurs.
<b>Cause</b>	This problem occurs if you disable the IP address of the network adapters for the SMVI log viewer.
<b>Corrective action</b>	You must enable the IP address of the network adapters for the SMVI log viewer in one of the following ways, depending on your operating system:

- Select **Control Panel > Network connections > Network and Sharing Center** in Windows 2003, Windows 2008, Windows 2008 R2, and Windows 7 environments.
- Select **Control Panel > Network connections** in Windows Vista and Windows XP environments.

## You may have reached the maximum number of NFS volumes configured in the vCenter

**Message** You may have reached the maximum number of NFS volumes configured in the vCenter. Check the vSphere Client for any error messages.

**Description** This message occurs when you attempt to mount a backup copy of an NFS datastore on a Vserver with the root volume in a load-sharing mirror relationship and the mount operation fails.

**Corrective action** When you add a storage system using clustered Data ONTAP to Backup and Recovery, use the server IP address instead of the Vserver IP address.

## Error writing backup metadata to repository\backups.xml: move failed

**Description** This message occurs when Backup and Recovery attempts to rename a temporary file during the process of creating of a new backups.xml file with updated backup information. The new file is saved to the Backup and Recovery repository folder.

**Corrective action** You must disable any antivirus programs that are running a scan of the Backup and Recovery repository folder. For a typical Windows installation, the repository can be found at the following location:

- C:\Program Files\NetApp\Virtual Storage Console\smvi\server\repository.

## Backup and Recovery unable to discover datastores on a Vserver without a management LIF

**Issue** Running a scheduled backup job fails.

**Cause** This problem occurs when a Vserver without a management LIF is added in Monitoring and Host Configuration. Backup and Recovery cannot resolve this Vserver and is unable to discover any datastores or volumes on the Vserver on which to perform Backup and Recovery tasks.

**Corrective action** You must add a Vserver with a management LIF before you can perform Backup and Recovery tasks.

## VMware vSphere does not remove all snapshot delta disks during a restore operation

<b>Issue</b>	When you delete a VMware snapshot, the snapshot delta disks that are created after quiescing the virtual machine are not removed.
<b>Cause</b>	This problem occurs during a backup of a virtual machine running on Windows 2008 or Windows 2008 R2. Virtual Storage Console for VMware vSphere creates a quiesced VMware snapshot, which results in the creation of snapshot delta disks. When you successfully restore the VM in VSC and then revert to the snapshot taken during the backup process and ultimately delete the VMware snapshot, not all the delta disk files are deleted.
<b>Corrective action</b>	None

## File already exists

<b>Description</b>	This message occurs when using Backup and Recovery to restore a virtual machine on certain versions of clustered Data ONTAP. The restore operation fails when the directory path does not contain <code>/vol/</code> . This issue was found in Data ONTAP 8.1.2P1.
<b>Corrective action</b>	You must specify the full directory path for an NFS volume, such as <code>/vol/XXX/XXX</code> .

## Using ESX hosts with NetApp storage

---

To use ESX hosts with NetApp storage systems, you need to correctly provision storage and configure ESX hosts.

The following sections provide general storage system and ESX host information.

### LUN type guidelines

LUNs must be created with the correct LUN type.

If the LUN will be configured with VMFS, then use the LUN type **vmware**.

If the LUN will be configured with RDM, then use the guest OS for the LUN type. For more information on the LUN type to use, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP software.

### Manually provisioning storage

To configure your storage systems to connect to virtual machines (VMs) running on VMware ESX or ESXi, you must create new volumes, LUNs, and igroups and then map the LUNs to the igroups.

#### Before you begin

You need the FC or iSCSI identifiers of the ESX or ESXi host.

For detailed instructions on how to perform the following steps, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP software.

#### About this task

The Provisioning and Cloning capability can also be used to provision storage.

#### Steps

1. Create an initiator group (igroup) for each VMware ESX or ESXi server using the **vmware** igroup type.

Use the WWPNs for all FC HBAs in the ESX or ESXi host or the iSCSI initiator node name of the ESX or ESXi host when creating the igroup.

For ESX 4.0, 4.1, and 5.0 configurations that support ALUA, enable the ALUA option on the igroup.

**Note: (Data ONTAP operating in 7-Mode)** If your storage systems are running Data ONTAP operating in 7-mode and you are using a Microsoft cluster (Windows Server 2003 MSCS or Server 2008 failover cluster) configuration, do not enable ALUA for LUNs used by guest operating systems. If ALUA is enabled, the cluster loses its persistent reservations during storage faults, causing the cluster service to be unavailable. In addition, you need to avoid ALUA mismatches so that no initiator participates in both ALUA and non-ALUA enabled groups. For more information, see Knowledge base article 2013316 *Microsoft failover clustering running on VMWare ESX 4.x Virtual Machines does not function properly when ALUA is enabled* at [kb.netapp.com/support/index?page=content&id=2013316](http://kb.netapp.com/support/index?page=content&id=2013316).

Check the Interoperability Matrix at [support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability) to see which versions of Data ONTAP are supported with Microsoft cluster (Windows Server 2003 MSCS and Server 2008 failover cluster) configurations.

2. Create the storage for each virtual machine.
  - a) Create one or more volumes to contain the LUNs.

FlexVol volumes are recommended in general, and are required if you are using Snapshot copies.

- b) Create a LUN for the VM's root disk.
- c) Create any additional LUNs needed for the VM's application data.
- d) Map all of the LUNs to the igroup for the ESX or ESXi host.

If you plan to use VMotion to move your guest operating systems from one VMware ESX or ESXi host to another, map the LUN to all hosts in the cluster. The LUN IDs must be identical.

3. Optionally, verify and if necessary correct the alignment of the VMDK partitions.

Sometimes partition misalignment problems can arise, which can lead to performance degradation under very heavy I/O. Depending on your configuration, you might need to align your VMDK partitions to avoid subsequent performance problems.

If you use RDM, and you use the correct guest OS for the LUN type, you should not experience alignment problems.

## How to set up VMware ESX

After creating the necessary LUNs and igroups, and after mapping the LUNs to the igroups, you must configure your host.

For guests in a Microsoft Windows cluster (MSCS) configuration only, you also need to change the path selection policy.

For more information about setting up ESX or ESXi, see the *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) and the VMware product documentation for your version of ESX.

### Related information

*VMware product documentation - [www.vmware.com/support/pubs/vi\\_pubs.html](http://www.vmware.com/support/pubs/vi_pubs.html)*

## Configuring the VMware ESX host

Configuring the VMware ESX or ESXi host requires rescanning the bus, creating a datastore, and creating a new VM.

### Before you begin

You need to create the required LUNs before starting this task.

More information on this task is available in the documentation provided by VMware.

### Steps

1. Rescan the SCSI bus to discover the new LUNs.
  - a) Open the VMware vSphere Client and connect to your ESX or ESXi host.
  - b) On the **Configuration** tab, select **Hardware > Storage Adapters**
  - c) Click **Rescan**.
2. Create a VMFS datastore on the LUN.
3. Create a new VM or add a new disk to an existing VM.

## (Data ONTAP operating in 7-Mode) Manually setting the path selection policy for Microsoft cluster configurations

If the storage system is running Data ONTAP operating in 7-mode and you are using a Microsoft cluster (Windows Server 2003 MSCS or Server 2008 failover cluster) configuration, you should disable ALUA on the igroup and change the path selection policy to FIXED for guest operating systems.

### About this task

For Microsoft Windows guest operating systems in a cluster configuration, always use the FIXED path selection policy and disable ALUA on the igroup for the LUNs. This might require you to manually set the path selection policy.

If ALUA is enabled, the Windows cluster loses its persistent reservations during storage faults, causing the cluster service to be unavailable. When ALUA is disabled, the FIXED path selection policy is required to avoid sending I/O over proxy paths.

The default path selection policy set by VSC for VMware vSphere should be used if the ESX or ESXi host does **NOT** have guest operating systems in a Windows cluster (MSCS or failover cluster) configuration. For ESX 4.0, 4.1, and 5.0 systems, the path selection policy is set to round robin (RR) for ALUA FC configurations and all iSCSI configurations, and set to FIXED for non-ALUA configurations. For ESX 3.5, the default policy is FIXED.

## Steps

1. To manually change the path selection policy, enter the following command on the ESX or ESXi host:

```
esxcli conn_options nmp device setpolicy --device device_name --psp
VMW_PSP_FIXED
```

For more information about the `esxcli` command, see the *vSphere Command-Line Interface Installation and Reference Guide*.

2. To disable ALUA for an igroup, enter the following command at a Data ONTAP command prompt:

```
igroup set igroup_name alua no
```

## Related information

*vSphere Command-Line Interface Installation and Reference Guide* - [www.vmware.com/pdf/vsphere4/r40/vsp\\_40\\_vcli.pdf](http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf)

## Timeout values for guest operating systems

The guest OS (GOS) timeout scripts set the SCSI I/O timeout values for supported Linux, Solaris, and Windows guest operating systems. The timeout values ensure correct failover behavior.

These scripts are provided as .ISO files in the Tools panel of the Monitoring and Host Configuration capability. There are two scripts for each operating system:

- A 60-second script
- A 190-second script

In most cases, the recommended value is 60 seconds. Knowledge base article 3013622, which is online at [kb.netapp.com/support/index?page=content&id=3013622](http://kb.netapp.com/support/index?page=content&id=3013622), contains information you can use when deciding which timeout value to use.

You can mount and run the script from the vSphere client. The Tools panel provides URLs for the scripts.

To get the script containing the timeout values you want for your operating system, you must copy the correct URL from the Tools panel and mount it as a virtual CD-ROM in the virtual machine using the vSphere client. Make sure you install the script from a copy of VSC for VMware vSphere that is registered to the vCenter Server that manages the virtual machine. After the script has been installed, you can run it from the console of the virtual machine.

The section *Installing GOS scripts* contains for detailed steps for performing this task.

## Running the GOS timeout scripts for Linux

The guest operating system timeout scripts set the SCSI I/O timeout settings for RHEL4, RHEL5, RHEL6, SLES9, SLES10, and SLES11. You can specify either a 60-second timeout or a 190-second timeout. You should always run the script each time you upgrade to a new version of Linux.

### Before you begin

You must mount the ISO image containing the Linux script before you can run it in the virtual machine.

### Steps

1. Open the console of the Linux virtual machine and log in to an account with root privileges.
2. Run the `linux_gos_timeout-install.sh` script.

### Result

For RHEL4 or SLES9, a message similar to the following is displayed:

```
Restarting udev... this may take a few seconds.
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For RHEL5 or RHEL6, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev.rules
Hunk #1 succeeded at 333 (offset 13 lines).
Restarting udev... this may take a few seconds.
Starting udev: [ OK ]
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For SLES10 or SLES11, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev-default.rules
Hunk #1 succeeded at 114 (offset 1 line).
Restarting udev ...this may take a few seconds.
Updating all available device nodes in /dev:           done
```

### After you finish

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from *filename.iso***.

## Running the GOS timeout scripts for Solaris

The timeout scripts set the SCSI I/O timeout settings for Solaris 10. You can specify either a 60-second timeout or a 190-second timeout.

### Before you begin

You must mount the ISO image containing the Solaris script before you can run it in the virtual machine.

### Steps

1. Open the console of the Solaris virtual machine and log in to an account with root privileges.
2. Run the `solaris_gos_timeout-install.sh` script.

### Result

For Solaris 10, a message similar to the following is displayed:

```
Setting I/O Timeout for /dev/s-a - SUCCESS!
```

### After you finish

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from *filename.iso***.

## Running the GOS timeout script for Windows

The timeout scripts set the SCSI I/O timeout settings for Windows guest operating systems. You can specify either a 60-second timeout or a 190-second timeout. You must reboot the Windows guest OS for the settings to take effect.

### Before you begin

You must mount the ISO image containing the Windows script before you can run it in the virtual machine.

### Steps

1. Open the console of the Windows virtual machine and log in to an account with Administrator privileges.
2. If the script does not automatically start, open the CD drive and run `windows_gos_timeout.reg`.

The Registry Editor dialog is displayed.

3. Click **Yes** to continue.

The following message is displayed: The keys and values contained in D:\windows\_gos\_timeout.reg have been successfully added to the registry.

4. Reboot the Windows guest OS.

**After you finish**

Unmount the ISO image by clicking the **CD/DVD Connections** icon in the vSphere Client and selecting **CD/DVD Drive 1 > Disconnect from filename.iso**.

## How to identify and fix VMDK partition alignment issues

In some cases, VMDK partitions can become misaligned, leading to performance degradation.

**Note:** This issue may be serious enough to warrant action, but the performance degradation depends on your I/O load and configuration. In many cases, the decrease in performance will be negligible.

Also note that this problem is not unique to NetApp storage platforms.

VMDK partitions need to be aligned at both the VMFS and guest OS levels. For example, you can align the partitions at the VMFS level by selecting the vmware LUN type when creating your LUNs. By doing so, the partitions are aligned to sector 128 or sector 0, depending on whether you use vCenter or vmkfstools to create the VMFS. Regardless, the partitions will be aligned as both are multiples of 4 KB, thereby fulfilling the WAFL read/write requirements.

However, because some operating systems implement sector offsets that are not aligned to 4 KB boundaries, the partitions might still not be aligned at the guest OS level. Therefore, you must manually align the .vmdk files at the guest OS level for VMFS and NFS datastores.

**Note:** If you use RDM and create the LUN with the correct guest OS for the LUN type, then you should not experience alignment issues with the RDM LUNs. The base VMDK might still have an alignment issue.

## Checking VMDK partition alignment with mbralign

You can use either the Optimization and Migration capability or the mbralign tool included with VSC for VMware vSphere to check VMDK partition alignment.

**About this task**

The mbralign tool replaces the mbrscan tool and is effective on -flat.vmdk and fixed .vhd files that are partitioned using a master boot record (MBR) partition table.

**Note:** There is a version mbralign for ESX hosts and one for ESXi hosts. When you download mbralign, you choose the software package based on whether you have an ESX host or an ESXi host. The example in these steps uses the standard ESX mbralign. The mbralign for ESXi does not support all the features used here, such as drive letter restoration.

If you do not want to power down the VM, take one of the following actions:

- Use the Optimization and Migration capability online alignment tool.
  - Note:** For information on using the Optimization and Migration capability tool to scan virtual machines (VM) and perform online alignments, see the section on Optimization and Migration.
- Use `mbralign` and take either a Data ONTAP Snapshot copy of the volume containing the datastore LUN or NFS datastore or a VMware snapshot of the VM in question. Then run `mbrscan` against the copy.

## Steps

1. On the ESX host console, change to the directory where `mbralign` is installed.
2. Enter the following command on the ESX host console:

```
mbralign { --scan all | filename }
```

The `--scan all` option scans all `-flat.vmdk` files.

*filename* specifies the name of a single file to scan.

The command displays information indicating whether the VMDK partition is correctly aligned.

```
# /opt/netapp/santools/mbralign --scan all
Building file list...
/vmfs/volumes/4c604abb-e41943c0-a81f-001b7845166c/
win2k3sp2_64v_esx-09/
  win2k3sp2_64v_esx-09_1-flat.vmdk P1 lba:63 Aligned: No

/vmfs/volumes/4c604abb-e41943c0-a81f-001b7845166c/
win2k3sp2_64v_esx-09/
  win2k3sp2_64v_esx-09_2-flat.vmdk P1 lba:63 Aligned: No
```

## VMDK partition alignment with `mbralign` overview

The `mbralign` tool enables you to correct misaligned VMDK partitions.

**Note:** The Optimization and Migration capability allows you to perform online alignments. In addition, you can also use the VMware vCenter Converter to perform offline alignments.

The `mbralign` tool works on primary VMDK partitions with a master boot record (MBR) partition table. If there are multiple partitions, the partitions must be in order on the disk.

Starting with the `mbralign` tool in VSC 2.1 for VMware vSphere, the 1-TB size limit on VMDKs has been removed. You can use `mbralign` with any VMDK size supported by VMware.

The `mbralign` tool has the following requirements:

- The destination datastore must have enough free space for a full copy of the `-flat.vmdk` file.

- GRUB-booted Linux guest operating systems need to have GRUB reinstalled after aligning the boot partition.
- The virtual machine using the VMDK must be shut down when running mbralign. If you use the feature to preserve Windows drive letter mapping, the mbralign program shuts down the VM after collecting drive letter information.
- For ESX clusters, you must run the mbralign program on the ESX or ESXi host where the VM is currently registered. For NFS datastores, the mbralign program cannot detect if the VM is powered down if the VM is running on another ESX or ESXi host.

The mbralign tool has the following limitations:

- Large NFS filesystems mounted on the ESX or ESXi host can greatly increase the time required to run mbralign, because mbralign scans them for VMDKs. Temporally unmounting large filesystems that do not contain VMDKs needing alignment should improve performance.
- VMDKs containing Windows dynamic disks or GPT partitions are not supported. The Windows operating system must be installed on the C: drive.
- VMDKs containing Linux LVM are not supported.
- The mbralign tool does not fix misalignment associated with linked clone or VMware delta files; it can only fix the base VMDK. All VMware snapshots and linked clones must be removed from the VM using the disk being aligned.
- If the storage system volume containing the VMDK is included in Snapshot copies, the alignment process can cause the Snapshot copies to grow very large (up to twice the size of the volume). The space used for the Snapshot copies is not released when the backup file is deleted. Deleting the Snapshot copies from before the alignment process releases the space used.
- Media devices, such as CD-ROM or DVD drives used by the VM, might not map to their Windows original drive letters after running the mbralign program. This can happen when there are multiple media drives or when the drive contains media.
- Do not use the `--force` option of the mbralign command on VMDKs for virtual machines running Windows 7, Windows Server 2008, or Windows Server 2008 R2. This can corrupt the boot LUN.
- Do not use mbralign with Solaris guest operating systems; it cannot correctly align them. Solaris ZFS file systems should not have alignment issues.

Starting with the mbralign tool in VSC 2.1 for VMware vSphere, you can now preserve the original drive mapping of Windows disks. Earlier versions of mbralign could only ensure that the C:\ drive mapped to the correct partition.

- The Windows virtual machine must be running when you start mbralign so that mbralign can collect the drive letter mapping information. The mbralign program prompts lead you through the process of shutting down the VM after collecting drive letter information and then starting the actual alignment process.
- The Windows operating system folder must be on the C:\ drive. For example, C:\Windows.
- The VMware tools package must be installed on the VM. Be sure to use the version of VMware tools that matches the ESX version on which the VM is running. See your VMware documentation for instructions on installing the VMware tools in the guest operating system.

- For Windows Server 2000, you must install the Windows 2000 Resource Kit Tools for administrative tasks, which includes the diskpart program. Be sure to take the default installation location.
- For 64-bit Windows Server 2003 and Windows XP guests, install Windows hotfix KB 942589 on the VM.
- Note that this drive letter mapping process does not apply to Linux virtual machines.

You can use the `--bs` option of the `mbralign` command to increase the default 8 KB block size to a larger value (16, 32, 64, 128, or 1024) using the `--bs` option of the `mbralign` command. Doing this improves performance of VMFS datastores. A larger block size generally means that there are fewer reads and writers.

To see all of the command options, you can use `mbralign --help` command, or you can refer to the `mbralign` man page.

### Related information

[Windows 2000 Resource Kit Tools \(KB 927229\) - support.microsoft.com/kb/927229](http://support.microsoft.com/kb/927229)  
[Hotfix KB 942589 - support.microsoft.com/kb/942589](http://support.microsoft.com/kb/942589)

## Fixing VMDK partition alignment using mbralign

If a VMDK partition is misaligned, you can align the partition using the `mbralign` tool included with VSC for VMware vSphere.

### Before you begin

Be aware that `mbralign` can take anywhere from 1 or 2 minutes to several minutes per gigabyte of storage in the affected files.

**Note:** The Optimization and Migration capability allows you to perform online alignments. In addition, you can use the VMware vCenter Converter to perform offline alignments.

### Steps

1. Remove any VMware snapshots from the VM that is to be realigned.
2. Temporarily unmount large NFS filesystems that do not contain VMDKs needing alignment from the ESX or ESXi host.
3. Shut down the VM.
4. For Linux VMs, and Windows VMs with only a C:\ drive, shut down the VM.

For a Windows VM with multiple drive letters mapped, the VM must be running so that `mbralign` can collect the drive letter information.

5. On the ESX or ESXi host console, change to the directory containing the `.vmdk` file for the VM.
6. Enter the following command:

```
path/mbralign name.vmdk
```

*path* is the path where the `mbralign` program is installed.

*name* is the name of the VMDK file being aligned.

7. If prompted, enter **yes** for a Windows VM to automatically collect and restore drive letters. Enter the Windows Administrator credentials for the VM.

The VM is automatically shut down after the drive letter information is collected.

8. When prompted Are you sure that no snapshots/linked clones exist for this vmdk? Enter **y**.

**Attention:** The use of `mbralign` on a VMDK file that has a snapshot or linked clone associated with it can result in unrecoverable data loss or data corruption.

9. For Windows guest operating systems for which you are not using the drive letter restore feature, restart the VM and verify that the guest operating system boots successfully.
10. For Linux guest operating systems using the GRUB boot loader, reinstall GRUB before restarting the VM.
11. After verifying the VM has booted and is operating correctly, delete the backup files created by `mbralign`.

These files are saved in the same directory as the `.vmdk` file and have names ending in `-mbralign-backup`.

In the following example, some output has been deleted for clarity and the lines have been truncated to fit the page.

```
[root@esxhost1 VM2]# /opt/netapp/santools/mbralign VM2_1.vmdk
The vmdk file looks like it belongs to a Windows Virtual Machine:
VM2.
Would you like to automatically align the vmdk and restore the
original drive letters?
If this is not a Windows Virtual Machine, or if this is a VM part
of a Microsoft Cluster,
select no (yes/no) yes
This VM also has the following vmdk files associated to it that
also need to be aligned.
  /vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2.vmdk
  /vmfs/volumes/4bb1f98a-a2c428cc-f253-001e4f2f3dd3/VM2/VM2_2.vmdk
Do you want to also align /vmfs/volumes/4bb1f98a-a2c428cc-
f253-001e4f2f3dd3/VM2/VM2.vmdk? (yes/no/all) yes
Do you want to also align /vmfs/volumes/4bb1f98a-a2c428cc-
f253-001e4f2f3dd3/VM2/VM2_2.vmdk? (yes/no/all) yes
Checking the power state of the VM.....ON
Please provide Administrator credentials or [enter] to skip:
ESX Server Username: root
ESX Server Password:
VM Domain:
VM Username: Administrator
VM Password:
...
```

```

Alignment complete for VM2_1.vmdk
The next vmdk to align is: /vmfs/volumes/4bb1f98a-a2c428cc-
f253-001e4f2f3dd3/VM2/VM2_2.vmdk
...
Press enter when you are ready to power on the VM [enter]
Powering on the VM.....DONE
Establishing connection to VM.....DONE
Collecting volume information.....DONE
Setting drive
letters.....DONE
Removing temporary files.....DONE

```

```

[root@esxhost1 V_M_2]# /opt/netapp/santools/mbralign V_M_2.vmdk
  Part  Type  old LBA  New Start LBA  New End LBA  Length in KB
    P1    83      63           64      208846      104391
    P2    8e   208845      208856      16771871      8281507

NOTICE:
This tool does not check for the existence of Virtual Machine
snapshots or linked clones.
The use of this tool on a vmdk file that has a snapshot or linked
clone associated with it
can result in unrecoverable data loss and/or data corruption.
Are you sure that no snapshots/linked clones exist for this vmdk?
(y/n)y

Creating a backup of V_M_2.vmdk
Creating a backup of ./V_M_2-flat.vmdk
Creating a copy the Master Boot Record
Working on partition P1 (2): Starting to migrate blocks from 32256
to 32768.
...

```

## Reinstalling GRUB for Linux guests after running mbralign

After running `mbralign` on disks for Linux guest operating systems using the GRUB boot loader, you must reinstall GRUB to ensure that the guest operating system boots correctly.

### Before you begin

The `mbralign` program has completed on the `.vmdk` file for the virtual machine.

### About this task

This topic applies only to Linux guest operating systems using the GRUB boot loader and SystemRescueCd.

### Steps

1. Mount the ISO image of Disk 1 of the installation CDs for the correct version of Linux for the virtual machine.

2. Check the box for **Connected** (if the VM is running) or **Connect at power on** (if the VM is not running).
3. Open the vSphere Client remote console for the VM.
4. If the VM is running and hung at the GRUB screen, click in the display area to make sure it is active, then press Ctrl-Alt-Insert to reboot the VM. If the VM is not running, start it, and then immediately click in the display area to make sure it is active.
5. As soon as you see the VMware BIOS splash screen, press the Esc key once.

The boot menu is displayed.

6. At the boot menu, select CD-ROM.
7. At the Linux boot screen, enter:  
**linux rescue**
8. Take the defaults for Anaconda (the blue/red configuration screens). Networking is optional.
9. Launch GRUB by entering:

**grub**

10. If there is only one virtual disk in this VM, or if there are multiple disks, but the first is the boot disk, run the following GRUB commands:

```
root (hd0,0)
```

```
setup (hd0)
```

```
quit
```

If you have multiple virtual disks in the VM, and the boot disk is not the first disk, or you are fixing GRUB by booting from the misaligned backup VMDK, enter the following command to identify the boot disk:

```
find /boot/grub/stage1
```

Then run the following commands:

```
root (boot_disk,0)
```

```
setup (boot_disk)
```

```
quit
```

*boot\_disk* is the disk identifier of the boot disk.

11. Press Ctrl-D to log out.

Linux rescue shuts down and then reboots.

## Copyright information

---

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**RESTRICTED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

---

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bypass, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

## How to send your comments

---

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

- .bat [196](#)
- .cmd [196](#)
- .sfr file
  - manually creating [129](#)
  - Restore Agent
    - creating a .sfr file for [129](#)
    - single file restore file
    - See .sfr file
- /etc/smvi.override file
  - changing the SnapManager for Virtual Infrastructure server IP address [130](#)

## A

- Access Management Console [79](#)
- accounts
  - configuring with RBAC [43](#)
- active backup
  - suspending [121](#)
- actual alignments [101](#)
- Add Single File Restore Session wizard
  - creating a limited self-service restore session [134](#)
  - creating a self-service restore session [131](#)
- Admin Assisted
  - restore sessions [130](#)
- administrator-assisted file restoration
  - of virtual disk files [128](#)
- administrator-assisted restore sessions, about [128](#)
- aggregate
  - adding to resources [80](#)
- alert status
  - getting more information [199](#)
- alignments
  - actual [101](#)
  - functional [101](#)
  - identifying and fixing partition [212](#)
  - offline [101](#)
  - online [101](#)
- ALUA
  - enabling [206](#)
- API
  - Provisioning and Cloning methods [158](#)
- APIs
  - VMware vCloud [156](#)
  - work with SOAP [156](#)

- architecture
  - VSC for VMware vSphere [13](#)

## B

- Backup and Recovery
  - configuring [113](#)
  - defined [112](#)
  - installing or uninstalling [21](#)
  - lock management [14](#)
  - SnapMirror requirements [113](#)
  - Snapshot and SnapRestore requirements [112](#)
  - upgrade considerations [26](#)
- Backup and Recovery capability
  - authentication credentials [113](#)
  - network connection to Restore Agent [127](#)
  - restoring data from backups [122](#)
- backup copies
  - finding [124](#)
  - locating [123](#)
  - restoring data from [123](#)
  - unmounting from ESX server [127](#)
- backup jobs
  - datastores or datacenters [118](#)
  - deleting [121](#)
  - for virtual machines [117](#)
  - specifications for adding [116](#)
  - starting a one-time backup [120](#)
  - suspending active [121](#)
  - troubleshooting [202](#)
- backup mount error [204](#)
- Backup Now window [120](#)
- backup operations
  - managing [116](#)
- backup retention
  - specifying [116](#)
- backup scripts
  - controlling the runtime of [202](#)
- Backup tab [133, 136](#)
- Backup wizard [117, 118](#)
- backups
  - editing a backup job [120](#)
  - failure [204](#)
  - mounting [126](#)
  - restoring data from [122](#)

- restoring data from backups with failed VMware consistency snapshots [123](#)
- resuming a suspended backup job [122](#)
- selecting the restore destination [123](#)
- baseline [78](#)
- boot loader
  - reinstalling for Linux after running mbralign [217](#)

## C

- C# [197](#)
- capabilities
  - accessing from VSC for VMware vSphere [14](#)
  - interaction [12](#)
  - VSC for VMware vSphere [11](#)
- Citrix [69, 75, 79, 80, 91, 168, 193](#)
- Citrix XenDesktop
  - adding connection brokers [75](#)
- clone
  - virtual machines [69](#)
- clone creation engine [156](#)
- clone data [75](#)
- clone operation [75](#)
- cluster [69, 161](#)
- communication
  - troubleshooting [201](#)
- Communities
  - provide information about VSC for VMware vSphere [198](#)
- concurrent VMware snapshots
  - controlling the number created or deleted [202](#)
- config\_mpath
  - setting path selection policy for MSCS guests [208](#)
- configuration
  - of Backup and Recovery [113](#)
  - troubleshooting [202](#)
- configuration files
  - for restore sessions [132, 135](#)
  - generated upon creation of restore sessions [128](#)
  - restore sessions [130](#)
  - used for single file restore [134](#)
- connection broker
  - adding [75](#)
- Connection brokers panel [75](#)
- connectivity
  - troubleshooting [201](#)
- consistency snapshots [123](#)
- controller
  - removing skipped or unmanaged [52](#)
  - supports direct connections to Vserver [50](#)

- supports vFiler unit, Vserver tunneling [50](#)
- controller name
  - correcting when unknown [52](#)
- controlling what is displayed
  - using Inventory panel [53](#)
- copy/clone offload engine [156](#)
- Create Rapid Clones Wizard [79](#)
- createClones [176, 177, 180, 181, 187](#)
- createDatastore [187](#)
- credentials
  - default for storage controllers [47](#)
  - setting for storage controller [47](#)
  - upgrade considerations [26](#)
  - using RBAC [43](#)
- csv [69, 75, 79, 98, 99](#)
- custom user accounts
  - configuring using RBAC [43](#)

## D

- data collection
  - log files [200](#)
- Data Collection panel
  - collecting diagnostic data [61](#)
- data collection programs
  - changing service credentials [62](#)
- datacenter [69, 161](#)
- datacenters
  - creating backup copies [118](#)
- datastore [98, 157, 158, 161, 163, 164, 169, 171, 181, 185, 187](#)
- datastore management engine [156](#)
- datastores
  - cloning [69](#)
  - creating [208](#)
  - creating backup copies [118](#)
  - mounting [89](#)
  - NFS indirect path [56](#)
  - reclaiming space [78](#)
  - resizing [90](#)
  - restoring [123, 124](#)
  - searching for backup copies of [124](#)
  - where to restore a backup [123](#)
- deduplication
  - enabling [89](#)
- default credentials
  - for storage controllers [47](#)
- default values in smvi.override
  - modifying [202](#)
- delta disks

- not removed during restore [205](#)
- Desktop Studio [80](#)
- destroy datastores [91](#)
- diagnostic data
  - collecting [61](#)
- discovery
  - correcting unknown controller name [52](#)
  - enabling for vFiler units [49](#)
  - storage system [46](#)
- disk files
  - restoring [123](#)
- disk partitions
  - checking and fixing alignment [63, 65](#)
- Disk tab [133, 136](#)
- Disk.QFullSampleSize [54](#)
- Disk.QFullThreshold [54](#)
- display
  - limiting using Inventory panel [53](#)
- drive letters
  - changing [133, 136](#)

**E**

- e-mails
  - restore session confirmations [131](#)
- Element Manager
  - launching [59](#)
- email alerts
  - specifying for backup jobs [116](#)
- email messages
  - link to Restore Agent download URL [134](#)
- email notification error [203](#)
- Emulex FC HBA timeouts [54](#)
- error logs
  - locations [203](#)
  - reviewing [203](#)
- error messages
  - file already exists [205](#)
  - mounting NFS datastore backup fails [204](#)
  - move failed [204](#)
- ESX [91, 157, 158, 160, 161, 163, 164, 166, 168, 169, 171, 176, 177, 179–181, 185](#)
- ESX host
  - displaying settings [199](#)
  - multipathing and timeout settings [54](#)
  - restoring virtual disks on VMFS datastores [125](#)
  - setting up [207](#)
  - timeout values [209](#)
- ESX, ESXi host settings [54, 56](#)
- ESXi

- enable ssh [63](#)
- event logs
  - locations [203](#)
  - reviewing [203](#)
- expiration times
  - of restore sessions, changing [129](#)

## F

- failed consistency snapshots [123](#)
- failover cluster
  - setting path selection policy for [208](#)
- FCP [185, 187](#)
- fileCloneOffload [173, 175](#)
- fileCopyOffload [173, 175](#)
- FilerView
  - launching [59](#)
- fixed
  - path selection policy [208](#)
- Flash Accel
  - used with VSC for VMware vSphere [58](#)
- FlexClone [81, 91, 157, 171](#)
- FlexClone license [132, 135](#)
- FQDN [193](#)
- functional alignments [101](#)

## G

- General tab
  - single file restore settings [129](#)
- getFileOpOffloadStatus [175](#)
- GRUB
  - reinstalling for Linux after running mbralign [217](#)
- guest OS
  - installing scripts [66](#)
  - setting timeouts for Linux [210](#)
  - setting timeouts for Solaris [211](#)
  - setting timeouts for Windows [211](#)
  - timeout values [209](#)

## H

- Help, online [15](#)
- host
  - multipathing and timeout settings [54](#)
- HTTP [190](#)
- httpd.admin.enable option [49](#)
- HTTPS [190](#)

- ## I
- igroup
    - creating [206](#)
  - in-place restore [123](#)
  - installation
    - of Restore Agent [132](#), [135](#)
    - of VSC for VMware vSphere using the installation wizard [19](#)
  - installing
    - guest operating system (GOS) scripts [66](#)
    - mbralign [63](#), [65](#)
    - VSC for VMware vSphere in silent mode [20](#)
  - Inventory panel
    - limiting what is displayed [53](#)
  - IPv6
    - Provisioning and Cloning issue [68](#)
  - iSCSI [185](#), [187](#)
  - issues
    - datastore discovery [204](#)
    - mounting NFS datastore backup fails [204](#)
    - VMware snapshot delta disks not removed during restore [205](#)
- ## J
- Job Properties dialog box
    - using to edit a backup job [120](#)
- ## K
- kaminoprefs file [91](#), [97](#)
- ## L
- license requirements
    - single file restore feature [132](#), [135](#)
  - limited self-service restore sessions
    - about [128](#)
  - limiting what is displayed
    - using Inventory panel [53](#)
  - Linux
    - reinstalling GRUB after running mbralign [217](#)
    - setting timeouts for guest OS [210](#)
  - linux\_gos\_timeout-install.iso
    - guest OS tool [210](#)
  - Load Configuration window
    - in Restore Agent [132](#)
  - locations of log files [203](#)
  - lock management [14](#)
  - Lock management
    - Optimization and Migration [102](#)
  - log configuration file
    - modifying [97](#)
  - log files
    - collecting [200](#)
  - log4j [97](#)
  - log4j.properties
    - log configuration file [97](#)
  - logs
    - collecting [200](#)
    - errors [203](#)
    - events [203](#)
    - locations [203](#)
    - message [203](#)
    - reviewing [203](#)
    - troubleshooting [203](#)
- ## LUN
- creating [206](#)
  - type [206](#)
- ## M
- MAC address [191](#)
  - Managed Object Browser [156](#)
  - management
    - of backup operations [116](#)
  - MBR tools
    - enable ssh for ESXi [63](#)
  - mbralign
    - checking partition alignment [212](#)
    - fixing partition alignment [215](#)
    - installing [63](#), [65](#)
    - overview [213](#)
  - mbrscan
    - replaced by mbralign [212](#)
  - memory requirements
    - VSC for VMware vSphere [18](#)
  - message logs
    - locations [203](#)
    - reviewing [203](#)
  - Monitoring and Host Configuration
    - upgrade considerations [26](#)
  - mount expiration
    - changing the default time [131](#), [134](#)
  - MSCS
    - setting path selection policy for [208](#)
  - multihost configurations
    - backing up [112](#)
    - recovering [112](#)

- multipathing
  - configuring ESX host [54](#)
- MultiStore
  - display differences with vFiler units [60](#)
  - enabling discovery of vFiler units on private networks [49](#)

## N

- Net.TcpipHeapMax [54](#)
- Net.TcpipHeapSize [54](#)
- NetApp Communities
  - API questions [156](#)
  - See also* Communities
- network configuration file [98](#)
- NFS
  - indirect path [56](#)
  - reclaiming space [78](#)
- NFS datastore backup
  - mounting a backup fails [204](#)
- NFS datastores
  - Optimization and Migration considerations [103](#)
- NFS paths
  - changing to direct access [57](#)
- NFS VAAI Plug-in
  - installing [58](#)
- NFS.HeartbeatFrequency [54](#)
- NFS.HeartbeatMaxFailures [54](#)
- NFS.HeartbeatTimeout [54](#)
- NFS.MaxVolumes [54](#)
- NTFS
  - reclaiming space [78](#)
- nvram [91](#)

## O

- object
  - object
    - vSphere [33, 34](#)
    - storage system [33, 34](#)
    - vSphere [33, 34](#)
- OnCommand System Manager
  - See* System Manager
- one-time backup jobs
  - starting [120](#)
- online Help [15](#)
- Optimization and Migration
  - actual alignments [101](#)
  - considerations when using NFS datastores [103](#)
  - features [100](#)

- functional alignments [101](#)
- important notes [102](#)
- lock management [14](#)
- Lock management [102](#)
- migrating multiple virtual machines [102](#)
- offline alignments [101](#)
- online alignments [101](#)
- SDRS [102](#)
- VAAI extended copy [102](#)
- VMware Storage vMotion [102](#)
- Windows 2008 R2 SP1 [102](#)
- out-of-place restore [123](#)
- Overview panel
  - limiting what is displayed [53](#)

## P

- parameters
  - ESX, ESXi host [54](#)
  - UNMAP [56](#)
- partition alignment
  - checking with mbralign [212](#)
  - fixing with mbralign [215](#)
  - identifying and fixing [212](#)
- partitions
  - checking and fixing alignment [63, 65](#)
- path selection policy
  - setting for MSCS and failover cluster [208](#)
- paths
  - changing to direct NFS paths [57](#)
- performing online alignments [107](#)
- permission
  - vCenter Server [33, 34](#)
- port groups
  - changing the network connection for [130](#)
  - management of [127, 128](#)
- preferences file [91](#)
- privileges
  - example of assigning privileges [35](#)
  - example of View privilege [40](#)
  - native vCenter Server [33, 34](#)
  - product level [40](#)
  - vCenter Server [39](#)
  - Virtual Storage Console [40](#)
  - VSC [39](#)
  - VSC specific [33, 34](#)
- Provisioning and Cloning
  - about [68](#)
  - accessing features [68](#)
  - cloning virtual machines [69](#)

- IPv6 not supported [68](#)
- lock management [14](#)
- mounting datastores [89](#)
- reclaiming space [78](#)
- upgrade considerations [26](#)
- using the wsimport tool [196](#)
- Provisioning and Cloning log files [97](#)
- Provisioning and Cloning methods [158](#)
- provisioning storage [206](#)
- PXE [91](#)

## Q

- QLogic
  - FC HBA timeouts [54](#)
  - iSCSI HBA IP\_ARP\_Redirect [54](#)
  - iSCSI HBA timeouts [54](#)

## R

- rapid cloning
  - fails if permission incorrectly assigned [202](#)
- RBAC
  - about [31](#)
  - configuring [43](#)
  - Data ONTAP [41](#)
  - Data ONTAP privileges [31](#)
  - Data ONTAP roles [42](#)
  - guidelines for standard VSC roles [38](#)
  - standard VSC roles [37](#)
  - upgrade considerations [26](#)
  - vCenter privileges [31](#)
  - vCenter Server [32](#)
- RDM [160](#)
- reclaiming space
  - using Provisioning and Cloning [78](#)
- redeploying clones
  - fails if permission incorrectly assigned [202](#)
- redeployVMs [176](#), [180](#), [181](#)
- registering
  - Virtual Storage Console with vCenter Server [23](#)
- Release Notes
  - checking [198](#)
- Remove Controller command [52](#)
- required ports
  - firewall requirements [24](#)
  - Virtual Storage Console [24](#)
- rescan SCSI bus [208](#)
- resizing datastores [90](#)
- resource pool [69](#)

- resources
  - discovering and adding [51](#)
- restoration
  - datastore [124](#)
- Restore Agent
  - changing the download URL [129](#)
  - clearing the configuration [133](#), [136](#)
  - disk partitions [133](#), [136](#)
  - downloading the application software [134](#)
  - installing [132](#), [135](#)
  - link to software download [131](#)
  - loading the configuration file [132](#)
  - network connection to Backup and Recovery capability [127](#)
  - post-installation configuration changes [127](#)
  - restoring virtual disk files [133](#), [136](#)
  - software requirements of [132](#), [135](#)
- restore operations
  - example of limited self-service [134](#)
  - from VMDKs [125](#)
  - restoring from backups with failed VMware consistency snapshots [123](#)
  - troubleshooting [202](#)
  - types of [123](#)
  - where to restore a backup [123](#)
- Restore panel
  - purpose [123](#)
  - searching for backup copies [124](#)
- restore sessions
  - Admin Assisted [130](#)
  - changing expiration times [129](#)
  - configuration files [130](#)
  - creating a new session after changing port group settings [127](#)
  - example of a limited self-service restore session [134](#)
  - types of [128](#)
- restore sessions,
  - example of a self-service restore session [131](#)
- Restore wizard
  - restoring a virtual machine or its disk files [125](#)
- restoring virtual data
  - of virtual disk files [127](#)
- resumption
  - of suspended backup job [122](#)
- role-based access control
  - .See RBAC
- roles
  - configuring with RBAC [43](#)

- S**
- scanning datastores [105](#)
- scanning with the Optimization and Migration capability [105](#)
- scheduled backup jobs
  - deleting [121](#)
- scheduled backups
  - editing [120](#)
- scripts
  - choosing for backup jobs [116](#)
- scripts, guest operating system (GOS)
  - installing [66](#)
- SCSI bus
  - rescan [208](#)
- SDRS
  - Optimization and Migration [102](#)
- Search field
  - using [124](#)
- searching for backup copies
  - steps [124](#)
- security
  - configuring using RBAC [43](#)
- self-service restore sessions
  - about [128](#)
- service account
  - changing credentials [62](#)
- settings
  - ESX, ESXi host [54, 56](#)
- Setup panel
  - single file restore settings [129](#)
- SFR files
  - clearing [133, 136](#)
  - loading with Restore Agent [132](#)
  - See also* configuration files
- show details
  - displaying ESX host settings [199](#)
- single file restore
  - recovery of virtual disk files [127](#)
- single file restore feature
  - license requirements [132, 135](#)
  - prerequisites to using [129](#)
- Single File Restore panel
  - changing the type of restore session [130](#)
  - checking the port group settings [127](#)
  - viewing the restore sessions list [130, 131, 134](#)
- skipped controller
  - removing [52](#)
- SMVI log viewer
  - enabling network adapters [203](#)
  - smvi.override file
    - modifying [202](#)
    - purpose [202](#)
- SnapManager for Virtual Infrastructure server
  - manually changing the IP address [130](#)
- SnapMirror requirements
  - Backup and Recovery [113](#)
- snapshot [91, 158, 160, 181](#)
- Snapshot and SnapRestore requirements
  - Backup and Recovery [112](#)
- snapshot autodelete [81](#)
- snapshot delta disks
  - not removed during restore [205](#)
- snapshots [123](#)
- SOAP
  - exposes APIs [156](#)
- software requirements
  - for Restore Agent [132, 135](#)
- Solaris
  - setting timeouts for guest OS [211](#)
- solaris\_gos\_timeout-install.iso
  - guest OS tool [211](#)
- spanned entities
  - backing up [116](#)
- ssh
  - enabling for ESXi [63](#)
- SSL certificate
  - regenerating [24](#)
- status
  - getting more information for alert [199](#)
- status reason
  - displaying column [199](#)
- storage controller
  - removing skipped or unmanaged [52](#)
  - setting default credentials [47](#)
  - using default credentials [47](#)
- storage controller exports file [91](#)
- storage controllers panel [80](#)
- Storage Distributed Resource Scheduler
  - See* SDRS
- storage resources
  - discovering and adding [51](#)
- storage system
  - assigning permissions [33, 34](#)
  - discovery [46](#)
  - managing with Element Manager [59](#)
  - managing with FilerView [59](#)
  - managing with System Manager [59](#)
- storage systems
  - adding [113](#)

- configuring using RBAC [43](#)
  - RBAC privileges [31](#)
- supported configurations
  - memory requirements [18](#)
- SUSE Linux
  - reinstalling GRUB after running mbralign [217](#)
- suspended backup job
  - resuming [122](#)
- suspension of backup jobs [121](#)
- sysprep answer file [76](#)
- System Manager
  - managing storage controllers [59](#)

## T

- the managed object reference of the requested object based on name and type.vCenter server [179](#)
- thick volume [81](#)
- thin provisioned LUN [81](#)
- timeout settings
  - configuring ESX host [54](#)
- timeout values
  - recommended values [209](#)
  - setting for guest OS [209](#)
- timeouts
  - ESX, ESXi host [54, 56](#)
- tools
  - partition alignment [63, 65](#)
  - setting Linux guest OS timeouts [210](#)
  - setting Solaris guest OS timeouts [211](#)
  - setting Windows guest OS timeouts [211](#)
- troubleshooting
  - cannot find file [205](#)
  - checking Release Notes [198](#)
  - collecting diagnostic data [61](#)
  - collecting log files [200](#)
  - email notification error [203](#)
  - error writing backup metadata to repository [204](#)
  - file already exists [205](#)
  - logs [203](#)
  - mounting NFS datastore backup fails [204](#)
  - move failure [204](#)
  - partition alignment [212](#)
  - unable to communicate with the server [201](#)
  - unable to discover datastores [204](#)
  - VMware snapshot delta disks not removed during restore [205](#)
- tunneling
  - supported for vFiler units [115](#)
  - supported for vFiler units, Vservers [50](#)

## U

- unable to communicate with the server
  - troubleshooting [201](#)
- unable to discover datastores
  - on a Vserver [204](#)
- uninstalling
  - VSC for VMware vSphere using Add/Remove Programs [29](#)
  - VSC for VMware vSphere using the CLI [30](#)
- unknown
  - controller name [52](#)
- unmanaged controller
  - removing [52](#)
- update command [46](#)
- updating
  - resource information [51](#)
- upgrading VSC for VMware vSphere
  - considerations [26](#)
- URL
  - Restore Agent downloads
    - changing [129](#)
- user accounts
  - creating custom users [114](#)
  - creating roles, groups, and users [114](#)
- user accounts, creating roles, groups, and users [115](#)
- user credentials
  - backup job [116](#)
- user name
  - configuring custom with RBAC [43](#)
- user privileges
  - controlling using RBAC [31](#)
- UUID [91, 98](#)

## V

- VAAI extended copy
  - Optimization and Migration [102](#)
- vApp [69](#)
- vApps
  - provisioning and cloning [156](#)
- vCenter inventory
  - RBAC privileges [31](#)
- vCenter Inventory [156](#)
- vCenter object
  - RBAC privileges [31](#)
- vCenter server [158, 160, 161, 163, 164, 166, 168, 169, 171, 176, 177](#)
- vCenter Server
  - permission [33, 34](#)

- privileges [39](#)
- registering Virtual Storage Console with [23](#)
- standard VSC roles [37, 38](#)
- vCenter Servers, configuration of [114](#)
- vCenter task [158, 160, 163, 164, 171](#)
- vCloud tenants
  - VMware vCloud [156](#)
- vFiler unit
  - display differences with physical storage controllers [60](#)
  - enabling discovery [49](#)
  - tunneling supported [50](#)
- vFiler units
  - discovering on private networks [49](#)
  - discovery [115](#)
  - tunneling [115](#)
- VI SDK [158](#)
- View privilege
  - example [35, 40](#)
- Virtual Center [158, 160](#)
- Virtual Desktops page [79](#)
- virtual disk files
  - administrator-assisted restoration of [128](#)
  - recovery using Restore Agent [133, 136](#)
- virtual entities
  - backing up [116](#)
- virtual hard disk [185](#)
- virtual hard drive [157, 160](#)
- virtual hard drives [181](#)
- virtual machine
  - cloning [69](#)
- virtual machine disk files
  - where to restore a backup [123](#)
- virtual machines
  - creating backup copies [117](#)
  - migrating in a group [109](#)
  - migrating with Optimization and Migration [102](#)
  - restarting [125](#)
  - restoring [123, 125](#)
  - searching for backup copies of [124](#)
  - where to restore a backup [123](#)
- virtual networks
  - connection to [128](#)
  - relationship of port groups to network connectivity [127](#)
- Virtual Storage Console
  - Backup and Recovery [112](#)
  - CLI [136](#)
  - configuration [46](#)
  - example of product privilege [35, 40](#)
  - firewall port requirements [24](#)
  - launching the VSC CLI [137](#)
  - memory requirements [18](#)
  - privileges [39, 40](#)
  - registering with vCenter Server [23](#)
  - required ports [24](#)
  - standard roles [37, 38](#)
- Virtual Storage Console for VMware vSphere service
  - changing credentials [62](#)
- VM BIOS file [91](#)
- vmdk [171, 185](#)
- VMDK [69](#)
- VMDK partition alignment
  - checking with mbralign [212](#)
  - fixing with mbralign [215](#)
  - identifying and fixing [212](#)
- VMDKs
  - restoring [125](#)
- VMFS [69, 157, 163, 171, 185, 187](#)
- VMFS datastores
  - restoring virtual disks on [125](#)
- VMkernel [91](#)
- VMware [191](#)
- VMware Session ID [180](#)
- VMware snapshot delta disks
  - not removed during restore [205](#)
- VMware snapshots [123](#)
- VMware Storage vMotion
  - Optimization and Migration [102](#)
- VMware vCenter SDK [180](#)
- VMware vCloud
  - accessing API documentation [156](#)
  - APIs [156](#)
- VMware VI API [181, 185](#)
- VMware VI SDK [156](#)
- VMware View [76, 193](#)
- VMware View Server
  - adding connection brokers [75](#)
- vmx [185](#)
- VMX Path [98](#)
- volume
  - adding to resources [80](#)
  - creating [206](#)
- volume autogrow [81](#)
- VSC CLI
  - launching [137](#)
- VSC for VMware vSphere
  - accessing capabilities [14](#)
  - additional installation checks [17](#)
  - architecture [13](#)

- capabilities [11](#)
  - Data ONTAP RBAC roles [42](#)
  - installation overview [16](#)
  - installing in silent mode [20](#)
  - installing or uninstalling Backup and Recovery [21](#)
  - installing using installation wizard [19](#)
  - interaction between capabilities [12](#)
  - lifecycle management for VMware environments [11](#)
  - limitations with direct connections to Vservers [50](#)
  - lock management [14](#)
  - management console for Flash Accel [58](#)
  - Monitoring and Host Configuration [45](#)
  - Optimization and Migration capability [100](#)
  - plug-in functions [11](#)
  - Provisioning and Cloning capability [68](#)
  - regenerating an SSL certificate [24](#)
  - support for Data ONTAP RBAC [41](#)
  - support for vCenter Server RBAC [32](#)
  - supported configurations [18](#)
  - uninstalling using Add/Remove Programs [29](#)
  - uninstalling using the CLI [30](#)
  - upgrade considerations [26](#)
  - upgrading the software [28](#)
  - VSC roles
    - guidelines [38](#)
  - Vserver
    - limitations when connecting directly to storage [50](#)
    - management LIF [204](#)
    - supports direct connections to storage [50](#)
    - tunneling supported [50](#)
  - vSphere
    - object [33, 34](#)
- ## W
- Windows
    - setting timeouts for guest OS [211](#)
  - Windows 2008 R2 SP1
    - Optimization and Migration [102](#)
  - Windows failover cluster
    - setting path selection policy for [208](#)
  - windows\_gos\_timeout.iso
    - guest OS tool [211](#)
  - workflow [104](#)
  - WSDL [196](#)
  - wsimport certificate [196](#)
- ## X
- XenDesktop
    - adding connection brokers [75](#)
- ## Z
- ZAPI [190](#)