



## Virtual Storage Console 6.0 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: 215-09831\_A0  
March 2015



# Contents

<b>Changes to this document: March 2015 .....</b>	<b>7</b>
<b>How to use this guide .....</b>	<b>9</b>
<b>Virtual Storage Console for VMware vSphere overview .....</b>	<b>11</b>
How VSC for VMware features work with optional plug-ins, virtual appliances ....	12
Provisioning and cloning datastores and virtual machines .....	13
How VSC for VMware vSphere optimizes I/O performance of misaligned virtual machines .....	13
Methods for migrating virtual machines .....	14
Backing up and restoring virtual machines and datastores .....	14
VSC for VMware vSphere protects system resources by using lock management .....	15
VSC for VMware vSphere architecture .....	16
<b>Plan your VSC for VMware vSphere installation .....</b>	<b>17</b>
<b>VSC for VMware vSphere installation overview .....</b>	<b>22</b>
VSC for VMware vSphere supported configurations .....	24
Initial installation of VSC for VMware vSphere .....	26
Installing VSC for VMware vSphere using the installation wizard .....	26
Installing VSC for VMware vSphere using silent mode .....	27
Installing or uninstalling the VSC backup and restore features .....	28
Upgrade installation of VSC for VMware vSphere .....	29
Performing a standard VSC for VMware vSphere upgrade installation .....	31
Upgrading VSC for VMware vSphere from a 32-bit installation to a 64- bit installation .....	32
Uninstalling VSC for VMware vSphere using Add/Remove Programs .....	34
Uninstalling VSC for VMware vSphere using silent mode .....	35
<b>Installing plug-ins, virtual appliances supported by VSC for     VMware vSphere .....</b>	<b>36</b>
NetApp NFS Plug-in for VAAI installation .....	36
VASA Provider for clustered Data ONTAP installation and registration .....	36
<b>Configuring your VSC for VMware vSphere environment .....</b>	<b>38</b>
ESX server and guest operating system setup .....	38
Configuring ESX server multipathing and timeout settings .....	38

Timeout values for guest operating systems .....	42
VSC for VMware vSphere configuration .....	47
Registering VSC for VMware vSphere with vCenter Server .....	47
Regenerating an SSL certificate for VSC for VMware vSphere .....	49
VSC for VMware vSphere port requirements .....	51
Performing VSC for VMware vSphere tasks across multiple vCenter Servers .....	51
Managing connection brokers .....	53
Adding connection brokers .....	53
Removing connection brokers .....	55
Configuring AutoSupport messages for backup jobs .....	55
Configuring email alerts for backup jobs .....	56
The preferences files .....	57
Enabling datastore mounting across different subnets .....	57
Setting the frequency of NFS path optimization checks .....	58
<b>Authentication and user management with vCenter RBAC and Data</b>	
<b>ONTAP RBAC .....</b>	<b>59</b>
vCenter Server role-based access control features in VSC for VMware vSphere ....	60
Components that make up vCenter Server permissions .....	61
Key points about assigning and modifying permissions .....	63
Advanced example of using vCenter Server permissions .....	64
Standard roles packaged with VSC for VMware vSphere .....	65
Product-level privilege required by VSC for VMware vSphere .....	68
Example of how the View privilege affects tasks in VSC for VMware	
vSphere .....	69
Data ONTAP role-based access control features in VSC for VMware vSphere .....	69
Recommended Data ONTAP roles when using VSC for VMware	
vSphere .....	70
How to configure Data ONTAP role-based access control for VSC for VMware	
vSphere .....	73
<b>Requirements for performing tasks in VSC for VMware vSphere .....</b>	<b>74</b>
<b>Navigating VSC for VMware vSphere .....</b>	<b>75</b>
<b>Working with storage systems .....</b>	<b>77</b>
Storage system discovery and credentials overview .....	77
Default credentials simplify administrating storage systems .....	78
Tunneled vFiler units and SVMs discovered automatically .....	80
Enabling discovery and management of vFiler units .....	81

Enabling discovery of vFiler units on private networks .....	81
Discovering storage systems and hosts .....	82
Manually adding storage systems .....	84
Refreshing the storage system display .....	85
Removing storage systems from VSC .....	86
Correcting storage system names displayed as "unknown" .....	87
Managing settings for volumes .....	88
MultiStore vFiler units are displayed differently .....	89
Differences between direct connections to SVMs and to cluster-management LIFs .....	90
Direct path access and NFS datastores .....	90
Changing NFS data paths to direct access .....	91
<b>Deploying virtual machines on NetApp storage .....</b>	<b>93</b>
Provisioning datastores .....	94
Cloning virtual machines from a template .....	97
Increasing storage efficiency by enabling deduplication .....	101
<b>Optimizing performance by aligning the I/O of misaligned virtual machines non-disruptively .....</b>	<b>103</b>
Scanning datastores to determine the alignment status of virtual machines .....	103
Checking the alignment status of virtual machines .....	105
Aligning the I/O of misaligned virtual machines non-disruptively .....	107
<b>Maintaining your VMware environment .....</b>	<b>110</b>
Migrating virtual machines to a new or existing datastore .....	110
Redeploying NFS-based virtual machine clones from a template .....	113
Reclaiming space from NFS-based virtual machines .....	114
Mounting datastores on hosts .....	115
Resizing datastores .....	116
Destroying datastores .....	118
<b>Backing up virtual machines and datastores .....</b>	<b>119</b>
Backup job specifications .....	121
Backup job requirements .....	121
Performing an on-demand backup of a virtual machine or datastore .....	122
Scheduling backup jobs .....	124
Adding a virtual machine or datastore to an existing backup job .....	126
Modifying the job properties of a scheduled backup job .....	127
Suspending an active backup job .....	127

Resuming a suspended backup job .....	128
Deleting a scheduled backup job .....	129
<b>Restoring virtual machines and datastores from backup copies .....</b>	<b>130</b>
Considerations for restore operations using data that was backed up with failed VMware consistency snapshots .....	130
Searching for backup copies .....	131
Mounting a backup copy .....	131
Unmounting a backup copy .....	132
Restoring data from backup copies .....	133
<b>Troubleshooting .....</b>	<b>135</b>
Information at NetApp VSC Communities Forum .....	135
Check the Release Notes .....	135
Uninstall does not remove standard VSC roles .....	136
Collecting the VSC for VMware vSphere log files .....	136
Updating vCenter credentials for background discovery .....	137
Resolution of issues that apply to backup and restore operations .....	137
Values that you can override for backup jobs .....	137
Location of backup event and error logs .....	138
Email notification for scheduled backup contains a broken link .....	139
You may have reached the maximum number of NFS volumes configured in the vCenter .....	139
Error writing backup metadata to repository\backups.xml: move failed ....	139
Virtual Storage Console unable to discover datastores on an SVM (Vserver) without a management LIF .....	140
VMware vSphere does not remove all snapshot delta disks during a restore operation .....	140
<b>Copyright information .....</b>	<b>141</b>
<b>Trademark information .....</b>	<b>142</b>
<b>How to send comments about documentation and receive update     notification .....</b>	<b>143</b>
<b>Index .....</b>	<b>144</b>

## Changes to this document: March 2015

---

This section contains information about the changes made to this guide for the 6.0 version of Virtual Storage Console for VMware vSphere. Previously, this guide was released with the 5.0 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 this VSC release, then this guide is updated with the new information, and this section lists what has 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 located on [NetApp Support](#).

### March 2015 update

In March 2015, this document was updated to add support for the 6.0 release of VSC.

- Using the VMware vSphere Web Client to manage multiple vCenter Servers that have VSC installed on them  
VSC 6.0 adds support for using the VMware vSphere Web Client to manage multiple vCenter Servers that have VSC installed on them. There is a one-to-one pairing between VSC and the vCenter Server, which means that one VSC Server can be registered with one vCenter Server. If you are performing a task where you must specify the vCenter Server where the storage is located, you will see a drop-down list of the available vCenter Servers.  
For more information, see [Performing VSC for VMware vSphere tasks across multiple vCenter Servers](#) on page 51.
- Connecting with storage systems using Transport Layer Security (TLS)  
Previously VSC used Secure Socket Layer (SSL) to communicate with storage systems.  
For more information, see [VSC for VMware vSphere supported configurations](#) on page 24.
- Automatically setting ESXi host values for NFS.MaxQueueDepth and NFS41.MaxVolumes  
NFS41.MaxVolumes, which applies to vSphere 6.0 is set to 256, and NFS.MaxQueueDepth, which applies to vSphere 5.0 and later, is set to 64.  
For more information, see [Configuring ESX server multipathing and timeout settings](#) on page 38.

VSC also supports VMware virtual volumes (VVOLs) for provisioning operations. You use VASA Provider for clustered Data ONTAP to create and manage VVOLs. VVOLs accept storage capability profiles, which you create and manage from the VASA Provider section of the VSC GUI.

For information about working with VVOLs, see [VASA Provider 6.0 for Clustered Data ONTAP User's Guide](#).

**Note:** At this time, VSC does not support VVOLs for the following operations: backing up, restoring, cloning, optimizing, or migrating. For the most current information about backing up VVOLs, see the *Virtual Storage Console 6.0 for VMware vSphere Release Notes*.



# How to use this guide

---

This guide has been arranged to make it easy for you to quickly get to the information you need.

The guide is organized into parts based on information that different types of users might need. These include the following broad categories:

- **Conceptional information about VSC, its key components, and how it works with optional plug-ins**

Target audience: Administrators, executives, and anyone who needs a high-level understanding of how VSC works

*Virtual storage Console for VMware vSphere overview* contains this information.

- **What is involved in setting up VSC for your environment**

Target audience: Administrators

*Plan your VSC for VMware Sphere Virtual storage Console overview* provides you with a checklist of issues you should consider before you install VSC. These issues can range from making sure that you have the correct system requirements to whether you want to also install optional features such as VASA Provider for clustered Data ONTAP.

- **Installation and configuration instructions**

Target audience: Administrators and VSC users

- *VSC for VMware vSphere installation overview* contains details you need to consider before starting your VSC installation as well as instructions for installing VSC the first time or upgrading an existing version of VSC.
- *VASA Provider for clustered Data ONTAP installation* tells you how to get VASA Provider so that you can use it with VSC.
- *NetApp NFS Plug-in for VAAI installation* tells you how to get that plug-in so that you can use it with VSC.
- *Configuring your VSC for VMware vSphere environment* provides instructions for configuring ESX server settings, guest operating system timeouts, regenerating VSC SSL certificates, setting up connection brokers, configuring AutoSupport and email alerts for backup jobs, enabling datastores to be mounted across subnets, and configuring the vCenter Server Heartbeat feature.

- **Role-based access control (RBAC)**

Target audience: Administrators

*Authentication and user management with vCenter RBAC and Data ONTAP RBAC* explains how VSC works with RBAC and describes ways that administrators can manage both vCenter Server and Data ONTAP RBAC when using VSC.

- **Prerequisites for VSC tasks**

Target audience: Administrators and VSC users

*Requirements for performing tasks in VSC for VMware vSphere* lists some of the requirements that must be in place before you perform VSC tasks. *Navigating VSC for VMWare vSphere* provides information about how to access tasks within the new VSC GUI.

- **VSC tasks**

Target audience: Administrators and VSC users

These sections provide the steps you need to perform to VSC tasks.

- *Working with storage systems* contains details about the tools VSC provides that enable you to manage storage systems.
- *Deploying virtual machines on NetApp storage* presents a workflow about provisioning datastores as well as instructions for performing provisioning and cloning operations.
- *Optimizing performance by aligning the I/O of misaligned virtual machines non-disruptively* contains instructions for tasks that affect checking the alignment of virtual machines and correcting misalignments.
- *Maintaining your VMware environment* contains tasks such as migrating virtual machines, redeploying virtual machines, reclaiming space from virtual machines, and managing datastores by mounting, resizing, and destroying them.
- *Backing up virtual machines and datastores* provides tasks that enable you to create and work with backup copies.
- *Restoring virtual machines and datastores from backup copies* contains the tasks you must perform when you need to restore a backup copy.

- **Troubleshooting**

Target audience: Administrators and VSC users

This section identifies possible issues that could affect your VSC installation.

If you are currently a VSC user, you should check *Changes to this document*, which summarizes some of the key differences from previous releases and how they affect this documentation.

# Virtual Storage Console 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 systems.

VSC integrates smoothly with the VMware vSphere Web Client and enables you to use single sign-on (SSO) services. If you have multiple vCenter Servers and each one has a VSC server registered to it, you can manage them from a single Web Client. In addition, the VSC Summary page allows you to quickly check the overall status of your vSphere environment.

**Note:** The NetApp blue "N" icon in the screens and portlets lets you easily distinguish the NetApp features from the VMware ones.

By running VSC, you can perform the following tasks:

- **Manage storage and configure the ESX host**

You can use VSC to add, remove, assign credentials, and set up permissions for storage controllers within your VMware environment. In addition, you can 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.

**Note:** You can also add Storage Virtual Machines (SVMs, formerly known as Vservers).

- **Create storage capability profiles and set alarms**

When you install VASA Provider for clustered Data ONTAP and register it with VSC, you can create and use storage capability profiles and VMware virtual volumes (VVOLs). You can also set alarms to warn you when the thresholds for volumes and aggregates are approaching full.

- **Provision datastores and clone virtual machines**

VSC 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.

- **Perform online alignments and migrate virtual machines singularly and in groups into new or existing datastores**

You can use VSC to quickly check the alignment status of virtual machines. If there are alignment issues with the virtual machines, you can, in most cases, resolve those issues without having to power down the virtual machines.

- **Back up and restore virtual machines and datastores**

VSC allows you to rapidly back up and restore virtual entities such as virtual machines or datastores on NetApp storage.

To provide security while performing tasks, VSC supports role-based access control (RBAC) at two levels:

- vSphere objects, such as virtual machines and datastores. These objects are managed using vCenter 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.

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

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.

### Related concepts

*[Provisioning and cloning datastores and virtual machines](#) on page 13*

*[How VSC for VMware vSphere optimizes I/O performance of misaligned virtual machines](#) on page 13*

*[Backing up and restoring virtual machines and datastores](#) on page 14*

*[Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 59*

## How VSC for VMware features work with optional plug-ins, virtual appliances

Virtual Storage Console for VMware vSphere supports optional plug-ins and virtual appliances that work with VSC features. You can enhance your experience with VSC by installing NFS Plug-in for VAAI and VASA Provider for clustered Data ONTAP.

VSC cloning and provisioning operations benefit from the NFS Plug-in for VMware VAAI. The plug-in integrates with VMware Virtual Disk Libraries to provide VMware vStorage APIs for Array Integration (VAAI) features, including copy offload and space reservations. These features can improve the performance of cloning operations because they do not need to go through the ESXi host.

VASA Provider is a virtual appliance that improves storage management. It provides information to the vCenter Server about the NetApp storage systems being used in the VMware environment. Integrating with the vCenter Server this way enables you to make more informed decisions. For example, you can create storage capability profiles that define different storage Service Level Objectives (SLOs) for your environment. You can then use these SLOs to select a datastore with the correct storage attributes when provisioning virtual machines. You can also set up alarms to notify you when a volume or aggregate is nearing full capacity or a datastore is no longer in compliance with its associated SLO.

## Provisioning and cloning datastores and virtual machines

Virtual Storage Console for VMware vSphere enables you to provision datastores and quickly create multiple clones of virtual machines in the VMware environment.

VSC's Create Rapid Clones wizard lets you create multiple clones from one virtual machine template. Cloning from a template saves time and enables you to set up virtual machines that all have the same configuration.

If you have VASA Provider for clustered Data ONTAP installed, you can use existing storage capability profiles to ensure that the new storage is configured consistently.

You can also use connection brokers, such as VMware View Server or Citrix XenDesktop, to import virtual machines into a virtual desktop infrastructure.

In addition, it is a good practice to have the NFS Plug-in for VMware VAAI installed before you perform provisioning and cloning operations. The plug-in can improve performance during the operations.

**Note:** VSC does not support IPv6. If you have IPv6 configured on a LIF, VSC cannot use any Storage Virtual Machines (SVMs, formerly known as Vservers) from that cluster.

### Related concepts

[Deploying virtual machines on NetApp storage](#) on page 93

### Related tasks

[Provisioning datastores](#) on page 94

## How VSC for VMware vSphere optimizes I/O performance of misaligned virtual machines

Virtual Storage Console for VMware vSphere provides a non-disruptive, interim solution for the performance penalty introduced by misaligned virtual machines. Rather than align the misaligned VMDKs, which requires downtime, VSC aligns the I/O without requiring downtime by offsetting the VMDKs within optimized datastores.

A virtual machine is misaligned when VMDK partitions do not align with the block boundaries of the storage system. As a result, the storage system might read or write to twice as many blocks of storage than is necessary.

VSC can scan datastores to determine which virtual machines are misaligned and, if possible, perform an online alignment by non-disruptively migrating the misaligned virtual machines to a datastore that is optimized for the VMDK layout. VSC optimizes the datastore by functionally aligning I/O to the offset of the largest partition.

Online alignment is a good choice for virtual machines that you cannot take offline. When possible, you should take the virtual machine offline and physically align the VMDK using a tool such as VMware vCenter Converter.

## Methods for migrating virtual machines

Virtual Storage Console for VMware vSphere provides two options for virtual machine migration: optimizing I/O performance for a misaligned virtual machine and moving virtual machines from one datastore to another.

Goal	What the migration does	Location in the vSphere Web Client
Align I/O of misaligned virtual machines non-disruptively	Performs an online alignment by non-disruptively migrating the misaligned virtual machines to a datastore that is optimized for the VMDK layout.	<b>Virtual Storage Console &gt; Optimization and Migration</b>
Migrate virtual machines to another datastore	<p>Migrates virtual machines to a new or existing datastore.</p> <p><b>Note:</b> If the selected virtual machines do not have the same offset group, the target datastore will not be optimized for all virtual machines. VSC creates a datastore optimized for the offset group of the last virtual machine that it migrates.</p>	<b>vCenter &gt; Inventory Lists &gt; Virtual Machines</b>

### Related tasks

[Optimizing performance by aligning the I/O of misaligned virtual machines non-disruptively](#) on page 103

[Migrating virtual machines to a new or existing datastore](#) on page 110

## Backing up and restoring virtual machines and datastores

Virtual Storage Console for VMware vSphere provides backup and restore features that enable you to create backup copies of virtual machines and datastores and later restore them.

To use these features, you must select the **Backup and Recovery** option included in the VSC installation program. Unlike other VSC features, the backup and restore features require that you have an SMVI license.

**Note:** The VSC GUI displays the backup and restore interface even if you have not installed it or purchased the SMVI license. But, if you attempt to perform a backup, restore, mount, or unmount operation, you receive an error message.

The VSC backup feature provides you with several options, including the following:

- Using SnapMirror or SnapVault when performing backups
- Performing a one-time, on-demand backup
- Scheduling backups to occur on a regular basis
- Specifying a retention policy for the backups
- Adding a virtual machine or datastore to an existing backup
- Modifying the job properties of an existing backup

You can restore a backup copy whenever you need to. VSC enables you to restore a datastore, an entire virtual machine, or particular disks from a virtual machine. To confirm that the backup copy is correct, you can mount the backup copy to a different host, verify the backup contents, unmount the backup from that host, and then restore the backup copy to the original location.

#### Related concepts

[\*Backing up virtual machines and datastores\*](#) on page 119

## VSC for VMware vSphere protects system resources by using lock management

Virtual Storage Console for VMware vSphere uses lock management to avoid having simultaneous tasks performed on the same target datastores or virtual machines. As a result, certain alignment, migration, provisioning or cloning, and backup and recovery features that could impact each other become temporarily unavailable if another task is being performed on the target datastore or virtual machine.

For example, if you are migrating virtual machines, you cannot clone one of the virtual machines until the migration operation completes. Or, if you provisioning storage, you cannot perform the following backup or restore operations on the target datastore or virtual machine until the provisioning operation completes:

- 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 during 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.

Before you start an operation, it is a good practice to make sure the target datastore or virtual machine is not being used by another operation.

## VSC for VMware vSphere architecture

The Virtual Storage Console for VMware vSphere architecture includes the storage system running Data ONTAP, the vCenter Server, the VMware vSphere Web client, and the ESX and ESXi hosts.

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 Web 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 and ESXi hosts.

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

- A single VSC plug-in with one user interface and help file
- The VSC server
- The SnapManager for Virtual Interface (SMVI) server

You can also write applications that communicate with the VSC server. For example, you can create this type of application using the PowerShell cmdlets that VSC supports.

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.

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.



## Plan your VSC for VMware vSphere installation

---

Before you install Virtual Storage Console for VMware vSphere, it is a good practice to plan your installation and decide how you want to configure your VSC environment, including whether you want to install other plug-ins or virtual appliances that work with VSC, such as NFS Plug-in for VAAI or VASA Provider.

The following is a high-level overview of what you need to consider when you install and configure VSC:

Consider...	Explanation...
What are the requirements for installing VSC?	<p>You must install the VSC software on a 64-bit Windows server with at least 4 GB of RAM. Do not install it on a client computer. Also, the vCenter Server must be running vSphere 5.5 or later.</p> <p>In addition, some of the VSC features use products that require you purchase a software license.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"><li>• Interoperability Matrix, which is available online at <a href="http://mysupport.netapp.com/matrix">mysupport.netapp.com/matrix</a></li><li>• <a href="#">Installation overview</a> on page 22</li><li>• <a href="#">VSC for VMware vSphere supported configurations</a> on page 47</li></ul>

Consider...	Explanation...
<p>What sort of role-based access control (RBAC) do you need?</p>	<p>VSC supports both vCenter Server RBAC and Data ONTAP RBAC.</p> <p>If you plan to run VSC as an administrator, you will have all the necessary permissions and privileges for all the tasks.</p> <p>If your company requires that you restrict access to vSphere objects, you can assign users to the standard VSC roles to meet the vCenter Server requirements.</p> <p>You can create the recommended Data ONTAP roles using the "RBAC User Creator for Data ONTAP" tool, which is available on the NetApp ToolChest.</p> <p>If a user attempts to perform a task without the correct privileges and permissions, the task options are grayed out.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#"><i>Installation overview</i></a> on page 22</li> <li>• <a href="#"><i>Standard roles packaged with VSC for VMware vSphere</i></a> on page 65</li> <li>• <a href="#"><i>Recommended Data ONTAP roles when using VSC for VMware vSphere</i></a> on page 70</li> <li>• <a href="#"><i>Authentication and user management with vCenter RBAC and Data ONTAP RBAC</i></a> on page 59</li> </ul>

Consider...	Explanation...
<p>Is this the first time you have installed VSC or is this an upgrade?</p>	<p><b>(Initial installation)</b> The VSC installation wizard automatically installs most of the VSC features, but it gives you the option of installing backup and restore features. You can also install these features later.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <i>Installing VSC for VMware vSphere using the installation wizard</i> on page 26</li> <li>• <i>Installing VSC for VMware vSphere using silent mode</i> on page 27</li> </ul> <p><b>(Upgrade installation)</b> You can only upgrade to the following:</p> <ul style="list-style-type: none"> <li>• 64-bit Windows systems</li> <li>• vSphere 5.5 or later environments</li> </ul> <p>Best practices before an upgrade include the following:</p> <ul style="list-style-type: none"> <li>• Record information about the storage systems being used and their credentials, especially those storage systems being used for backup and restore operations. After the upgrade, verify that all the storage systems were automatically discovered and they have the correct credentials.</li> <li>• If you modified any of the standard VSC roles, you should clone those roles in order to save your changes. VSC overwrites the standard roles with the current defaults each time you restart the VSC service.</li> <li>• If you made any changes to the VSC preferences file, you should record those changes. Each time you install VSC, it overwrites the current preferences files.</li> </ul> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <i>Upgrade installation of VSC for VMware vSphere</i> on page 29</li> <li>• <i>Performing a standard VSC for VMware vSphere upgrade installation</i> on page 31</li> <li>• <i>Upgrading from a 32-bit installation to a 64-bit installation of VSC for VMware vSphere</i> on page 32</li> </ul>

Consider...	Explanation...
Have you registered your VSC installation?	<p>After you install VSC, you must register it with the vCenter Server.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Registering VSC for VMware vSphere with vCenter Server</a> on page 47</li> </ul>
Are your ports set up correctly for VSC?	<p>VSC uses designated ports to enable communication between its components. If you have firewalls enabled, you might need to manually grant access to specific ports for VSC.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">VSC for VMware vSphere port requirements</a> on page 51</li> </ul>
Do you need to regenerate an SSL certificate for VSC?	<p>The SSL certificate is automatically generated when you install VSC. You might need to regenerate it to create a site-specific certificate.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Regenerating an SSL certificate for VSC</a> on page 49</li> </ul>
Were your ESX server values set correctly?	<p>While most of your ESX server values are set by default, it is a good practice to verify them.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">ESX server and guest operating system setup</a> on page 38</li> </ul>
Do you need to set up the guest operating system timeout values?	<p>The guest operating system (GOS) timeout scripts set the SCSI I/O timeout values for supported Linux, Solaris, and Windows guest operating systems to ensure correct failover behavior.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Timeout values for guest operating systems</a> on page 42</li> </ul>
Will you be performing provisioning and cloning tasks using connection brokers?	<p>You can add connection brokers to your system and use them to import virtual machines into a virtual desktop infrastructure.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Managing connection brokers</a> on page 53</li> </ul>
Do you plan to use storage capability profiles?	<p>To use storage capability profiles, you must install VASA Provider. This virtual appliance is installed separately from VSC.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">VASA Provider for clustered Data ONTAP installation</a> on page 36</li> </ul>

Consider...	Explanation...
Do you plan to use NFS Plug-in for VAAI?	<p>The plug-in provides VAAI features, such as copy offload and space reservations, which can improve the performance some provisioning and cloning operations.</p> <p><b>More information:</b></p> <ul style="list-style-type: none"><li>• <a href="#"><i>NFS Plug-in for VAAI installation</i></a> on page 36</li></ul>

## VSC for VMware vSphere installation overview

---

While the basic installation of the Virtual Storage Console for VMware vSphere software is simple, you must decide how you want to handle certain options, including whether you want to install the backup and restore features, whether you need a special software license, and, if you use role-based access control (RBAC), whether you have the correct RBAC privileges.

### Installation guidelines

To install VSC, you must have 64-bit Windows server and the vCenter Server must be running vSphere 5.5 or later. For information about which versions of Windows and other features are supported, see the Interoperability Matrix, which is available online at [mysupport.netapp.com/matrix](http://mysupport.netapp.com/matrix).

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

The following are some guidelines for installing the VSC software:

- 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 plug-in.
- At a minimum, the display must be set to 1,280 by 1,024 pixels to view VSC pages correctly.

### Backup and restore features require manual selection during installation

VSC automatically installs the software modules that support all its features except the backup and restore features. You must purchase a license for SnapManager for Virtual Infrastructure in order to perform backup and restore operations. As a result, VSC requires that you explicitly select the backup and restore features in the VSC installation menu.

The VSC GUI displays information about the backup and restore options regardless of whether you installed those features. However, if you attempt to perform a backup, restore, mount, or unmount operation, you will receive an error message. You can only perform those operations when the software supporting the backup and restore features is installed.

### Software licenses that might be needed

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

- The required protocol license (NFS, FCP, iSCSI)
- SnapManager for Virtual Infrastructure (if performing backup and restore operations)

- SnapMirror (required for the provisioning and cloning template distribution feature and for the backup SnapMirror update option)
- SnapRestore (if performing backup and restore operations)
- A\_SIS (if using provisioning and cloning features when configuring deduplication settings)
- MultiStore (if using provisioning and cloning features and working with vFiler units)
- FlexClone  
The FlexClone license is required in the following situations:
  - You are using VSC to clone virtual machines.
  - You are performing backup and restore operations in NFS environments and running a version of Data ONTAP prior to 8.1.

You do not need a FlexClone license if you are performing backup and restore operations 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

## RBAC access to VSC required for users

VSC supports vCenter Server and Data ONTAP role-based access control (RBAC). As a result you must provide users with the appropriate RBAC permissions. The vCenter Server permissions determine whether a user has or does not have access to certain VSC tasks for certain vSphere objects. The Data ONTAP privileges 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 an 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. However, 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.

### **All VSC users must 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. This privilege must be included in all VSC roles, or the user will not be able to view the VSC interface.

When you are working in an environment that has multiple VSC-vCenter Server instances and you are not an administrator, you must have the View privilege across all the vCenter Servers in that environment. Otherwise, the VMware vSphere Web Client will not load VSC.

The view privilege is used in the standard VSC Read-only role to the user. If you want to limit a user to read-only access to VSC, you can assign that user the VSCRead-only role.

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

### **Confirm that the storage systems and their credentials are available**

VSC provides centralized management for storage discovery and credentials. All VSC features use the credentials entered during the initial setup and discovery process.

If you are upgrading from VSC 4.1 or earlier, you should check the storage systems to ensure that there are no issues with credentials, especially in regard to backup and recovery operations, after you install the backup and restore features.

### **Make sure the times are synchronized when installing VSC on a different server**

If you are not installing VSC on the vCenter Server, you must make sure that the times on the VSC installation server and the vCenter Server are synchronized.

The vCenter Server will not accept the certificate of the VSC installation server when the times differ on the servers.

For information on synchronizing the server times, see your operating system documentation.

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

Virtual Storage Console for VMware vSphere requires the VMware vSphere Web Client and is supported on specific releases of ESX and ESXi, vSphere, and Data ONTAP software. See the Interoperability Matrix, which is online at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix), for details.

### **vSphere server configuration**

VSC requires that you have vCenter Server 5.5 or later.

If you are installing VSC on a vCenter Server that has a large number of ESX or ESXi hosts, make sure that there is sufficient CPU and memory to support VSC in your environment. If the vCenter Server is consuming all the resources to manage the hosts, the VSC service cannot respond to the requests it receives. The number of resources needed varies based on your system setup.



## Windows 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. Memory requirements for 64-bit environments where VSC is installed on a separate machine are currently 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. 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

**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 reenabled after VSC is installed.

## VMware vSphere Web Client configuration

The client computer that runs the vSphere Web Client software must have a web browser installed.

VSC supports all browsers that the vSphere Web Client supports.

VSC supports having the vSphere Web Client manage multiple vCenter Servers when there is a unique one-to-one pairing between VSC and a vCenter Server.

## VASA Provider and VSC configuration

VASA Provider for clustered Data ONTAP is designed to work with VSC as long as you install the same version of each product. For example, VASA Provider 6.0 does not work with VSC 5.0; it only works with VSC 6.0.

## Network security protocol configuration

VSC requires that you have Transport Layer Security (TLS) enabled on the storage systems. If your storage systems are running clustered Data ONTAP, TLS is automatically enabled. If your storage systems are running Data ONTAP operating in 7-mode, you must manually enable TLS using the

`options.tls.enable` on command. For TLS to take effect on storage systems running Data ONTAP operating in 7-mode, the option `httpd.admin.ssl.enable` must be set to `on`.

You can find additional information about TLS in the Data ONTAP documentation.

## Initial installation of VSC for VMware vSphere

If this is the first time you are installing Virtual Storage Console for VMware vSphere, you do not need to worry about upgrade issues. You can install VSC using either the installation wizard or silent mode.

The wizard steps you through the process and gives you the option of installing the backup and restore features. It installs the rest of the VSC features by default. Because the backup and restore features require the purchase of a SnapManager for Virtual Infrastructure license, some people choose not to install those features.

If you use silent mode to install VSC, you must enter a command line telling VSC which options to install.

### 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 all the features except the backup and restore features. You must manually select the option to install those features because they normally require that you purchase a license for SnapManager for Virtual Infrastructure.

#### Before you begin

- You must log in with administrator privileges for the machine on which you install VSC.
- You can only install VSC on a 64-bit Windows server.
- Before you install VSC, verify that your system meets the VSC requirements listed in the Interoperability Matrix, which is available online at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix).

#### 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 restore features, you must select that option. Otherwise, the VSC installer does not install it. Menu options for the features will appear in the installed VSC GUI, but you will not be able to use them.

4. Click **Finish** to complete the installation.

5. At the web page that appears after the installation completes, 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.

## 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 enter a command line that lets you automatically install all the VSC features at one time.

### Before you begin

- You must log on with administrator privileges for the machine on which you install VSC.
- You can only install VSC on a 64-bit Windows server.
- Before you install VSC, verify that your system meets the VSC requirements listed in the Interoperability Matrix, which is available online at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix).

### Steps

1. Download the VSC installer.
2. 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 features, including the backup and restore features.

### Example

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

```
VSC-6.0-win64.exe /s /v"/qn /Li install.log ADDLOCAL=ALL INSTALLDIR="C:
\Program Files\NetApp\Virtual Storage Console\""
```

3. 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.

## Installing or uninstalling the VSC backup and restore features

The installation wizard lets you install or uninstall the Virtual Storage Console for VMware vSphere backup and restore features separately. The other VSC features can be installed or uninstalled only as a group. You can install or uninstall the backup and restore features with the other features or at a different time.

### About this task

The steps that follow describe how to install the backup and restore features after you have already installed the other features.

You can also use this backup and restore check box to uninstall only those features. It has no effect on the other installed VSC features.

### Steps

1. On the Windows server where you installed the VSC software, go to the **Control Panel** and select the option for your version of Windows that lets you add programs.  
  
For example, if you are running Windows Server 2003, you would choose **Control Panel > Add/Remove Programs**, but for Windows 2008, you would choose **Control Panel > Programs and Features**.
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 the backup and restore features if you want to add them, or clear the check box if you want to remove them, and click **Next**.
5. If you are installing the backup and restore features, click **Install** to start the installation.
6. Click **Finish** to complete the installation or the removal of these features.

### After you finish

You must close the vSphere Client and restart it to activate the new features.

The VSC interface displays the menu options for the backup and restore features even if they are not installed; however, you can only use these menu options if the features are installed.

If you install the backup and restore features separately from the other features, you must restart the VSC service after the installation completes.

## Upgrade installation of VSC for VMware vSphere

If you are running a 4.x version of Virtual Storage Console for VMware vSphere, you can perform an upgrade installation as long as your environment meets the VSC requirements. Before you upgrade, make a record of the information that VSC does not update and the storage systems in use. Also, if you are using VASA Provider for clustered Data ONTAP, you must unregister it.

### Check for changes to VSC requirements

VSC 6.0 only works with 64-bit Windows servers. If you are currently running in a 32-bit Windows environment, you must perform special steps and manually move certain data in order to upgrade to a 64-bit Windows system. For information on doing this, see [Upgrading from a 32-bit installation to a 64-bit installation of VSC for VMware vSphere](#) on page 32.

In addition, VSC 6.0 requires that the vCenter Server be running vSphere 5.5 or later. If you attempt to install VSC 6.0 on an earlier version of vSphere, you will receive an error message. This message varies depending on whether you have thick provisioning enabled. With thick provisioning the message is similar to the following:

```
The VSC plugin is only available in the vSphere web client.
```

Without it, the error message is the following:

```
Registration failed with the following message:  
NVPF-00017: This version of vCenter (5.1.0) found at  
https://00.00.000.000:443/sdk is not supported. Please upgrade  
to vCenter 5.5 or later.
```

For complete information about the VSC requirements, see the [NetApp Interoperability Matrix Tool](#), which is available online at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix).

### Record storage system information

VSC automatically rediscovers your storage after you perform an upgrade. It is a good practice to record your storage system information before the upgrade so that you can confirm that all the storage systems were rediscovered after the upgrade.

If you are using backup and restore features, it is important to record the storage systems used for those operations and the credentials associated with them. Prior to VSC 4.2 for VMware vSphere, the Backup and Recovery plug-in managed its own storage system discovery and credentials. Because VSC did not discover those storage systems, they might not appear in the list of storage systems that VSC manages.

After an upgrade, when you check the storage systems used for backup and restore operations, you should also verify that the systems have at least the minimum credentials required to perform these operations. You can use VSC to update the credentials.

If any storage systems are missing after the upgrade, you can select the **Update All** icon to force VSC to discover storage systems. If that does not work, you can manually add the storage system using the **Add** icon found on the VSC Storage System page or the **Add storage system** option found in the vCenter Actions menu.

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

### Record any changes you made to standard VSC roles

You should not modify the VSC standard roles. If you make changes to these roles, you will lose those changes when you upgrade your VSC installation or restart the VSC Windows service. These roles return to the current default values each time you install VSC, restart the VSC Windows service, or modify your VSC installation.

If you made any changes to these roles, you should record the changes. After you upgrade your installation, you can create new roles that reflect those changes.

**Note:** Instead of editing the standard VSC roles, you should clone them and then edit the cloned roles.

### Record any changes you made to the preferences file

The upgrade process overwrites the existing preferences files with new preferences files for features that VSC uses. It is a good practice to record changes you made to preference files before an upgrade.

For provisioning and cloning tasks, VSC creates a backup of the preferences file, `etc/kamino/old_kaminoprefs.xml`. If you had modified the `etc/kamino/kaminoprefs.xml` preferences file, you can copy the changes from `etc/kamino/old_kaminoprefs.xml` to the new file that VSC creates during the update.

Instructions for registering VASA Provider

### Unregister VASA Provider

If you are using VASA Provider for clustered Data ONTAP, you must unregister it before you upgrade your VSC software. If you do not unregister it, you might not be able to see the VASA Provider section of the VSC GUI when you register the VASA Provider server in the upgraded installation.

Instructions for registering VASA Provider are in the [VASA Provider 6.0 for Clustered Data ONTAP User's Guide](#).

## Remove UI extensions from the vCenter Server

You must remove the user interface (UI) extensions that are cached on the vCenter Server before you upgrade VSC or VASA Provider. See the [Virtual Storage Console 6.0 for VMware vSphere Release Notes](#) for instructions about removing these directories.

## Performing a standard VSC for VMware vSphere upgrade installation

If you currently have version 4.x 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 VSC component to determine whether you are upgrading to a newer version.

### Before you begin

The VSC installer automatically upgrades all the installed VSC features to the newer versions. If you do not have backup and restore features installed, the VSC installer gives you the option of installing them.

The VSC installer does not support upgrades from the following:

- VSC 2.1.2 or earlier
- A standalone version of Rapid Cloning Utility (RCU)
- A standalone version of SnapManager for Virtual Infrastructure (SMVI)

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.

**Important:** If you are using VASA Provider for clustered Data ONTAP, unregister it from VSC before you install the upgrade.

### 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. Review your installation options.

The installation wizard gives you the option to select the backup and restore features if they are not currently installed.

5. Click **Next** to start the installation.

The wizard automatically selects all currently installed features and upgrades them.

The installation might take several minutes.

6. Click **Finish** to complete the installation.
7. 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.

## Upgrading VSC for VMware vSphere from a 32-bit installation to a 64-bit installation

If you are currently running Virtual Storage Console 4.x for VMware vSphere on a 32-bit Windows platform, you can upgrade to VSC 6.0, which requires a 64-bit platform. Going to a 64-bit platform requires that you manually move directories from the 4.x installation to the 6.0 installation in addition to performing the standard upgrade procedures.

### Before you begin

You must have administrator privileges on the system where you are installing VSC.

You should always check the Interoperability Matrix, which is online at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix), to confirm that your system meets the requirements for VSC.

### Steps

1. Download the VSC software package.
2. **(VSC 4.x system)** Stop the VSC Windows service.  
  
If you have been using the backup and restore features, make sure that you also stop the SnapManager for Virtual Infrastructure (SMVI) service.
3. **(VSC 4.x system)** Copy the following VSC 4.x directories and files.

**Note:** To make moving these files to your VSC 6.0 installation easier, you can create a .zip file to contain them.

These directories and files are all relative to the VSC installation directory.

- etc\keystore.properties
- etc\nvpf.keystore
- etc\nvpf.override
- etc\network-interface.properties
- etc\caster\casterprefs.xml



- etc\caster\derby\
- etc\caster\kaminosdkprefs.xml
- etc\kamino\baselines.ser
- etc\kamino\connectionBrokers.ser
- etc\kamino\vcenters.ser
- etc\kamino\kaminoprefs.xml
- etc\vsc\vsc.xml
- etc\vsc\vscPreferences.xml
- log\
- smvi\server\etc\cred
- smvi\server\etc\keystore
- smvi\server\etc\smvi.keystore
- smvi\server\repository\

**4. (VSC 6.0 system)** Run the VSC 6.0 installation program.

This program automatically installs most VSC features. If you want to use the VSC backup and restore features, you must manually select that option. The backup and restore features require the SMVI license.

See the section "Software licenses" for more information about licenses that you might need in order to use VSC.

**5. (VSC 6.0 system)** Register VSC with the vCenter Server when the registration web page opens.

For details about how to register VSC or what to do if the registration web page does not open, see the section "Registering VSC for VMware vSphere with vCenter Server."

**6. (VSC 6.0 system)** Stop the VSC Windows service.

**7. (VSC 6.0 system)** Place the VSC 4.x files into your VSC 6.0 installation.

If you created a .zip file to contain these directories and files, you need to paste it into the VSC 6.0 installation directory and then unzip it.

You should paste these files relative to the VSC 6.0 installation directory.

**8. (VSC 6.0 system)** Restart the VSC Windows service.

**9. (VSC 6.0 system)** Re-register VSC with the vCenter Server.

**10. (VSC 6.0 system)** Verify that the expected data (storage systems, backup jobs, and so on) appears in VSC after you complete upgrade from a 32-bit system to VSC 6.0.

11. **(VSC 4.x system)** When your VSC 6.0 installation is running and you have verified that it has the correct data, uninstall the VSC 4.x program.

You can use one of the following methods to uninstall VSC 4.x:

- Go to the Windows Add or Remove Programs list and remove VSC 4.x.
- Perform a “silent” uninstall by entering the following at the command line:

```
installer.exe /s /v"/qn /Li logfileREMOVE=ALL INSTALLDIR=
\"install_path\" " "
```

## 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 VSC features 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 features.

### Example

The following is an example of the command line you might use if you were uninstalling VSC 6.0 from a 64-bit host machine:

```
VSC-6.0-win64.exe /s /v"/qn /Li uninstall.log REMOVE=ALL
```

## Installing plug-ins, virtual appliances supported by VSC for VMware vSphere

---

You can install plug-ins and virtual appliances that work with Virtual Storage Console for VMware vSphere to enhance VSC tasks and storage management. VSC supports both NFS Plug-in for VMware VAAI and VASA Provider for clustered Data ONTAP.

### NetApp NFS Plug-in for VAAI installation

The NetApp NFS Plug-in for VMware VAAI is not shipped with Virtual Storage Console for VMware vSphere; however, you can get the plug-in installation package and instructions for installing it from the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com). You can then go to use the Virtual Storage Console **Tools > NFS VAAI** page to complete your installation.

The plug-in is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. It is a good practice to install the plug-in because VAAI (VMware vStorage APIs for Array Integration) features such as copy offload and space reservations can improve the performance of cloning operations.

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.

To download the plug-in, go to the Software Download page on the NetApp Support Site and log in. In the row "NetApp NFS Plug-in for VAAI," select ESXi and click **Go!**. Continue through the pages until you reach NetApp NFS Plug-in for VMware VAAI Download. This page provides links to both the installation package and the installation guide.

Follow the VSC installation instructions in *Installing the NetApp NFS 1.0.20 Plug-in for VMware VAAI*.

After you install the plug-in, you must reboot the host. VSC then automatically detects the plug-in and uses it. You do not need to perform additional tasks to enable it.

### VASA Provider for clustered Data ONTAP installation and registration

When you install VASA Provider for clustered Data ONTAP, you must deploy it on an ESXi host and then register it with Virtual Storage Console for VMware vSphere. VSC is the management console for VASA Provider and also provides a VASA Provider GUI for certain tasks.

After you register VASA Provider with VSC, you must log out of the VMware vSphere Web Client and then log back in. You cannot see the VASA Provider GUI that VSC provides until you log back in.

When you have VASA Provider installed and registered, you can perform tasks that pertain to VVOLs, storage capability profiles, alarms, and general VASA Provider maintenance. You must use the correct VASA Provider user interface for the task:

- To work with VVOLs, you must select the VASA Provider for clustered Data ONTAP option in the VMware vSphere Web Client Actions menu, not the VSC GUI.
- To register VASA Provider, manage storage capability profiles, map them to datastores, and set threshold alarms, you must select the VASA Provider for clustered Data ONTAP section of the VSC GUI.
- To adjust settings for VASA Provider and perform maintenance tasks, you must use the VASA Provider maintenance menus, which are accessible from the console of the virtual appliance. The Main Menu provides several options for configuring VASA Provider and performing diagnostic operations.

If you need to create a support bundle, you should use the Vendor Provider Control Panel screen located at [https://vm\\_ip:9083](https://vm_ip:9083). You can use the maintenance menu to create a support bundle, but the Vendor Provider Control Panel enables you to create a more complete bundle.

## Configuring your VSC for VMware vSphere environment

---

Virtual Storage Console for VMware vSphere supports numerous environments. Some of the features in these environments might require additional configuration. In some cases, you might need to perform maintenance operations.

Some of the configuration and maintenance work that you might need to perform includes the following:

- Verify your ESX host settings, including UNMAP
- Add timeout values for guest operating systems.
- Manually register VSC.
- Regenerate a VSC SSL certificate.
- Verify that you have the correct port settings.
- Set up connection brokers to work with the provisioning and cloning operations.
- Register VASA Provider and create storage capability profiles and threshold alarms.
- Work with the preferences file to enable datastore mounting across different subnets.

## ESX server and guest operating system setup

Most of your ESX server values should be set by default; however, it is a good practice to verify the values and make sure they are right for your system setup. Virtual Storage Console for VMware vSphere also provides ISO files to enable you to set the correct timeout values for guest operating systems.

## Configuring ESX server multipathing and timeout settings

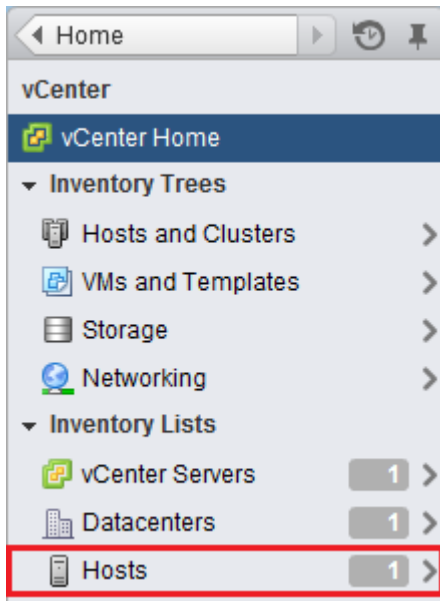
Virtual Storage Console for VMware vSphere checks and sets the ESX or ESXi host multipathing and HBA timeout settings that ensure proper behavior with NetApp storage systems.

### 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.

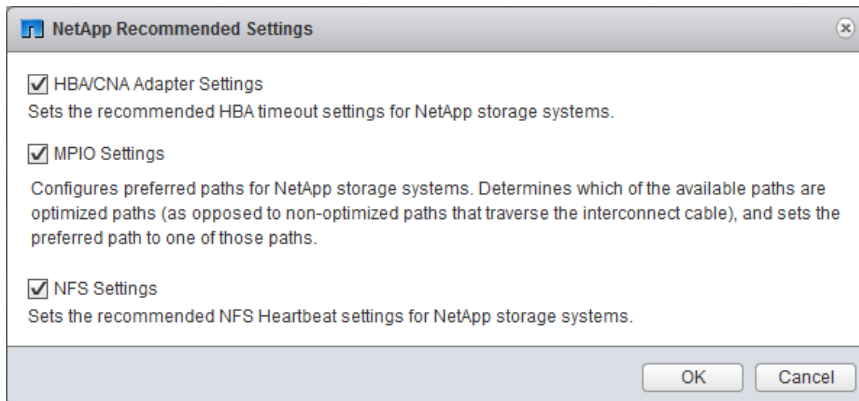
## Steps

1. From the VMware vSphere Web Client **Home** page, click **vCenter > Hosts**.



2. Right-click on a host and select **Actions > NetApp VSC > Set Recommended Values**.
3. In the Recommended Settings pop-up box, select the values that work best with your system.

The standard, recommended values are set by default.



4. Click **OK**.

## **ESX host values set by VSC for VMware vSphere**

Virtual Storage Console for VMware vSphere sets ESX or ESXi host timeouts and other values to ensure best performance and successful failover. The values that VSC sets are based on internal NetApp testing.

VSC sets the following values on an ESX or ESXi host:

### **ESX advanced configuration**

#### **/VMFS/hardwareacceleratedlocking**

Set to 1.

### **NFS settings**

#### **Net.TcpipHeapSize**

If you are using vSphere 5.0 or later, set to 32.

For all other NFS configurations, set to 30.

#### **Net.TcpipHeapMax**

If you are using vSphere 5.5 or later, set to 512.

If you are using vSphere 5.0 up to 5.5, set to 128.

For all other NFS configurations, set to 120.

#### **NFS.MaxVolumes**

If you are using vSphere 5.0 or later, set to 256.

For all other NFS configurations, set to 64.

#### **NFS41.MaxVolumes**

If you are using vSphere 6.0 or later, set to 256.

#### **NFS.MaxQueueDepth**

If you are using vSphere 5.0 or later, set to 64.

#### **NFS.HeartbeatMaxFailures**

Set to 10 for all NFS configurations.

#### **NFS.HeartbeatFrequency**

Set to 12 for all NFS configurations.

#### **NFS.HeartbeatTimeout**

Set to 5 for all NFS configurations.

### **FC/FCoE settings**

#### **Path selection policy**



Set to `RR` (round robin) for ESX 4.0 or 4.1 and ESXi 5.x, FC paths with ALUA enabled.  
Set to `FIXED` for all other configurations.

Setting this value to `RR` helps provide load balancing across all active/optimized paths.  
The value `FIXED` is for older, non-ALUA configurations and helps prevent proxy I/O. In other words, it helps keep I/O from going to the other node of a high availability pair (HA) in an environment that has Data ONTAP operating in 7-mode.

### **Disk.QFullSampleSize**

Set to 32 for all configurations. This setting is available with ESXi 5.x and ESX 4.x.  
Setting this value helps prevent I/O errors.

**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).

### **Disk.QFullThreshold**

Set to 8 for all configurations. This setting is available with ESXi 5.0 and ESX 4.x.  
Setting this value helps prevent I/O errors.

**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).

### **Emulex FC HBA timeouts**

For ESX 4.0 or 4.1 or ESXi 5.x, use the default value.

### **QLogic FC HBA timeouts**

For ESX 4.0 or 4.1 or ESXi 5.x, use the default value.

## **iSCSI settings**

### **Path selection policy**

Set to `RR` (round robin) for all iSCSI paths.

Setting this value to `RR` helps provide load balancing across all active/optimized paths.

### **Disk.QFullSampleSize**

Set to 32 for all configurations. This setting is available with ESX 4.x and ESXi 5.x.  
Setting this value helps prevent I/O errors.

**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).

### **Disk.QFullThreshold**

Set to 8 for all configurations. This setting is available with ESX 4.x and ESXi 5.x. Setting this value helps prevent I/O errors.

**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, Virtual Storage Console for VMware vSphere 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, VSC 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 Update 2. VSC ensures that this feature is disabled with all versions of ESXi 5.x.

## Timeout values for guest operating systems

The guest operating system (GOS) timeout scripts set the SCSI I/O timeout values for supported Linux, Solaris, and Windows guest operating systems. The timeout values help improve disk I/O behavior in a failover situation.

These scripts are provided as .ISO files. You can get a copy of the scripts by clicking **Tools > Guest OS Tools** from the Virtual Storage Console Home page. 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 Guest OS Tools page 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 Virtual Storage Console 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.

## Adding the CD-ROM to a virtual machine

To enable installing the guest operating system scripts, you need to add the CD-ROM to a virtual machine if it does not exist.

### Steps

1. In the vSphere Client, select the desired virtual machine and power it off.
2. Right-click the virtual machine and select **Manage > VM Hardware**.
3. Select **CD/DVD Drive** in the **New device** drop-down box and 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 virtual machine.

## Installing guest operating system scripts

The ISO images of the guest operating system scripts are loaded on the Virtual Storage Console for VMware vSphere server. To use them to set the storage timeouts for virtual machines, you must mount and run them from the vSphere Web Client.

### Before you begin

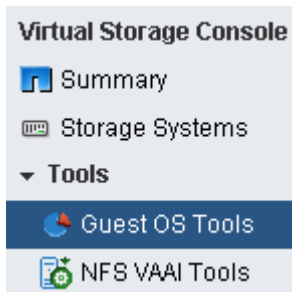
- The virtual machine must be running.
- The CD-ROM must already exist in the virtual machine, or it must have been added.
- The script must be installed from the copy of the VSC registered to the vCenter Server that manages the virtual machine.

### About this task

If your environment includes multiple vCenter Servers, you must select the server that contains the virtual machines for which you want to set the storage timeout values.

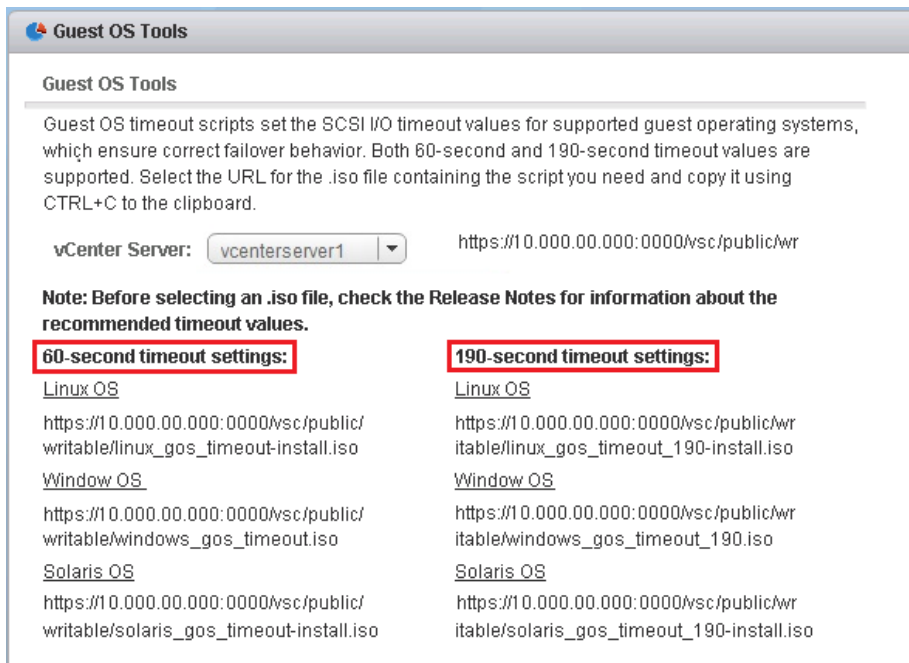
## Steps

1. From the Virtual Storage Console **Home** page, expand **Tools** and click **Guest OS Tools**:



2. Under **Guest OS Tools**, press Ctrl-C to copy the link to the ISO image for your guest operating system version to the clipboard.

VSC provides both 60-second timeout scripts and 190-second timeout scripts for Linux, Windows, and Solaris. Select the script for your operating system that provides the timeout value you want to use.



3. Return to the vSphere Web Client **Home** page and select **vCenter**.
4. Select the desired virtual machine and click the **Manage > VM Hardware**.
5. Select **CD/DVD Drive 1 > Connect to ISO image on local disk**.

6. Paste the link you copied into the **File Name** field and then click **Open**.

Be sure that the link you are using is from the copy of the VSC running on the vCenter Server that manages the virtual machine.

### After you finish

Log in to the virtual machine and run the script to set the storage timeout values.

## Running the GOS timeout scripts for Linux

The guest operating system timeout scripts set the SCSI I/O timeout settings for RHEL4, RHEL5, RHEL6, RHEL7, 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**.

## VSC for VMware vSphere configuration

Most of the Virtual Storage Console for VMware vSphere configuration happens automatically when you install the software. In some cases, you might need to register VSC with the vCenter Server or regenerate an SSL certificate.

It can also be helpful to know which VSC ports are available through the firewall.

## Registering VSC for VMware vSphere with vCenter Server

After installing the Virtual Storage Console for VMware vSphere software, you must register it with the vCenter Server. If your environment includes multiple vCenter Servers, you must register each

VSC instance with a single vCenter Server. By default, the registration web page opens when the VSC for VMware vSphere installation finishes.

### Before you begin

You must be logged in 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

If you register VSC to a vCenter Server that is part of a multi-vCenter Server environment, each instance of VSC must be registered individually with a vCenter Server.

If you change the credentials of the account that was used to register VSC, then you must register it again.

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 on which you installed VSC. If you are not performing this step from the computer on which you installed VSC, you must replace `localhost` with the host name 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.

The vSphere Plugin Registration web page appears.

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. Enter the vCenter Server host name or IP address and the administrative credentials to register the vCenter Server with this instance of VSC.

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. If you are registering additional vCenter Servers, repeat steps 3, 4, and 5 to register each instance of VSC with a vCenter Server.



7. Close the registration page after you complete the registration process, because the web page used to register with the vCenter Server is not automatically refreshed.

## Regenerating an SSL certificate for VSC for VMware vSphere

The SSL certificate is generated when you install Virtual Storage Console 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.

### About this task

You might also need to regenerate the certificate if you are using the VSC backup and restore features. The certificate that is automatically generated for the SMVI process, which supports the backup and restore features, uses a weak algorithm. The certificate that you regenerate uses a stronger algorithm.

### Steps

1. Stop the `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 installation directory and enter the following command:

```
bin\vsc ssl setup -cn <HOST>
```

For `<HOST>`, enter the host name of the system running VSC or a fully qualified domain name of the system running VSC.

### 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\vsc 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).

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 VSC provisioning and cloning or optimization and migration features, perform the following two steps:

- a. Change to the directory where `keytool.exe` is located. VSC installation directory:

This directory is under the VSC installation directory. While you can specify a different name for your installation directory, the default location for `keytool.exe` is the following.

```
C:\Program Files\NetApp\Virtual Storage Console\jre\bin\keytool.exe
```

- b. Enter the following command:

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

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

- c. Import the certificate to the local a Java keystore by entering the command:

```
c:\Program Files\NetApp\Virtual Storage Console\jre\bin\keytool.exe -import -alias npvf -file npvf.cer -keystore "C:\Program Files\NetApp\Virtual Storage Console\jre\lib\security\cacerts"
```

6. If you are using the backup and restore features, perform the following steps to ensure that certificate is not using a weak algorithm:

- a. Copy the following entries from the `keystore.properties` file that you just created in the installation directory.

The encrypted passwords might look similar but they are not an exact match to the following:

```
http.ssl.key.password=1LKPQWaJIEvANxOqsF3ILKCKThDZv1F5
```

```
http.ssl.keystore.file=C:\\Program Files\\NetApp\\Virtual Storage Console\\etc\\npvf.keystore
```

```
http.ssl.keystore.password=u4SIS8KXhzMDq5JebwyF0AXT36YVCfvX
```

- b. If the `smvi.override` file exists, add these entries to the file, which is located at `http.ssl.keystore.file=C:\\Program Files\\NetApp\\Virtual Storage Console\\smvi\\server\\etc`

If the `smvi.override` file does not exist, you must create it at that location and then add the entries.

- c. Restart the SMVI service
- d. Verify that the new certificate, which was generated with the `sha1WithRSAEncryption` algorithm, is being used.

You can use a tool such as SSLScan to verify SSL certificates.

**7. Secure the `etc\keystore.properties` file.**

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.

**8. Restart the `vsc` service.**

**9. You can review and accept the SSL certificate after the vSphere Client receives the certificate by clicking the NetApp icon in the vSphere Client.**

**10. Now 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.

## VSC for VMware vSphere port requirements

By default, Virtual Storage Console 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 uses. If you do not grant access to these ports, an error message such as `Unable to communicate with the server` is displayed.

VSC uses the following default ports:

Default port number	Description
8143	VSC listens for secure communication on this port.
8043	The VSC backup and restore features and the VSC Restore Agent and CLI listen for secure communication on this port.
9060	VSC listens for communication from the VSC PowerShell Toolkit on this port.

## Performing VSC for VMware vSphere tasks across multiple vCenter Servers

To use Virtual Storage Console for VMware vSphere in an environment with multiple vCenter Servers and to ensure that more operations are completed simultaneously within one maintenance

window, you can add multiple VSC and vCenter Server instances. You can manage all the servers running vCenter 5.5 or later from a single vSphere Web Client in both linked mode and nonlinked mode.

Linked mode uses Microsoft Active Directory Application Mode (ADAM), which is installed automatically as part of vCenter Server installation, to store and synchronize data across multiple vCenter Server systems.

Unlike VSC, VASA Provider currently requires a single vCenter Server environment.

If your environment includes multiple vCenter Servers, you can use the VMware vSphere Web Client to perform VSC tasks across multiple vCenter Servers in the following instances:

- Multiple instances of VSC have been installed.
- Each VSC server is registered with a single vCenter Server.  
This must be a unique one-to-one pairing. For example, VSC server A is registered to vCenter Server A, VSC server B is registered to vCenter Server B, VSC server C is registered to vCenter Server C and so on.
- To use VSC as a non-administrator user, all vCenter Servers in the VMware inventory must have a VSC server registered with them.  
If the VMware inventory includes a vCenter Server that does not have a VSC server registered to it and you log in to the vCenter Server as anything other than an administrator, even if it has one or more vCenter Servers that have VSC registered, it does not display VSC at all.
- The non-administrator must have the correct RBAC permissions and VSC-specific View privilege across all the vCenter Servers registered to the single sign-on (SSO), along with a registered instance of VSC.
- The **vCenter Server** box displays the vCenter Servers in alphanumeric order with the default being the first server in the list.
- When the location of the storage is known (for example, when using the Provisioning wizard and the datastore is on a host managed by a specific vCenter Server), the vCenter Server is displayed, but it is read-only.  
This applies only when you are using the **Action** menu by right-clicking an item in the vSphere Web Client.

Just as VSC does in a single vCenter Server environment, it warns you in a multi-vCenter Server environment when you attempt to select an object that it does not manage. You can filter storage systems based on a specific vCenter Server on the VSC summary page. A summary page appears for every VSC instance registered with a selected vCenter Server. You can manage storage systems associated with a specific VSC instance and vCenter Server but keep the registration information for each storage system separate if you are running multiple instances of VSC.

## 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 [Cloning virtual machines](#) on page 97 for details.

To work with connection brokers in Virtual Storage Console for VMware vSphere, you must have .Net 3.5 available on the system where you have VSC 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

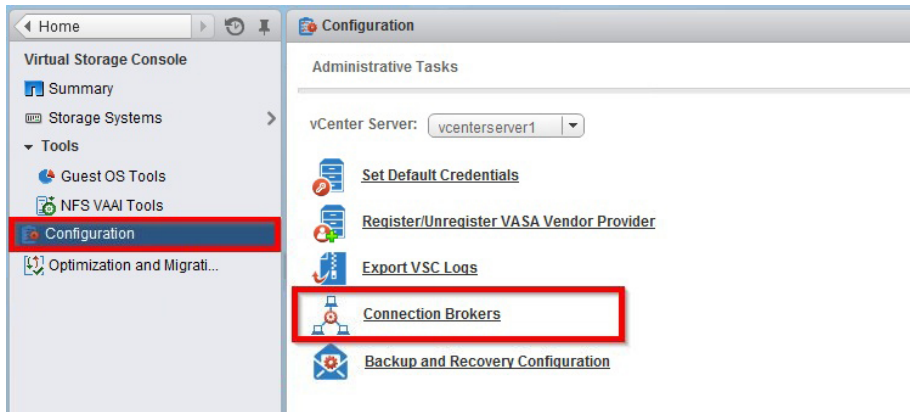
You can add connection brokers to Virtual Storage Console for VMware vSphere.

### Before you begin

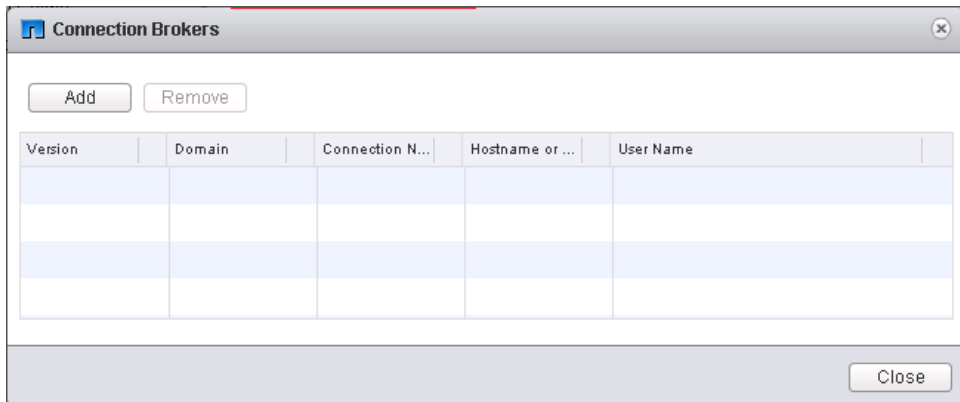
- You must have .Net 3.5 available on the system where you have VSC installed.  
For some versions of Windows, such as Windows 2008, .Net 3.5 is included as part of the installation. For other versions, such as Windows 2003, it is not part of the base install, so you must manually install it.
- You must ensure that your configuration is supported.  
You can do this by checking the Interoperability Matrix, which is available online at [mysupport.netapp.com/matrix](http://mysupport.netapp.com/matrix).
- You must have selected the vCenter Server that you want to use for this task.

### Steps

1. From the Virtual Storage Console **Home** page, select **Configuration > Connection brokers**.



2. In the pop-up **Connection Brokers** box, click **Add**.



3. In the **Add Connection Broker** window, specify the following information:
  - a. For **Connection Broker Version**, select the connection broker name and version from the drop-down list.
  - b. For **Domain**, enter the domain containing the connection broker.
  - c. For **Connection name** (XenDesktop 5.0 only), enter the name given the Citrix XenDesktop 5.0 connection.
  - d. For **Hostname or IP Address** (VMware View Server only), enter the connection broker host name or IP address.
  - e. For **Username** (VMware View Server only), enter the domain user name.
  - f. For **Password** (VMware View Server only), enter the domain password.

## Removing connection brokers

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

### Steps

1. From the Virtual Storage Console **Home** page, select **Configuration > Connection brokers**.
2. In the pop-up **Connection Brokers** box, click **Remove**.
3. Click **Yes** to confirm.

## Configuring AutoSupport messages for backup jobs

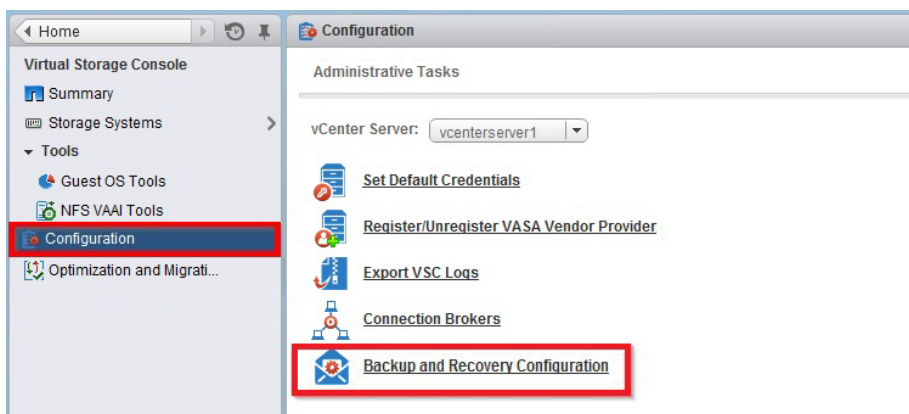
Information about specific errors that occur during backup and restore operations is automatically sent to the EMS log files. Enabling AutoSupport messages in Virtual Storage Console for VMware vSphere ensures that events such as a backup failure because of an error in taking a VMware snapshot or making a Snapshot copy are also sent to the EMS logs.

### Before you begin

You must have selected the vCenter Server that you want to use for this task.

### Steps

1. From the Virtual Storage Console **Home** page, click **Configuration > Backup and Recovery Configuration**.



2. In the **Backup and Recovery Configuration** dialog box, select the **Enable AutoSupport** check box to enable AutoSupport messages.

**Related tasks**

*Collecting the VSC for VMware vSphere log files* on page 136

## Configuring email alerts for backup jobs

You can configure in advance the email alert notifications that you want the Virtual Storage Console server to send when a backup job runs. The information you provide in the Backup and Recovery Configuration dialog box is then populated in the Schedule Backup wizard so that you do not have to type this information for every backup job.

**Before you begin**

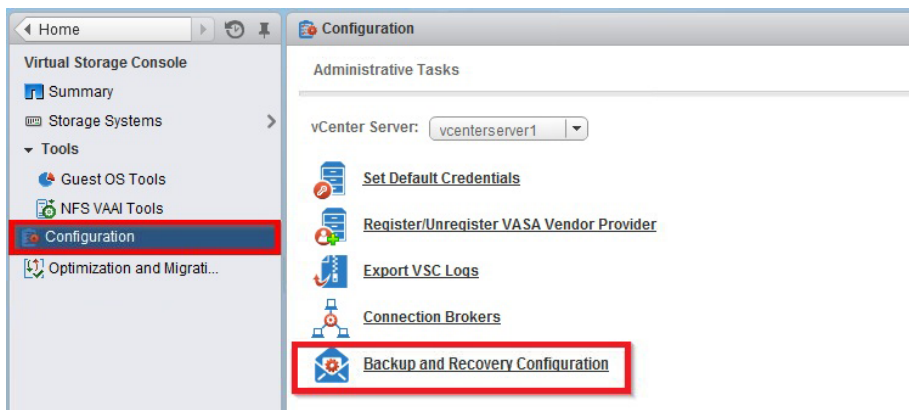
You must have selected the vCenter Server that you want to use for this task.

The following information must be available:

- Email address from which the alert notifications are sent
- Email address to which the alert notifications are sent
- Host name to configure the SMTP server

**Steps**

1. From the Virtual Storage Console **Home** page, click **Configuration > Backup and Recovery Configuration**.



2. In the **Backup and Recovery Configuration** dialog box, specify the email address and SMTP server from which the alert notifications are sent, as well as the email server to which the alert notifications are to be sent.
3. Optional: Click **Send Test Email** to verify that the outgoing email server to which the alert notifications are to be sent is working correctly.



## The preferences files

The preferences files contain settings that control Virtual Storage Console for VMware vSphere operations. You should not need to modify the settings in these files, except in some rare situations.

VSC has several preference files. These files include entry keys and values that determine how VSC performs operations. The following are some preference files that can be useful:

```
VSC_install_dir\etc\kamino\kaminoprefs.xml
```

```
VSC_install_dir\etc\caster\kaminosdkprefs.xml
```

```
VSC_install_dir\vsc\vscPreferences.xml
```

One reason to modify the preferences files is when you use iSCSI or NFS and the subnet is different between your ESX hosts and your storage system. In that situation, if you do not modify settings in the preferences files, provisioning a datastore fails because VSC cannot mount the datastore.

## Enabling datastore mounting across different subnets

If you use iSCSI or NFS and the subnet is different between your ESX hosts and your storage system, you need to modify two settings in the Virtual Storage Console for VMware vSphere preferences files. Otherwise, provisioning will fail because VSC cannot mount the datastore.

### About this task

When datastore provisioning fails, VSC logs the following error:

```
Unable to continue. No ip addresses found when cross-referencing kernel ip
addresses and addresses on the controller.
```

```
Unable to find a matching network to NFS mount volume to these hosts.
```

### Steps

1. Open the following files in a text editor:

```
VSC_install_dir\etc\kamino\kaminoprefs.xml
```

```
VSC_install_dir\etc\kamino\kaminosdkprefs.xml
```

2. Update both files as follows:

If you use...	Do this...
iSCSI	Change the value of the entry key <code>default.allow.iscsi.mount.networks</code> from ALL to your ESX host subnet masks.

If you use...	Do this...
NFS	Change the value of the entry key <code>default.allow.nfs.mount.networks</code> from ALL to your ESX host subnet masks.

The preferences files include example values for these entry keys.

**Note:** ALL does not mean all networks. It means that all matching networks between the host and storage system can be used to mount datastores. Specifying subnet masks enables mounting across the specified subnets only.

3. Save and close the files.

## Setting the frequency of NFS path optimization checks

By default, Virtual Storage Console for VMware vSphere checks the optimization of NFS paths every five minutes. You can change the frequency of these checks by modifying the `vscPreferences.xml` file.

### Steps

1. In the `etc/vsc/vscPreferences.xml` file, add the following entry:

```
<entry key="default.server.path.optimization.sleep"
value="value_in_seconds"/>
```

`value_in_seconds` is the amount of time, specified in seconds, that you want VSC to wait before rechecking the optimization of the NFS paths.

2. To confirm that the optimization checks are now occurring at the new interval, check the entries in the `log/vsc.log` log file.

The log file contains statements, similar to the following, that enable you to determine how often the path checks are occurring:

```
2013-10-17 16:26:58,665 [Thread-16] DEBUG (PathOptimizationManager)
- Refreshing controllers and recalculating path optimization
details...
2013-10-17 16:26:58,667 [Thread-16] DEBUG (PathOptimizationManager) -
Path optimization detection complete.
```

# 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.

The administrator handles setting up the RBAC roles. Depending on your system setup, you might have different administrators handling these two 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.

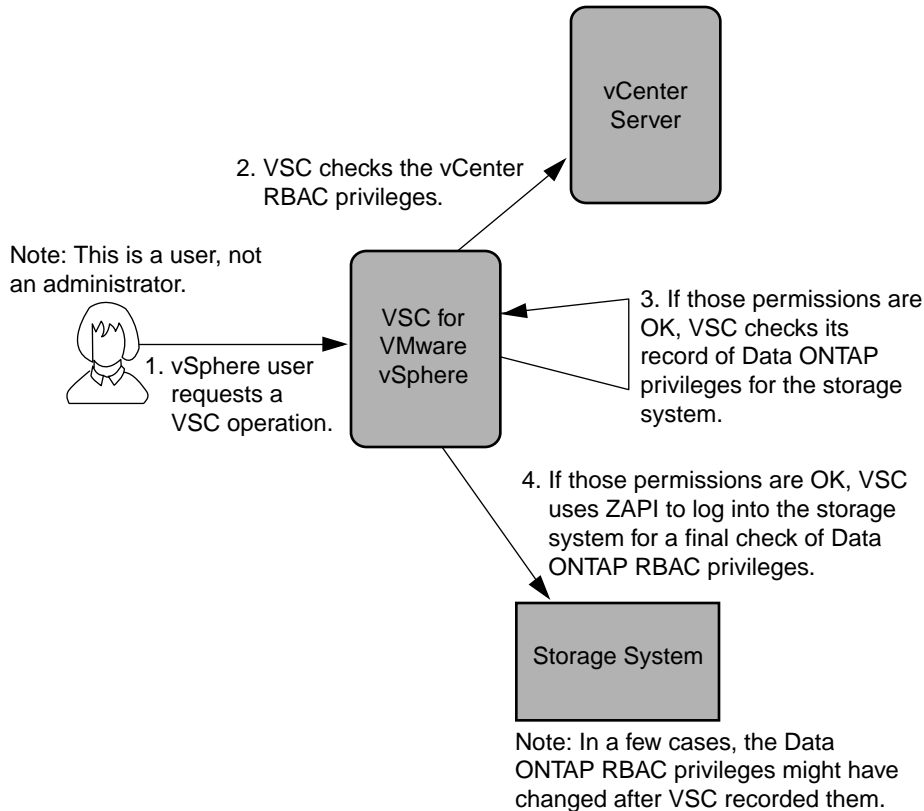
- **Data ONTAP RBAC**

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 VSC for each storage system that VSC manages. 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):



### Related concepts

[vCenter Server role-based access control features in VSC for VMware vSphere](#) on page 60

[Standard roles packaged with VSC for VMware vSphere](#) on page 65

[Data ONTAP role-based access control features in VSC for VMware vSphere](#) on page 69

## 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.

To successfully complete a task, you must have the appropriate vCenter Server RBAC permissions. During 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

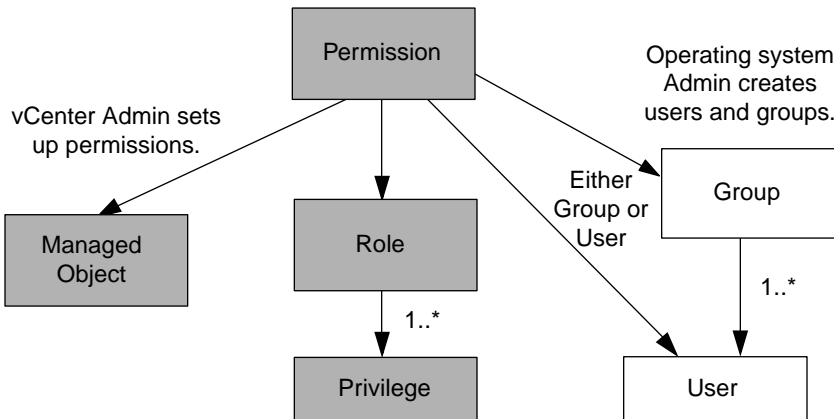
The vCenter Server recognizes permissions, not privileges. Each vCenter Server permission consists of three components.

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 *Virtual Storage Console for VMware vSphere Advanced RBAC Configuration*. 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 the 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 determines 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 to 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 users.

## Key points about assigning and modifying permissions

There are several key points to keep in mind when you are working with vCenter Server permissions. Whether 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 VSC 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-vcenter-server-51-security-guide.pdf>. NetApp follows the VMware recommendations for creating and using permissions.

## Standard roles packaged with VSC for VMware vSphere

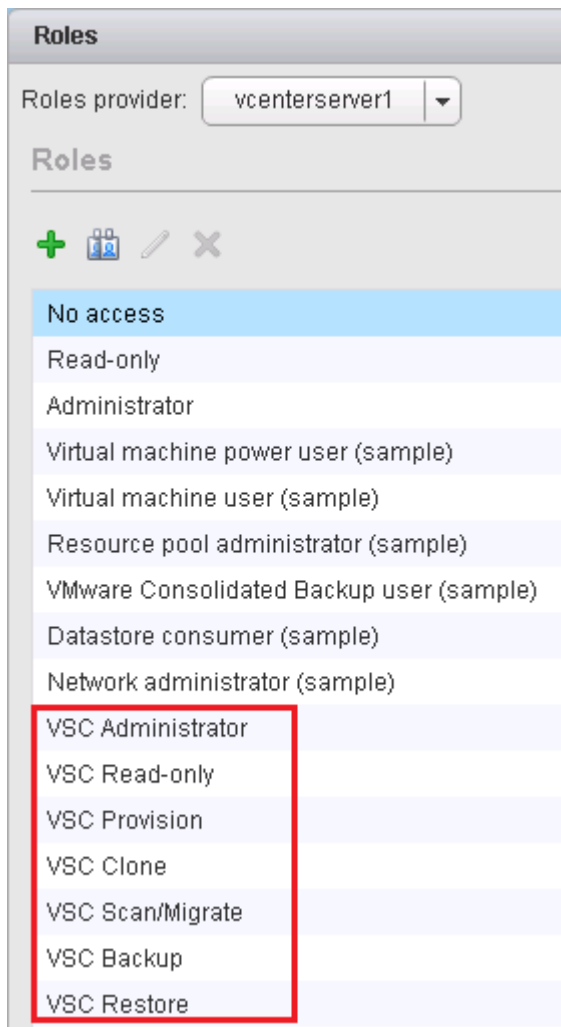
To simplify working with vCenter Server privileges and role-based access control (RBAC), Virtual Storage Console for VMware vSphere provides a set of standard VSC roles that enables users to perform key VSC tasks. There is also a read-only role that allows users to view VSC information, but not perform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



The roles VSC provides allow you to perform the following tasks:

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.

Role	Description
VSC Provision	<p>Provides all native vCenter Server and VSC-specific privileges necessary to provision storage.</p> <p>The user can perform the following tasks:</p> <ul style="list-style-type: none"> <li>• Create new datastores</li> <li>• Destroy datastores</li> <li>• View information about storage capability profiles</li> </ul> <p>The user cannot perform the following tasks:</p> <ul style="list-style-type: none"> <li>• Create clones</li> <li>• Reclaim space</li> <li>• Distribute templates</li> </ul>
VSC Clone	<p>Provides all native vCenter Server and VSC-specific privileges necessary to clone storage and view information about storage capability profiles.</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 and view information about storage capability profiles.</p> <p>With this role, the user can perform all tasks involving optimizing and migrating storage.</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) and view information about storage capability profiles.</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>

Role	Description
VSC Restore	<p>Provides all native vCenter Server and VSC-specific privileges necessary to restore vSphere objects that have been backed up and view information about storage capability profiles.</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>

Details about the privileges needed for these roles are included in *Virtual Storage Console for VMware vSphere Advanced RBAC Configuration Guide*.

## 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 the credentials (user name and password) that you set up within it 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 for a specific storage system, you cannot perform any tasks on a vSphere object hosted on

that storage system. 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 vSphere objects residing on storage
- Scanning, optimizing, and migrating vSphere objects residing on storage
- Backing up and recovering vSphere objects residing on storage

Using 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 (the user name and 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.

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.


There are several ways to create Data ONTAP roles:

- RBAC User Creator for Data ONTAP  
*[NetApp Community Document: RBAC User Creator for Data ONTAP](#)*
- OnCommand System Manager, which can be downloaded for either Windows or Linux platforms
- The CLI (command-line interface), using the `security login` set of commands  
The *System Administrator's Guide for Clustered Data ONTAP Administrators* contains information about using this command.

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 some of 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. **Discovery**  
The Discovery 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 Discovery 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. **Destroy 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. Discovery**

The Discovery role enables you to add storage systems.

**2. Backup-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 Discovery role.

If you are using VASA Provider for clustered Data ONTAP, you should also set up a PBM (policy-based management) role. That role will allow you to manage storage using storage policies. This role requires that you also set up the Discovery 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.

Details about the privileges needed for these roles are included in *Virtual Storage Console for VMware vSphere Advanced RBAC Configuration Guide*.

## 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.
- 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 VSC. Each time you log in to a storage system using these credentials, you are presented with the set of VSC functions that you set up in Data ONTAP when you created the credentials.

VSC performs an upfront privilege validation for Data ONTAP RBAC when you log in. VSC does not perform the upfront validation if the storage system is directly connected to a Storage Virtual Machine (SVM, formerly known as 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.

## Requirements for performing tasks in VSC for VMware vSphere

---

A few requirements impact many of the tasks that you can perform in Virtual Storage Console for VMware vSphere. You should understand requirements for privileges, running simultaneous tasks, and provisioning and cloning operations.

- Unless you log in as administrator, you must have the appropriate RBAC privileges correctly assigned to complete tasks.
- VSC must not be performing another operation on the target virtual machine or datastore that can have a negative impact on the currently executing operation.  
If VSC is performing a task on the target virtual machine or datastore, other tasks are temporarily unavailable.
- (NFS only) Before performing provisioning or cloning operations, you should have enabled the NFS Plug-in for VMware VAAI.  
While not required, installing the plug-in is a best practice because it reduces load from the host and places it on the storage system, which increases cloning efficiency.

### Related concepts

*[Authentication and user management with vCenter RBAC and Data ONTAP RBAC](#) on page 59*

*[VSC for VMware vSphere protects system resources by using lock management](#) on page 15*

*[How VSC for VMware features work with optional plug-ins, virtual appliances](#) on page 12*

*[NetApp NFS Plug-in for VAAI installation](#) on page 36*

## Navigating VSC for VMware vSphere

---

Virtual Storage Console for VMware vSphere integrates smoothly into the VMware vSphere Web Client and vCenter Server.

To help you distinguish between the VSC features and the vSphere Web Client features, the NetApp blue "N" icon appears in the screens and portlets associated with NetApp features.

You can access VSC tasks from both the vCenter and the Virtual Storage Console portions of the vSphere Web Client. Many of the Actions menus includes a **NetApp VSC** menu.

The Actions menu displays the VSC tasks available based on the object you are interacting with in the vSphere Web Client. For example, if you right-click on a vSphere cluster, the **NetApp VSC** menu displays an option to Provision Datastore.

You can initiate a task by doing any of the following:

- Right-clicking the object to display the Actions menu.
- Selecting the Actions menu from the menu bar.
- Clicking the icon in the menu bar on that page that is associated with the task you want to perform.

You can access data, such as datastores or virtual machines, either by using the navigation panel on the left side of the screen or by clicking an icon. For example, if you want to clone virtual machines, you can either click the VMs and Templates icon on the vSphere Web Client Home page or you can select **vCenter > VMs and Templates**. Both actions take you to the same place.

Many VSC pages provide a filter option. You can use the drop-down list in this option to organize the display to show only the columns with which you want to work. You can also use a text string to filter information. VSC performs a search based on the string and displays the results in a new window. For example, if you enter **aa** in the filter box and there is a datastore named "AA\_tester," then VSC displays information on that datastore. To return to your previous window, which lists all the datastores, clear the information in the filter box and press Enter.

If you are new to the vSphere Web Client, here are some tips for navigating through it:

- The Home icon takes you back to the Home view.
- The vCenter icon takes you to the vCenter.
- You can use the back arrow in the navigation pane to return to your previous location.
- The Recent Tasks pane on the right side of the screen lets you monitor the progress of a task that is under way.
- The Work In Progress pane shows you tasks that have been started and paused. To resume a paused task, click its name.

- The Alarms pane lists the Alarms that have occurred.

## Working with storage systems

---

Virtual Storage Console for VMware vSphere provides tools you can use to work with storage systems.

Using VSC, you perform tasks such as the following:

- Have VSC automatically discover storage systems
- Manually add and remove storage systems
- Set up default credentials for VSC to use when it adds storage systems
- Modify the credentials associated with a storage system
- Use VSC's interface to get a quick view of the storage system details

## Storage system discovery and credentials overview

Virtual Storage Console for VMware vSphere provides a single mechanism to discover storage systems and to set the storage credentials. The credentials provide the necessary Data ONTAP permissions to enable VSC users to perform tasks using the storage systems.

Before VSC can display and manage storage resources, it must discover the storage systems. As part of the discovery process, you must supply Data ONTAP credentials for your storage systems. These are the privileges (or roles) associated with the user name/password pair assigned to each storage system. These username and password pairs use Data ONTAP role-based access control (RBAC) and must be set up from within Data ONTAP. You cannot change their credentials from within VSC. You can, however, define Data ONTAP RBAC roles using a tool such as "RBAC User Creator for Data ONTAP".

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

When you add a storage system to VSC, you must supply an IP address for the storage system and the user name and password pair associated with that system. You can either set up default credentials that VSC can use during its storage system discovery process, or you can 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 `httpd.admin.enable` for the vFiler unit to enable discovery.

If your environment includes multiple vCenter Servers, when you add a storage system to VSC from the Storage Systems page, the **Add Storage System** dialog box displays a **vCenter Server** box where you can specify which vCenter Server the storage system is to be added to. If you add a

storage system by right-clicking a datacenter name, you do not have the option to specify a vCenter Server, because the server is already associated with that datacenter.

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, VSC begins its automatic background discovery process.
- You click the **Update All** icon or select it from the Actions menu ( **Actions > Netapp VSC > Update All**).

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

All of the VSC features require specific permissions to perform tasks. You can limit what users can do based on the credentials associated with the Data ONTAP role. All users with the same storage system user name and password pair share the same set of storage system credentials and can perform the same operations.

## Default credentials simplify administrating storage systems

You can use Virtual Storage Console for VMware vSphere to set default credentials for storage systems, hosts, and virtual machines. Setting default credentials that are valid for the storage system means that you do not need to manually enter credentials each time VSC adds a storage system.

For example, when VSC discovers a new storage system, it attempts to log in using the default credentials. If the login fails, the storage system status is set to Authentication Failure, and you must enter credentials manually.

You can set the default credentials by clicking **Configuration > Set Default Credentials** from the Virtual Storage Console Home page. VSC displays a pop-up box with tabs to let you set credentials for the different objects.

If you change the default credentials and select **Update All**, VSC uses the new credentials and attempts to log in to any storage system that has a status of either "Authentication Failure" or "TLS is not configured".

## Specifying default credentials

You can use Virtual Storage Console for VMware vSphere to create default credentials for a storage system, host, and virtual machine.

### Before you begin

You must have selected the vCenter Server that you want to use for this task.

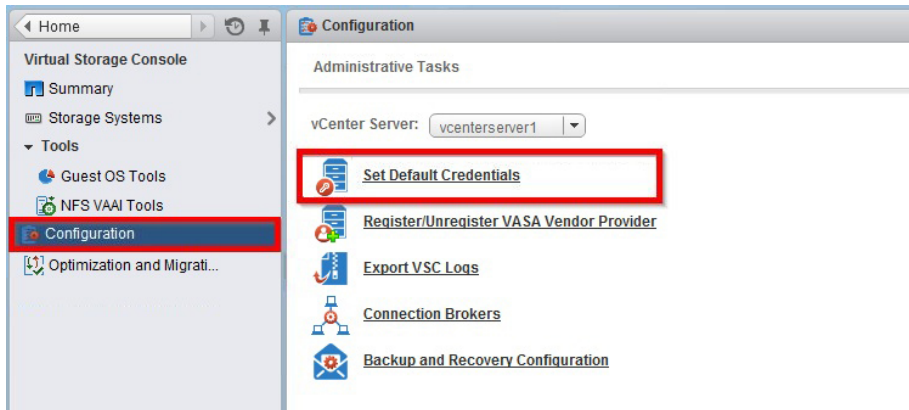
### About this task

When you set up default credentials, VSC uses them to attempt to log in to a storage system it has just discovered.

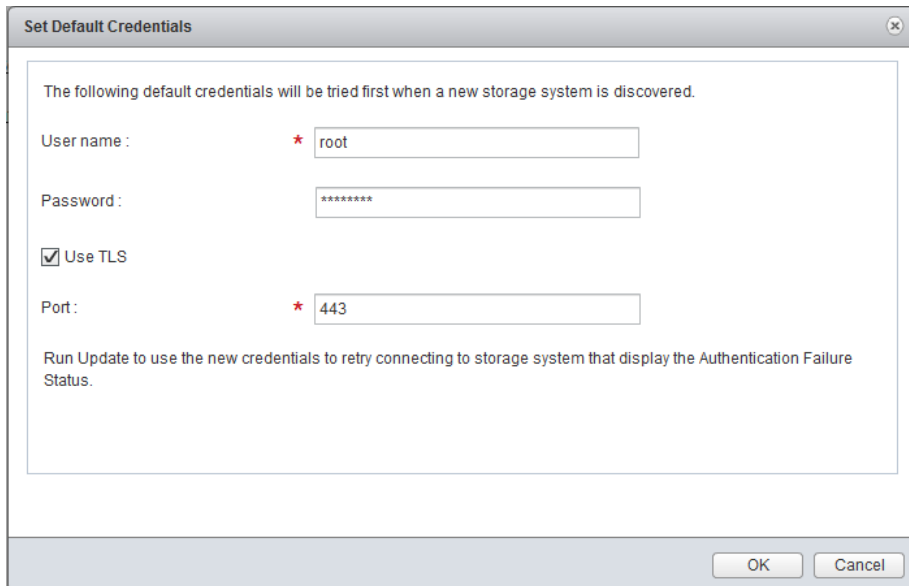
If the default credentials do not work, you must manually log in.

## Steps

1. From the Virtual Storage Console **Home** page, click **Configuration > Set Default Credentials**.



2. In the **Set Default Credentials** pop-up box, enter the credentials for the storage system.



Storage system field	Explanation
User name/password	Storage controller credentials are assigned in Data ONTAP 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 use VSC to change the roles associated with that user name/password pair. To do that, you must use a tool such as "RBAC User Creator for Data ONTAP."
Use TLS	Select this box to enable Transport Layer Security (TLS).
Port	The default management port number is 443 if the TLS box is selected and 80 if it is not selected.  These are the Data ONTAP defaults. If you toggle the TLS check box, the port number switches between 443 and 80. You can specify a different port number. If you do that, then toggling the TLS check box only changes the TLS state in the dialog box.

3. After you enter the information, click **OK**.

#### After you finish

If you updated the storage system credentials because a storage system reported Authentication Failure Status, select **Update Hosts and Storage Systems**, which is available from the **Actions > NetApp VSC** menu. When you do this, VSC tries to connect to the storage system using the new credentials.

#### Related tasks

[Discovering storage systems and hosts](#) on page 82

## Tunneled vFiler units and SVMs discovered automatically

Virtual Storage Console for VMware vSphere automatically supports vFiler and Storage Virtual Machine (SVM, formerly known as Vserver) tunneling for the storage systems. You do not need to manually add these vFiler units and SVMs.

When you enter information for a cluster administrative LIF or vfiler0, VSC discovers all the subordinate vFiler units and SVMs.

If you are using VSC's backup and restore features, there are certain environments where these features can only access the physical storage system, not the vFiler unit, for communication on a storage network. You need to have vFiler tunneling enabled in order to use the backup and restore features to create Snapshot copies.



## 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.

### About this task

You do not need to perform this task if your vFiler units were 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 these steps for each vFiler unit you want to manage using VSC.

## Enabling discovery of vFiler units on private networks

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

### Before you begin

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

### Steps

1. You can add a storage system by using either the **Add** icon or the **Add Storage System** menu option:

Starting location	Action
Virtual Storage Console Home page	<ol style="list-style-type: none"> <li>a. Click <b>Storage System</b>.</li> <li>b. Click the <b>Add</b> icon.</li> </ol>
VMware vSphere Web Client Home page	<ol style="list-style-type: none"> <li>a. Click the <b>Storage</b> icon.</li> <li>b. Select a datacenter.</li> <li>c. Click the <b>Actions &gt; NetApp VSC &gt; Add Storage System</b>.</li> </ol>

2. In the **Add Storage System** dialog box, enter the management IP address and credentials for the pFiler and then click **OK**.

When you add storage from the **VSC Storage System** page, you must also specify the vCenter Server where the storage will be located. The Add Storage System pop-up box provides a drop-down list of the available vCenter Servers. VSC does not display this option if you are adding storage to a datacenter that is already associated with a vCenter Server.

### Result

VSC discovers any vFiler units belonging to the parent of the vFiler that provide storage to ESX hosts.

## Discovering storage systems and hosts

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

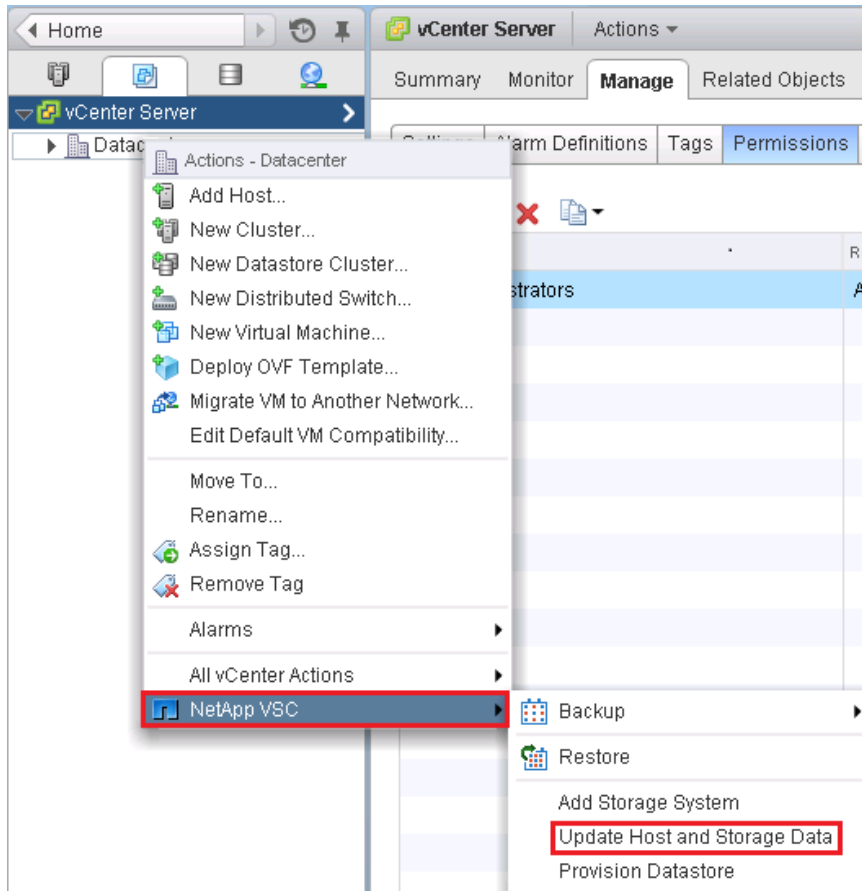
### About this task

You can discover new storage systems or update information on them to get the latest capacity and configuration information at any time. You can also modify the credential that VSC uses to log into the storage system.

The discovery process also 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.

### Steps

1. From the the vSphere Web Client **Home** page, select **vCenter**.
2. Right-click a datacenter and select **Actions > NetApp VSC > Update Host and Storage Data**.



3. VSC displays a **Confirm** dialog box that warns you that this operation can take a long time. Click **OK**.
4. Right-click any discovered storage controllers that have the status **Authentication Failure** and select **Modify**.
5. Fill in the information in the **Modify Storage System** dialog box.

#### After you finish

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

## Manually adding storage systems

Each time you start the VSC Windows service or select **Update All**, Virtual Storage Console for VMware vSphere automatically discovers the available storage systems. You can also manually add storage systems to VSC.

### About this task

If you have a large number of storage systems, manually adding a new one might be faster than using **Update All** to discover it.

### Steps

1. You can manually add a storage system to VSC using either the **Add** icon or the **Add Storage System** menu option:

Starting location	Action
Virtual Storage Console Home page	<ol style="list-style-type: none"><li>a. Click <b>Storage System</b>.</li><li>b. Click the <b>Add</b> icon.</li></ol>
VMware vSphere Web Client Home page	<ol style="list-style-type: none"><li>a. Click the <b>Storage</b> icon.</li><li>b. Select a datacenter.</li><li>c. Click the <b>Actions &gt; NetApp VSC &gt; Add Storage System</b> .</li></ol>

2. In the **Add Storage System** dialog box, enter the management IP address and credentials for that storage system.

You can also change the defaults for TLS and the port number in this dialog box.

When you add storage from the VSC **Storage System** page, you must also specify the vCenter Server where the storage will be located. The Add Storage System dialog box provides a drop-down list of the available vCenter Servers. VSC does not display this option if you are adding storage to a datacenter that is already associated with a vCenter Server.

**Add Storage System**

vCenter Server: vcenterserver1

Name or IP Address : \*

User name : \*

Password :

☒ Use TLS

Port : \*

OK Cancel <<Options

3. Click **OK** when you have added all the necessary information.

## Refreshing the storage system display

You can use the update feature provided by Virtual Storage Console for VMware vSphere to refresh the information about storage systems and force VSC to discover storage systems. This can be especially useful if you changed the default credentials for the storage systems after receiving an authentication error.

### About this task

You should always perform an update operation if you changed the storage system credentials after a storage system reported an Authentication Failure Status. During the update operation, VSC tries to connect to the storage system using the new credentials.

Depending on your system setup, this task can take a long time to complete.

### Steps

1. Go to the **Storage** page by clicking **Storage** from either the navigation pane of the Virtual Storage Console **Storage** page or the icon on the VMware vSphere Web Client **Home** page.

2. Start the update:

Location	Click ...
Virtual Storage Console	The <b>Update All</b> icon.
vCenter	<b>Actions &gt; NetApp VSC &gt; Update Host and Storage Data</b>

3. Click OK at the Confirm dialog box.

4. Click OK at the Success Message dialog box.

This operation works in the background.

## Removing storage systems from VSC

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

### About this task

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

### Step

1. You can remove a storage system by clicking **Storage** from either the VSC **Home** page or the VMware vSphere Web Client **Home** page.

Starting location	Action
Virtual Storage Console Home page	<b>a.</b> Click <b>Storage</b>  <b>b.</b> Right-click a storage system and select <b>Actions &gt; NetApp VSC &gt; Delete</b>
VMware vSphere Web Client Home page	<b>a.</b> Click the <b>Storage</b> icon  <b>b.</b> Right-click a datastore and select <b>Actions &gt; NetApp VSC &gt; Destroy</b>

The action VSC takes depends on whether the storage system is attached to a host:

If the storage system is ...	VSC...
Not attached to host	Removes the storage system

If the storage system is ...	VSC...
Attached to a host	Displays an error message and does not change the storage system

## Correcting storage system names displayed as "unknown"

If Virtual Storage Console for VMware vSphere displays a storage system name as "unknown," you can modify the credentials and add the management IP address of the storage system in the Modify Storage System pop-up box.

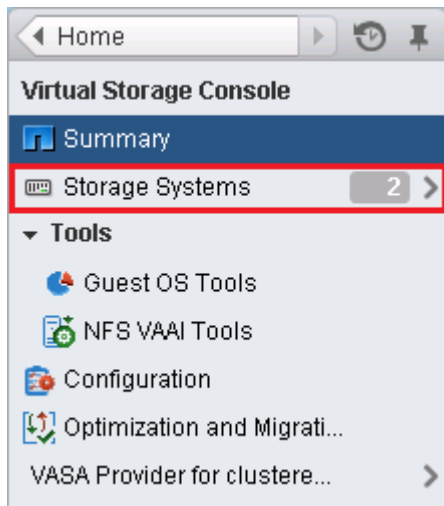
### About this task

An "unknown" name 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. From the Virtual Storage Console **Home** page, click **Storage Systems**.



2. Right-click the storage system shown as "unknown" and select **Modify Storage System**.



3. Enter the management IP address of the storage system and the storage system credentials in the **Modify storage system -unknown-** pop-up box.


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

## Managing settings for volumes

You can specify advanced settings that apply when you provision new volumes on the storage system. These settings include features such as enabling thin provisioning and deleting a volume when the last LUN contained in it is deleted.

### Steps

1. From the Virtual Storage Console **Home** page, click **Storage Systems**.
2. Right-click on a the storage system and select **Modify** from the Actions menu.



3. In the **Modify Storage System** pop-up box, click the **Provisioning Options** tab.

If storage system properties have been locked to prevent changes, click **Enable Editing**. In the resulting dialog box, enter the username and password for that storage system. Virtual Storage Console for VMware vSphere will not let you change the settings until you enter this information.

4. After you have enabled editing, check the boxes next to the advanced options and enter information to modify the FlexVol efficiency settings to match those on the LUN being deployed.

**Note:** 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.



Option	Explanation
Create a new volume for a new LUN	Selecting this option creates a FlexVol with the same name as the LUN. If a volume with that name already exists, VSC appends a number to the volume name; for example: Volname01
Reserve space for volumes that contain thin provisioned LUNs	Checking this results in having a thin LUN in a thick volume when a thin LUN is chosen.
Thin provision volume clones	Sets the 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 the last LUN on it is deleted.
Buffer space between volume and LUN (GB)	Specifies the amount of additional capacity in a volume that contains a LUN-based datastore.

5. Click **OK** when you finish.

## MultiStore vFiler units are displayed differently

Virtual Storage Console for VMware vSphere displays vFiler units differently than physical storage controllers.

When you have a vFiler unit created with the optional MultiStore feature of Data ONTAP software, VSC 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 VSC. The **Normal** icon means that VSC 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.`
- No aggregate information is displayed for vFiler units.  
Direct vFiler units and direct Storage Virtual Machines (SVMs) do not have aggregates to display.

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

Virtual Storage Console for VMware vSphere supports connecting a storage system directly to either a Storage Virtual Machine (SVM, formerly known as Vserver) or a cluster-management LIF. When the storage connects directly to an SVM, 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 system connects directly to an SVM:

- **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 an SVM.
- **NFS path checking**  
When you are using clustered Data ONTAP and a cluster-management LIF, VSC can query the storage system to determine whether that storage system is using a direct or indirect path. VSC then reports this information and supplies information that you can use to set up a direct path. Better performance is normally seen when direct paths are used. If a storage system connects directly to an SVM, VSC cannot query the storage system 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 system is directly attached to an SVM.
- **Load-sharing mirrors updates**  
VSC cannot update the mirror if you use load-sharing mirrors for your SVM's root volume.
- **EMS logging**  
VSC cannot perform EMS logging when the storage system is directly attached to an SVM.

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

## 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 negatively affect I/O performance and should be corrected. Virtual Storage Console for VMware vSphere provides tools to scan for direct and indirect NFS paths and provide you with the information you need to manually correct paths.

An indirect path 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.

VSC 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 clicking the **NAS** tab in the Related Objects page for a storage system 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:** VSC 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, VSC writes the path information to a log file.

If a direct Storage Virtual Machine (SVM) connection is made, VSC 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. Virtual Storage Console for VMware vSphere provides information about the paths; however, the task of moving the path must be performed by a storage administrator using either the storage system console or a NetApp tool such as System Manager.

### Before you begin

Only a storage administrator should change the path.

### About this task

The need to change paths only occurs when your clustered Data ONTAP configuration has 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. From the Virtual Storage Console **Home** page, click **Storage** > *<storage system name>* > **Related Objects** > **NAS**.
2. 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.

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

VSC 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.

4. 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.

The **Direct Data Path Choices** pop-up contains all the 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.

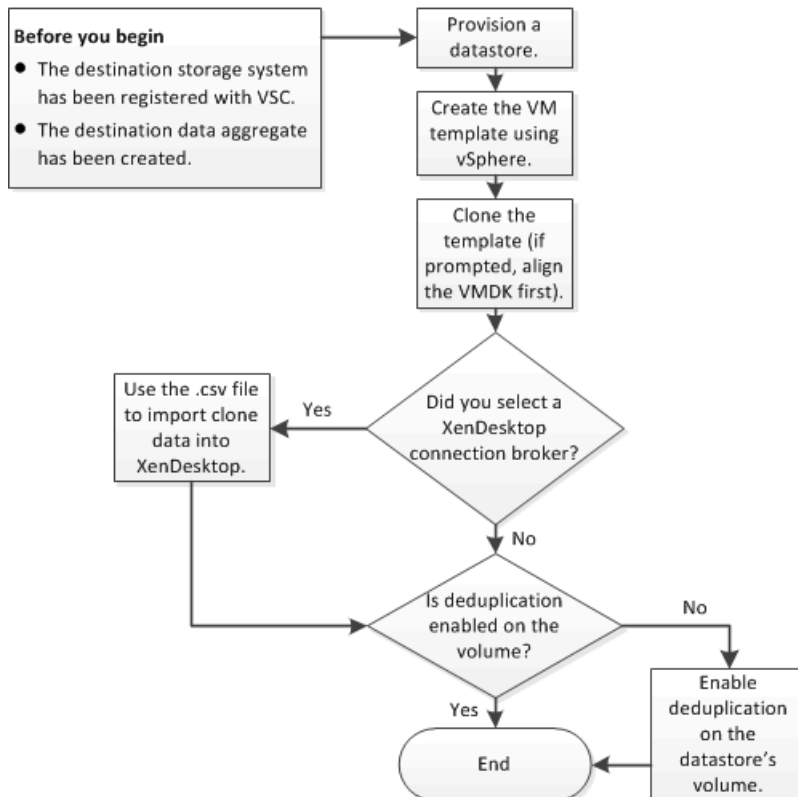
5. Use either the storage system console or a NetApp tool such as System Manager to change the path.

**Note:** Whenever the path value changes, VSC writes the information to a log file.

## Deploying virtual machines on NetApp storage

You can use Virtual Storage Console for VMware vSphere to deploy virtual machines by provisioning datastores and then rapidly cloning the virtual machines from a template into the provisioned datastores.

The following workflow shows how you can provision datastores using the Datastore Provisioning wizard before using the Create Rapid Clones wizard to clone virtual machines. This workflow is beneficial because the Datastore Provisioning wizard allows you to specify a storage capability profile, which ensures that consistent Service Level Objectives (SLOs) are maintained and simplifies the provisioning process, if you use VASA Provider for clustered Data ONTAP.



## Provisioning datastores

Provisioning a datastore creates a logical container for your virtual machines and their VMDKs. You can provision a datastore and attach it to a single host, to the hosts in a cluster, or to the hosts in a datacenter by using the Datastore Provisioning wizard.

### Before you begin

- To provision NFS datastores to vFiler units, you must have added the default vFiler unit (vFiler0) to Virtual Storage Console for VMware vSphere.
- To provision a datastore to a Storage Virtual Machine (SVM, formerly known as Vserver) that is directly connected to VSC, you must have added the SVM to VSC using a user account that has the appropriate privileges, not the default vsadmin user account or vsadmin role.
- If you use NFS or iSCSI and the subnet is different between your ESX hosts and your storage system, NFS or iSCSI settings in the VSC preferences file must include the ESX host subnet masks.

For more information, see [Enabling datastore mounting across different subnets](#) on page 57.

### About this task

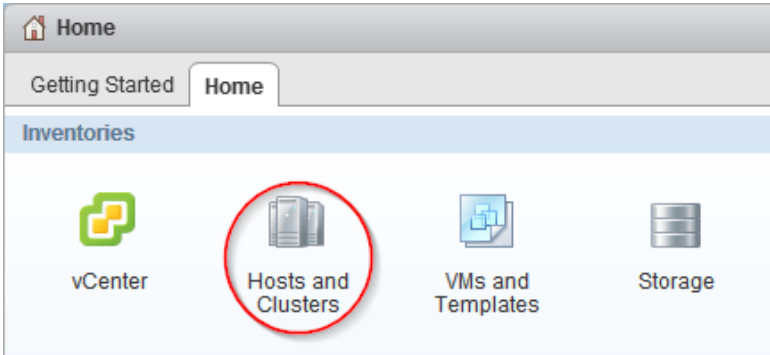
VSC enables you to provision datastores from wizards other than the Datastore Provisioning wizard (for example, from the Create Rapid Clones wizard). Using the Datastore Provisioning wizard, though, is beneficial because it allows you to specify a storage capability profile, which ensures consistent Service Level Objectives (SLOs) and simplifies the provisioning process.

VSC creates a datastore on an NFS volume or a LUN. The following happens when you provision a datastore:


- For an NFS datastore, VSC creates an NFS volume on the storage system and updates export policies.
- For a VMFS datastore, VSC creates a new volume (or uses an existing volume, if you selected that option), and creates a LUN and an igroup.

### Steps

1. From the vSphere Web Client **Home** page, click **Hosts and Clusters**.



2. In the navigation pane, expand the datacenter where you want to provision the datastore.



3. Specify the hosts to which you want to mount the datastore:

To make the datastore available to...	Do this...
All hosts in a datacenter	Right-click the datacenter and select <b>NetApp VSC &gt; Provision Datastore</b> .
All hosts in a cluster	Right-click a cluster and select <b>NetApp VSC &gt; Provision Datastore</b> .
A single host	Right-click a host and select <b>NetApp VSC &gt; Provision Datastore</b> .

4. Complete the pages in the **Datastore Provisioning** wizard to create the datastore.

- a. In the **Name and type** page, specify a datastore name, datastore type, and select a storage capability profile, if desired.

You can specify an existing storage capability profile that the wizard will use when defining the type of storage that you need for your virtual machines. The storage capability profile determines the following storage features: availability, disaster recovery, performance, protocol, and space efficiency. Storage capability profiles are available only if you installed and registered the VASA Provider for clustered Data ONTAP. You can select a default storage capability profile, which ships with the VASA Provider, a profile that you created, or a profile that was auto-generated. To provision a datastore without a storage capability profile, select **None**.

- b. In the **Storage system** page, specify the storage system that you want to use for the datastore.

**Note:** When connecting directly to an SVM, the provisioning operation might begin, but later fail due to insufficient privileges for the SVM user. SVM privileges are not visible to VSC prior to the operation. If the operation fails, you need to modify the privileges for that SVM user. You can check VSC logs for messages that identify the failed command. An alternative is to use the "RBAC User Creator for Data ONTAP" tool.

- c. In the **Details** page, specify details about the datastore that you want to create.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Thin provision	Allocates space on the volume when data is written, which allows you to provision more storage than is currently available. If disabled, space is reserved immediately.  You must closely monitor the available space in the containing aggregate because thin provisioning can oversubscribe the available space. In an NFS configuration, you can enable auto grow to automatically expand the datastore when space is needed. Make sure that the value you specify for auto grow is larger than the size of the datastore.
Aggregate	Defines the aggregate on which you want to create a new volume. If you selected an SVM that is directly connected to VSC, striped aggregates appear as available; however, they are not supported. Provisioning to a striped aggregate will fail.
Volume	Specifies the volume on which you want to create the datastore. For clustered Data ONTAP, you should not create a datastore in the Storage Virtual Machine (SVM) root volume.
Auto grow	(NFS only) Automatically expands the datastore by the specified increment when space is needed, up to the size limit. This size limit you specify must be larger than the existing datastore.
Datastore cluster	Adds the datastore to a cluster if the Storage Distributed Resource Scheduler (SDRS) feature is enabled on the vCenter Server. Do not mix datastores with varying offsets in the same cluster and do not mix optimized and non-optimized datastores.

- d. In the **Ready to complete** page, review the summary of your selections and click **Finish**.

## Result

VSC creates the datastore.



**After you finish**

Add virtual machines to the datastore.

**Related concepts**

[Storage system discovery and credentials overview](#) on page 77

[How to configure Data ONTAP role-based access control for VSC for VMware vSphere](#) on page 73

## Cloning virtual machines from a template

Setting up virtual machines can be a lengthy process. If you need to deploy multiple identical virtual machines, you can save time by setting up a single virtual machine as the template and then rapidly cloning virtual machines from that template.

**Before you begin**

- You should have created a virtual machine template using VMware vSphere.
- You should have installed the NFS Plug-in for VMware VAAI.  
While not required, installing the plug-in is a best practice because it reduces load from the host and places it on the storage system, which increases cloning efficiency.

**About this task**

Cloning performance is affected by many factors, including the vCenter Server hardware configuration, the number and hardware configuration of the ESX/ESXi hosts, and the current load on the vCenter Server and the hosts.

Performance can degrade if you request a large number of clones in a single operation. If you need to create a large number of clones, consider whether you should perform two cloning operations instead of one. For example, instead of requesting 2,000 clones in each operation, you might perform two operations that each request 1,000 clones each.

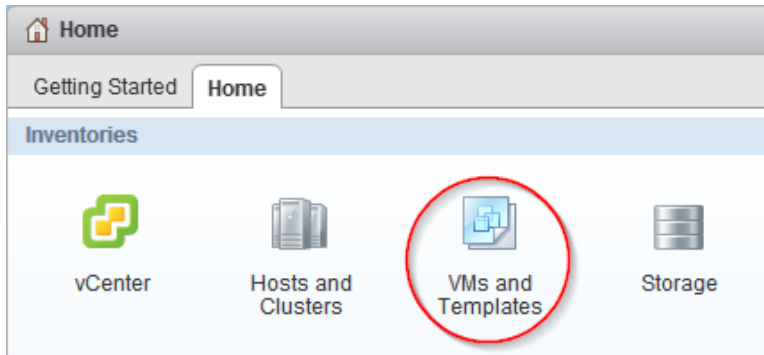
**Steps**

1. Power down the virtual machine template.


Powering down the virtual machine is recommended because it enables VSC to check the virtual machine's alignment and perform the cloning process faster. Checking the alignment is important because you should not clone a functionally aligned or misaligned virtual machine. Doing so can result in misaligned clones.

For more information about functionally aligned and misaligned virtual machines, see [Optimizing I/O performance with online alignment and migration of virtual machines](#) on page 103.

2. From the vSphere Web Client **Home** page, click **VMs and Templates**.



3. In the navigation pane, expand the datacenter that contains the virtual machine template.



4. Right-click the virtual machine template and select **NetApp VSC > Create Rapid Clones**.

If VSC warns you that the virtual machine is misaligned or functionally aligned, take the virtual machine offline and use a tool like VMware vCenter Converter to fix the VMDK alignment before you proceed.

**Note:** If you do not fix the alignment of a functionally aligned virtual machine, the clones can be misaligned if the destination datastores are not optimized for the VMDK layout of the clones.

5. Complete the pages in the **Create Rapid Clones** wizard to clone the virtual machines.

- a. In the **Clone destination** page, select a destination for the clones (a host, host cluster, or datacenter) and a folder to hold the clones (the default is no folder).

If you choose a cluster or datacenter, VSC spreads the virtual machines evenly across the hosts.

- b. In the **Clone folder** page, select a folder for the clones.

**Tip:** You can create folders on the VMs and Templates page using vCenter actions.

**Note:** This page appears if you chose the **Select a folder** option in the Clone destination page.

- c. In the **Disk format** page, select a disk format for the clones.

If you choose the thin provisioned format or the thick format, the wizard warns you that a vSphere clone operation might be required, which can take longer.

- d. In the **Virtual machine details** page, specify details about the virtual machine clones.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Number of virtual processors	Specifies the number of virtual CPUs for the virtual machines.
Upgrade hardware version?	Upgrades the hardware version of the virtual machine clone if the destination host supports a later version.
Connection broker version	Automatically imports clone data into a VMware View Server or creates a .csv file that you can import into Citrix XenDesktop.
Customization specification	Applies a VMware specification to the new virtual machines. Refer to your VMware documentation for information about customization specifications.
Stagger powering on the virtual machines	<p>Stagger the start up of virtual machines to avoid overwhelming your system. You should select this option if you have a large number of virtual machines. The number of virtual machines to start per minute depends on your system environment.</p> <p><b>Note:</b> If a problem prevents VSC from starting some of the virtual machines, the delay could result in VSC powering on a large number of virtual machines at once. For example, if you specify 10 virtual machines per minute and the start is delayed by five minutes, VSC starts 50 virtual machines at once. After the delay, VSC starts the specified number of virtual machines per minute.</p>

- e. In the **Storage system details** page, select the storage system where you want to provision the clones.
- f. In the **Datastore options** page, choose basic mode or advanced mode to specify the datastore options.

The advanced mode is a good choice if you want to distribute configuration files and VMDK files across multiple datastores.

- g. In the **Datastore details** page, select existing datastores or create new datastores for the clones.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance when you create new datastores:

Field	Description
Number of datastores	Specifies the number of datastores to create for the clones. The maximum is 256. The number of clones must be evenly divisible by the number of datastores.
Thin provision	<p>Allocates space on the volume when data is written, which allows you to provision more storage than is currently available. If disabled, space is reserved immediately.</p> <p>You must closely monitor the available space in the containing aggregate because thin provisioning can oversubscribe the available space. In an NFS configuration, you can enable auto grow to automatically expand the datastore when space is needed. Make sure that the value you specify for auto grow is larger than the size of the datastore.</p>
Size (GB)	Specifies the size per datastore.
Aggregate	Defines the aggregate on which you want to create a new volume. If you selected an SVM that is directly connected to VSC, striped aggregates appear as available; however, they are not supported. Provisioning to a striped aggregate will fail.
Volume	Specifies the volume on which you want to create the datastore. For clustered Data ONTAP, you should not create a datastore in the Storage Virtual Machine (SVM) root volume.
Auto grow	(NFS only) Automatically expands the datastore by the specified increment when space is needed, up to the size limit. This size limit you specify must be larger than the existing datastore.
Datastore cluster	Adds the datastore to a cluster if the Storage Distributed Resource Scheduler (SDRS) feature is enabled on the vCenter Server. Do not mix datastores with varying offsets in the same cluster and do not mix optimized and non-optimized datastores.

- h. In the **Connection broker** page, specify the VMware view or Citrix XenDesktop connection broker to which you want to import clone data.

If your connection broker does not appear, you must first add it by going to **Virtual Storage Console > Configuration > Connection Brokers**.

**Note:** This page appears if you chose a connection broker version in the Virtual machine details page.

- i. In the **Ready to complete** page, Review the summary of your selections and click **Finish**.

**Result**

VSC creates the virtual machine clones and creates a `.csv` file that includes details about the cloning process. The file, named `import_generic_timestamp.csv`, is created here: `VSC_install_dir\etc\kamino\exports`

If you chose a VMware View connection broker, VSC automatically imports clone data into the VMware View Server.

If you chose a XenDesktop connection broker, VSC creates a `.csv` file that you can use to import into XenDesktop. The file, named `xenDesktop_timestamp.csv`, is created here: `VSC_install_dir\etc\kamino\exports`

**After you finish**

If you chose a XenDesktop connection broker, use the `.csv` file with the Citrix Access Management Console (XenDesktop 4) to create a new desktop group or with Desktop Studio (XenDesktop 5) to create or modify an existing catalog.

**Related concepts**

[How VSC for VMware features work with optional plug-ins, virtual appliances](#) on page 12

[NetApp NFS Plug-in for VAAI installation](#) on page 36

## Increasing storage efficiency by enabling deduplication

Deduplication is a Data ONTAP feature that reduces physical storage space by eliminating duplicate data within a volume. Deduplication enables virtual machines to share the same common data in a datastore, similar to how they share system memory. You should enable deduplication if it is not enabled.

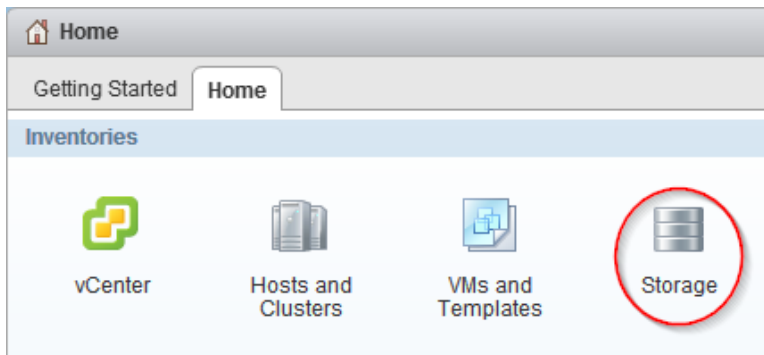
**About this task**

When Virtual Storage Console for VMware vSphere creates a new volume for a datastore, it enables deduplication by default. If you created the volume through Data ONTAP or OnCommand System Manager, deduplication is not enabled by default. Enabling deduplication is a best practice.


For more information about deduplication, refer to the *Storage Management Guide* for your version of Data ONTAP.

**Steps**


1. From the vSphere Web Client **Home** page, click **Storage**.



2. In the navigation pane, expand the datacenter that contains the datastore.



3. Select the datastore.
4. Click the **Summary** tab if it does not automatically display.
5. In the **Deduplication Details** pane, view the **State** field to determine whether deduplication is enabled or disabled.



6. If deduplication is disabled, at the bottom of the **Deduplication Details** pane, click **Enable**.  
VSC enables deduplication on the volume. Deduplication runs daily at midnight.
7. Optional: To start deduplication immediately, click **Start**.

### After you finish

You can view the Volume Space Saving field to identify the percentage and amount of storage that deduplication saved and the Volume Space Shared field to identify the amount of shared data.

# Optimizing performance by aligning the I/O of misaligned virtual machines non-disruptively

---

Virtual Storage Console for VMware vSphere can scan your datastores to determine the alignment status of virtual machines. You can then use VSC to functionally align the I/O to certain misaligned virtual machines without having to power them down.

## About this task

Online alignment is a good choice for virtual machines that you cannot take offline. When possible, you should take the virtual machine offline and physically align the VMDK using a tool such as VMware vCenter Converter.

## Steps

1. [Scan datastores to determine the alignment status of virtual machines](#) on page 103
2. [Check the alignment status of virtual machines](#) on page 105
3. [Align the I/O to any misaligned virtual machines](#) on page 107

## Related concepts

*[Methods for migrating virtual machines](#) on page 14*

*[How VSC for VMware vSphere optimizes I/O performance of misaligned virtual machines](#) on page 13*

# Scanning datastores to determine the alignment status of virtual machines

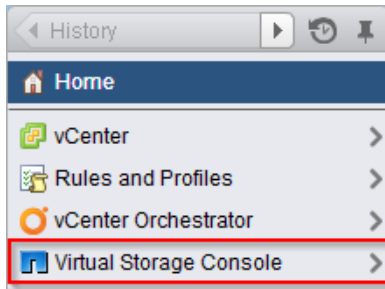
You should periodically scan datastores to identify whether any of your virtual machines are misaligned. A misaligned virtual machine can negatively affect I/O performance.

## About this task

- It is a good practice to scan datastores during noncritical production times.  
The time required to perform a scan can increase as more virtual machines are scanned.
- Virtual Storage Console for VMware vSphere uses VMware snapshots to scan virtual machines that reside in VMFS datastores and then deletes the snapshots when they are no longer needed.

## Steps

1. From the vSphere Web Client **Home** page, click **Virtual Storage Console**.



2. In the navigation pane, click **Optimization and Migration**.
3. If your environment includes multiple vCenter Servers, you must select the server that contains the datastores for which you want to scan.
4. Initiate a scan of the datastores:

To...	Do this...
Schedule recurring scans	<ol style="list-style-type: none"><li>a. Click <b>Global Scan Schedule</b> and set a schedule for the scan. You should schedule the scans during noncritical production times.</li><li>b. Click <b>OK</b>.</li><li>c. To add or remove datastores from the global scan, select a datastore in the datastores table (the table at the top of the Optimization and Migration page) and click <b>Exclude</b> or <b>Include</b>.</li></ol>
Initiate a one-time scan	<ol style="list-style-type: none"><li>a. Choose to scan all datastores or specific datastores:<ul style="list-style-type: none"><li>• To scan all datastores, click <b>Scan All</b>.</li><li>• To scan one or more datastores, select the datastores and click <b>Scan selected</b>.</li></ul></li><li>b. Click <b>OK</b> to confirm the scan.</li></ol>

## Result

VSC scans the datastores according to the options that you selected.

## After you finish

Check the alignment status of the virtual machines after VSC finishes scanning them.



# Checking the alignment status of virtual machines


Check the alignment status of a virtual machine to determine whether it is aligned or misaligned. You should fix the alignment of a misaligned virtual machine to optimize I/O performance.

## Before you begin


You should have scanned your datastores.

## Steps


1. In the **Optimization and Migration** page, select a datastore from the datastores table (the table at the top of the page).




Global Scan Schedule



Exclude



Include



Scan selected

Datastore Name	Host(s)	Type	Optimized	Excluded
datastore1	dilbert	VMFS	Unknown	No
datastore2	dilbert	VMFS	No	No

After you select a datastore, the virtual machines running on that datastore appear in the virtual machines table (the table at the bottom of the Optimization and Migration page).

2. In the virtual machines table, view the **Alignment State** column to identify the alignment state of the virtual machines on that datastore.

**Tip:** To quickly identify any misaligned virtual machines, select **Show only misaligned virtual machines**.

**Note:** The View Disk Details button shows you the alignment details of a virtual machine's VMDKs.

Migrate

View Disk Details

☒ Show only misaligned virtual machines

Virtua...	Alignment State	Offset Group	Operation St...	Number of VMDKs	
om-prefix-	Misaligned, Online migration capable	7	Ready	1	...

The following states can appear in the Alignment State column:

For this state...	Do this...
Actually aligned	Nothing. The partitions of the virtual machine's hard disk are aligned to the storage system and start at the correct offset.

For this state...	Do this...
Functionally aligned	<p>Nothing.</p> <p>The partitions of the virtual machine's hard disk are misaligned; however, when residing on an optimized datastore, they align on correct boundaries. As a result, the virtual machine performs as though it is aligned.</p> <p>If you want to clone a functionally aligned virtual machine, you should take it offline and fix the VMDK alignment before you clone it. Otherwise, the clones can be misaligned if the destination datastores are not optimized for the VMDK layout of the clones.</p>
Misaligned, Offline alignment only	<p>Power down the virtual machine and align the VMDK using a tool such as VMware vCenter Converter.</p> <p>The virtual machine is misaligned; however, VSC cannot align it due to either of the following:</p> <ul style="list-style-type: none"> <li>• The virtual machine has more than one disk with different offsets</li> <li>• The virtual machine has multiple disks spanning multiple datastores</li> </ul>
Misaligned, Online migration capable	Use VSC to perform an online alignment.
Other	<p>Review the following reasons why VSC cannot determine the alignment state of the virtual machine:</p> <ul style="list-style-type: none"> <li>• It is inaccessible or reports an error during read attempts</li> <li>• It has a disk size of 0</li> <li>• It does not have any partitions</li> <li>• It has independent disks or dynamic disks</li> </ul>

3. Repeat steps 1 and 2 for each datastore.

## Aligning the I/O of misaligned virtual machines non-disruptively

A misaligned virtual machine can negatively affect I/O performance. To correct the misalignment, Virtual Storage Console for VMware vSphere can migrate a virtual machine to a new or existing optimized datastore so that the I/O is functionally aligned. This process does not require downtime.

### Before you begin


- You should have scanned your datastores and found virtual machines with the alignment state "Misaligned, Online migration capable."
- For NFS datastores, the storage system must be running Data ONTAP 8.1.3 or later.
- The volume on which the datastore resides must not be a SnapLock volume.
- You should be aware of the following caveats and limitations with optimized datastores:
  - You cannot use the vStorage APIs for Array Integration (VAAI) extended copy feature with an optimized datastore.
  - For NFS-optimized datastores, using Data ONTAP to perform an NDMP copy, NDMP restore, or dump restore to the volume can be slower.
  - Migrating a virtual machine from an optimized datastore to a non-optimized datastore will result in misaligned I/O to the virtual machine.


### About this task


- Migrating multiple virtual machines at one time is I/O intensive.  
You should limit the number of virtual machines that you migrate at one time to avoid over-stressing your system.
- The alignment might require more space because deduplication is temporarily turned off when VSC aligns the virtual machine.


### Steps

1. In the **Optimization and Migration** page, select a datastore on which a misaligned virtual machine is running.


Global Scan Schedule


Exclude


Include


Scan selected

Datastore Name	Host(s)	Type	Optimized	Excluded
datastore1	dilbert	VMFS	Unknown	No
datastore2	dilbert	VMFS	No	No

After you select a datastore, the virtual machines that reside on that datastore appear in the virtual machines table (the table at the bottom of the Optimization and Migration page).

- In the virtual machines table, select one or more virtual machines with the status "Misaligned, Online migration capable."

If you select multiple virtual machines, they must have the same offset group (the offset of the largest disk partition).

You should limit the number of virtual machines that you migrate at one time to avoid over-stressing your system.

Migrate

View Disk Details

☒ Show only misaligned virtual machines

Virtua...	Alignment State	Offset Group	Operation St...	Number of VMDKs	
om-prefix-	Misaligned, Online migration capable	7	Ready	1	⋮

- Click **Migrate**.
- Complete the pages in the **Migrate Virtual Machines** wizard to migrate the virtual machine to an optimized datastore.

- In the **Destination datastore** page, specify whether you want to use an existing datastore or a new datastore.

For existing datastores, you will be able to choose from datastores that are optimized for the offset group of the virtual machine.

- (New datastore) In the **Name and type** page, specify a datastore name, datastore type (NFS or VMFS), and for a VMFS datastore, the VMFS protocol (FC/FCoE or iSCSI).
- In the **Storage system** page, specify the storage system that you want to use for the datastore.
- (Existing datastore) In the **Datastore selection** page, select the destination datastore.

VSC lists the datastores that are optimized for the VMDK layout of the virtual machine. If no datastores are listed, go back and select the new datastore option.

- (New datastore) In the **New datastore details** page, specify details about the datastore that you want to create.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Thin provision	<p>Allocates space on the volume when data is written, which allows you to provision more storage than is currently available. If disabled, space is reserved immediately.</p> <p>You must closely monitor the available space in the containing aggregate because thin provisioning can oversubscribe the available space. In an NFS configuration, you can enable auto grow to automatically expand the datastore when space is needed. Make sure that the value you specify for auto grow is larger than the size of the datastore.</p>
Aggregate	Defines the aggregate on which you want to create a new volume. If you selected an SVM that is directly connected to VSC, striped aggregates appear as available; however, they are not supported. Provisioning to a striped aggregate will fail.
Volume	Specifies the volume on which you want to create the datastore. For clustered Data ONTAP, you should not create a datastore in the Storage Virtual Machine (SVM) root volume.
Auto grow	(NFS only) Automatically expands the datastore by the specified increment when space is needed, up to the size limit. This size limit you specify must be larger than the existing datastore.
Datastore cluster	Adds the datastore to a cluster if the Storage Distributed Resource Scheduler (SDRS) feature is enabled on the vCenter Server. Do not mix datastores with varying offsets in the same cluster and do not mix optimized and non-optimized datastores.

- f. In the **Ready to complete** page, Review the summary of your selections and click **Finish**.

## Result

VSC starts the migration task. You cannot cancel this task.

## After you finish

If the old datastore is empty, use it for other virtual machines or destroy it.

## Maintaining your VMware environment

---

You can use Virtual Storage Console for VMware vSphere to maintain your VMware environment by migrating virtual machines, redeploying virtual machines, reclaiming space from virtual machines, and managing datastores by mounting, resizing, and destroying them.

### Migrating virtual machines to a new or existing datastore

Migrating virtual machines moves them from one datastore to another. For example, you might need to migrate a virtual machine to a new datastore to balance disk space usage.

#### Before you begin

- For NFS datastores, the storage system must be running Data ONTAP 8.1.3 or later.
- The volume on which the datastore resides must not be a SnapLock volume.
- To avoid migration errors, the virtual machines must be part of datastores that have been scanned by Virtual Storage Console for VMware vSphere.

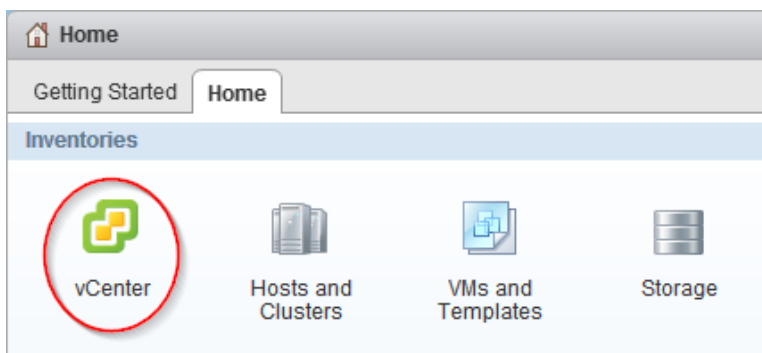
**Note:** The Optimization and Migration page lists the offset group of a virtual machine, if its containing datastore was scanned.

#### About this task

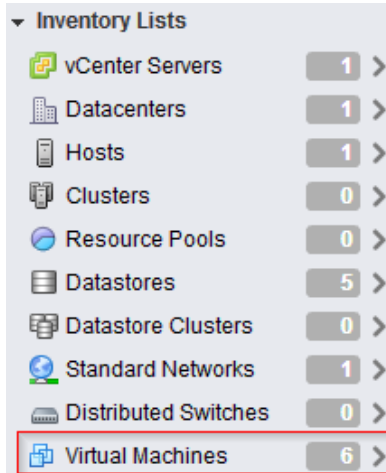
If the selected virtual machines do not have the same offset group, the target datastore will not be optimized for all virtual machines. VSC creates a datastore optimized for the offset group of the last virtual machine that it migrates.

#### Steps

1. From the vSphere Web Client **Home** page, click **vCenter Inventory Lists**.



2. In the navigation pane, under **Inventory Lists**, click **Virtual Machines**.



3. In the **Objects** table, select the virtual machines that you want to migrate.

Migrating multiple virtual machines at one time is I/O intensive. You should limit the number of virtual machines that VSC migrates at one time to avoid over-stressing your system.

4. Click **Actions > NetApp VSC > Migrate**.
5. Click **Yes** to confirm the action.
6. Complete the pages in the **Migrate Virtual Machines** wizard to migrate the virtual machines to a new or existing datastore.
  - a. In the **Destination datastore** page, specify whether you want to migrate the virtual machines to an existing datastore or a new datastore.
  - b. (New datastore) In the **Name and type** page, specify a datastore name, datastore type (NFS or VMFS), and for a VMFS datastore, the VMFS protocol (FC/FCoE or iSCSI).
  - c. In the **Storage system** page, specify the storage system that you want to use for the datastore.
  - d. (Existing datastore) In the **Datastore selection** page, select the destination datastore.
  - e. (New datastore) In the **New datastore details** page, specify details about the datastore that you want to create.

Most of the fields on this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Thin provision	<p>Allocates space on the volume when data is written, which allows you to provision more storage than is currently available. If disabled, space is reserved immediately.</p> <p>You must closely monitor the available space in the containing aggregate because thin provisioning can oversubscribe the available space. In an NFS configuration, you can enable auto grow to automatically expand the datastore when space is needed. Make sure that the value you specify for auto grow is larger than the size of the datastore.</p>
Aggregate	Defines the aggregate on which you want to create a new volume. If you selected an SVM that is directly connected to VSC, striped aggregates appear as available; however, they are not supported. Provisioning to a striped aggregate will fail.
Volume	Specifies the volume on which you want to create the datastore. For clustered Data ONTAP, you should not create a datastore in the Storage Virtual Machine (SVM) root volume.
Auto grow	(NFS only) Automatically expands the datastore by the specified increment when space is needed, up to the size limit. This size limit you specify must be larger than the existing datastore.
Datastore cluster	Adds the datastore to a cluster if the Storage Distributed Resource Scheduler (SDRS) feature is enabled on the vCenter Server. Do not mix datastores with varying offsets in the same cluster and do not mix optimized and non-optimized datastores.

- f. In the **Ready to complete** page, Review the summary of your selections and click **Finish**.

## Result

VSC starts the migration task. You cannot cancel this task.

## After you finish

If the old datastore is empty, use it for other virtual machines or destroy it.

## Related concepts

*[Methods for migrating virtual machines](#)* on page 14



## Redeploying NFS-based virtual machine clones from a template

After you clone virtual machines from a template, you might need to patch or update the cloned virtual machines. You can use Virtual Storage Console for VMware vSphere to redeploy NFS-based virtual machine clones from an updated template. Redeploying VMFS-based clones is not supported.

### Before you begin

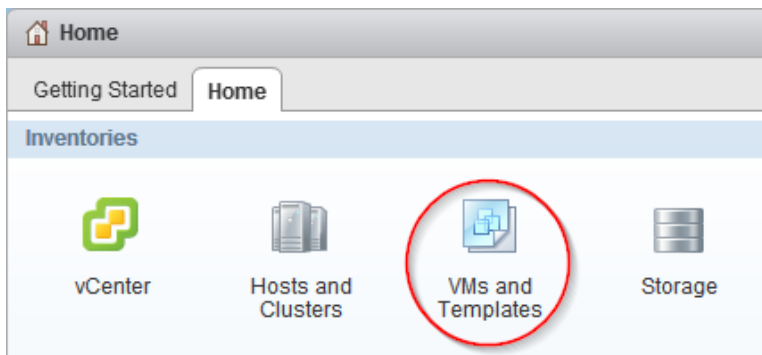
- You must have used VSC when you originally cloned the virtual machines from the template.
- Because redeploying resets the clone to the state of the template, you should first have backed up any needed data.

### About this task


To redeploy a cloned virtual machine, VSC powers off the virtual machine. Make sure this is an acceptable time to take the virtual machine offline. VSC can power on the virtual machine after the redeployment is complete, if you select that option.

### Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates**.



2. In the navigation pane, expand the datacenter that contains the virtual machine template.



3. Right-click the virtual machine and select **NetApp VSC > Redeploy Clones**.
4. In the **Redeploy Clones** dialog box, select the clones, choose their settings, and click **OK**.

The clones can inherit the VMware specification and power state from the template, or you can define new settings.

For a large number of virtual machines, stagger the start up of those virtual machines so you do not overwhelm your system.

**Note:** If a problem prevents VSC from starting some of the virtual machines, the delay could result in VSC powering on a large number of virtual machines at once. For example, if you specify 10 virtual machines per minute and the start is delayed by five minutes, VSC starts 50 virtual machines at once. After the delay, VSC starts the specified number of virtual machines per minute.

5. To confirm that you want to power off and redeploy the selected virtual machines, click **OK**.

### Result

VSC powers off the virtual machines and redeploys them based on the new template. After the redeployment, VSC powers on the virtual machines, if you chose that option.

### Related tasks

[Cloning virtual machines from a template](#) on page 97

## Reclaiming space from NFS-based virtual machines

When users delete data from a virtual machine, the storage space from NTFS partitions is not immediately returned to the NFS datastore. You can reclaim the space to return it to the datastore. Reclaiming space from VMFS-based virtual machines is not supported.

### Before you begin

- Virtual machine files must be on NFS datastores that are not backed by a qtree on a vFiler unit.

- VMDKs must have NTFS partitions.
- VMware Tools must be installed on the virtual machine.
- ISOs mounted to the virtual machine must be contained in an NFS datastore.

**About this task**

To reclaim the space, VSC powers off a virtual machine by using VMware Tools. Make sure this is an acceptable time to take the virtual machine offline. After the process completes, VSC returns the virtual machine to its previous state.

You can perform this task on an individual virtual machine or on a datastore, which reclaims space from all virtual machine disks in a datastore. If you do not want to take all of the virtual machines in a datastore offline, reclaim the space from one virtual machine at a time.

**Steps**

1. Reclaim space from one or more virtual machines:

To...	Do the following...
Reclaim space from all virtual machines in a datastore	From the vSphere Web Client Home page, click <b>Storage</b> .
Reclaim space from one virtual machine	From the vSphere Web Client Home page, click <b>VMs and Templates</b> .

2. In the navigation pane, expand the datacenter that contains the datastore or the virtual machine.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Reclaim Space**.
4. In the **Reclaim Space** dialog box, click **OK**.

**Result**

VSC powers off the virtual machines and starts reclaiming the space.

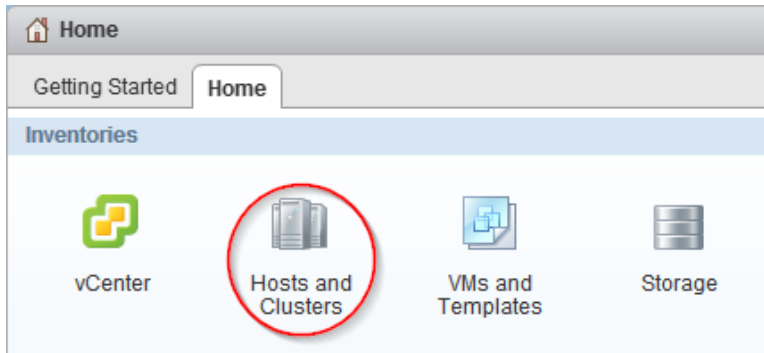
Do not power on the virtual machines while space reclamation is in progress.

## Mounting datastores on hosts


Mounting a datastore gives a host access to storage. You might need to mount a datastore on a host after you add the host to your VMware environment.

**Steps**

1. From the vSphere Web Client **Home** page, click **Hosts and Clusters**.



2. In the navigation pane, expand the datacenter that contains the host.



3. Right-click the host and select **NetApp VSC > Mount Datastores**.
4. Select the datastores that you want to mount and click **OK**.

### Result

VSC mounts the datastores on the host.

## Resizing datastores

Resizing a datastore gives you more or less storage for your virtual machine files. You might need to change the size of a datastore as your infrastructure requirements change.

### Before you begin

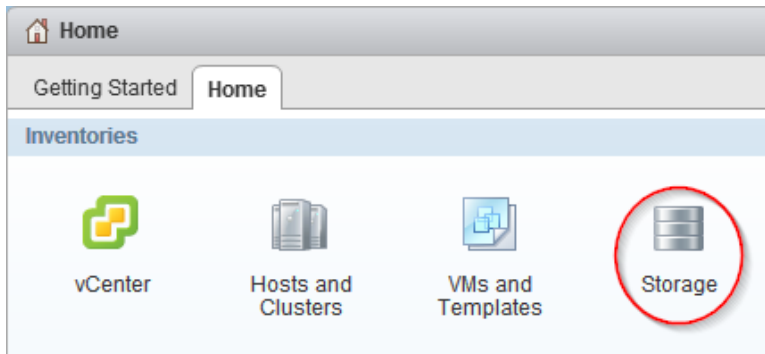
If you want VSC to resize the containing volume when it resizes the VMFS datastore, you should have enabled the **Create new volume for each new LUN** option in the VSC provisioning options for the storage system.

### About this task


You can increase or decrease the size of an NFS datastore. You can only increase the size of a VMFS datastore.

## Steps

1. From the vSphere Web Client **Home** page, click **Storage**.



2. In the navigation pane, expand the datacenter that contains the datastore.



3. Right-click the datastore and select **NetApp VSC > Resize**.
4. In the **Resize** dialog box, specify a new size for the datastore and click **OK**.

## Result

VSC resizes the datastore.

## Related tasks

*[Managing settings for volumes](#)* on page 88

## Destroying datastores

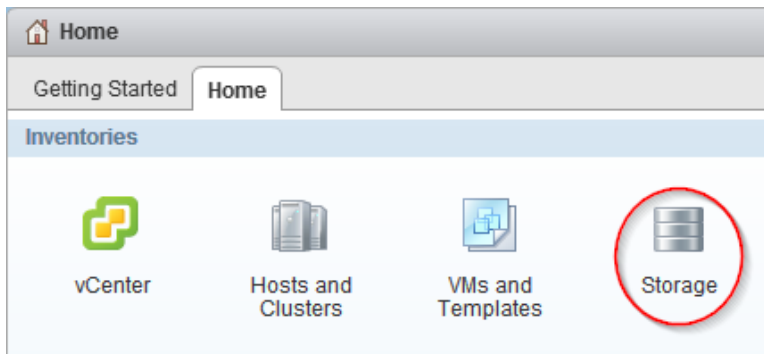
Destroying a datastore returns storage to your storage system and deletes associated objects such as export policies and igroups. You might need to destroy a datastore when you decommission your virtual machines.

### About this task

When you destroy a datastore, the virtual machines within that datastore are also destroyed. Virtual Storage Console for VMware vSphere displays a list of the affected virtual machines before you destroy the datastore.

### Steps

1. From the vSphere Web Client **Home** page, click **Storage**.



2. In the navigator pane, right-click the datastore and select **NetApp VSC > Destroy**.
3. Click **OK**.

### Result

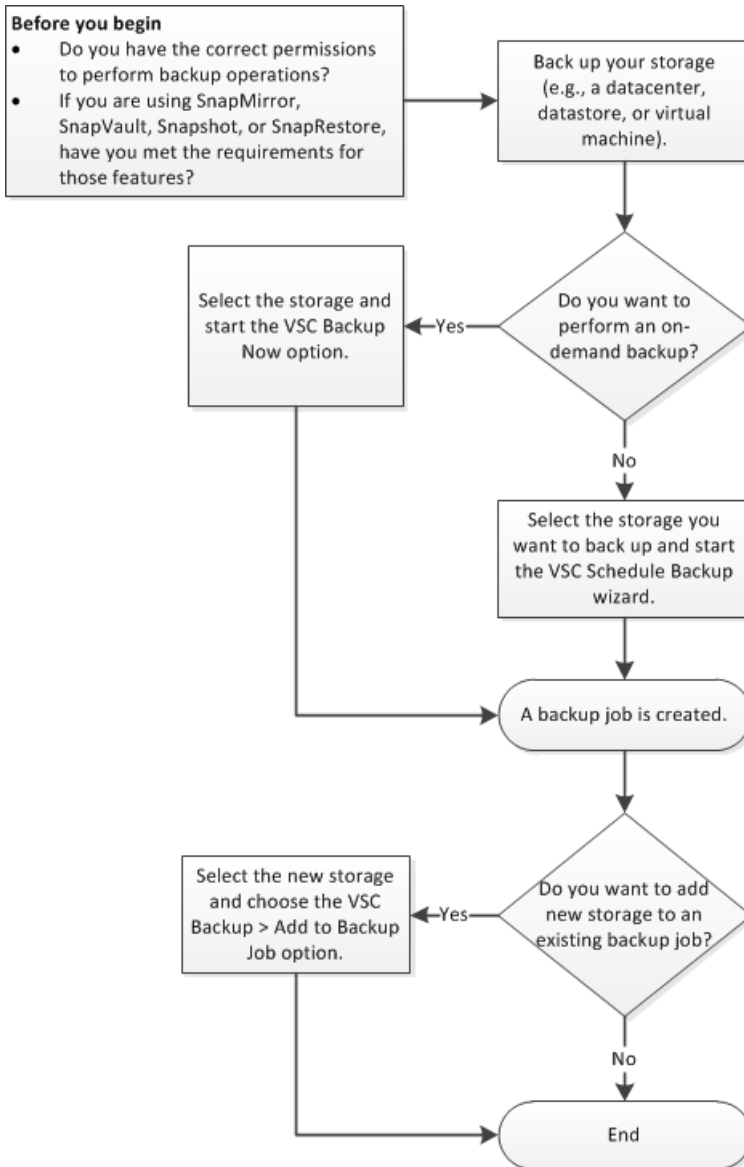
VSC destroys the datastore.

## **Backing up virtual machines and datastores**

---

You can back up individual virtual machines or datastores on demand or on an automated schedule using Virtual Storage Console for VMware vSphere.

You can set a backup job schedule and specify a retention policy for the backup copy when you create a new backup job using the backup wizard. You can also change the schedule and retention policy, as well as suspend and resume or delete a backup job.



## Related tasks

[Performing an on-demand backup of a virtual machine or datastore](#) on page 122

[Scheduling backup jobs](#) on page 124

[Adding a virtual machine or datastore to an existing backup job](#) on page 126



## Backup job specifications

Before you create a backup job, 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 or SnapVault update on the virtual entities that are to be backed up or create VMware snapshots for every backup. If you select virtual machines that span multiple datastores, you can specify one or more of the 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 prebackup or postbackup script that results in an output file when the script is run, the output file is saved to the same directory to which you initially installed the prebackup or postbackup 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 or the maximum number of backup copies and 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 Schedule Backup wizard include the following:

### **Always**

An alert notification is always sent.

### **Errors or Warnings**

A backup failure or a warning triggers an alert notification.

### **Errors**

A backup failure or partial failure triggers an alert notification.

### **Never**

An alert notification is never sent.

## Backup job requirements

A SnapRestore license is required to perform restore operations and a SnapMirror and SnapVault license is required if you are going to use the SnapMirror or SnapVault options for a backup job.

The following is a SnapRestore requirement:

- SnapRestore technology must be licensed for the storage systems where the datastore and virtual machine system images reside.

Following are some SnapMirror and SnapVault requirements:

- The volumes containing the active datastore and virtual machine images must be configured as SnapMirror or SnapVault source volumes.
- The SnapVault policy must have a rule that specifies labels for the VSC backup schedule:

Schedule type	Required label
Hourly	VSC_JOB_HOURLY
Daily	VSC_JOB_DAILY
Weekly	VSC_JOB_WEEKLY
Monthly	VSC_JOB_MONTHLY
One-time only	VSC_ONDEMAND

To use values other than the default values, you must specify the following labels in the `smvi.config` file and specify the same labels when you create the SnapVault protection policy.

```
"snapvault.job.hourly.label"="VSC_XXXX"
"snapvault.job.daily.label"="VSC_XXXX"
"snapvault.job.weekly.label"="VSC_XXXX"
"snapvault.job.monthly.label"="VSC_XXXX"
"snapvault.ondemand.label"="VSC_XXXX"
```

- The source volumes must have a SnapMirror or SnapVault relationship with target volumes on a second storage system that is licensed for SnapMirror or SnapVault.
- The host names and IP addresses of the SnapMirror or SnapVault 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 Storage Virtual Machine (SVM) administrators must create node-management or cluster-management LIFs, which are required to update SnapMirror or SnapVault relationships for storage systems running clustered Data ONTAP 8.2 or later. The cluster-management LIF is required for storage systems running a version of clustered Data ONTAP prior to 8.1.

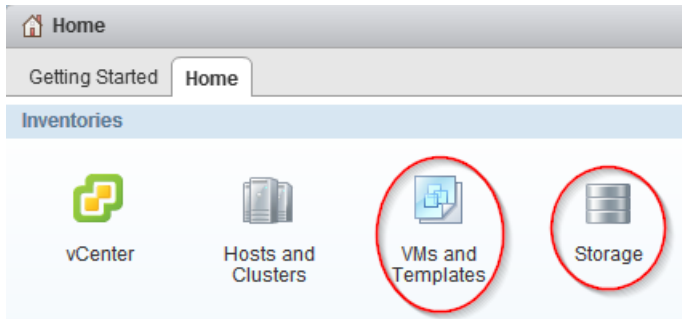
## Performing an on-demand backup of a virtual machine or datastore

You can launch a one-time backup operation for a virtual machine or for an entire datastore. This type of backup is useful if you do not want to schedule regular backups for a particular virtual

machine or selected datastores or if you need to create a one time, non-scheduled backup to retain important changes.

## Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.



2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Backup Now**.
4. In the **Backup Now** dialog box, type a name for the backup job.
5. Optional: Check the box next to the options you want this backup job to use:

- If you want to start a SnapVault update on the selected entities concurrent with every backup copy, select **Initiate SnapVault update**.

**Note:** The SnapVault option is only supported on clustered Data ONTAP 8.2 or later.

For this option to execute successfully, the selected entities must reside in volumes that are configured as SnapVault source volumes and a destination volume must also exist.

- If you want to start a SnapMirror update on the selected entities concurrent with every backup copy, select **Initiate SnapMirror update**.

For this option to execute successfully, the selected entities must reside in volumes that are configured as SnapMirror source volumes and a destination volume must also exist.

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 every time the backup job runs, select **Perform VMware consistency snapshot**.
- If you want to include independent disks from datastores that contain temporary data, select **Include datastores with independent disks**.

6. Click **OK**.

# Scheduling backup jobs

You can create scheduled backup jobs by selecting an entire datacenter, a datastore, or a virtual machine. You can also view all backup jobs on the Backup Jobs page in the vSphere Web Client and create backup jobs from this page using the backup job wizard to select a virtual entity.

## Before you begin

vSphere Web Client must be connected to a vCenter Server to create backup copies.

## About this task

You can set the schedule for your backup jobs, specify a retention policy, and create an automated policy for email alerts.

## Steps

1. To create and schedule a backup job, take one of the following actions:

To...	Do this...
Create a backup job for a specific datastore or virtual machine	<p><b>a.</b> From the vSphere Web Client Home page, click <b>VMs and Templates</b> to view virtual machines or click <b>Storage</b> to view datastores.</p> <p><b>b.</b> In the navigation pane, expand the datacenter that contains the virtual machine or datastore, depending on whether you are in the <b>VMs and Templates</b> view or the <b>Storage</b> view.</p> <p><b>c.</b> Right-click the datastore or virtual machine and select <b>NetApp VSC &gt; Backup &gt; Schedule Backup</b>. You can create a scheduled backup job for an entire datacenter by right-clicking the datacenter and selecting <b>NetApp &gt; Backup &gt; Schedule Backup</b>.</p>
Create a backup job using the backup job wizard to select a datastore or virtual machine	<p><b>a.</b> From the vSphere Web Client Home page, click <b>vCenter Inventory Lists</b>.</p> <p><b>b.</b> In the navigation pane, under NetApp, click <b>Backup Jobs</b>.</p> <p><b>c.</b> Click the <b>Add</b> icon on the Backup Jobs page in the vSphere Web Client.</p>

2. Type a name for the backup job and add a description.
3. If your environment includes multiple vCenter Servers, you must select the server that contains the datastores or virtual machines for which you want to back up.

4. Optional: Select the box next to the options you want for this backup job:

- If you want to start a SnapVault update on the selected entities concurrent with every backup copy, select **Initiate SnapVault update**.

**Note:** The SnapVault option is only supported on clustered Data ONTAP 8.2 or later.

For this option to execute successfully, the selected entities must reside in volumes that are configured as SnapVault source volumes and a destination volume must also exist.

- If you want to start a SnapMirror update on the selected entities concurrent with every backup copy, select **Initiate SnapMirror update**.  
For this option to execute successfully, the selected entities must reside in volumes that are configured as SnapMirror source volumes and a destination volume must also exist.  
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 every time the backup job runs, select **Perform VMware consistency snapshot**.
- If you want to include independent disks from datastores that contain temporary data, select **Include datastores with independent disks**.

5. Click **Next**.

6. If you have not selected a datastore or virtual machine, select the virtual entities available for this backup job and click **Next**.

7. Select the spanned entities available for this backup job and click **Next**.

Spanned entities might be a VM with multiple VMDKs across multiple datastores.

8. 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.

9. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job.

10. Specify either a maximum number of days or a maximum number of backup copies for this backup job and click **Next**.

**Attention:** If you select **A Maximum of Backups**, you can specify a number up to 254; you cannot specify the maximum of 255 backup copies per backup job.

11. Select **Use default vCenter credentials** or type the user name and password for a specific vCenter Server user.

12. Select the frequency for receiving email alerts and click **Next**.

You can add multiple email addresses by using semicolons to separate each email address.

13. Optional: Click **Send Test Email** to verify that the outgoing email server where the alert notifications are to be sent is working correctly.
14. Review the summary of your selections and click **Finish**.
15. Optional: Select the **Run Job Now** check box to immediately run the job.

#### Related tasks

*[Modifying the job properties of a scheduled backup job](#) on page 127*

*[Suspending an active backup job](#) on page 127*

*[Resuming a suspended backup job](#) on page 128*

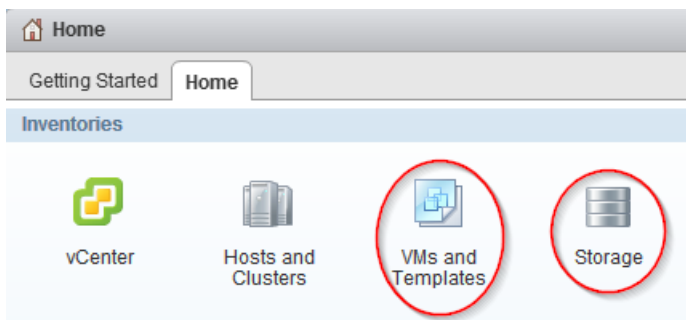
*[Deleting a scheduled backup job](#) on page 129*

## Adding a virtual machine or datastore to an existing backup job

You can add a new virtual machine or datastore to an existing backup job. If you have already created a backup job with specific schedule and retention properties, you can then add a new datastore or virtual machine to the existing backup job.

#### Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.



2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Add to Backup Job**.
4. In the **Add to Backup Job** dialog box, select the backup job to which you want to add the datastore or virtual machine.

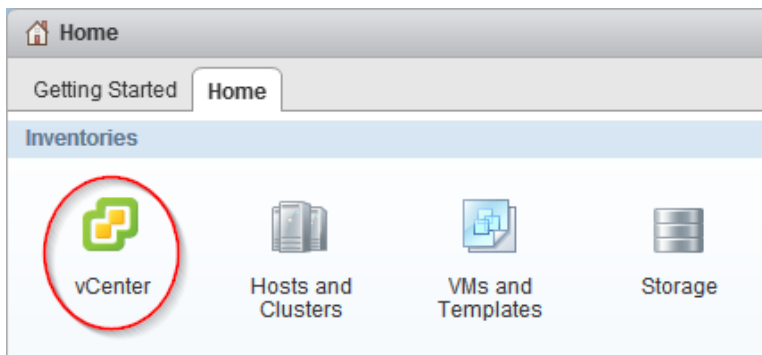
5. Click **OK**.

## Modifying the job properties of a scheduled backup job


You can modify the name and description, the datastores and virtual machines that are assigned, the backup scripts, the user credentials, the schedule, the retention policy, and the email alerts for an existing backup job using the Modify Backup Job dialog box.

### Steps

1. From the vSphere Web Client **Home** page, click **vCenter Inventory Lists**.



2. In the navigation pane, under **NetApp**, click **Backup Jobs**.



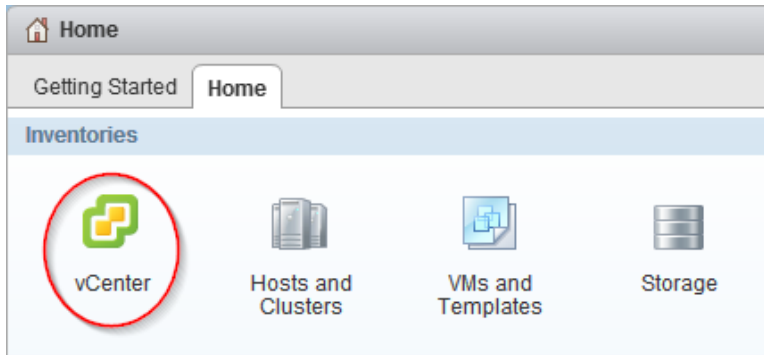
3. Right-click the backup job whose properties you want to modify and select **Modify**.
4. Click the appropriate tab for the properties that you want to modify for this backup job.
5. Modify backup job properties as necessary, and then click **OK** to change the properties.

## 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, during periods of high activity, or for other reasons.

### Steps

1. From the vSphere Web Client **Home** page, click **vCenter Inventory Lists**.



2. In the navigation pane, under **NetApp**, click **Backup Jobs**.



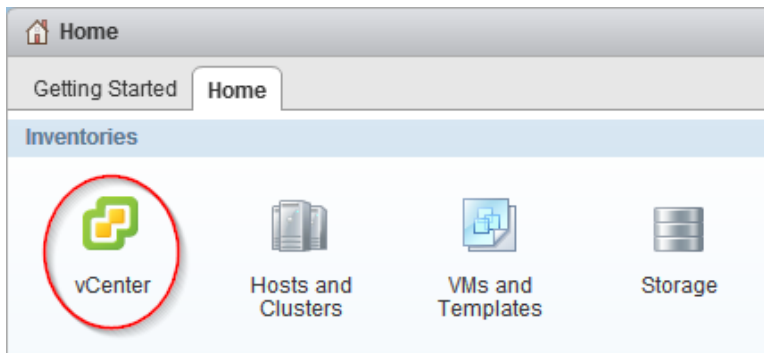
3. Right-click the active backup job that you want to suspend and select **Suspend**.
4. Click **OK** when you receive the confirmation prompt to suspend the active backup job.

## Resuming a suspended backup job

You can resume and run a suspended backup job at any time after you temporarily halt the backup job.


### Steps

1. From the vSphere Web Client **Home** page, click **vCenter Inventory Lists**.



2. In the navigation pane, under **NetApp**, click **Backup Jobs**.





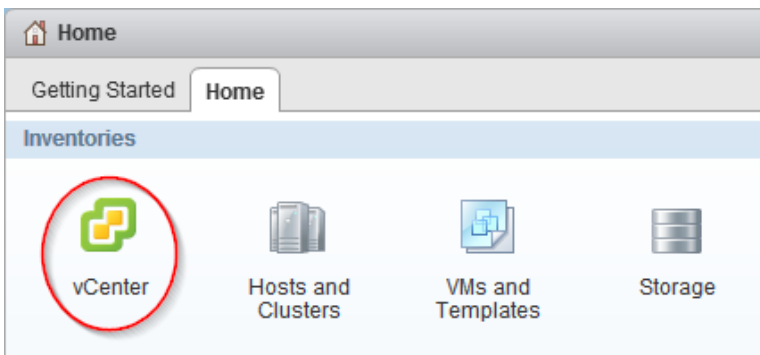
3. Right-click the suspended backup job that you want to resume and select **Resume**.
4. Click **OK** when you receive the confirmation prompt to resume the suspended backup job.

## 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.

### Steps

1. From the vSphere Web Client **Home** page, click **vCenter Inventory Lists**.



2. In the navigation pane, under **NetApp**, click **Backup Jobs**.



3. Select one or more backup jobs that you want to delete.
4. Right-click each selected backup job, and then select **Delete**.
5. Click **OK** at the confirmation prompt to delete the scheduled backup job.

## Restoring virtual machines and datastores from backup copies

---

You can restore your virtual machines and datastores from backup copies using Virtual Storage Console for VMware vSphere. Virtual machines are always restored to the most current datastore; only VMDKs can be restored to an alternate datastore.

### Related tasks

[\*Restoring data from backup copies\*](#) on page 133

[\*Mounting a backup copy\*](#) on page 131

[\*Unmounting a backup copy\*](#) on page 132

## Considerations for 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 backed up entities contained in the backup copy in the Restore wizard and use it for restore operations.

When creating a VMware 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 Data ONTAP Snapshot copy is taken. Despite the VMware snapshot failure, the virtual machine is still included in the Data ONTAP Snapshot copy.

The Quiesced column can display the following values:

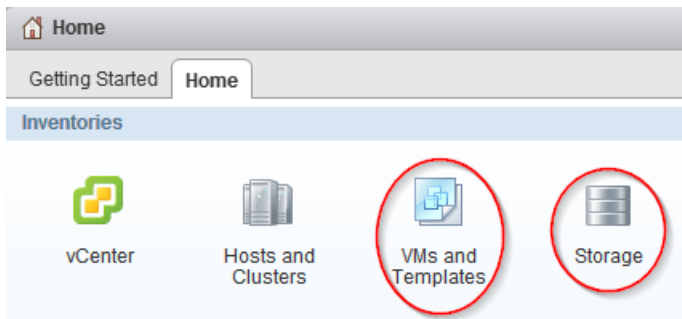
- Yes, if a VMware snapshot operation was successful and the guest operating system was quiesced.
- No, if a VMware snapshot was not selected or the operation failed because the guest operating system could not be quiesced.
- Not Applicable, for entities that are not virtual machines.

## Searching for backup copies

You can search for and find a specific backup copy of a virtual machine or datastore using the Restore wizard. After you locate a backup copy, you can then restore it.

### Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.



2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore, depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Restore**.
4. Click **Advanced Filter** in the **Restore** wizard.
5. Type one or more search terms, and then click **OK**.

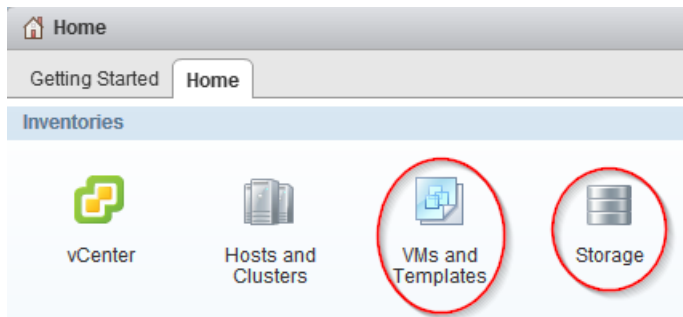
Available criteria for search are the name of the backup job, time range of the backup job, whether the backup job contains a VMware snapshot, or whether the backup job has been mounted.

## Mounting a backup copy

You can mount an existing backup copy onto an ESX server to verify the contents of the backup copy prior to completing a restore operation.

### Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.



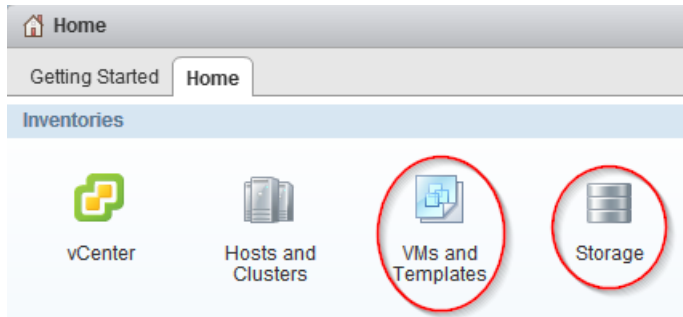
2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Mount Backup**.
4. In the **Mount Backup** dialog box, select the name of an unmounted backup copy that you want to mount.
5. Select the name of the ESX server to which you want to mount the backup copy.  
 You can only mount one backup copy at a time, and you cannot mount a backup that is already mounted. All datastores residing in the backup copy, even ones that were added because of spanned VMs, are mounted.
6. Click **OK**.

## Unmounting a backup copy

After you have verified 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.

### Steps

1. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.



2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Unmount Backup**.
4. In the **Unmount Backup** dialog box, select the name of a mounted backup that you want to unmount.
5. Click **OK**.

## Restoring data from backup copies

You can restore a datastore, 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

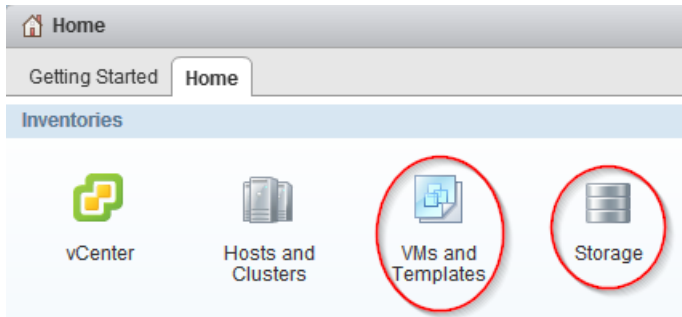
You must have already backed up a virtual machine before you can restore either the entire VM or its individual VMDKs.

### 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. From the vSphere Web Client **Home** page, click **VMs and Templates** to view virtual machines or click **Storage** to view datastores.




2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Restore**.
4. In the **Restore** wizard, select the backup copy that you want to restore from and click **Next**.
5. Select one of the following restore options:

Option	Description
<b>The entire virtual machine</b>	Restores the contents to the last datastore in which it resided from a Snapshot copy with a particular time and date. The <b>Restart VM</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 individual VMDKs to the most current or alternate datastore. This option is enabled when you clear the <b>The entire virtual machine</b> option.

6. Click **Next**.
7. Review the summary of your selections and click **Finish**.

You can track the progress of the restore operation from the Recent Tasks pane in the vSphere Web Client.



## Troubleshooting

---

If you encounter unexpected behavior during the installation or configuration of VSC for VMware vSphere, you can follow specific troubleshooting procedures to identify and resolve the cause of such issues.

### Information at NetApp VSC Communities Forum

The NetApp Communities Forum provides information about Virtual Storage Console for VMware vSphere. When you join the forum you can ask questions and talk with other VSC users.

The NetApp Communities Forum also provides information about tools you can use with VSC news. At the forum, you can do the following:

- Get links to tools, such as the "RBAC User Creator for Data ONTAP."
- See video blogs created by NetApp VSC team members.
- See the latest NetApp news about VSC, such as when a beta testing program might be available.

It is good practice to check the NetApp Communities Forum periodically.

The NetApp VSC Communities Forum is online at:

<http://communities.netapp.com/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 [mysupport.netapp.com](https://mysupport.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.

## Collecting the VSC for VMware vSphere log files

You can collect the Virtual Storage Console for VMware vSphere log files using the Export VSC Logs page. Technical support might ask you to collect the log files to help troubleshoot a problem.

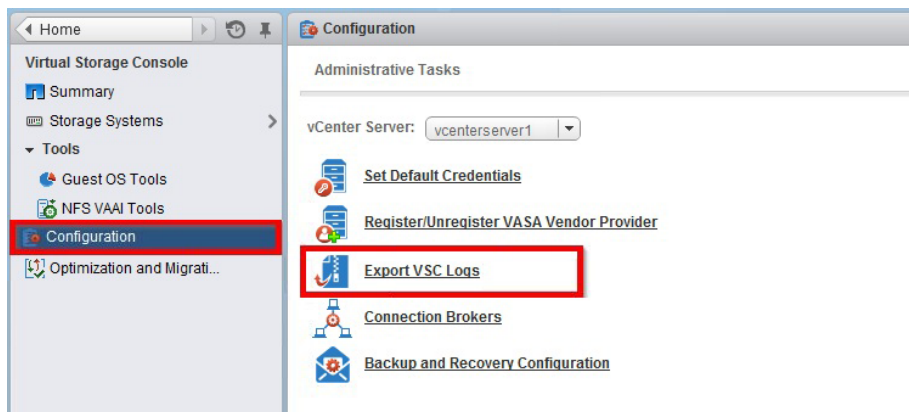
### Before you begin

You must have selected the vCenter Server that you want to use for this task.

### Steps

1. From the Virtual Storage Console **Home** page, click **Configuration > Export VSC Logs**.

This operation can take several minutes.





2. When prompted, save the file to your local computer.

#### After you finish

Send the .zip file to technical support.

## 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.

## Resolution of issues that apply to backup and restore operations

If you encounter unexpected behavior during a backup or restore operation, you can follow specific troubleshooting procedures to identify and resolve the cause of such issues.

### Values that you can override for backup jobs

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. If you modify any of these entries, you must restart the server for the changes to take effect.

**`vmware.max.concurrent.snapshots=6`**

Specifies six as the default maximum number of VMware snapshots created or deleted per datastore during a backup.

**`vmware.quiesce.retry.count=0`**

Specifies zero as the maximum number of retry attempts for VMware snapshots.

**`vmware.quiesce.retry.interval=5`**

Specifies the amount of time, in seconds, between retry attempts for VMware snapshots.

**`vim.client.log.verbose=true`**

When the value is true, logs the interactions between the SMVI server and the vCenter server.

**`smvi.script.timeout.seconds=120`**

Specifies the SMVI timeout value for a prebackup or postbackup script, which is when the SMVI server stops waiting for the script to finish running.

**`smvi.snapshot.recent.naming`**

Disables the snapshot naming convention that adds the *\_recent* suffix to the latest Snapshot copy. To disable the *\_recent* naming convention, you must add the `smvi.snapshot.recent.naming` property, with the value set to false, to the `smvi.override` configuration file.

## Location of backup event and error logs

Virtual Storage Console 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

### Issue

When you click the link to view the log files in the email notification for a backup job, an error occurs.

### Cause

This problem occurs if you disable the IP address of the network adapters for the SMVI log viewer.

### Corrective action

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

This message occurs when you attempt to mount a backup copy of an NFS datastore on a Storage Virtual Machine (SVM, formerly known as Vserver) with the root volume in a load-sharing mirror relationship.

The mount operation fails with the following message:

You may have reached the maximum number of NFS volumes configured in the vCenter. Check the vSphere Client for any error messages.

To prevent this problem, use the server IP address instead of the SVM IP address when you add a storage system running clustered Data ONTAP to Virtual Storage Console.

## 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.

## **Virtual Storage Console unable to discover datastores on an SVM (Vserver) without a management LIF**

Running a scheduled backup job fails when a Storage Virtual Machine (SVM, formerly known as Vserver) without a management LIF is added in Virtual Storage Console. VSC cannot resolve this SVM and is unable to discover any datastores or volumes on the SVM on which to perform backup or restore operations.

You must add an SVM with a management LIF before you can perform backup or restore operations.

## **VMware vSphere does not remove all snapshot delta disks during a restore operation**

### **Issue**

When you delete a VMware snapshot, the snapshot delta disks that are created after quiescing the virtual machine are not removed.

### **Cause**

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.

### **Corrective action**

None

## Copyright information

---

Copyright © 1994–2015 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, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

## How to send comments about documentation and receive update notification

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

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.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

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

32-bit installations

    upgrading to 64-bit installations [32](#)

## A

accounts

    configuring with RBAC [73](#)

active backup jobs

    suspending [127](#)

alignments

    checking the state of virtual machines [105](#)

    performing an online alignment [107](#)

    scanning datastores to determine for virtual machines [103](#)

architecture

    VSC [16](#)

AutoSupport messages

    enabling for backup jobs [55](#)

## B

backing up

    about [14](#)

backup and restore features

    installing or uninstalling [28](#)

backup copies

    finding [131](#)

    mounting [131](#)

    restoring data from [130](#)

    searching for [131](#)

    unmounting from ESX servers [132](#)

backup jobs

    considerations for restoring data from backups with failed VMware consistency snapshots [130](#)

    deleting scheduled [129](#)

    enabling AutoSupport messages for [55](#)

    modifying job properties of scheduled [127](#)

    resuming suspended [128](#)

    setting up alert notifications [56](#)

    SnapMirror and SnapVault requirements [121](#)

    Snapshot and SnapRestore requirements [121](#)

    specifications for adding [121](#)

    suspending [127](#)

    troubleshooting [137](#)

backup mount error [139](#)

backup operations

    on-demand [122](#)

    performing on virtual machines, datastores, and datacenters [119](#)

backup retention

    specifying [121](#)

backup scripts

    controlling the runtime of [137](#)

backups

    failure [139](#)

## C

CD-ROM

    adding to virtual machine [43](#)

Citrix [53](#)

Citrix XenDesktop

    adding connection brokers [53](#)

    importing clone data into [97](#)

clone data [53](#)

clone operation [53](#)

clones

    creating from a template [97](#)

    importing clone data into connection broker [97](#)

    redeploying from a template [113](#)

comments

    how to send feedback about documentation [143](#)

Communities

    provide information about VSC [135](#)

concurrent VMware snapshots

    controlling the number created or deleted [137](#)

configuration

    troubleshooting [137](#)

    troubleshooting VSC for VMware vSphere [135](#)

    VSC for VMware vSphere [47](#)

connection broker [55](#)

connection brokers

    adding [53](#)

    importing clone data into [97](#)

Connection brokers panel [53](#)

controller

    removing skipped or unmanaged [86](#)

    supports vFiler unit, SVM tunneling [80](#)

copies

    unmounting backup, from ESX servers [132](#)

credentials



- default for storage controllers [78](#)
- modifying storage system credentials [82](#)
- setting for storage system [78](#)
- upgrade considerations [29](#)
- using RBAC [73](#)

csv [53](#)

- custom user accounts
  - configuring using RBAC [73](#)

## D

datastores

- adding to existing backup jobs [126](#)
- destroying [118](#)
- enabling deduplication on [101](#)
- enabling mounting across subnets [57](#)
- migrating virtual machines to [110](#)
- mounting on hosts [115](#)
- NFS indirect path [90](#)
- on-demand backup operations [122](#)
- provisioning [94](#)
- resizing [116](#)
- restoring [133](#)
- returning space to (NFS only) [114](#)
- scanning to determine virtual machine alignment [103](#)
- scheduling backup jobs [124](#)
- searching for backup copies of [131](#)
- unable to discover on a Storage Virtual Machine [140](#)

deduplication

- checking the state of [101](#)
- enabling [101](#)

default credentials

- for storage controllers [78](#)
- setting for storage system [78](#)

default values in smvi.override

- modifying [137](#)

delta disks

- not removed during restore [140](#)

discovering

- hosts [82](#)
- storage systems [82](#)

discovery

- correcting unknown storage system name [87](#)
- enabling for vFiler units [81](#)
- manually adding storage systems [84](#)
- storage system [77](#)

Disk.QFullSampleSize [40](#)

Disk.QFullThreshold [40](#)

documentation

- how to receive automatic notification of changes to [143](#)

- how to send feedback about [143](#)

## E

email alert notifications

- configuring [56](#)

email alerts

- specifying for backup jobs [121](#)

email notification error [139](#)

Emulex FC HBA timeouts [40](#)

error logs

- locations [138](#)
- reviewing [138](#)

error messages

- mounting NFS datastore backup fails [139](#)
- move failed [139](#)

ESX host

- restoring virtual disks on VMFS datastores [133](#)

ESX hosts

- multipathing and timeout settings [38](#)
- settings [40](#)
- timeout values [42](#)

ESX, ESXi host settings [42](#)

event logs

- locations [138](#)
- reviewing [138](#)

## F

failed consistency snapshots

- considerations for restore operations backed up with VMware [130](#)

feedback

- how to send comments about documentation [143](#)

files, log

- collecting [136](#)

## G

guest OS

- installing scripts [43](#)
- setting timeouts for Linux [45](#)
- setting timeouts for Solaris [46](#)
- setting timeouts for Windows [47](#)
- timeout values [42](#)

## H

### hosts

- discovering [82](#)
- mounting datastores on [115](#)
- multipathing and timeout settings [38](#)

httpd.admin.enable option [81](#)

## I

### information

- how to send feedback about improving documentation [143](#)
- locating in this guide [9](#)

### installation

- of VSC installation wizard [26](#)
- overview of [22](#)
- planning [17](#)
- troubleshooting VSC for VMware vSphere [135](#)
- upgrading from 32-bit to 64-bit [32](#)

### installing

- guest operating system (GOS) scripts [43](#)
- VSC [27](#)

### iSCSI

- enabling datastore mounting across subnets [57](#)

### issues

- datastore discovery [140](#)
- mounting NFS datastore backup fails [139](#)
- VMware snapshot delta disks not removed during restore [140](#)

## J

### jobs

- modifying job properties of scheduled backup [127](#)

## K

### kaminoprefs.xml

- modifying to enable datastore mounting across subnets [57](#)

### kaminosdkprefs.xml

- modifying to enable datastore mounting across subnets [57](#)

## L

### Linux

- setting timeouts for guest OS [45](#)

### linux\_gos\_timeout-install.iso

- guest OS tool [45](#)

### locations of log files

- listed [138](#)

### lock management

- VSC [15](#)

### log files

- collecting [136](#)

### logs

- errors [138](#)
- events [138](#)
- locations [138](#)
- message [138](#)
- reviewing [138](#)
- troubleshooting [138](#)

## M

### memory requirements

- VSC [24](#)

### message logs

- locations [138](#)
- reviewing [138](#)

### MetroCluster

- manually adding storage systems to VSC [84](#)

### misaligned virtual machines

- how VSC optimizes them [13](#)

### Modify Backup Job dialog box

- using to modify backup job properties [127](#)

### multi-vCenter environment

- for VSC [51](#)

### multipathing

- configuring ESX host [38](#)

### multiple vCenter Servers

- deploying [51](#)

### MultiStore

- display differences with vFiler units [89](#)
- enabling discovery of vFiler units on private networks [81](#)

## N

### Net.TcpipHeapMax [40](#)

### Net.TcpipHeapSize [40](#)

### NetApp Communities

- See* Communities

### NFS

- enabling datastore mounting across subnets [57](#)

### NFS datastore backup

- mounting a backup fails [139](#)

- NFS paths
  - changing frequency of optimization checks [58](#)
  - changing to direct access [91](#)
  - indirect path [90](#)
- NFS VAAI Plug-in
  - installing [36](#)
- NFS.HeartbeatFrequency [40](#)
- NFS.HeartbeatMaxFailures [40](#)
- NFS.HeartbeatTimeout [40](#)
- NFS.MaxVolumes [40](#)

## O

- object
  - storage system [61, 63](#)
  - vSphere [61, 63](#)
- on-demand backups
  - starting [122](#)
  - using VSC for [119](#)
- OnCommand System Manager
  - See* System Manager
- organization
  - using this guide [9](#)

## P

- parameters
  - ESX hosts [40](#)
  - UNMAP [42](#)
- path selection policy [40](#)
- paths
  - changing to direct NFS paths [91](#)
- permission
  - vCenter Server [61, 63](#)
- plug-ins
  - supported with VSC [12](#)
- preferences files
  - what they are [57](#)
- privileges
  - example of assigning privileges [64](#)
  - example of View privilege [69](#)
  - native vCenter Server [61, 63](#)
  - product level [68](#)
  - Virtual Storage Console [68](#)
  - VSC specific [61, 63](#)
- provisioning
  - datastores [94](#)
- provisioning and cloning
  - about [13](#)

## Q

- QLogic
  - FC HBA timeouts [40](#)
  - iSCSI HBA IP\_ARP\_Redirect [40](#)
  - iSCSI HBA timeouts [40](#)

## R

- RBAC
  - about [59](#)
  - configuring [73](#)
  - Data ONTAP [69](#)
  - Data ONTAP privileges [59](#)
  - Data ONTAP roles [70](#)
  - standard VSC roles [65](#)
  - upgrade considerations [29](#)
  - vCenter privileges [59](#)
  - vCenter Server [60](#)
- registering
  - Virtual Storage Console with vCenter Server [47](#)
- Release Notes
  - checking [135](#)
- Remove Controller command [86](#)
- required ports
  - firewall requirements [51](#)
  - VSC [51](#)
- requirements
  - backup job [121](#)
- resources
  - discovering and adding [82](#)
- restore operations
  - considerations when using data backed up with failed VMware consistency snapshots [130](#)
  - enabling AutoSupport messages for [55](#)
  - from backup copies [130](#)
  - from VMDKs [133](#)
  - mounting backup copies [131](#)
  - troubleshooting [137](#)
- Restore wizard
  - restoring a virtual machine or its disk files [133](#)
- restoring
  - about [14](#)
- resumption
  - of suspended backup jobs [128](#)
- role-based access control
  - See* RBAC
- roles
  - configuring with RBAC [73](#)

## S

- Schedule Backup wizard
    - configuring alert notifications [56](#)
    - using to back up virtual machines or datastores [124](#)
  - scheduled backup jobs
    - adding a virtual machine or datastore [126](#)
    - deleting [129](#)
    - modifying job properties [127](#)
    - of virtual machines or datastores [124](#)
    - using VSC for [119](#)
  - scripts
    - choosing for backup jobs [121](#)
  - scripts, guest operating system (GOS)
    - installing [43](#)
  - security
    - configuring using RBAC [73](#)
  - settings
    - ESX hosts [40](#)
    - ESX, ESXi host [42](#)
  - silent mode
    - using to uninstall VSC for VMware vSphere [35](#)
  - skipped controller
    - removing [86](#)
  - SMTP server
    - specifying for alert notifications [56](#)
  - SMVI
    - regenerating SSL certificate [49](#)
  - SMVI log viewer
    - enabling network adapters [139](#)
  - smvi.override file
    - modifying [137](#)
    - purpose [137](#)
  - SnapMirror requirements
    - for backing up datastores and virtual machines [121](#)
  - SnapRestore requirements
    - for backing up datastores and virtual machines [121](#)
  - snapshot delta disks
    - not removed during restore [140](#)
  - Snapshot requirements
    - for backing up datastores and virtual machines [121](#)
  - snapshots
    - considerations for restore operations backed up with failed VMware consistency [130](#)
  - SnapVault requirements
    - for backing up datastores and virtual machines [121](#)
  - Solaris
    - setting timeouts for guest OS [46](#)
  - solaris\_gos\_timeout-install.iso
    - guest OS tool [46](#)
  - spanned entities
    - backing up [121](#)
  - SSL certificate
    - regenerating [49](#)
    - regenerating for SMVI [49](#)
  - storage controller
    - removing skipped or unmanaged [86](#)
    - using default credentials [78](#)
  - Storage Distributed Resource Scheduler
    - See* SDRS
  - storage resources
    - discovering and adding [82](#)
  - storage system name
    - correcting when unknown [87](#)
  - storage systems
    - adding manually [84](#)
    - assigning permissions [61](#), [63](#)
    - configuring using RBAC [73](#)
    - discovering and adding [82](#)
    - discovery [77](#)
    - modifying credentials [82](#)
    - RBAC privileges [59](#)
    - setting default credentials [78](#)
    - specifying a vCenter Server [51](#)
    - specifying volume settings [88](#)
    - supports direct connections to SVMs [90](#)
    - updating [85](#)
  - Storage Virtual Machine
    - management LIF [140](#)
  - suggestions
    - how to send feedback about documentation [143](#)
  - supported configurations
    - memory requirements [24](#)
  - suspended backup jobs
    - resuming [128](#)
  - suspension
    - of active backup jobs [127](#)
  - SVMs
    - limitations when connecting directly to storage [90](#)
    - support direct connections to storage [90](#)
- ## T
- tasks
    - accessing from VSC [75](#)
  - templates
    - cloning new virtual machines from [97](#)
    - redeploying virtual machines from [113](#)
  - timeout settings
    - configuring ESX hosts [38](#)

- timeout values
  - ESX hosts [40](#)
  - ESX, ESXi host [42](#)
  - recommended values [42](#)
  - setting for guest OS [42](#)
- tools
  - setting Linux guest OS timeouts [45](#)
  - setting Solaris guest OS timeouts [46](#)
  - setting Windows guest OS timeouts [47](#)
- troubleshooting
  - backup and restore operations [137](#)
  - checking Release Notes [135](#)
  - collecting log files [136](#)
  - email notification error [139](#)
  - error writing backup metadata to repository [139](#)
  - logs [138](#)
  - mounting NFS datastore backup fails [139](#)
  - move failure [139](#)
  - unable to discover datastores [140](#)
  - VMware snapshot delta disks not removed during restore [140](#)
- tunneling
  - supported for vFiler units, SVMs [80](#)
- twitter
  - how to receive automatic notification of documentation changes [143](#)

## U

- uninstalling
  - VSC for VMware vSphere using Add/Remove Programs [34](#)
  - VSC using a command line [35](#)
- unknown
  - storage system name [87](#)
- unmanaged controller
  - removing [86](#)
- update command
  - forces storage system discovery [77](#)
- updating
  - resource information [82](#)
- upgrading VSC
  - considerations [29](#)
- user credentials
  - backup job [121](#)
- user interfaces
  - installation software [36](#)
  - VASA Provider [36](#)
- user name
  - configuring custom with RBAC [73](#)

- user privileges
  - controlling using RBAC [59](#)

## V

- VASA Provider
  - installation software [36](#)
  - storage capability profiles [36](#)
  - supported with VSC [12](#)
  - user interfaces [36](#)
  - VSC GUI [36](#)
  - VSC requirement [24](#)
- vCenter inventory
  - RBAC privileges [59](#)
- vCenter object
  - RBAC privileges [59](#)
- vCenter Server
  - permission [61](#), [63](#)
  - registering Virtual Storage Console with [47](#)
  - standard VSC roles [65](#)
- vFiler unit
  - display differences with physical storage controllers [89](#)
  - enabling discovery [81](#)
  - tunneling supported [80](#)
- vFiler units
  - discovering on private networks [81](#)
- View privilege
  - example [64](#), [69](#)
- virtual appliances
  - supported with VSC [12](#)
- virtual entities
  - backing up [121](#)
- virtual machine
  - adding CD-ROM [43](#)
- virtual machines
  - adding to existing backup jobs [126](#)
  - aligning I/O non-disruptively [107](#)
  - checking the alignment of [105](#)
  - cloning from a template [97](#)
  - deploying [93](#)
  - how VSC optimizes misaligned virtual machines [13](#)
  - methods for migrating [14](#)
  - migrating a group [110](#)
  - on-demand backup operations [122](#)
  - reclaiming space from (NFS only) [114](#)
  - redeploying from a template [113](#)
  - restarting [133](#)
  - restoring [133](#)
  - scanning to determine alignment [103](#)

- scheduling backup jobs [124](#)
- searching for backup copies of [131](#)
- Virtual Storage Console
  - example of product privilege [64](#), [69](#)
  - privileges [68](#)
  - registering with vCenter Server [47](#)
  - standard roles [65](#)
- VMDKs
  - restoring [133](#)
- VMFS datastores
  - restoring virtual disks on [133](#)
- VMware consistency snapshots
  - considerations for restore operations backed up with failed [130](#)
- VMware snapshot delta disks
  - not removed during restore [140](#)
- VMware View Server
  - adding connection brokers [53](#)
  - importing clone data into [97](#)
- volumes
  - enabling deduplication [101](#)
  - specifying settings [88](#)
- VSC
  - accessing tasks [75](#)
  - architecture [16](#)
  - backing up and restoring [14](#)
  - checking optimization of NFS paths [58](#)
  - configuration tasks [38](#)
  - Data ONTAP RBAC roles [70](#)
  - firewall port requirements [51](#)
  - how it optimizes misaligned virtual machines [13](#)
  - installation overview [22](#)
  - installing in silent mode [27](#)
  - installing or uninstalling backup and restore features [28](#)
  - installing using installation wizard [26](#)
  - lifecycle management for VMware environments [11](#)
  - limitations with direct connections to SVMs [90](#)
  - lock management [15](#)
  - memory requirements [24](#)
  - overview [11](#)

- planning installation [17](#)
- provisioning and cloning [13](#)
- regenerating an SSL certificate [49](#)
- required ports [51](#)
- requirements for performing tasks [74](#)
- support for Data ONTAP RBAC [69](#)
- support for vCenter Server RBAC [60](#)
- supported configurations [24](#)
- supported plug-ins [12](#)
- uninstalling using a command line [35](#)
- upgrade considerations [29](#)
- upgrading from 32-bit installations [32](#)
- upgrading the software [31](#)
- using multiple vCenter Servers [51](#)
- using this guide [9](#)
- VASA Provider requirement [24](#)
- VSC Communities
  - See* Communities
- VSC for VMware vSphere
  - configuration [47](#)
  - uninstalling using Add/Remove Programs [34](#)
- Vserver
  - tunneling supported [80](#)
- Vservers
  - See* SVMs
- vSphere
  - object [61](#), [63](#)

## W

- Windows
  - setting timeouts for guest OS [47](#)
- windows\_gos\_timeout.iso
  - guest OS tool [47](#)

## X

- XenDesktop
  - adding connection brokers [53](#)
  - See also* Citrix XenDesktop