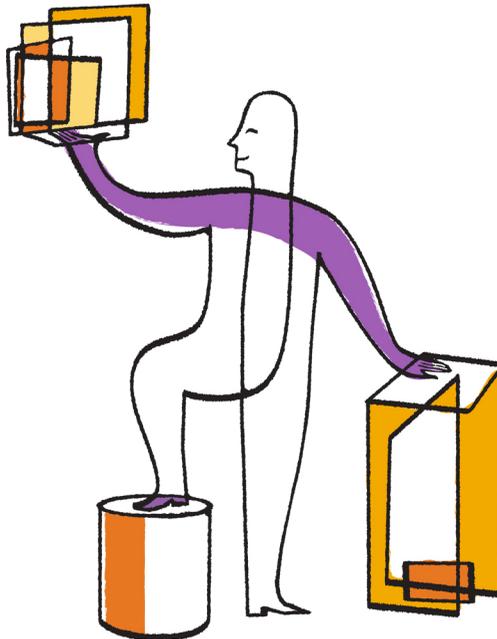




NetApp®

Virtual Storage Console 5.0 for VMware® vSphere®

Performing Basic Tasks



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
Feedback: doccomments@netapp.com

Part number: 215-08617_A0
April 2014

Contents

Deciding whether to use this guide	5
Deploying virtual machines on NetApp storage	6
Provisioning datastores	7
Cloning virtual machines from a template	9
Increasing storage efficiency by enabling deduplication	12
Optimizing performance by aligning the I/O of misaligned virtual machines non-disruptively	14
Scanning datastores to determine the alignment status of virtual machines	14
Checking the alignment status of virtual machines	15
Aligning the I/O of misaligned virtual machines non-disruptively	17
Maintaining your VMware environment	19
Migrating virtual machines to a new or existing datastore	19
Redeploying NFS-based virtual machine clones from a template	21
Reclaiming space from NFS-based virtual machines	22
Mounting datastores on hosts	23
Resizing datastores	24
Destroying datastores	25
Backing up virtual machines and datastores	27
Backup job requirements	28
Performing an on-demand backup of a virtual machine or datastore	29
Scheduling backup jobs	31
Adding a virtual machine or datastore to an existing backup job	34
Modifying the job properties of a scheduled backup job	34
Suspending an active backup job	35
Resuming a suspended backup job	36
Deleting a scheduled backup job	37
Restoring virtual machines and datastores from backup copies	38
Considerations for restore operations using data that was backed up with failed VMware consistency snapshots	38
Searching for backup copies	38
Mounting a backup copy	39
Unmounting a backup copy	40

Restoring data from backup copies	41
Copyright information	43
Trademark information	44
How to send your comments	45
Index	46

Deciding whether to use this guide

This guide provides high-level steps for performing basic tasks in Virtual Storage Console for VMware vSphere. If you are experienced with VSC, you can use this guide to quickly find information about tasks involving provisioning and optimizing storage and backing up and restoring storage.

The information in this guide is for the average VSC user, not the system administrator. It does not provide information about all the task options or about all the conceptual background for the tasks. This guide shares information with the *Virtual Storage Console for VMware vSphere Installation and Administration Guide*, which contains complete details on installing, configuring, and using VSC. If you are a new VSC user or a system administrator, you should use the *Virtual Storage Console for VMware vSphere Installation and Administration Guide*.

Requirements for using this guide and successfully completing tasks

To use this guide, you must have a basic understanding of VSC and what you can do with it.

This guide is based on the assumption that your system administrator has installed VSC and configured it for your system environment. This includes the following:

- The features you need for the tasks you are performing have been installed. These features might include the VSC backup and restore options, the NFS Plug-in for VMware VAAI, and/or the VASA Provider for clustered Data ONTAP.
- VSC has discovered all the storage systems that you are using.
- Role-based access control (RBAC) for both vCenter Server and Data ONTAP has been set up. If you do not have all of the necessary privileges and permissions, the task you are performing might appear to work as you go through the wizard, but still fail to complete successfully.

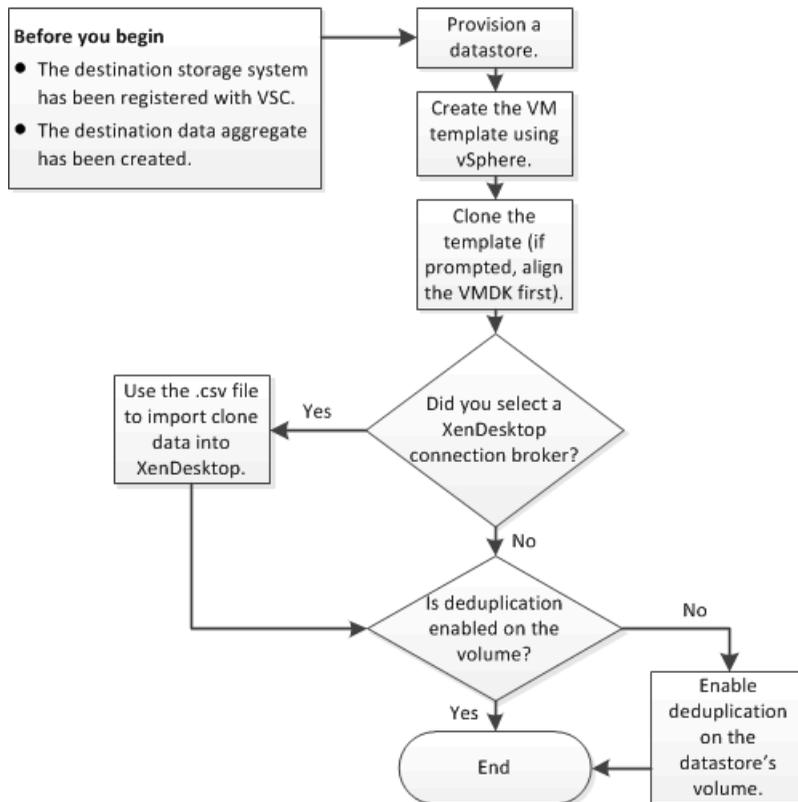
Most of the tasks presented in this guide also require that only one operation at a time be performed on the target datastore or virtual machine.

If you have questions about VSC or want share information with other VSC users, you can go to the NetApp Communities Forum, which is online at <http://communities.netapp.com/vsc>.

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

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.

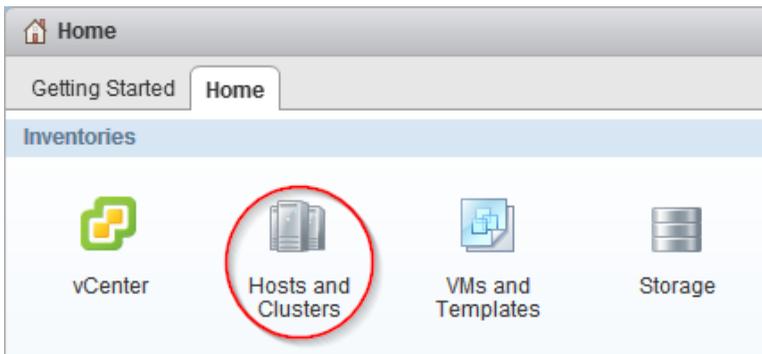
Before you begin

- 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.

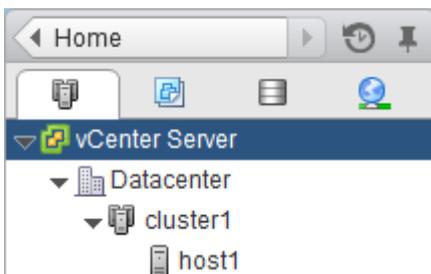
For more information, see the *Virtual Storage Console for VMware vSphere Installation and Administration Guide*.

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 NetApp VSC > Provision Datastore .
All hosts in a cluster	Right-click a cluster and select NetApp VSC > Provision Datastore .
A single host	Right-click a host and select NetApp VSC > Provision Datastore .

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

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

Field	Description
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.
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.
Storage capability profile	(Recommended) Specifies 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 .

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>
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.

Cloning virtual machines from a template

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

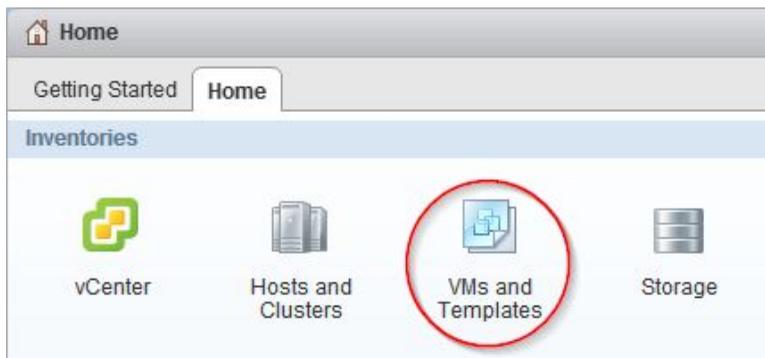
1. Power down the virtual machine template.

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.

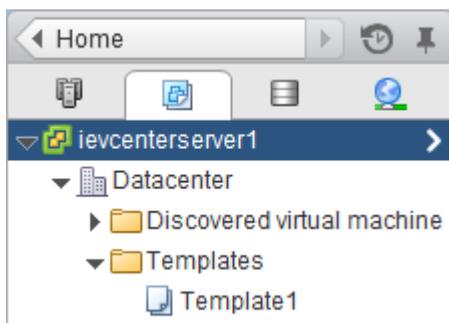
10 | VSC 5.0 for VMware vSphere Performing Basic Tasks

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 14.

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.

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

Field	Description
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.
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.
Connection broker version	Automatically imports clone data into a VMware View Server or creates a <code>.csv</code> 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.
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.
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.
Number of virtual processors	Specifies the number of virtual CPUs for the virtual machines.
Size (GB)	Specifies the size per datastore.
Stagger powering on the virtual machines	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.
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>
Upgrade hardware version?	Upgrades the hardware version of the virtual machine clone if the destination host supports a later version.

Field	Description
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.

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.

Increasing storage efficiency by enabling deduplication

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.

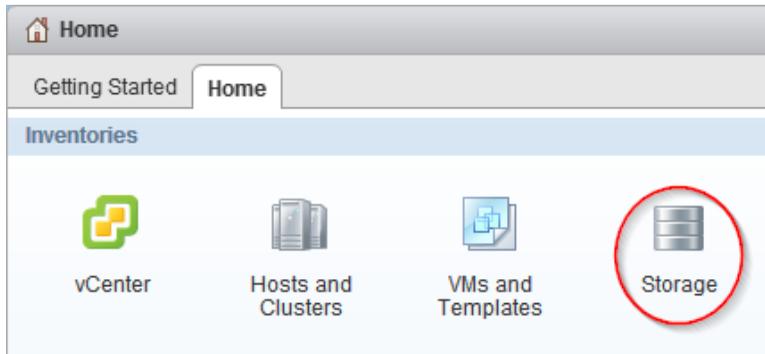
About this task

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.

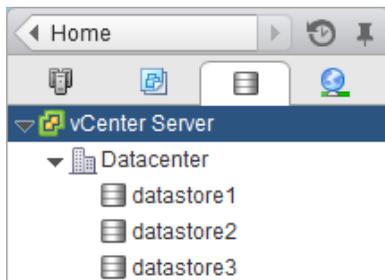
For more information about deduplication, refer to the *Storage Management Guide* for your version of Data ONTAP.

Steps

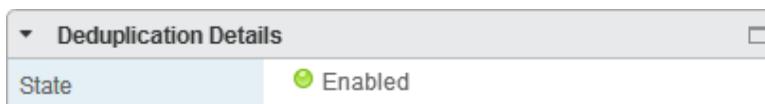
1. From the vSphere Web Client **Home** page, click **Storage**.



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

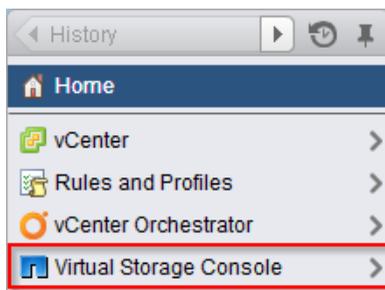
1. [Scan datastores to determine the alignment status of virtual machines](#) on page 14
2. [Check the alignment status of virtual machines](#) on page 15
3. [Align the I/O to any misaligned virtual machines](#) on page 17

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.

Steps

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



2. In the navigator pane, click **Optimization and Migration**.
3. Initiate a scan of the datastores:

To...	Do this...
Schedule recurring scans	<ol style="list-style-type: none"> a. Click Global Scan Schedule and set a schedule for the scan. You should schedule the scans during non-critical production times. b. Click OK. 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 Exclude or Include.
Initiate a one-time scan	<ol style="list-style-type: none"> a. Choose to scan all datastores or specific datastores: <ul style="list-style-type: none"> • To scan all datastores, click Scan All. • To scan one or more datastores, select the datastores and click Scan selected. b. Click OK to confirm the scan.

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).

The screenshot shows a row of four buttons: "Global Scan Schedule" (with a calendar icon), "Exclude" (with a globe and red X icon), "Include" (with a globe and blue plus icon), and "Scan selected" (with a magnifying glass icon). Below the buttons is a table with the following data:

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.

16 | VSC 5.0 for VMware vSphere Performing Basic Tasks

Migrate	View Disk Details	<input checked="" type="checkbox"/> 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.
Functionally aligned	Nothing. 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. 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.
Misaligned, Offline alignment only	Power down the virtual machine and align the VMDK using a tool such as VMware vCenter Converter. The virtual machine is misaligned; however, VSC cannot align it due to either of the following: <ul style="list-style-type: none"> • The virtual machine has more than one disk with different offsets • The virtual machine has multiple disks spanning multiple datastores
Misaligned, Online migration capable	Use VSC to perform an online alignment.
Other	Review the following reasons why VSC cannot determine the alignment state of the virtual machine: <ul style="list-style-type: none"> • It is inaccessible or reports an error during read attempts • It has a disk size of 0 • It does not have any partitions • It has independent disks or dynamic disks

3. Repeat steps 1 and 2 for each datastore.

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

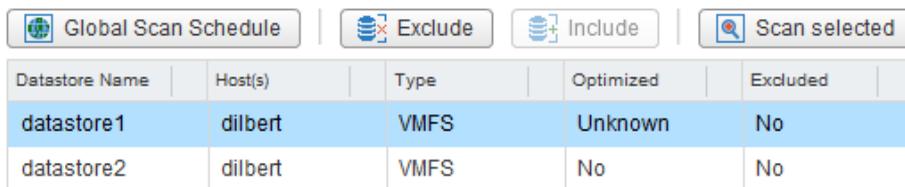
To correct the misalignment, Virtual Storage Console for VMware vSphere can migrate a virtual machine to a new or existing optimized datastore so that the I/O is functionally aligned. This process does not require downtime.

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.

Steps

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



The screenshot shows a user interface with four buttons: "Global Scan Schedule", "Exclude", "Include", and "Scan selected". Below the buttons is a table with the following data:

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).

2. 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	<input checked="" type="checkbox"/> Show only misaligned virtual machines			
Virtual Machine	Alignment State	Offset Group	Operation Status	Number of VMDKs	
om-prefix-	Misaligned, Online migration capable	7	Ready	1	⌵

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

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

Field	Description
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.
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.
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>
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.

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.

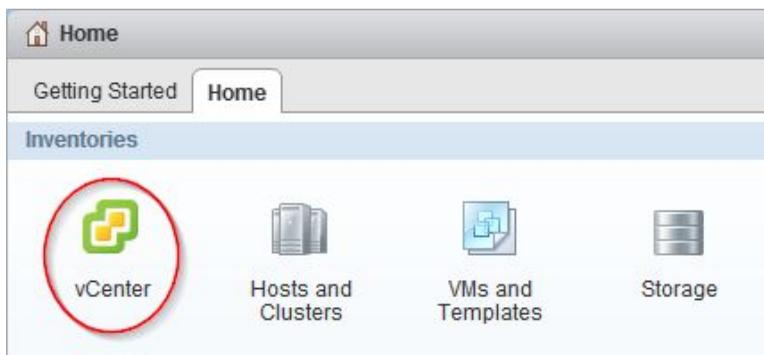
Note: The Optimization and Migration page lists the offset group of a virtual machine, if its containing datastore was scanned.

About this task

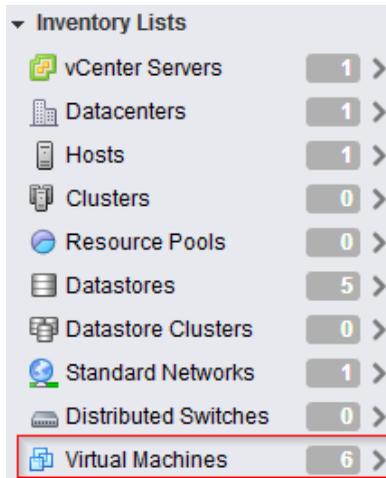
If the selected virtual machines do not have the same offset group, the target datastore will not be optimized for all virtual machines. VSC creates a datastore optimized for the offset group of the last virtual machine that it migrates.

Steps

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



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.

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

Field	Description
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.
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.

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>
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.

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.

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

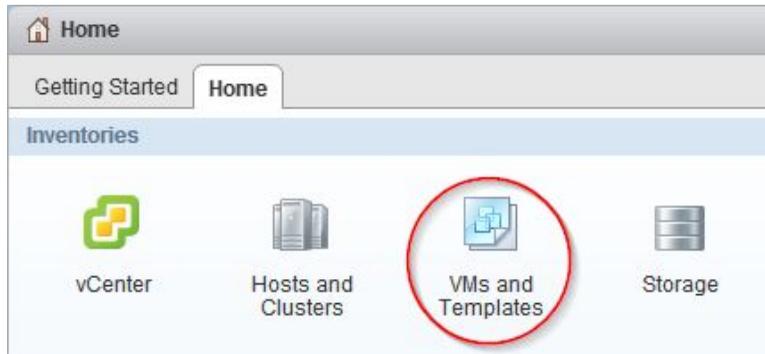
- You must have used VSC when you originally cloned the virtual machines from the template.
- Because redeploying resets the clone to the state of the template, you should first have backed up any needed data.

About this task

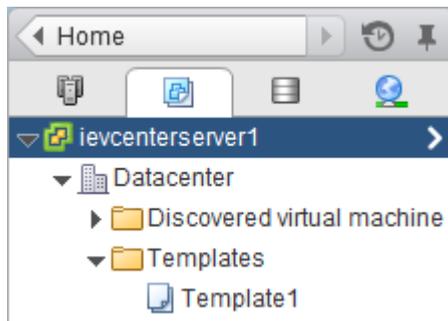
To redeploy a cloned virtual machine, VSC powers off the virtual machine. Make sure this is an acceptable time to take the virtual machine offline. VSC can power on the virtual machine after the redeployment is complete, if you select that option.

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**.
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.

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 Storage .
Reclaim space from one virtual machine	From the vSphere Web Client Home page, click VMs and Templates .

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.

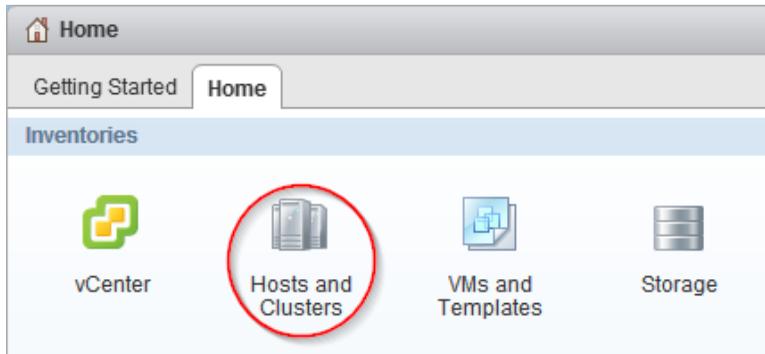
Do not power on the virtual machines while space reclamation is in progress.

Mounting datastores on hosts

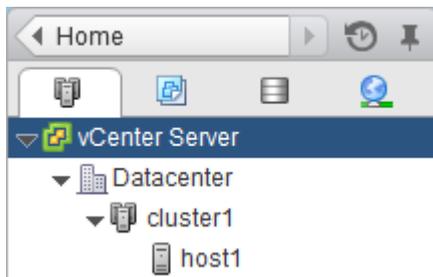
Mounting a datastore gives a host access to storage. You might need to mount a datastore on a host after you add the host to your VMware environment.

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**.

Resizing datastores

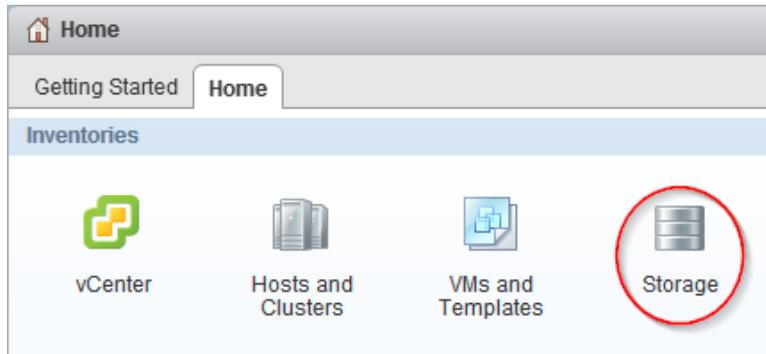
You might need to change the size of a datastore as your infrastructure requirements change.

About this task

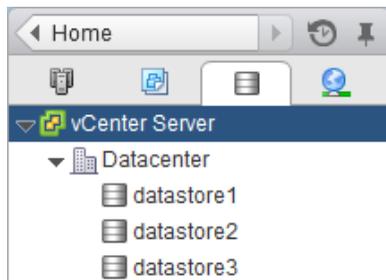
You can increase or decrease the size of an NFS datastore. You can only increase the size of a VMFS datastore.

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**.

Destroying datastores

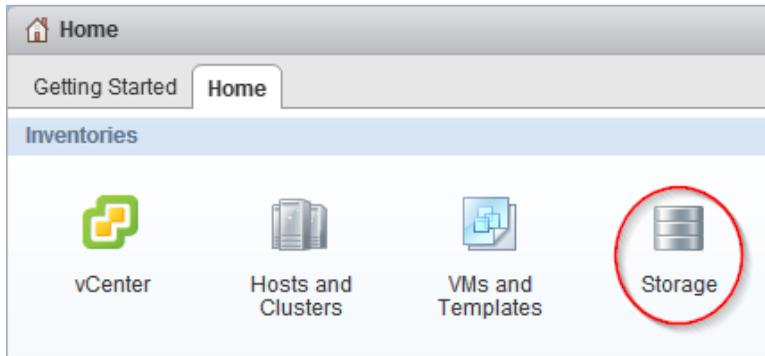
You might need to destroy a datastore when you decommission your virtual machines.

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**.

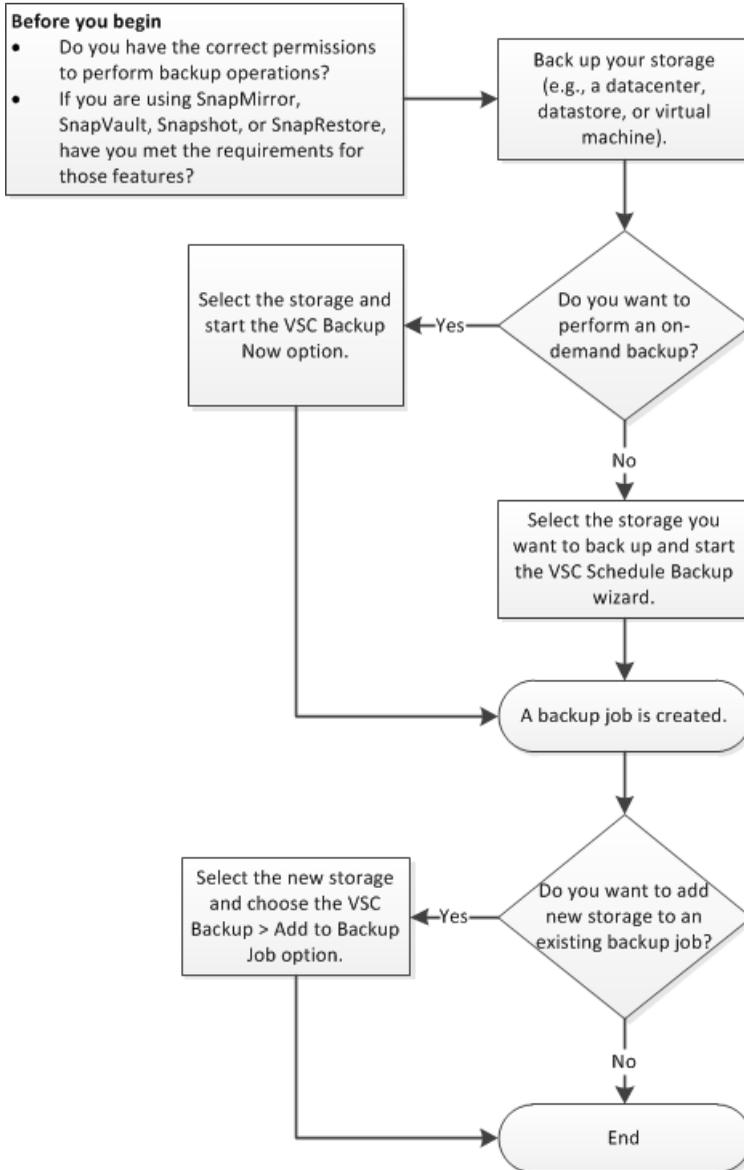


2. In the navigator pane, right-click the datastore and select **NetApp VSC > Destroy**.
3. Click **OK**.

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.



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.

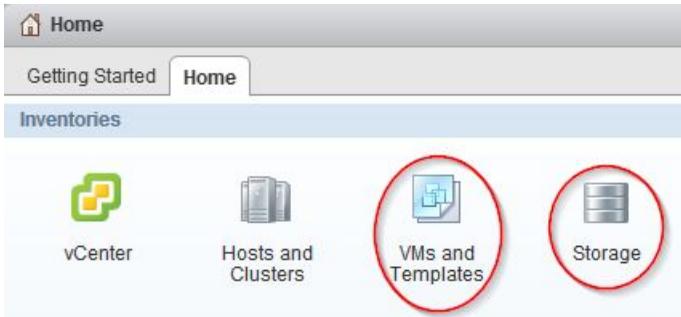
Performing an on-demand backup of a virtual machine or datastore

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. Check the box next to the options you want this backup job to use:

Option...	Action that occurs...
Initiate SnapVault update	<p>Starts a SnapVault update on the selected entities concurrent with every backup copy.</p> <p>Note: The SnapVault option is only supported on clustered Data ONTAP 8.2 or later.</p> <p>The selected entities must reside in volumes that are configured as SnapVault source volumes and a destination volume must also exist.</p>
Initiate SnapMirror update	<p>Starts a SnapMirror update on the selected entities concurrent with every backup copy.</p> <p>The selected entities must reside in volumes that are configured as SnapMirror source volumes and a destination volume must also exist.</p> <p>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 <code>snapmirror.conf</code> file.</p>

Option...	Action that occurs...
Perform VMware consistency snapshot	Creates a VMware snapshot each time the backup job runs.
Include datastores with independent disks	Includes independent disks from datastores that contain temporary data.

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

The 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	<ol style="list-style-type: none"> a. From the vSphere Web Client Home page, click VMs and Templates to view virtual machines or click Storage to view datastores. b. 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. c. Right-click the datastore or virtual machine and select NetApp VSC > Backup > Schedule Backup. You can create a scheduled backup job for an entire datacenter by right-clicking the datacenter and selecting NetApp > Backup > Schedule Backup.

To...	Do this...
Create a backup job using the backup job wizard to select a datastore or virtual machine	<ol style="list-style-type: none"> a. From the vSphere Web Client Home page, click vCenter. b. In the navigation pane, under NetApp, click Backup Jobs. c. Click the Add icon on the Backup Jobs page in the vSphere Web Client.

2. Type a name for the backup job and add a description.
3. Check the box next to the options you want this backup job to use:

Option...	Action that occurs...
Initiate SnapVault update	<p>Starts a SnapVault update on the selected entities concurrent with every backup copy.</p> <p>Note: The SnapVault option is only supported on clustered Data ONTAP 8.2 or later.</p> <p>The selected entities must reside in volumes that are configured as SnapVault source volumes and a destination volume must also exist.</p>
Initiate SnapMirror update	<p>Starts a SnapMirror update on the selected entities concurrent with every backup copy.</p> <p>The selected entities must reside in volumes that are configured as SnapMirror source volumes and a destination volume must also exist.</p> <p>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 <code>snapmirror.conf</code> file.</p>
Perform VMware consistency snapshot	Creates a VMware snapshot each time the backup job runs.
Include datastores with independent disks	Includes independent disks from datastores that contain temporary data.

4. Click **Next** to proceed with the **Schedule Backup** wizard.

Most of the fields in the wizard are self-explanatory. The following table describes the information you need to provide:

Option	Description
Virtual entities (only for backup jobs in which a datastore or VM has not been selected)	Select the virtual entities available for this backup job.

Option	Description
Spanned entities	Spanned entities might be a VM with multiple VMDK's across multiple datastores.
Backup scripts	Select the backup scripts you want to use with this job. If one or more of the selected scripts has been deleted, an error message appears. You can save the backup job without using any of the scripts in the selected scripts list, thereby removing the deleted script from the job. Otherwise, the backup job continues to use the deleted script.
Schedule the time for the backup job	Specify whether this backup job should occur hourly, daily, weekly, or monthly.
Limits on how long a backup job is kept	Specify either a maximum number of days, maximum number of backup copies, or backup indefinitely to retain your backup job. 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.
vCenter Server credentials	Either select Use default vCenter credentials or enter a user name and password for a specific vCenter Server user.
Email alerts	Select the frequency for receiving email alerts. You can add multiple email addresses by using semicolons to separate each email address.
Email server check	Optional: Click Send Test Email to verify that the outgoing email server where the alert notifications are to be sent is working correctly.

5. Review the summary of your selections and click **Finish**.
6. Optional: Select the **Run Job Now** check box to immediately run the job.

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.



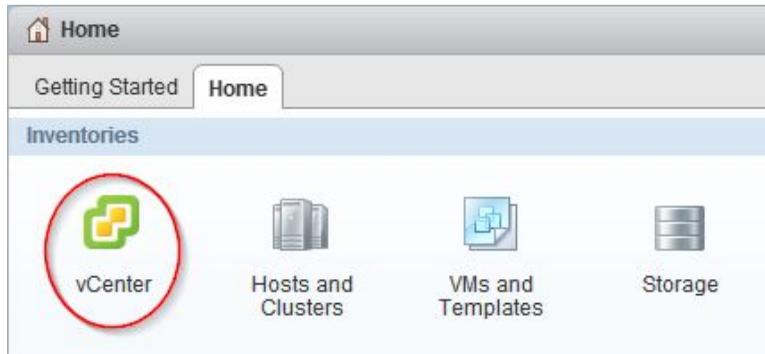
2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Add to Backup Job**.
4. In the **Add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

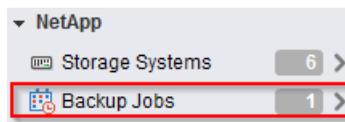
You can modify the name and description, the datastores and virtual machines that are assigned, the backup scripts, the user credentials, the schedule, the retention policy, and the email alerts for an existing backup job using the Modify Backup Job dialog box.

Steps

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



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



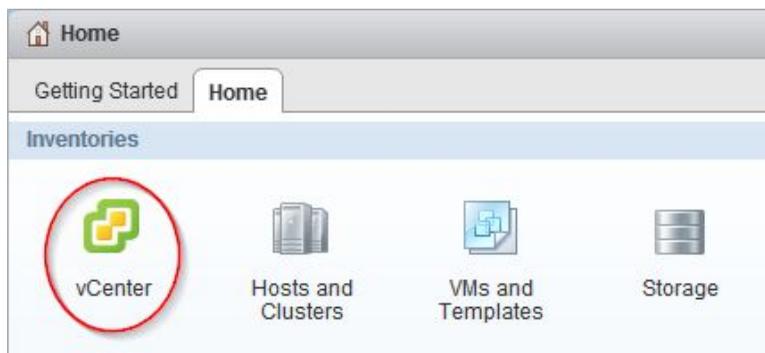
3. Right-click the backup job whose properties you want to modify and select **Modify**.
4. Click the appropriate tab for the properties that you want to modify for this backup job.
5. Modify backup job properties as necessary, and then click **OK** to change the properties.

Suspending an active backup job

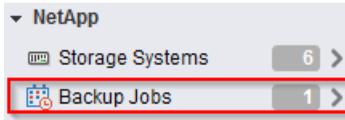
You can suspend the scheduled operations of an active backup job without deleting the job. This gives you the ability to temporarily halt backup jobs in case of planned maintenance, during periods of high activity, or for other reasons.

Steps

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



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



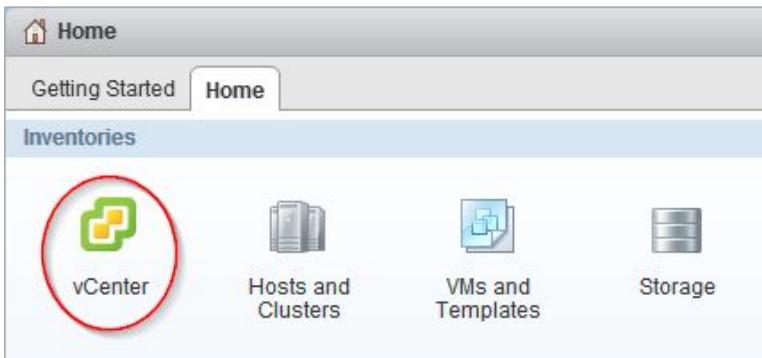
3. Right-click the active backup job that you want to suspend and select **Suspend**.
4. Click **OK** when you receive the confirmation prompt to suspend the active backup job.

Resuming a suspended backup job

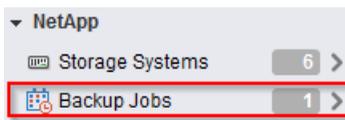
You can resume and run a suspended backup job at any time after you temporarily halt the backup job.

Steps

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



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



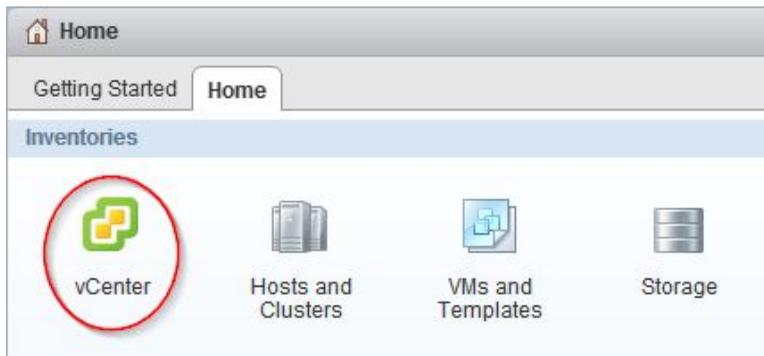
3. Right-click the suspended backup job that you want to resume and select **Resume**.
4. Click **OK** when you receive the confirmation prompt to resume the suspended backup job.

Deleting a scheduled backup job

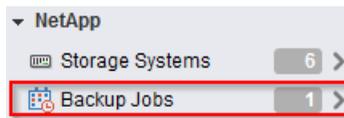
You can select and delete one or more backup jobs from the list of scheduled jobs, but you cannot delete any backup jobs that are running.

Steps

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



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.

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:

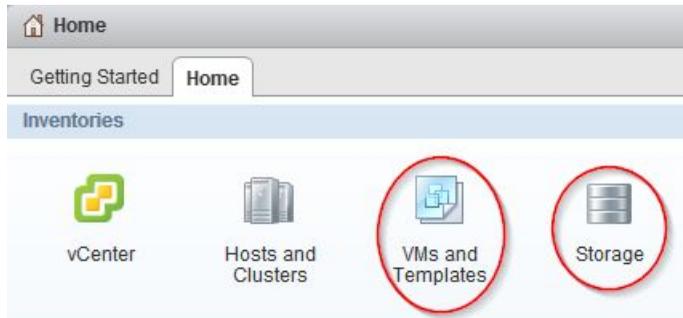
- Yes, if a VMware snapshot operation was successful and the guest operating system was quiesced.
- No, if a VMware snapshot was not selected or the operation failed because the guest operating system could not be quiesced.
- Not Applicable, for entities that are not virtual machines.

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.

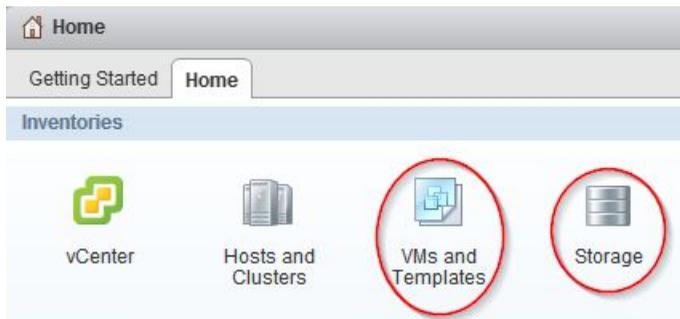
3. Right-click the datastore or virtual machine and select **NetApp VSC > Backup > Mount Backup**.
4. In the **Mount Backup** dialog box, select the name of an unmounted backup copy that you want to mount.
5. Select the name of the ESX server to which you want to mount the backup copy.
You can only mount 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.



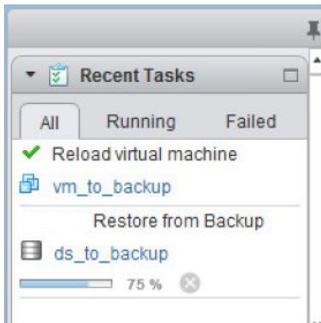
2. In the navigation pane, expand the datacenter that contains the virtual machine or datastore depending on whether you are in the **VMs and Templates** view or the **Storage** view.
3. Right-click the datastore or virtual machine and select **NetApp VSC > Restore**.
4. In the **Restore** wizard, select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

Option	Description
The entire virtual machine	Restores the contents to the last datastore in which it resided from a Snapshot copy with a particular time and date. The Restart VM check box is enabled if you select this option and the virtual machine is registered.

Option	Description
Particular virtual disks	Restores the contents of individual VMDKs to the most current or alternate datastore. This option is enabled when you clear the The entire virtual machine 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.



Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- A**
- active backup jobs
 - suspending [35](#)
 - alignments
 - checking the state of virtual machines [15](#)
 - performing an online alignment [17](#)
 - scanning datastores to determine virtual machine alignment [14](#)
 - audience
 - experienced users [5](#)
- B**
- backup copies
 - finding [38](#)
 - mounting [39](#)
 - restoring data from [38](#)
 - searching for [38](#)
 - unmounting from ESX servers [40](#)
 - backup jobs
 - considerations for restoring data from backups with failed VMware consistency snapshots [38](#)
 - deleting scheduled [37](#)
 - modifying job properties of scheduled [34](#)
 - resuming suspended [36](#)
 - SnapMirror and SnapVault requirements [28](#)
 - Snapshot and SnapRestore requirements [28](#)
 - suspending [35](#)
 - backup operations
 - on-demand [29](#)
 - performing on virtual machines, datastores, and datacenters [27](#)
- C**
- Citrix XenDesktop
 - importing clone data into [9](#)
 - clones
 - creating from a template [9](#)
 - importing clone data into connection broker [9](#)
 - redeploying from a template [21](#)
 - connection brokers
 - importing clone data into [9](#)
 - copies
 - unmounting backup, from ESX servers [40](#)
- D**
- datastores
 - adding to existing backup jobs [34](#)
 - destroying [25](#)
 - enabling deduplication on [12](#)
 - migrating virtual machines to [19](#)
 - mounting on hosts [23](#)
 - on-demand backup operations [29](#)
 - provisioning [7](#)
 - resizing [24](#)
 - restoring [41](#)
 - returning space to (NFS only) [22](#)
 - scanning to determine virtual machine alignment [14](#)
 - scheduling backup jobs [31](#)
 - searching for backup copies of [38](#)
 - deduplication
 - checking the state of [12](#)
 - enabling [12](#)
- E**
- ESX host
 - restoring virtual disks on VMFS datastores [41](#)
- F**
- failed consistency snapshots
 - considerations for restore operations backed up with VMware [38](#)
- H**
- hosts
 - mounting datastores on [23](#)
- J**
- jobs
 - modifying job properties of scheduled backup [34](#)
- M**
- Modify Backup Job dialog box
 - using to modify backup job properties [34](#)

- O**
- on-demand backups
 - starting [29](#)
 - using VSC for [27](#)
- R**
- requirements
 - backup job [28](#)
 - restore operations
 - considerations when using data backed up with failed VMware consistency snapshots [38](#)
 - from backup copies [38](#)
 - from VMDKs [41](#)
 - mounting backup copies [39](#)
 - Restore wizard
 - restoring a virtual machine or its disk files [41](#)
 - resumption
 - of suspended backup jobs [36](#)
- S**
- Schedule Backup wizard
 - using to back up virtual machines or datastores [31](#)
 - scheduled backup jobs
 - adding a virtual machine or datastore [34](#)
 - deleting [37](#)
 - modifying job properties [34](#)
 - of virtual machines or datastores [31](#)
 - using VSC for [27](#)
 - SnapMirror requirements
 - for backing up datastores and virtual machines [28](#)
 - SnapRestore requirements
 - for backing up datastores and virtual machines [28](#)
 - Snapshot requirements
 - for backing up datastores and virtual machines [28](#)
 - snapshots
 - considerations for restore operations backed up with failed VMware consistency [38](#)
 - SnapVault requirements
 - for backing up datastores and virtual machines [28](#)
 - suspended backup jobs
 - resuming [36](#)
 - suspension
 - of active backup jobs [35](#)
- T**
- templates
 - cloning new virtual machines from [9](#)
 - redeploying virtual machines from [21](#)
- V**
- virtual machines
 - adding to existing backup jobs [34](#)
 - aligning I/O non-disruptively [17](#)
 - checking the alignment of [15](#)
 - cloning from a template [9](#)
 - deploying [6](#)
 - migrating a group [19](#)
 - on-demand backup operations [29](#)
 - reclaiming space from (NFS only) [22](#)
 - redeploying from a template [21](#)
 - restarting [41](#)
 - restoring [41](#)
 - scanning to determine alignment [14](#)
 - scheduling backup jobs [31](#)
 - searching for backup copies of [38](#)
 - VMDKs
 - restoring [41](#)
 - VMFS datastores
 - restoring virtual disks on [41](#)
 - VMware consistency snapshots
 - considerations for restore operations backed up with failed [38](#)
 - VMware View Server
 - importing clone data into [9](#)
 - volumes
 - enabling deduplication [12](#)
 - VSC
 - experienced users [5](#)
 - Vservers
 - See* SVMs
- X**
- XenDesktop
 - See* Citrix XenDesktop