Dell Virtual Storage Manager: Installation Considerations and Local Data Protection

Virtual Storage Manager (VSM) end-to-end array management, local data protection and recovery, and VMware vSphere Virtual Volumes (vVols) with Dell PS Series storage

Dell Storage Engineering
November 2019

A Dell Deployment and Configuration Guide

Revisions

Date	Description
June 2015	Initial release
November 2019	vVols branding update

Table of contents

Executive summary

This document provides guidance for VMware® and Dell PS Series SAN administrators on the installation and usage of the Dell Virtual Storage Manager (VSM) versions 4.0 or 4.5. It focuses on installing VSM, highlights the functionality provided by the Datastore Manager component of the VSM plugin, and discusses adding array-based snapshots to data-protection capabilities. Finally, it introduces VMware® vSphere® Virtual Volumes™ (vVols) and discusses the role they play in the data center today and tomorrow.

1 Introduction

Data centers today use VMware virtualization solutions and Dell PS Series SAN storage to consolidate servers and storage for efficient utilization and ease of management. The encapsulation of a virtual machine (VM) into a set of files increases both the flexibility of data protection as well as the challenges of managing the protection of virtualized assets. VMware uses a snapshot technology within VMware vCenter® that can quiesce and help protect the VMs. Dell has combined the intelligence of native point-in-time PS Series SAN snapshots with vCenter snapshots to provide a scalable and automated data protection package for the virtual environment.

The Dell Virtual Storage Manager (VSM) is a next-generation VMware vCenter plugin that allows administrators to coordinate data protection and recovery within their virtual environment, and perform many day-to-day storage administration tasks directly from the vSphere Web Client GUI. The Dell VSM is a virtual appliance that is downloaded as part of the all-inclusive Dell PS Series software support and installed into an existing VMware vCenter environment. VSM contains many tools and capabilities that help VMware administrators gain better control and functionality in their PS Series environment including:

- Datastore management: Provision, expand, delete, and monitor VMFS and vVol datastores across multiple PS Series groups from within vCenter
- Data protection: Create hypervisor-consistent smart copies and replicas for local and remote data protection and disaster recovery
- Dell PS Series vSphere APIs for Storage Awareness (VASA) Provider: A protocol that allows vCenter and the PS Series SAN to communicate, provide better storage awareness within vCenter, and to enable the VMware Virtual Volumes feature
- Seamless vSphere integration: Enables PS Series storage management directly from the VMware new vCenter Web Client

Note: For more information on the VMware snapshot process (which is invoked before the datastore volume is snapped at the SAN level) refer to VMware KB article 1015180, “Understanding virtual machine snapshots in VMware ESXi and ESX” at http://kb.vmware.com/kb/1015180.

1.1 Introduction to Virtual Volumes

VMware vSphere 6.0 introduces Virtual Volumes (vVols), a significant change in how storage is utilized in a virtualized environment. Enabled by the second-generation VASA Provider included with Virtual Storage Manager 4.5, this feature enables storage to be virtual-machine aware, and for virtual machines to be first-class citizens in the storage array.

vVols does not significantly change day-to-day activities for a vSphere administrator; a virtual machine is still a virtual machine, and the workflows within vCenter do not change. What changes on the storage side is that a virtual machine consists of a grouping of volumes on the array (explained in section 1.1.2). This change results in storage-centric tasks being the domain of the array.

1.1.1 Benefits of Virtual Volumes

Cloning a virtual machine, or deploying a virtual machine from a template without vVols, is a large file-copy operation that is accelerated with the VAAI primitive Full Copy. With vVols, the cloning operation (which is manipulating block pointers and reserving space) is completed within a matter of seconds.

VMware recommends that when executing virtual machine snapshots, to limit them to 24 or 72 hours and 2-3 delta files in a chain because performance may be decreased. The workflow remains unchanged, but the old delta file snapshots become efficient pointer-based snapshots on the array. This results in a rapid creation of snapshots that can be used for a quick restore and kept for an indefinite period of time.

Note: While the array firmware permits a volume to have 512 snapshots, the current vSphere vVol implementation is limited to 32 snapshots. Even with this limitation, vSphere administrators are able to complement their current backup strategy with more frequent and rapidly restorable snapshots.

With vVols, a virtual machine is a group of volumes on the array, this enables the existing Dell SAN Headquarters (SAN HQ) array performance-monitoring tool to provide a detailed I/O analysis on a per-virtual-machine and per-virtual -disk level. While similar performance metrics can be seen in vCenter, these are generated from the host side, and cannot show the same level and detail that can be seen on the array side, such as the impact of I/Os, latency, and block size on the underlying physical disks. SAN HQ, coupled with the PS Series vCenter Operation Manager™ adapter, makes this detailed information available within vCenter Operations, and enables both the vSphere administrator and PS Series array administrator to see the same information from their respective preferred interfaces.

1.1.2 Understanding how Virtual Volumes changes storage

Traditionally, when storage has been deployed to a vSphere environment, a volume is created on the array. This becomes a datastore within vCenter where virtual machines are placed. With Virtual Volumes, some of this remains the same; virtual machines are placed within datastores and work flows that depend on this remain unchanged. What has changed is that the object backing up the datastore is a storage container. While it does contain virtual machines and backs up a datastore, it should not be thought of as a volume. Rather, think of the storage container as a reservation of space on the array where virtual machines consume storage. A typical virtual environment design involves using multiple volumes due to queue depth concerns, SCSI-2 Reservation concerns, and differing data-protection needs. With vVol storage containers, this changes to potentially having a single storage container for the entire environment, as Storage Policy Based Management and data-protection policies can be applied at the individual virtual-machine level.

The size of a storage container is limited only by the size of the pool it exists in on a PS Series group. However, non-technical issues, such as a preference for keeping a virtual machine isolated to a specific department or project, may drive a preference for multiple storage containers. In an environment where a PS Series group has multiple pools, multiple storage containers are needed as a storage container does not span pools.

Previously, a virtual machine consisted of a VMX file (configuration file), one or more VMDK files (virtual disk), a VSWP file (memory swap file), and other miscellaneous files including log files. With vVols, a virtual machine consists of a grouping of volumes on the array consuming space from the storage container space reservation. A vVol-based virtual machine consists of the following types of Virtual Volumes:

- **Config**: A small, VMFS-formatted 4GB volume that hosts the VMX and other miscellaneous files, including log files
- **Data**: The equivalent of a VMDK; one exists for each virtual disk attached to the virtual machine
- **Swap**: The equivalent of the VSWP file; it exists only when the virtual machine is powered on

If a VMware snapshot is taken of the virtual machine, two more virtual volume types will exist:

- **Snapshot:** One of these hidden Virtual Volumes exists for each data virtual volume included in each virtual machine snapshot that has been taken, and stores the delta of changes since the previous snapshot was taken.
- **Memory:** This virtual volume is created if the option to include a memory dump with the snapshot is selected.

2 Installation considerations

To ensure a smooth installation of Virtual Storage Manager v4.x, a few considerations should be taken into account and prerequisites completed. The installation process for VSM can be completed using either the vCenter Web Client or the legacy vCenter Client, however, it is recommended that the vSphere Web Client is used.

VSM 4.x can only be used with the vCenter Web Client. For users of the legacy vCenter Client, VSM 3.5 provides similar functionality. For information on VSM 3.5, see the technical reports, *Dell Virtual Storage Manager 3.5: Installation Considerations and Datastores Manager* and *Virtual Machine Protection with Dell Virtual Storage Manager 3.5*.

Note: Minor workflow and GUI differences between the vSphere Web Client under vSphere 5.5 and 6.0 are not called out, and only significant differences are identified.

2.1 vCenter Server managed IP requirement

The installation process for Virtual Storage Manager v4.x has been significantly streamlined over prior versions, with most of the configuration completed during the import of the VSM virtual appliance.

One setting that is populated automatically is the vCenter IP address. This is pulled from querying the vCenter where the VSM virtual appliance is being installed and is dependent on the **vCenter Server managed IP address** field being populated, as shown in Figure 1. Depending on the version of vCenter in use, this may or may not be already populated.

2.1.1 To verify or set the vCenter Server managed IP

1. From the vSphere Web Client landing homepage, click **vCenter**, and then under **Inventory Lists,** click **vCenter Servers**.
2. Select the vCenter Server that manages the environment where the VSM will be installed. In the **Manage** tab, click **Edit**.
3. Under **Runtime settings,** set the **vCenter Server managed address** to the IP address assigned to the vCenter Server.

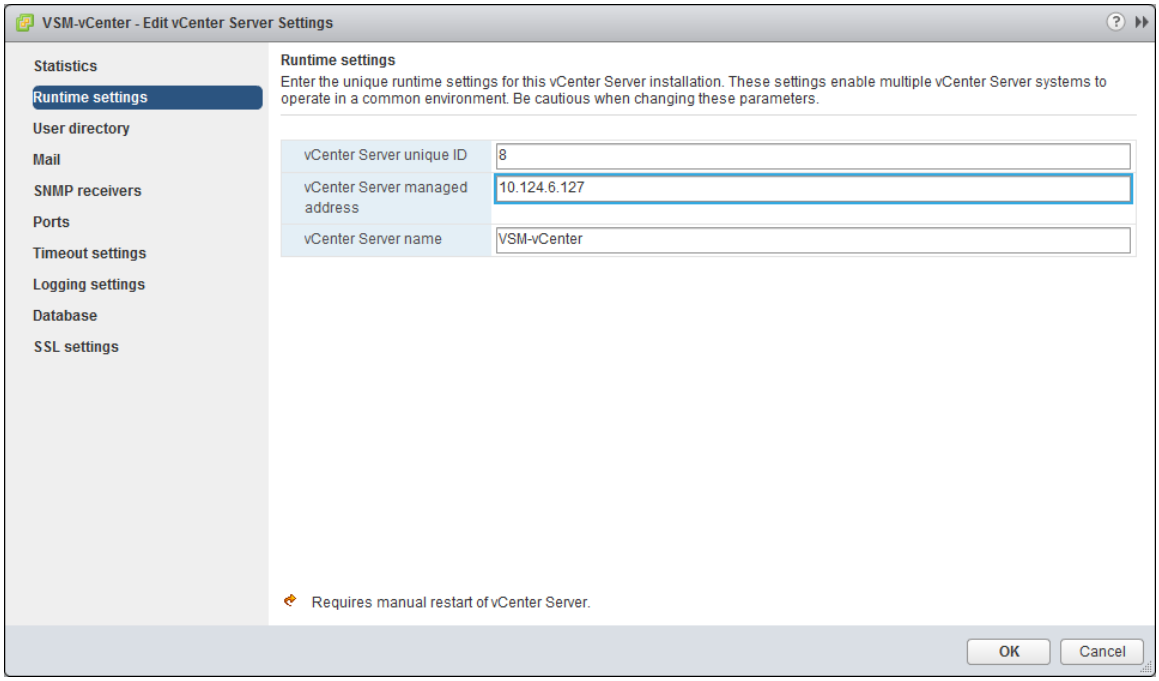

Figure 1 Setting vCenter Server managed address

4. Click **OK** to apply.

2.2 Protecting the VSM appliance

Exercise care when selecting the datastore to host the VSM appliance VM. VSM should not be placed on local storage datastores, or on a vVol datastore. While VSM can be placed on a datastore that is being backed up with replications or snapshots, it should not be placed on a datastore where the replication is managed or Smart Copies are taken by the VSM. This is because the quiescing that is performed as part of these operations may disrupt the VSM while it is managing the replication or Smart Copy. For Virtual Volume environments, do not migrate the VSM appliance to a vVol datastore or storage container.

Because VSM is an integral part of a virtualized environment, especially in environments utilizing Virtual Volumes, protect the VSM virtual appliance and data. VSM must be run on a vSphere High Availability cluster to provide continued availability of VSM in the event of host failure.

For failures that are more serious, it may be necessary to re-install the VSM. The critical data that VSM contains exists in an internal database where a locally stored backup is automatically created each day. This database backup can be accessed through the VSM CIFS share, located at: \\<*IP Address or Hostname*>\database\dbbackup<*hostname*><*date and time stamp*>.sql. This should be copied or backed up to another location. Using the instructions below, this database backup file, coupled with a newly installed instance of VSM, quickly restores a corrupted or accidently deleted VSM to a working state.

1. To install VSM, follow the process for importing the VSM virtual appliance, which is similar to that of other virtual appliances. Refer to the product documentation for further details.
2. Browse to the VSM CIFS share, \\<*IP Address or Hostname*>\database\, and copy the backup of the database to the CIFS share. The backup filename must begin with dbbackup and have an extension of .sql.
3. Launch the VM console for the VSM virtual appliance and log in to the VSM console.
4. From the menu, select **Maintenance**, and then **Database Restore**.
5. When the console displays the database backups in the backup folder, select the appropriate one and press [Enter]. VSM will then begin the process of restoring the database backup.

Note: Depending on the amount of data in the backup, it can take several minutes for the restore to complete.

6. Once the database restore operation is completed, the state of VSM is restored, including the information about the snapshots and replicas that it created on the PS Series array.

   In a Virtual Volume environment, while the VSM appliance is being recreated, do not power on or off, snapshot, or migrate Virtual-Volume-based VMs.

2.3 Connecting to the storage network

Virtual Storage Manager communicates with vCenter and by default, communicates with the PS Series arrays using their Group IP. However, the Group IP exists on the iSCSI network, which is often kept isolated from the rest of the networking environment.

There are three options for enabling communications with the PS Series group:

1. Create an exception in the firewall, isolating the iSCSI network to permit traffic from the VSM to be passed.
2. Frequently in virtualized environments, a Guest OS of the VMs directly accesses the storage to use tools such as EqualLogic Auto Snapshot Manager for Microsoft® Windows and Linux, and offers data protection to applications like Microsoft Exchange and Microsoft SQL Server®. This requires the design of the virtualized environment networking to include a VM network with access to the iSCSI network. VSM can also use this network.
3. Enablement of the dedicated PS Series management network provides access to the PS Series array management functions while maintaining the preferred isolation of the iSCSI network. For details on enabling the management network, see the section, “About Dedicated Management Networks” in the *Group Manager Administrator’s Manual* for the PS Series firmware for your system (available on eqlsupport.dell.com; login required).

Note: Option 1 does not require a second NIC in the VSM. Option 2 requires a second NIC in the VSM. Option 3 may require a second NIC in the VSM depending on the network configuration of the environment. See section 2.4.3, “Adding a second NIC to VSM”.

2.4 Post-install configuration

The installation process of importing the VSM appliance into vCenter configures the network settings of VSM. However, a few important steps remain:

1. Configure the VASA Provider credentials (section 2.4.1). This is required for the Virtual Volumes to function since the VASA Provider is the out-of-band communication channel between the vCenter server and the PS Series array.
2. Change the VSM appliance root password from the default (section 2.4.2).
3. In some environments, as discussed in section 2.3, a second NIC may be required for communication with the PS Series array (section 2.4.3).
4. Finally, register PS Series groups with VSM.

2.4.1 Configuring the VASA Provider

Note: This is a required step for Virtual Volumes.

1. Open the virtual machine console to the VSM appliance.
2. Log in to the VSM console using the default credentials: username: **root** and password: **eql**.
3. From the **Setup** menu, select **Configuration**.
4. From the **Configuration** menu, select **Register VMware vSphere Storage APIs for Storage Awareness**.
5. Provide a **username** for the credentials of the account to be created for the vCenter VASA Service and the PS Series VASA Provider to communicate with, and then press [Enter].

Note: This is a set of credentials unique to the VASA communication between the vCenter VASA Service and the PS Series VASA Provider.

6. Enter a **password** for this account, press **[**Enter] and then re-enter the password for verification.

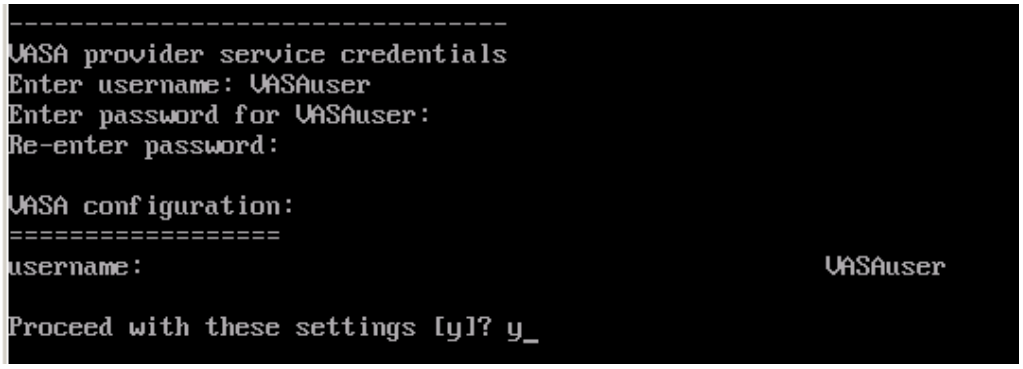

Figure 2 Creating the VASA provider service account

7. Enter **y** to proceed with these settings.
8. The VSM PS Series VASA Provider will then be registered with the vCenter VASA Service. This process takes approximately two minutes. Once complete, the PS Series VASA Provider is listed under **Storage Providers.**

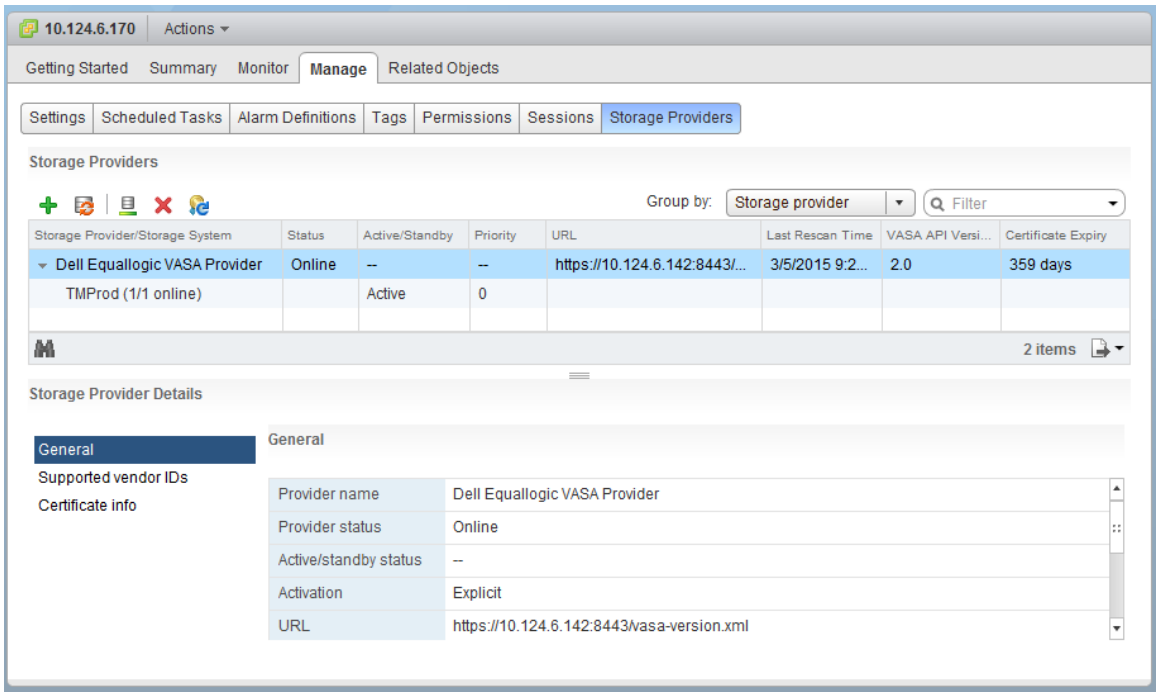

Figure 3 Listing of vSphere Storage Providers – vSphere 6.0 example

2.4.2 Changing the root account password

Change the root password from the default using the following steps.

1. Open a virtual machine console to the VSM virtual appliance.
2. Log into the VSM console using the default credentials: username: **root** and password: **eql**.
3. From the **Setup** menu, select **Configuration**.
4. From the **Configuration** menu, select **Change root password**.
5. At the prompt enter the new password, and then press [Enter]**.**
6. Re-enter the password to verify it, and then press [Enter].
7. Press [Enter] again to return to the main setup menu.

2.4.3 Adding a second NIC to VSM

In some environments, a second NIC is necessary due to the subnet layout or security requirements. Use the steps below for adding a NIC with the vSphere client.

1. From the vCenter Web Client **Home** screen, click the **Dell Virtual Storage Manager** icon.
2. In **Manage** tab, select **Storage Network**, and then click **Edit…**

3. Select the **Enabled** checkbox. From the **Network Connection** dropdown menu, select the appropriate VM Network for communicating with the PS Series array. If using DHCP, select the **Use DHCP** checkbox. Otherwise, enter the static **IP Address** and appropriate **Netmask**.

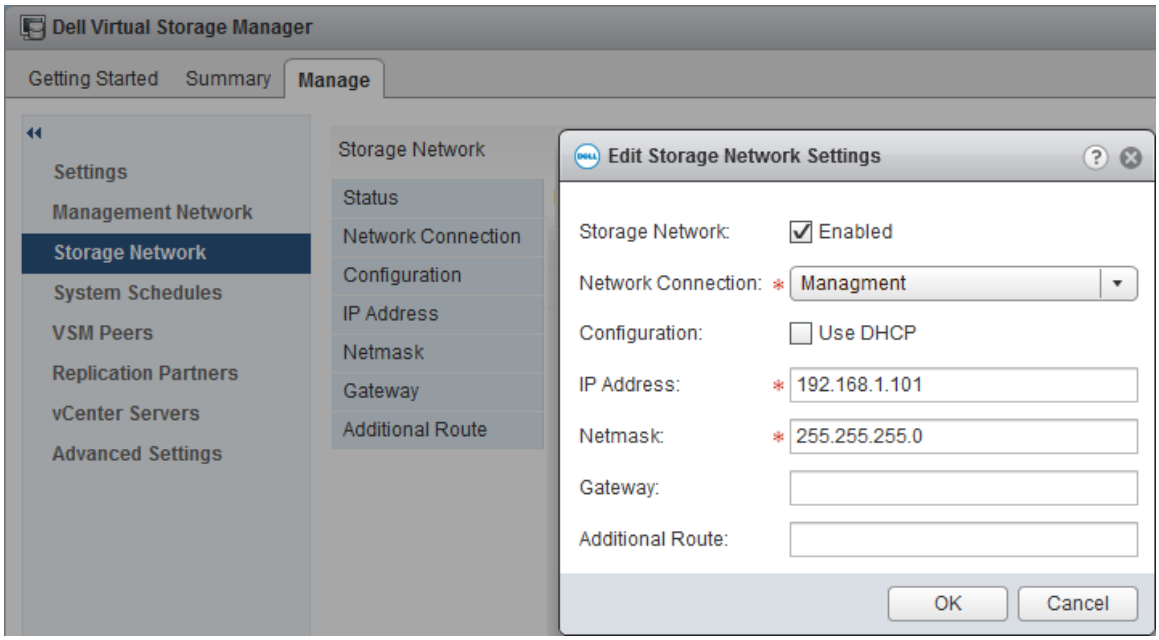

Figure 4 Configuring the optional Storage Network in VSM 4.0

Note: Depending on the network configuration, some environments may require the use of an additional route or a gateway IP address.

4. Click **OK** to continue.

   VSM then reconfigures the virtual appliance and requests a reboot in order to complete the reconfiguration.

5. Once the reboot is completed, VSM can communicate with the PS Series arrays on the newly added storage network.

2.5 Attaching PS Series storage

Configuring VSM to connect with PS Series storage is detailed in "Connecting PS Series storage" below. Configuring access controls for a VMFS volume can be done if a few different ways depending on preferences. Access controls for Virtual Volumes is handled differently from VMFS volumes, and is also explained below.

2.5.1 Connecting PS Series storage

Once the installation and configuration of VSM is completed, the PS Series storage array(s) supporting the virtualized environment need to be registered with VSM.

Use these steps to connect to a PS Series group.

1. From the vSphere Web Client **Home** screen, click the **Dell VSM** icon to view the VSM Home page.
2. On the **PS Groups** menu, select the **Objects** tab, and click the **Add Groups** icon.

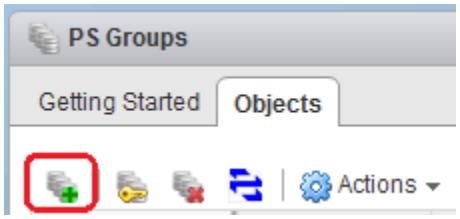

3. In the **Add PS Series Group** dialog box, enter the group name or IP address and credential with grpadmin privileges, and then click **Add**.

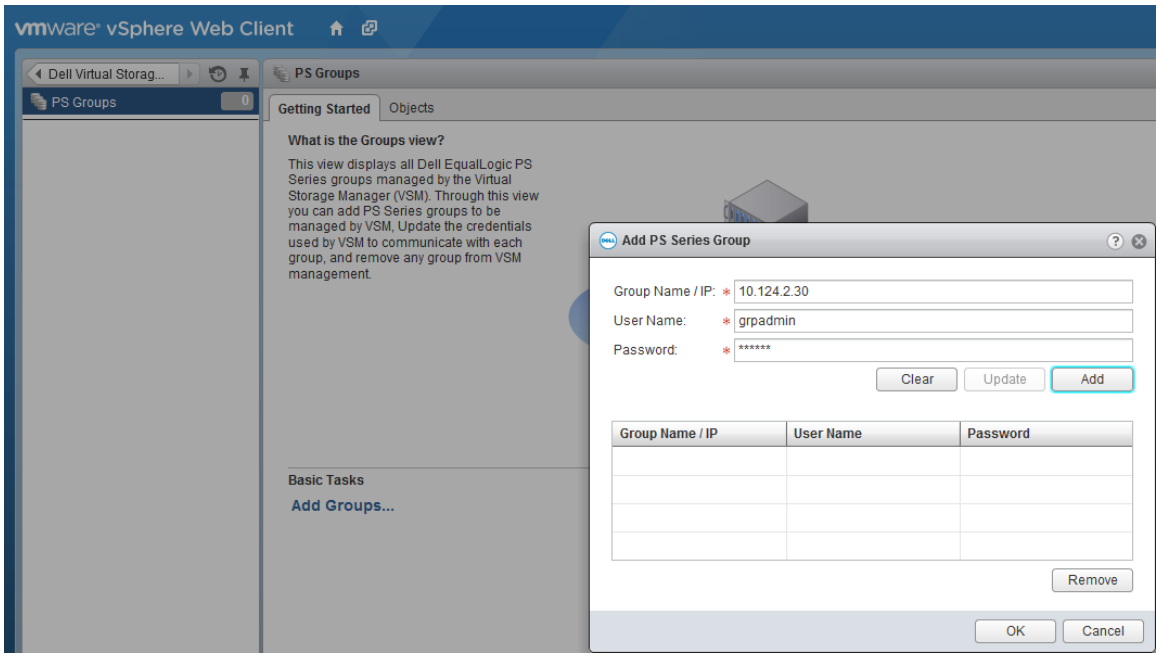

4. Repeat step 3 for each group that will be managed by this VSM, then click **OK** to continue.
5. VSM connects to each of the groups, and begins to populate VSM with information about the groups. The time it takes for this to occur depends on the configuration of the groups, and the number of volumes and VMs that are stored on them.

VSM can now be used to manage these PS Series groups, enabling the creation of VMFS and vVol datastores, and the protection of VMs through SmartCopy Snapshots and SmartCopy Replication.

2.5.2 Access control for VMFS datastores

The ability for a VMware ESXi™ host to access a particular volume on a PS Series SAN is restricted by the volume access controls. Controls are vital for the integrity of the data because serious issues can arise if any server can access any volume. VSM enables the creation of datastores, and their underlying volumes, directly from the vSphere Web Client, and will create the access controls based on user input.

The access controls for volumes created with VSM can be done in two ways:

- Have VSM auto-generate an access control policy or access control list (ACL).

  For PS Series groups running firmware 7.0 and above, VSM will assign an Access Control Policy that contains the IQNs of the servers contained in the cluster or server selected on the Hardware Resource step. If no existing Access Control Policy meets these requirements, a new Access Control Policy will be created. For more information on Access Control Policies and Access Control Lists, see the PS Series technical report titled, *Access Control Policies*.

  For PS Series groups running firmware prior to 7.0, VSM will create a traditional Access Control List (ACL) on the volume that contains the IQNs of the servers contained in the cluster or server selected on the Hardware Resource step. It is important to note the ACL are limited to sixteen entries.

- Where auto-generating an access controls is not preferred, an ACL can be manually created.

  Manually created ACLs can consist of up to sixteen entries containing the CHAP user name, IP address, or IQN to reflect servers that should be permitted access, and can be saved as a template for future re-use.

2.5.3 Access controls for Virtual Volumes datastores

vSphere 6.0 introduces the vVol datastore which is a logical space on the array known as a Storage Container where the virtual machines are placed. Storage Containers differ from traditional VMFS datastores in that a file system is not used. Instead, virtual machines are directly using volumes on the SAN.

Access controls for storage containers are handled at the protocol endpoint using the steps below.

1. Click **PS Groups**, and then select the group with the protocol endpoint that a host or cluster will access.
2. In the **Manage** tab, click **Protocol Endpoint**.
3. Select the host that needs to be accessed and click **Allow**. Alternatively, select a cluster and click **Allow All**.

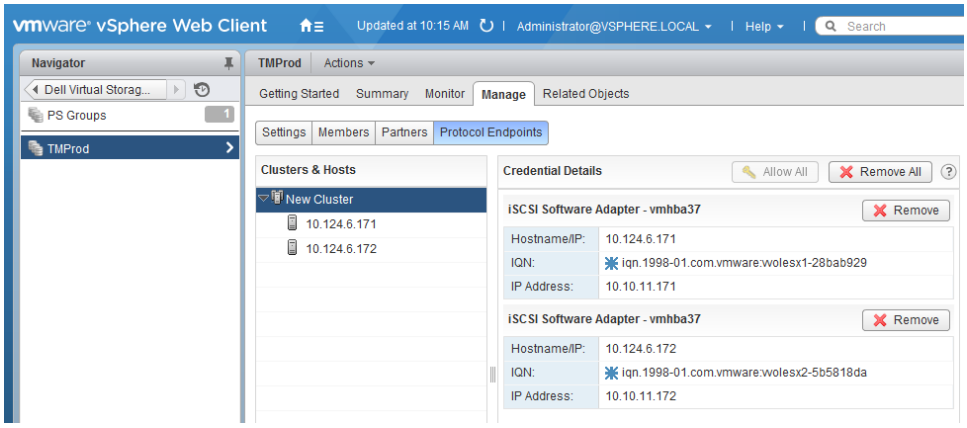

The selected host or cluster is given access to the protocol endpoint, and in turn, access any storage containers created on that group.

Note: See appendix C for a glossary of Virtual Volume terms. See the document, *VMware vSphere Virtual Volumes on Dell PS Series Storage* for information on Virtual Volumes.

3 VSM datastore management

The design principles behind VSM provide the vSphere administrator with:

- A view into the storage that supports the datastore presented to the vSphere virtualized environment
- The ability to perform normal day-to-day storage administration tasks
- The ability to use all the data protection capabilities of a PS Series array

For details on the data protection capabilities of VSM, see sections 6 and 7.

The next section gives examples of the information views that VSM provides. These views supply the vSphere administrator with the current status of the storage.

3.1 Dell Storage view

This first view, accessed from **VSM** > **PS Groups**, shows a high level overview of the status of all the PS Series groups registered with the VSM. The vSphere administrator can quickly see the overall array storage capacity used, and any events of concern on the storage that may need attention.

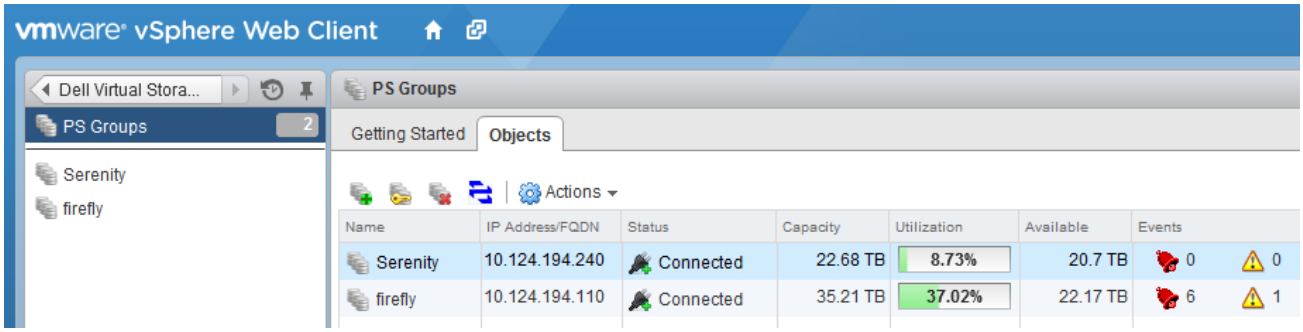

Figure 5 Listing of PS Series groups registered with VSM

In the example shown, the second array named **firefly**, which is utilizing 37% of its 35TB capacity, has six errors and one warning event. Double clicking on the entry allows the administrator to view additional information about the errors and warnings as well as on the array as a whole.

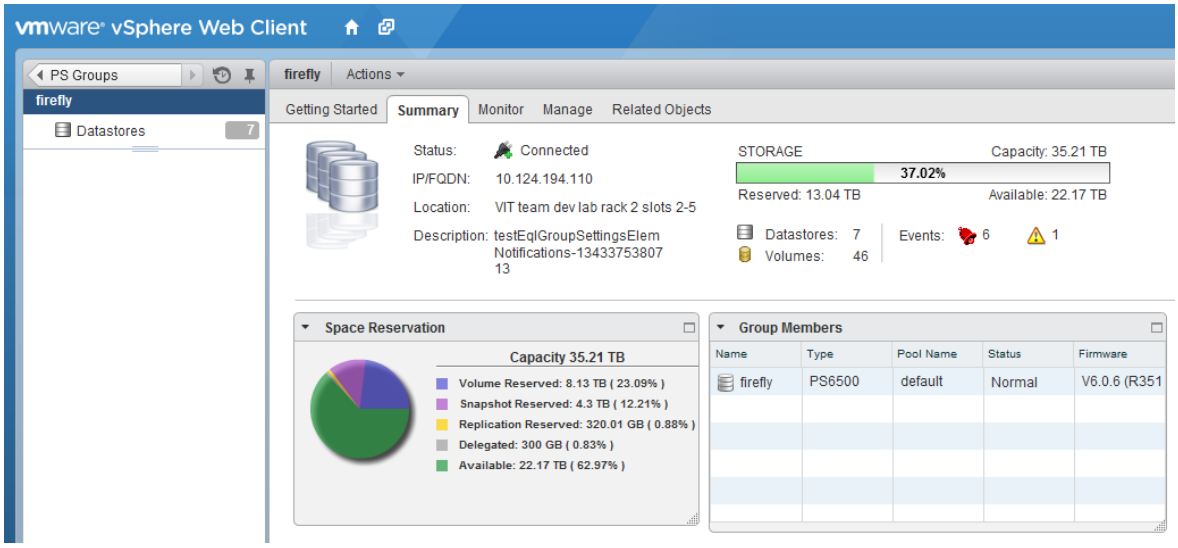

Figure 6 Summary screen of the **firefly** Group registered with VSM

The detailed view displays several tabs with additional information about the selected **firefly** group. The **Summary** tab displays details about the individual members that make up the group, a breakdown of the storage consumption, information on the total number of volumes on the array, and how many datastores are used in the vSphere environment.

Additional information on the members that make up the group can be found on the **Manage** tab under **Members**. Also, under the **Manage** tab is the **Partners** section where information on replication partners can be found and configured.

Details about the error and warning events called out in Figure 5 can be found in the **Monitor** tab. More information on these events can be found in the *VSM Users Guide* and the PS Series *Group Manager Administrator's Manual* on eqlsupport.dell.com (login required).

The **Related Objects** tab has a detailed view of the datastores backed by this group. This view displays the same information as seen and described below, but it shows all the groups registered to the VSM.

3.2 VSM Inventory Datastores view

The **Datastores** view under **VSM Inventory** provides a lot of useful information at a glance. The vSphere administrator can quickly see the capacity, snapshot and replication status of all the datastores in the environment.

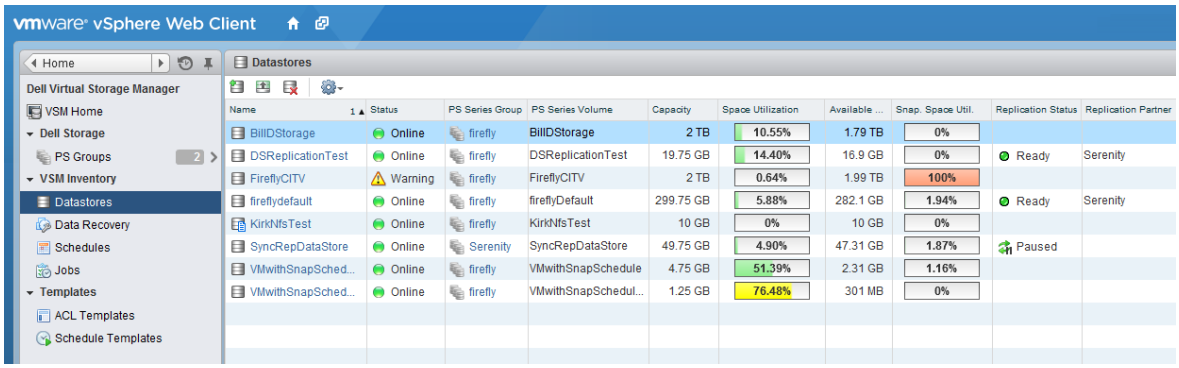

Figure 7 VSM Datastores view

This Datastores screen displays the following states:

- The datastore **FireflyCITV** has a warning status because its **Snapshot Space Utilization** is exceeding 90%.
- The datastores **DSReplicationTest** and **firefly** are both being successfully replicated from the **firefly** group to the **Serenity** group.
- The datastore **SyncRepDataStore**, which is hosted on the Serenity group, is configured for synchronous replication, but its **Replication Status** is currently **Paused**.
- The datastore **VMwithSnapSchedule-clone-11-25-2013-1:26-PM** Space Utilization is in a **Yellow** state because it is consuming more than 60% of its capacity.

Combining all these statistics from vCenter and the PS Series arrays on one screen provides the vSphere administrator with a way to quickly check the status of their datastores, snapshots and replicas. This enables them to stay more informed about their overall environment and to manage it more efficiently.

3.3 Datastore management

Virtualized environments have made it easier to meet the dynamic needs of businesses. VSM enables you to complete many storage tasks directly from the vSphere Web Client.

A toolbar at the top of the **VSM Inventory** > **Datastores** screen provides easy access to starting points for the common storage tasks such as: creating, resizing, and deleting datastores.

3.3.1 Creating a VMFS datastore

VSM can be used to create a single datastore or multiple datastores, backed by either PS Series block or file storage. See section 3.3.2, “Creating a vVol datastore”.

1. From the **VSM Inventory** > **Datastores** page click the **Create Dell Datastore** icon 
2. Enter a **Name** for the datastore, and the volume on the array backing the datastore will be created with the same label. Select the **Type** of datastore to be created (VMFS-5, VMFS-3 or NFS), and select the **Inventory Location** (for example, Datacenter, Datastore Folder, or Datastore Cluster) where the datastore should be placed.
3. Select the cluster or individual host where the datastore should be mounted.
4. From the dropdown menus, select the **PS Group** and the **Storage Pool** where the volume backing the datastore will be created.
5. Enter the **Number of Datastores** to be created as well as the **Volume Size**. Select if the volume is to be **Thin Provisioned** and if **Enable Thin Provision Stun** is to be enabled. (For more information on VMware Thin Provision Stun, see the document, *Dell PS Series Arrays: Advanced Storage Features in VMware vSphere*.) Select the percentage of additional allocated space for the **Snapshot Reserve** and **Snapshot Reserve Warning** threshold. Finally, select whether to **Enable Snapshot Space Borrowing**.

   At the bottom of the dialog, there is a pie chart representing the **Storage Pool Capacity** and the impact the selections will have.

Note: vSphere 5.0 and above support 64TB datastores, however the maximum PS Series volume size is 15TB. If multiple datastores are being created, the datastore name will be suffixed with a two digit number starting with 01. If a volume of that name already exists, the next higher number is selected.

6. The **ACL Source** has two options: **Auto-generate ACL** or **Specify ACL**. If Auto-generate ACL is selected, it creates an ACL on the volume based on the IQNs of the hosts to be granted volume access.

Note: If the array is running firmware version 7.0 or later it creates or uses an Access Control Policy that reflects the IQNs of the hosts.

7. If **Specify ACL** is selected, an existing **ACL Template** can be selected from the dropdown menu, or a new ACL can be created, and optionally saved as an ACL Template.
8. Steps 5, 6, and 7 in the GUI enable a volume to be assigned a Snapshot Schedule, and configured for Replication or Synchronous Replication. For more information on these data protection strategies see section 6.
9. The final dialog in Create Datastore is a summary page detailing the chosen options. Click **Finish** to continue. VSM will then perform all the steps required, and the datastore will become available within the vSphere environment for use.

3.3.2 Creating a vVol datastore

Creating a vVol datastore from VSM is similar to creating a VMFS datastore. From vCenter both look the same, however on the SAN there are significantly different. A VMFS datastore is backed by a traditional volume, and formatted with the VMFS file system. A vVol datastore is backed by a Storage Container, where a file system is not used. Instead, virtual machines consist of a group of virtual volumes placed within a Storage Container.

1. From the **VSM Inventory** > **Datastores** page, click the **Create Dell Datastore** icon 

2. Enter a **Name** for the datastore, select the datastore **Type VVol**, and select the **Inventory Location** (for example, Datacenter, Datastore Folder or Datastore Cluster) where the datastore should be placed.
3. Select the cluster or individual host where the vVol datastore should be mounted.
4. From the dropdown menus, select the **PS Group** and the **Storage Pool** where the Storage Container backing the datastore will be created.
5. Enter the **Number of Datastores** to be created as well as the **Container Size**. At the bottom of the dialog, there is a pie chart representing the **Storage Pool Capacity** and the impact the selections will have on it. Click **Next** to continue.
6. The final dialog in Create Datastore is a summary page detailing all the chosen options. Click **Finish** to continue. VSM will perform all the steps required, and the vVol datastore will become available within the vSphere environment.

3.3.3 Resizing a datastore

Virtual environments are in continuous flux, reflecting the ever-changing demands of the businesses they support. What is expected from storage often shifts from day to day, and the PS Series array architecture is designed to continuously balance the workload among several arrays. For more information on the PS Series load balancers, see *EqualLogic PS Series Architecture: Load Balancers*.

Performance is not the only storage requirement that shifts over time. Often the capacity requirements of a datastore can change over time as well. VSM enables you to increase the size of a datastore in a minimal number of steps.

1. From the **VSM Inventory** > **Datastores** page, highlight the datastore to be resized, and then click the **Resize Dell Datastore** icon  . Alternatively, right click a datastore and select **Resize Dell Datastore** from the context menu.
2. Select a larger **Volume Size** for the datastore. Optionally, the volume **Snapshot Reserve** of a datastore can be altered.
3. Click **OK** to start the resize task.

Note: It is only possible to increase the size of a datastore.

When VSM finishes increasing the size of the volume on the array and VMFS partition on the volume, additional capacity on the datastore is available.

vVol datastores can be resized in the same manner.

3.3.4 Deleting a datastore

Occasionally, datastores need to be deleted from the virtual infrastructure. Without VSM, this involved a number of steps in vCenter, followed by some steps on the array. VSM reduces this procedure to just a few clicks from within vCenter.

Note: VSM will not delete a datastore that has registered VMs on it.

1. From the **VSM Inventory** > **Datastores** page, highlight a datastore to resize, and then click the **Delete Dell Datastore** icon . Alternatively, right click a datastore and select **Delete Dell Datastore** from the context menu.
2. A verification dialog is displayed, with the warning ALL DATA WILL BE LOST! Select the checkbox to acknowledge the deletion and enable the **OK** button allowing you to proceed.

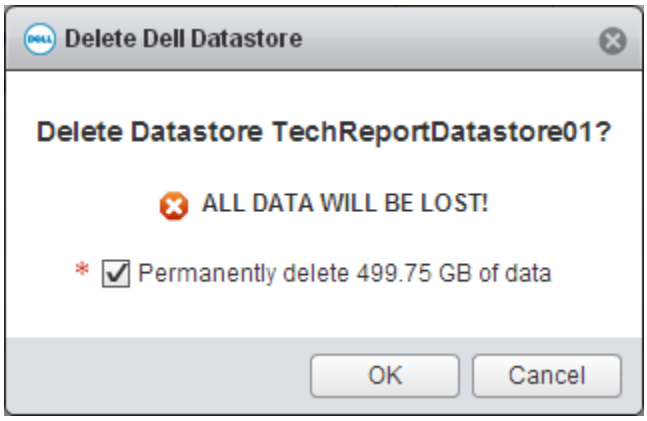

Figure 8 Datastore deletion warning

3. When you click **OK**, VSM performs all the necessary tasks within vCenter to unmount the datastore from all of the hosts, log the iSCSI initiators out, take the datastore offline and delete the datastore volume from the PS Series array.

   vVol datastores can be deleted in the same manner.

3.3.5 VMFS datastore access policy

Access to iSCSI volumes is restricted by an Access Control List (ACL) entry. When creating a datastore through VSM, two options are presented for restricting access to the volume in the ACL Source section: **Auto-generate ACL** and **Specify ACL**.

Auto-generate ACL behaves differently depending on the version of array firmware. For all versions of firmware 6.0 and prior, Auto-generate ACL creates an ACL entry or entries consisting of the selected host(s) IQN. For arrays running firmware 7.0 and above, Auto-generate ACL uses the new Access Policies feature. If an existing Access Control Policy or Access Control Policy Group meets the access needs of the selected cluster or hosts, it will be used. If not, a new Access Control Policy or Access Control Policy Group is created to meet the access needs of the cluster or hosts. For more information on Access Policies, see the document, *Access Control Policies*.

3.3.5.1 Creating a new ACL template

ACL templates provide a way to predefine the access policy used for a volume on an iSCSI array. Typically, in a virtualized environment, all hosts in a cluster access the same volumes, so a single ACL template would be created to reflect this. This results in one ACL template for each vSphere cluster in the environment. Occasionally, there may be exceptions where access to a datastore is limited to only a few hosts in a cluster or where access to a datastore is required from hosts in multiple clusters. ACL templates can be created to reflect these requirements.

1. From the VSM home, under **Templates**, click **ACL Templates**.
2. Click the **Create ACL Template** icon .
3. Provide a name for the new ACL Template and optionally add a description.
4. Click **Add**, and enter an ACL entry or entries. An ACL entry can consist of either a CHAP user name, IP address, or iSCSI initiator name or a combination of these. Click **OK** to save the ACL entry.
5. If the ACL Template requires multiple ACL entries, click **Add** again and repeat step 4 until the policy has all the ACL entries required.

Note: If IP-address ACLs are used, include the IP address of all the VMkernel Ports bound to the iSCSI initiator on each host. If IQN ACLs are used on hosts with dependent or independent HBAs, include the IQN of each HBA used for the iSCSI from each host.

6. Once the ACL Policy has completed, click **OK**.

3.3.5.2 Creating an ACL policy from an existing datastore volume

If deploying VSM into an existing environment, there may be existing datastore volumes that have an ACL suitable for creating an ACL Template.

1. Click **VSM Inventory** > **Datastores**, and then right click on the datastore that has a suitable ACL.
2. From the context menu, select **All Dell VSM Action** > **More Uncategorized Actions** > **Clone ACL from Volume**.
3. Provide a suitable name for the cloned ACL Template and optionally add a description.
4. Click **OK**.

After this process, datastores created with VSM will have the option of selecting the new ACL Template as the ACL Source in the Access Policy section.

4 Role-based access controls

Dell VSM is a powerful tool that enables vSphere administrators to complete storage management and data protection tasks from within vCenter. However, in many vSphere environments, there will be multiple vSphere users with varying degrees of knowledge, skills, and responsibilities. VSM 4.0 introduces over 50 additional privileges enabling the creation of very granular vCenter roles that reflect the skills and business needs of the individual user.

From a high level, the vCenter predefined roles range from limited read-only access to full administrative access, with a number of additional role possibilities. VSM builds on the defined roles and helps create other roles that meet the business access requirements.

Some businesses clearly define teams and roles (for example, a network team, server team, storage team, and application team). Because hypervisors cross over several core data center roles, more-granular roles were needed. The additional privileges that VSM provides allows vSphere administrators to create a vCenter storage administrator role that enables members of the storage team to add PS Series groups, deploy additional volumes, and present those volumes as VMFS formatted datastores to the vSphere environment. At the same time, the storage team members' ability to interact with VMs can be limited since that is a task for the hypervisor team which has not been given this permission.

Even in businesses that do not have distinct teams, there are tasks that require distinct access permissions. For example, incorrectly recovering a VM from a snapshot can have undesired consequences if a *Rollback Restore* is performed instead of a *Restore VM from Snapshot*. The role privileges that VSM provides permit that level of granularity. It also allows roles that are limited from all snapshot and replication data protection privileges.

For details on the more than 50 vCenter privileges that VSM enables in vSphere, see the *VSM Installation and User's Guide* on eqlsupport.dell.com (login required).

5 VASA Provider

The PS Series VSM virtual appliance includes the PS Series VASA Provider. The VASA Provider is a set of APIs that enable vCenter to communicate with the virtual environment's underlying storage. This non-SCSI communication with the PS Series array enables vSphere to learn the capabilities of each datastore volume presented to the virtual environment. These datastore volume capabilities are displayed in a number of locations in the vSphere Client interface, providing virtual administrators with valuable information about their storage infrastructure.

VSM displays these datastore volume capabilities in the Datastore section of the plugin, in the System Storage Capability column, as shown in Figure 9.

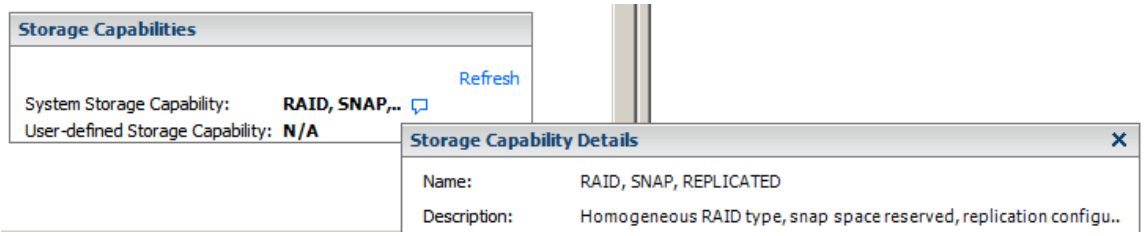

Figure 9 Storage capability of a datastore as shown in the legacy vCenter Client

While these capabilities of VASA do not require a particular vSphere license, there are three vSphere features that leverage this information from VASA: Storage Based Profile Management (formerly Profile-Driven Storage), Storage Distributed Resource Scheduler, and Virtual Volumes. Details on Virtual Volumes are provided in section 7 of this document.

Storage Based Profile Management uses the VASA-provided information to ensure that virtual machines reside on the datastores that meet the VM needs. An administrator can create various storage profiles that reflect particular data protection needs. When a VM is created, the administrator can select the storage profile that meets the requirements of the VM and place the VM on a datastore that provides these capabilities. If the VM is later migrated to a datastore that is not suitable for the storage profile, or should the capabilities of the datastore change and no longer fit the storage profile, the VM fails the storage profile compliance check. This compliance status can be seen on the individual VM summary page, and on the VM storage profile for all VMs assigned to a storage profile. Storage Based Profile Management with the PS Series VASA Provider enables administrators to place VMs on the right datastore and quickly ensure that VMs continue to reside on a datastore that meet their needs.

Storage Distributed Resource Scheduler (Storage DRS) uses the VMware CPU and memory resource management concepts and applies them to datastores. Similar to the VMware traditional DRS capability, Storage DRS groups datastores with like performance characteristics into a datastore cluster. When a VM is deployed, it is not deployed to a particular datastore but rather to a datastore cluster. Storage DRS determines where to place the VM, based on space utilization and I/O load. Like DRS, Storage DRS continuously monitors the cluster space utilization, and the I/O load using storage I/O control. Should space utilization or I/O response time thresholds be exceeded, or if there is a significant difference in space utilization among the datastores within the datastore cluster, Storage DRS seeks to relocate a VM using Storage vMotion®.

However, while Storage DRS is aware of the datastores (and volumes) involved, it is not aware of the volume location in the PS Series storage. Therefore, prior to initiating a Storage vMotion action on a VM, Storage DRS consults with the PS

Series VASA Provider to find out whether the migration of the VM and its workload would benefit the overall I/O workload distribution of the PS Series array. If the migration will not result in an improvement in the distribution of the I/O workload (for example, if the volumes involved reside on the same PS Series group members) the VASA Provider informs Storage DRS not to perform the migration. Conversely, if the VASA Provider agrees that the migration will result in an improvement in the distribution of I/O (for example, if the volumes involved reside on different PS Series group members) the Provider approves the migration request. In this case, Storage DRS leverages Storage vMotion to move the VM and its I/O workload to the selected datastore.

In another parallel to the VMware classic DRS feature, Storage DRS has the ability to operate in maintenance mode. When a datastore in a cluster is placed in maintenance mode, the VMs and VMDKs residing on the datastore are moved to other datastores within the datastore cluster by Storage vMotion. Storage DRS ensures that the I/O workload and space utilization remains balanced across the remaining datastores that are not in maintenance mode.

Storage DRS also has a placement constraint rule that is enforced during migrations. The first option, enabled by default, is the **Intra-VM VMDK affinity rule** which keeps all of a specific VM VMDKs together on the same datastore. The inverse of that rule, the **VMDK anti-affinity rule**, keeps the VMDKs of a specific VM on separate datastores within the datastore cluster. Finally, there is the **VM anti-affinity rule** which prevents certain VMs from sharing the same datastore.

6 Local data protection strategies with VMFS datastores

Once the VSM appliance is installed and running in the environment, a **Dell Virtual Storage Manager** icon appears in the vCenter Web UI.

Throughout this document, the term VSM snapshot refers to the coordinated protection process of VMware snapshots and PS Series SAN snapshots being used together to create a hypervisor-aware array snapshot recovery point.

In addition to launching VSM from the Home screen icon, there are options available inside the vCenter Web UI. In the Hosts and Clusters view, right click an object in the left pane to reveal the **All Dell VSM Actions** menu with the available and relevant tasks. These PS Series menu options show up throughout the vCenter Web UI whenever a PS Series VSM-related task can be performed. There are multiple points where the VSM data-protection wizards are accessible. All achieve the same result, and all are based on ease of use and comfort with the tools.

To launch the VSM GUI and manage or monitor snapshots, click the **Dell Virtual Storage Manager** icon in the **Home** screen > **Inventories** section.

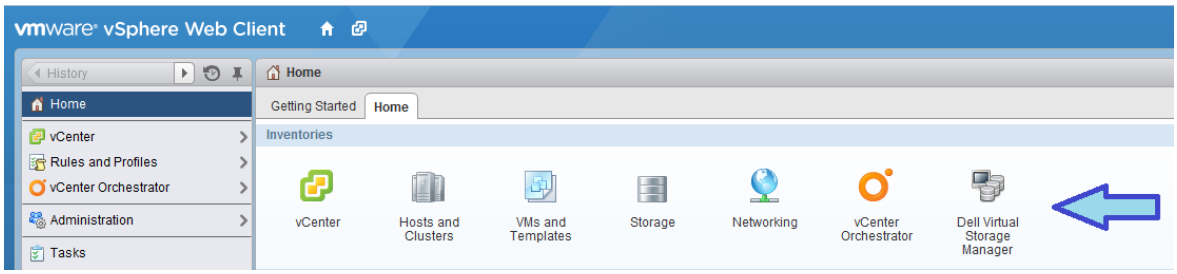

Figure 10 vCenter inventories

Options available from this screen include launching all of the views within VSM and managing or monitoring datastores, snapshots, and replicas. This document focuses on local virtual machine data protection with snapshots.

6.1 Protection with VSM snapshots

A VSM Snapshot is a hypervisor- or application-aware VMware snapshot combined with a PS Series SAN snapshot. When VMware puts the VM into snapshot mode, it quiesces the I/O to the virtual machine VMDK files and, if possible, quiesces the application inside the VM. The level of application consistency is based on the operating system of the VM, the VMware tools, and the application. There are multiple options, including the ability to save the memory to a disk, but once these VMs are quiesced, any new changes to the VM are stored in a separate delta VMDK. Once the VM is quiesced, VSM coordinates with the SAN to determine which PS Series volume(s) to snapshot. These datastore volumes have a PS Series snapshot created on them and then VSM coordinates with vCenter to release the VM snapshot. The benefit to this is that the same consistency is obtained without leaving the virtual machine in snapshot mode for an extended period of time, which could possibly lead to longer consolidation times for the snapshot and space consumption on the datastore.

There are multiple ways to launch the Create Snapshot wizard. This design flexibility allows each of the various features to be launched from a variety of places, including Hosts and Clusters view, VMs and Templates view, and Storage view.

The supported objects for a VSM Snapshot are: VM, VM Folder, Datastore, Datastore Folder, and Datastore Cluster.

6.1.1 Creating a snapshot

All of the available actions for a particular object appear in the right-click menu of any supported object. **All Dell VSM Actions** are listed at the bottom of the menu.

1. Select an object and click **All Dell VSM Actions** > **Create Snapshot** to take a VSM snapshot using the Create Snapshot wizard.

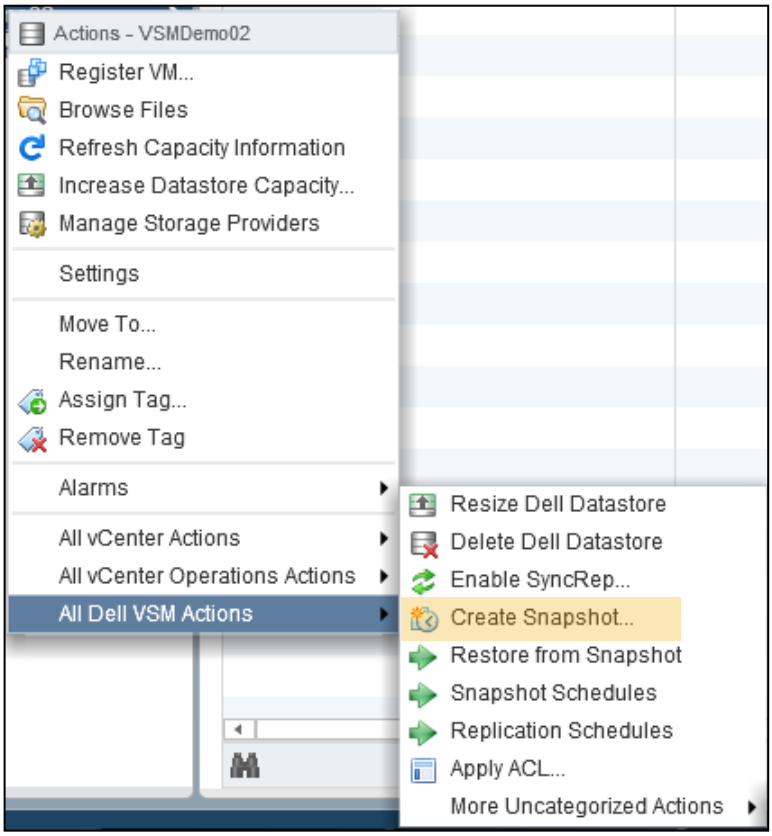

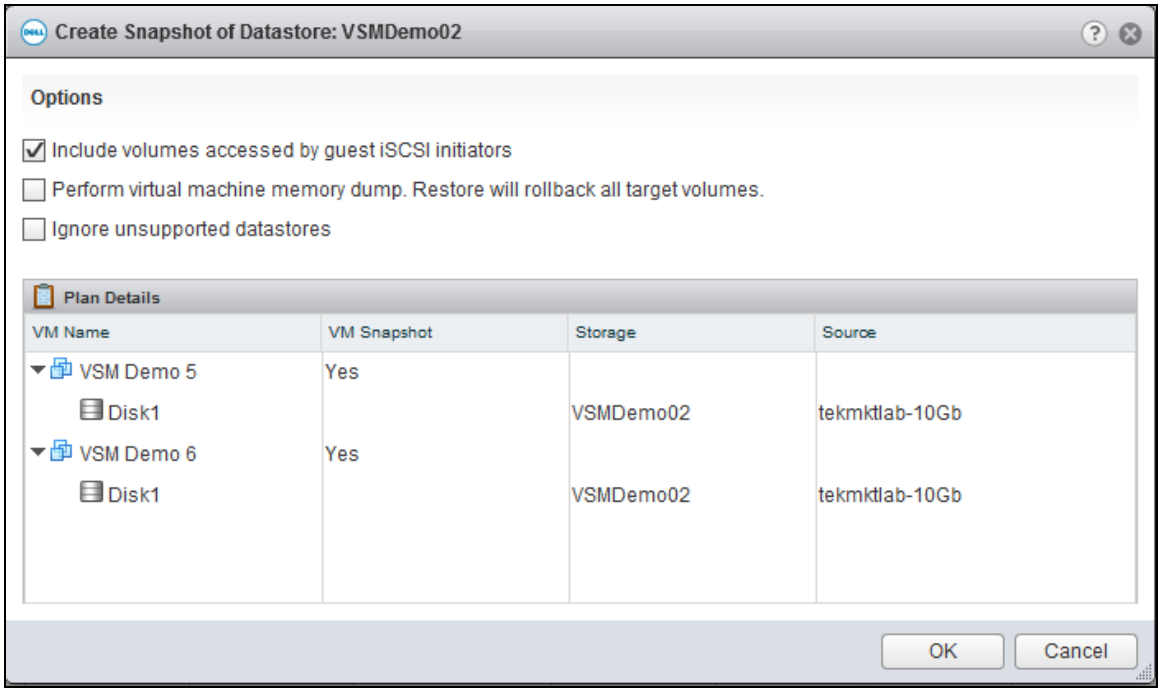

The snapshot options displayed are optional parameters that can be selected for the snapshot. These options apply to all VMs included in the Snapshot.

- **Include volumes accessed by guest iSCSI initiators:** This option requires the VMs to be powered on and the VMware tools to be installed. If these conditions are met, VSM will query the tools and any connected PS Series iSCSI initiated volumes and include them in the snapshot. These volumes must reside on a group that is also managed by the VSM.
- **Perform virtual machine memory dump:** This option requires the VMs to be powered on and the VMware tools to be installed. As part of the VMware initiated snapshot, the memory of the virtual machine is written to a disk.

Note: The virtual machine is stunned during the memory commit process. Depending on the size of memory and activity, the time it takes to stun the VM could pose a problem for applications and access. This problem is especially true if the VM is only being captured a few times a day, making the memory state almost useless. Consider this process and potential impact during the creation of snapshots with virtual machine memory dump option enabled.

- **Ignore unsupported datastores:** Choose this option to continue the snapshot operation regardless of any unsupported datastores. The job history log will indicate which VMs could potentially be affected. This is important because a VM that spans between supported and unsupported datastores would result in that VM becoming non-recoverable.
- **Plan Details:** The Plan Details pane lists all of the virtual machines that will be affected by the snapshot. Information such as the group and volume location as well as any discovered problems are listed here.

2. Make a selection and click **OK** to create the snapshot.

During the snapshot process, each VM is placed into VMware snapshot mode, quiescing the virtual machine (if VMware Tools are installed). Once the VM snapshots are created, VSM coordinates PS Series snapshots for each of the included PS Series volumes. When the PS Series snapshots are completed, VSM deletes the VMware snapshots associated with the snapshot. This does not delete existing VMware snapshots on the VMs, just the ones created for this snapshot.

3. The Job Details log lists all of the steps taken and the results.

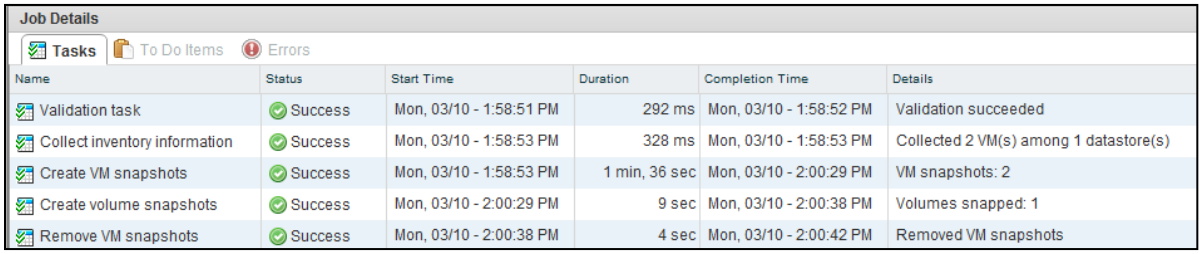

Name	Status	Start Time	Duration	Completion Time	Details
Validation task	Success	Mon, 03/10 - 1:58:51 PM	292 ms	Mon, 03/10 - 1:58:52 PM	Validation succeeded
Collect inventory information	Success	Mon, 03/10 - 1:58:53 PM	328 ms	Mon, 03/10 - 1:58:53 PM	Collected 2 VM(s) among 1 datastore(s)
Create VM snapshots	Success	Mon, 03/10 - 1:58:53 PM	1 min, 36 sec	Mon, 03/10 - 2:00:29 PM	VM snapshots: 2
Create volume snapshots	Success	Mon, 03/10 - 2:00:29 PM	9 sec	Mon, 03/10 - 2:00:38 PM	Volumes snapped: 1
Remove VM snapshots	Success	Mon, 03/10 - 2:00:38 PM	4 sec	Mon, 03/10 - 2:00:42 PM	Removed VM snapshots

Inside the Group Manager GUI, the associated PS Series volumes have a new snapshot created with the description: Created by Auto-Snapshot Manager/VMware Edition.

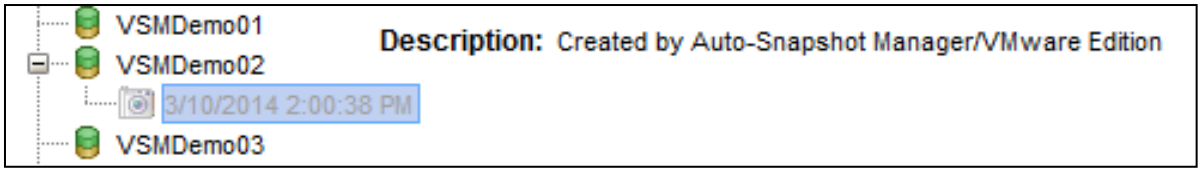

6.2 Scalability with folders and datastores

As virtual environments grow, it becomes increasingly important to be able to protect these environments. However, protecting these growing and changing environments can also be a challenge. VSM enables protecting folders of virtual machines and folders of datastores to allow scaling and adding protected objects without constantly having to adjust protection schemes. By utilizing the folder structure in vCenter Server to organize the VMs based on administrative roles or protection groups, administrators can select an entire folder of VMs or datastores and create a snapshot or snapshot schedule. VSM queries to see which VMs are in the folder, which PS Series volumes the VMs reside on, and then take a snapshot of the entire set. This keeps web server farms consistent or file servers coordinated in their protection.

This process also allows VMs to migrate from one datastore volume to another by either Storage vMotion or Migration; allows VMs to retain their protection strategy as it is assigned at the folder level; and includes multiple datastores.

These VM folders, datastore folders, and even datastore clusters can be selected as the object of a snapshot. More importantly, protection schemes can be scheduled around them.

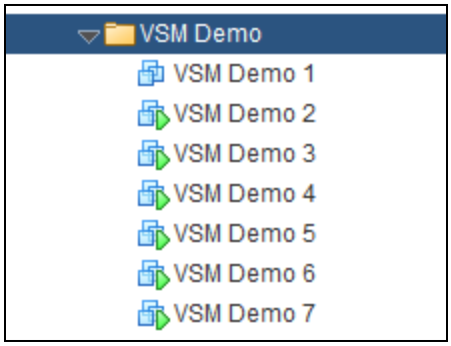

Figure 11 Example of folders in vCenter for protection

6.3 Automating protection with schedules

Individual snapshots are useful for one-off situations, such as testing a new patch or software build, but the real power from VSM comes from the built-in scheduling function. This provides a layer of protection that allows VMs to meet a better SLA for recoverability. Everything that can have a snapshot taken can also have a schedule created to automate the process. VMs, folders, datastores, datastore folders, and datastore clusters can be scheduled for snapshots.

Schedules allow tiering of protection levels for VMs. The administrator can have different schedules for different folders or different datastores, depending on the needs of those VMs. When a new VM is created, it can fall under a certain tier of protection and the administrator does not have to adjust the schedule since it inherits the protection scheme of the folder or datastore where it is located.

Creating a snapshot schedule is done in the same way that a standard snapshot is created. Two methods for completing this task are:

- Right click an object and then click **All Dell VSM Actions** > **Snapshot Schedules**.
- Click on the object, click the **Manage** tab, and then click the **Dell VSM** tab. Under **Snapshot Schedules**, click the green + symbol to add a new schedule and launch the **Add Snapshot Schedule** wizard.

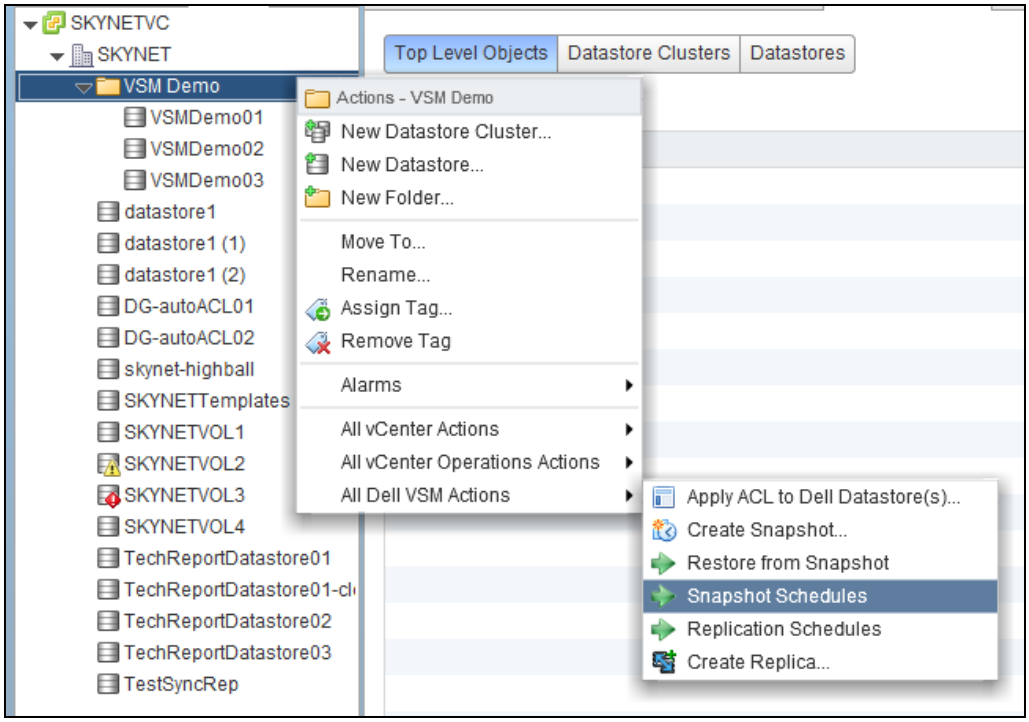

Either method for creating a snapshot schedule displays existing schedules for that particular object.

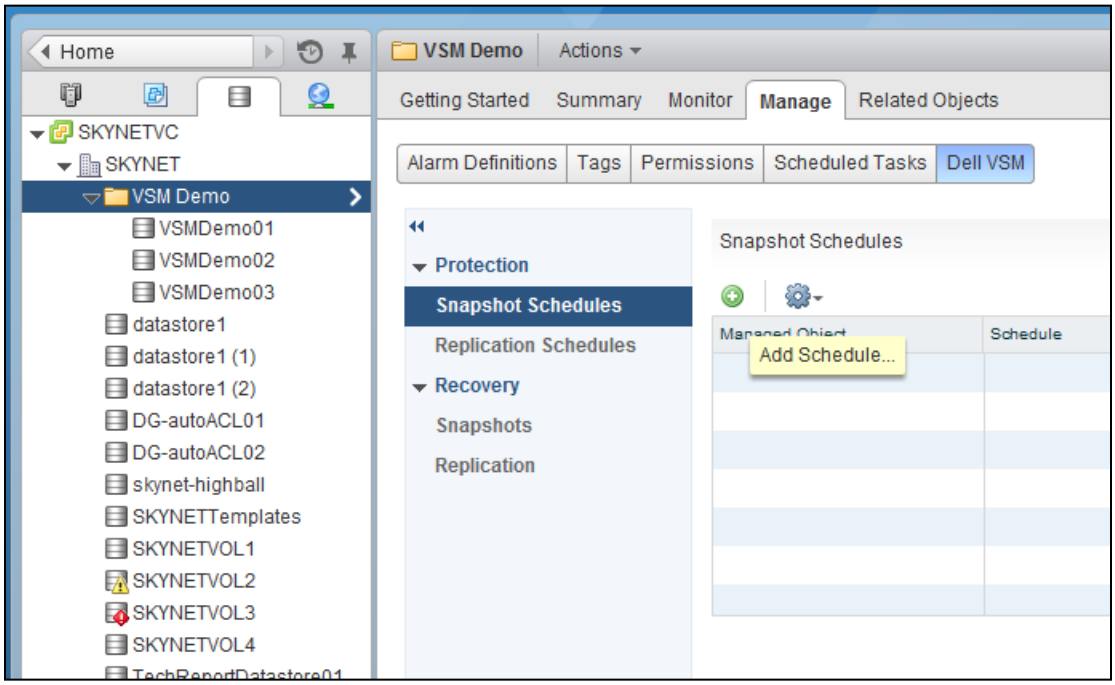

6.3.1 Adding a snapshot schedule

In this example, protection is established for the datastore folder, VSM Demo, that contains the three datastores (VSMDemo01, VSMDemo02, and VSMDemo03). A four hour schedule that keeps the past 8 copies is created, and it runs from 6 a.m. to 6 p.m. This means a snapshot is created at 6 a.m., 10 a.m., 2 p.m., and 6 p.m. and then kept for two days. The schedule is repeated Monday through Friday.

With the power of schedules, multiple layers of point-in-time protection for various objects can be created. The great thing about using the schedule on the datastore folder is that any new VM provisioned to those datastores automatically inherits the protection from the schedule without administrator intervention. In the advanced options, administrators get the same options for a standard snapshot.

1. Open the **Add Snapshot Schedule** wizard and give the snapshot schedule a meaningful name.

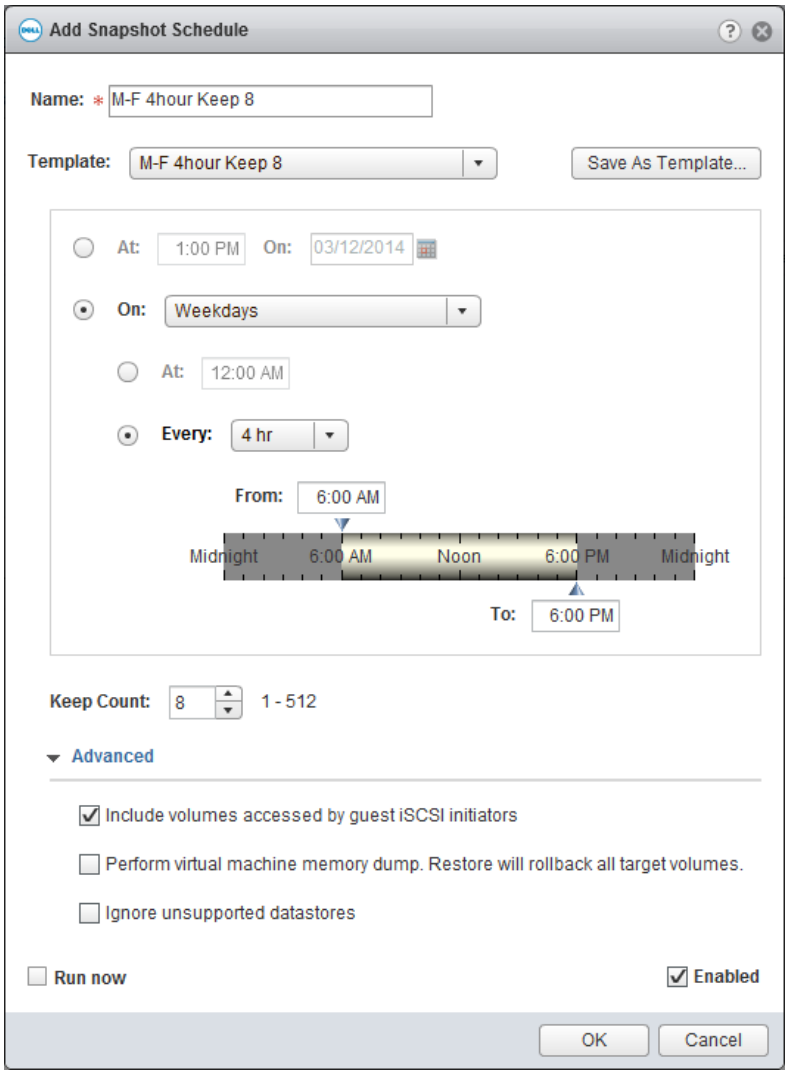

VSM automatically populates the **Name** field with the object description, but you can change this if you plan on applying it to other objects.

2. To apply the same schedule to mulitple types of objects, click **Save As Template...**.

   Templates are a useful tool for managing different types of schedules. VSM comes with a few example templates: Business Hours every hour, Weekly Snapshot, and Gold 2hr. These can be modified or used as examples for creating new protection schemes.

3. Create a schedule that fits your business needs, and then click **OK**.

   The schedule is placed in the list under **Snapshot Schedules**. It is immediatly ready to run and can be enabled or disabled.

4. At this time, more schedules can be applied to this object. In addition, the number of times the snapshot schedule has been run and the objects inside it can be displayed.

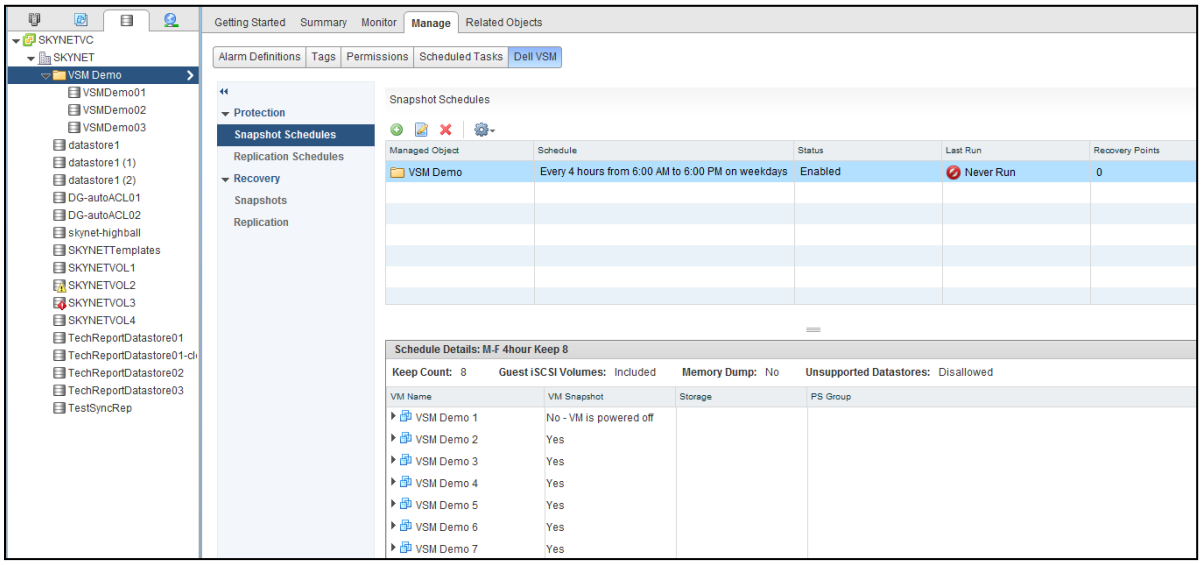

Note: With PS Series firmware version 6.x, snapshot borrowing on the volume can be enabled to allow borrowing unused snapshot space from another volume or from the free pool space. This allows the retention policy to be met in the event that the volume does not have sufficient snapshot reserve. For more information on snapshot borrowing, refer to *EqualLogic PS Series Architecture: Snapshot Space Borrowing Overview*.

6.3.2 Overlapping datastore schedules

It is important to have an understanding of what VSM is doing in the background during these schedules. Otherwise, overlapping datastore schedules could occur, preventing the data protection scheme from being achieved. Every object that has a snapshot is first put into VMware snapshot mode and then the underlying datastore volume on the PS Series SAN is snapped.

In this scenario, imagine that there is a folder with a VM in a two-hour reoccurring snapshot. Elsewhere in the cluster, another folder with a VM is in a six-hour reoccurring snapshot. The second VM resides on the same datastore volume as the first VM. This causes the PS Series SAN to create multiple snapshots of the same volume, but snapshots of the first VM only happen during its schedule. While the second snapshot is occurring on the same datastore volume, VSM is not placing the first VM in VMware snapshot mode. Therefore, the snapshot of the second folder is not usable as a consistent restore point for the VM in the first folder. The solution to this problem is to either place both VM folders in a higher-level folder, or move the folder of VMs to a different or new datastore. Because of this, proper VM placement for protection strategies is important.

When looking at objects in the **Snapshot Schedules** pane, if one is a part of a higher-level snapshot schedule, an alert is displayed.

6.4 Managing and monitoring snapshots

VSM includes a variety of tools to use for managing the snapshots. The Recovery section lists the snapshot schedules and all of the snapshots of an object. Snapshots can be deleted individually or as a whole. Objects are also displayed as part of a particular snapshot, which is useful for recovery purposes.

Another benefit to VSM is the ability to see the recently performed tasks. Completed snapshots and scheduled operations are displayed, along with any errors. For detailed information, click **Jobs** in the VSM window. Select a snapshot to list all of the tasks associated with running that job.

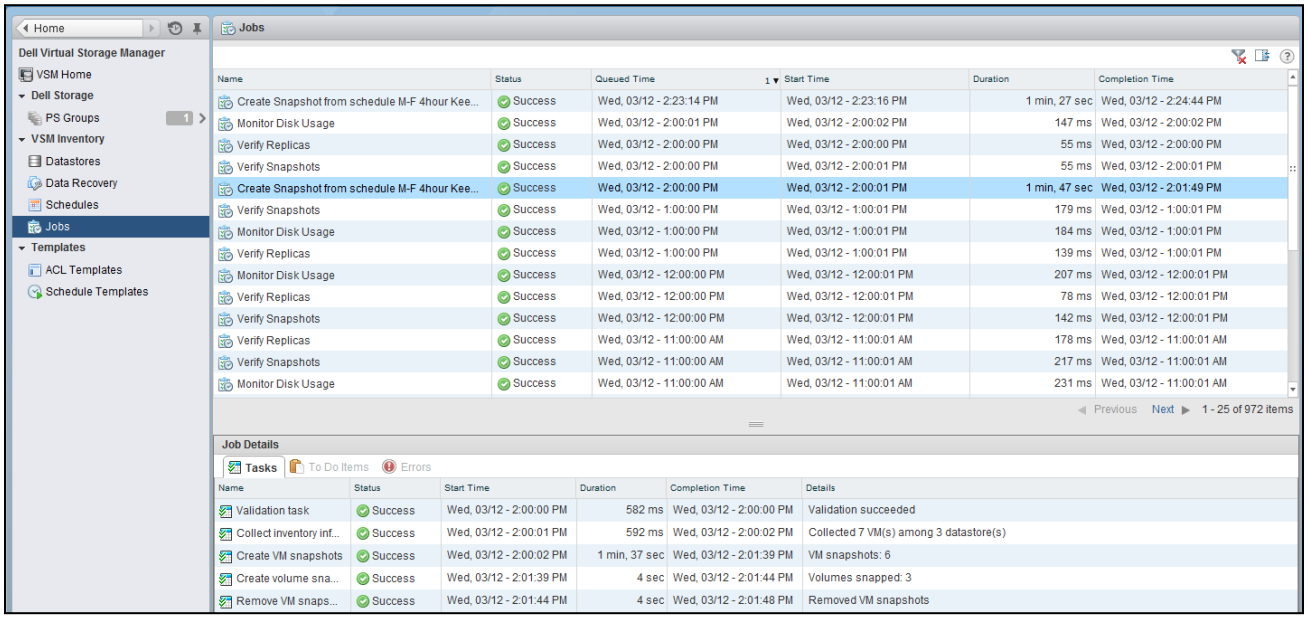

Figure 12 VSM task list

6.5 Recovering with snapshots

There are many reasons for recovering virtual machines: A bad patch or software build, corrupt file or virtual machine, or even a file that was deleted by accident. Creating snapshots on a standard schedule adds time-specific recovery

points of the virtual environment to a traditional backup schedule in case data needs to be restored. By utilizing the snapshots in addition to the traditional backup schemes, administrators gain a shorter recovery time objective. The act of deploying a new virtual machine, patching it, installing the applications, backing up the agent, and then recovering data results in the loss of hours or even days of work. Instead, snapshots can be utilized to rapidly roll a virtual machine back to a good point in time and work can continue with minimal disruption.

VSM has two different ways to start the process of recovery: From the **Data Recovery** option in the main VSM menu or from the object itself.

6.5.1 Data Recovery menu

The Data Recovery option is useful for viewing all of the available snapshots or replicas in one place. This screen also provides the ability to search for an object that may not exist in the environment but exists in a snapshot.

1. From the main VSM menu, click **Data Recovery**.
2. In the main pane, search for an object or search for all of a certain type of object.
3. Click **Search** to find the object to recover.

   It does not matter if the object has a single snapshot or dozens, a single object is displayed.

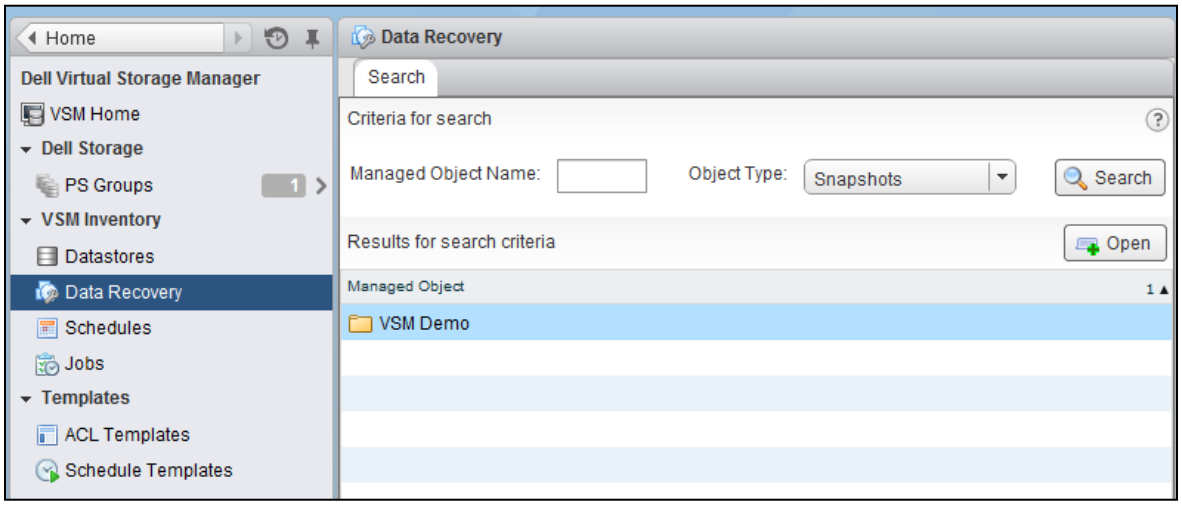

4. Double click an object to open a new tab in the **Data Recovery** pane.

   This will list all of the snapshots for that particular object.

5. Select each snapshot in the lower pane to list all of the virtual machines, volumes, and groups that are part of the snapshot.

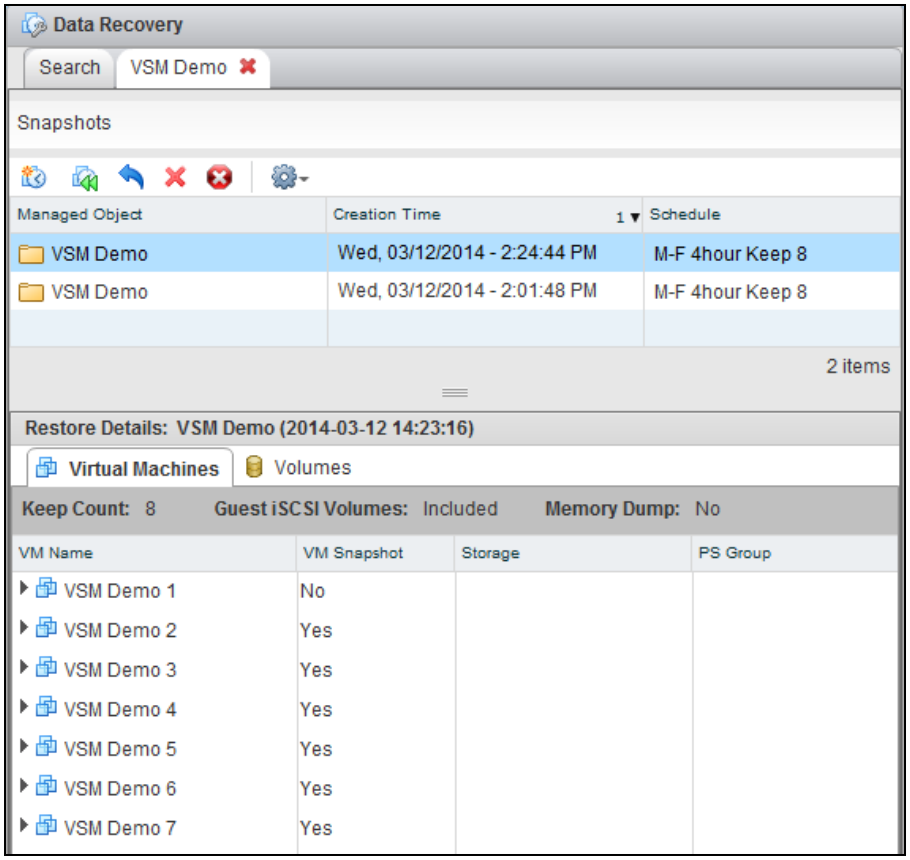

6. Choose an action to perform on the selected snapshot.

Icon	Action	Description
	Create Snapshot	An additional way to create a one-time snapshot on this object.
	Selective Restore	Use VSM to restore an individual VM or just a few VMs that are contained inside the snapshot without impacting other VMs.
	Rollback Restore	Revert everything in the snapshot to the point in time that the snapshot was taken. This affects every VM and every volume in the snapshot.
	Delete	Delete the highlighted snapshot.
	Delete All Snapshots	Delete all of the snapshots for this object.

In both cases of a restore, the VMware snapshot is reverted and deleted for all of the VMs affected to bring the VM back to the exact state it was in when the snapshot was created.

6.5.2 Selective restore

Note: This process takes longer than a restore by rollback but it does not impact other VMs on the datastore.

1. To perform a selective restore of any number of VMs in the snapshot, select a snapshot timestamp and click the selective restore icon in the menu bar.
2. **Inventory**: Check the VMs to recover. The **Current Status** displays whether the VM has been deleted.
3. Click **Next** to continue.

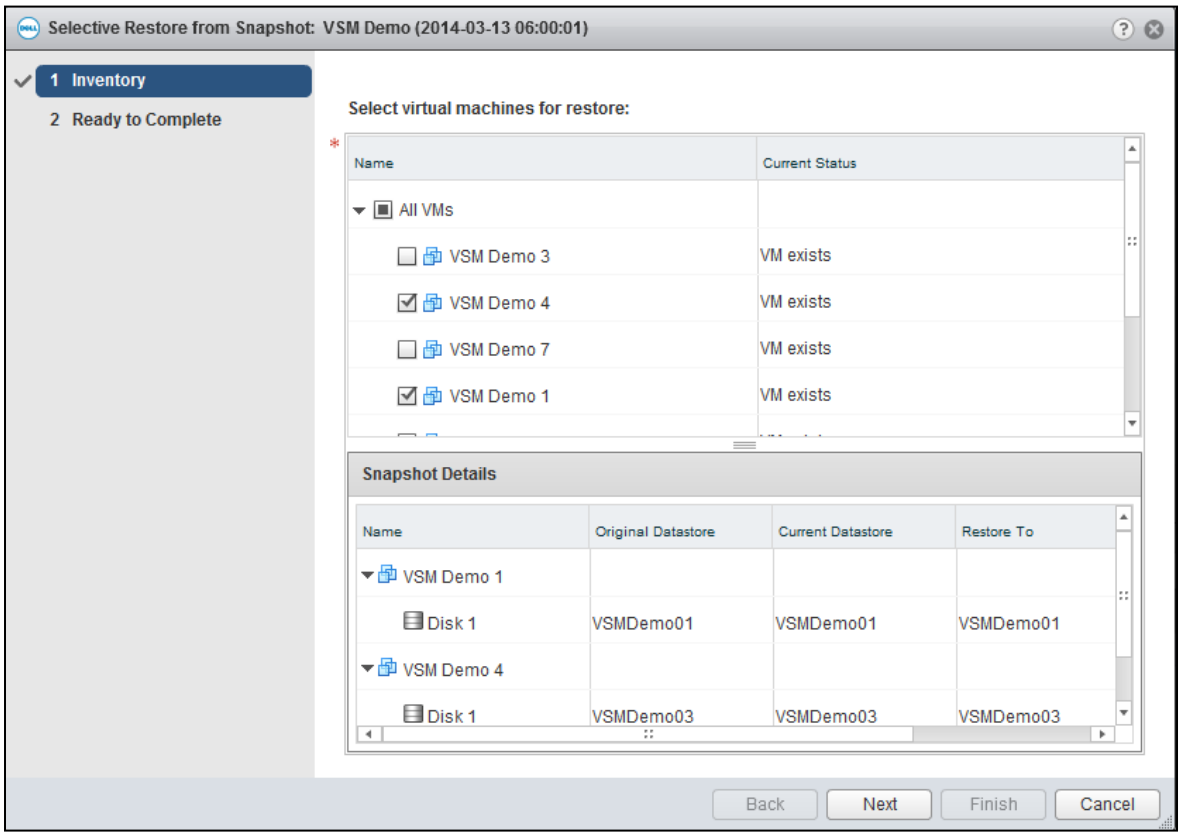

4. **Ready to Complete:** Verify the VMs that are being restored and correct any detected problems.
5. Click **Finish** to start the restore process.

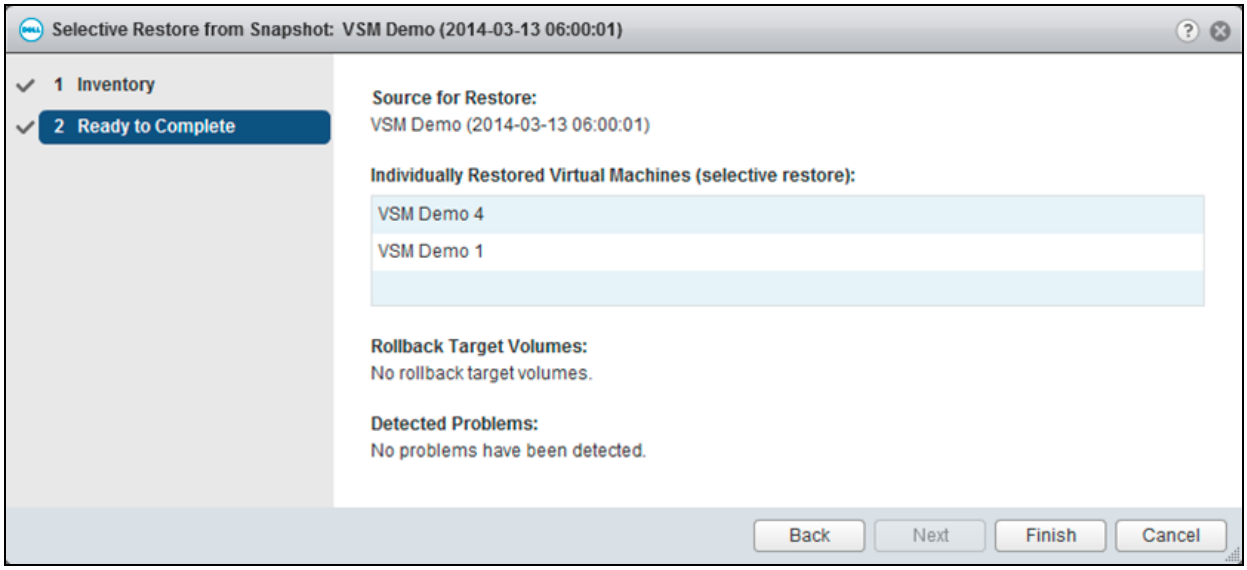

Monitor the restore task in the VSM **Job Details** pane. It may be necessary to occasionally refresh the vCenter web UI.

VSM performs the following steps:

1. Powers off the VMs that are being restored.
2. Creates clones of the datastore volumes named VSM-temp-*****.
3. Rescans the VMware ESX® cluster and registers the cloned volumes.
4. Deletes the VMs that are being restored.
5. Copies the VMs from the clone volume to the original datastore.
6. Registers and reverts the VM to the snapshot state.
7. Cleans up the clones and environment.

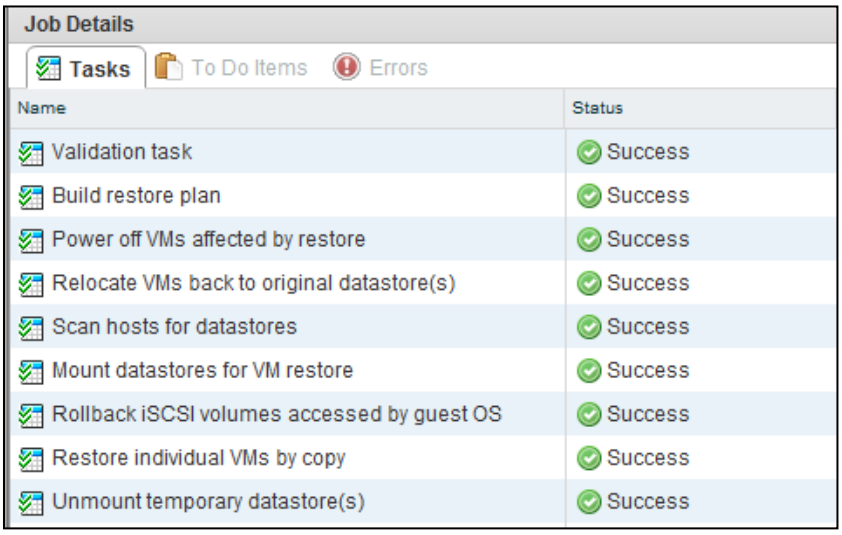

6.5.3 Rollback restore

Note: This method rolls back the entire datastore, and affects all VMs on the datastore including new VMs that might not be part of an older snapshot. VSM provides a warning if these impacts exist.

When all of the information in a snapshot needs to be rolled back to the point when it was created, use rollback restore. This reverts every single object in the snapshot including every VM and every volume.

1. Select the point in time to recover from and click the **Rollback Restore** icon in the menu bar.
2. **Inventory:** After the restore is complete, the Job Details list additional user intervention needed as well as any information or warnings.

6.6 Creating clones from snapshots

Creating clones in any environment is useful for a number of reasons. Clones allow quick deployment for multiple sets of virtual machines to test configurations and they can be used to create identical environments. VSM has the capability to clone running virtual machine environments and create clones from previous snapshots. This allows the administrator to bring online copies of virtual machines from a prior point in time. This can be helpful for troubleshooting or a side-by-side comparison of machines.

Another use of clones is the ability to test new software without impacting existing production machines. An administrator typically takes a snapshot of a set of virtual machines before upgrading software. If the upgrade results in an outage, or the software is incompatible, the administrator can roll back to the snapshot. This can be disruptive since the environment is unavailable during the restore. Creating brand new machines to test the software upgrade very rarely introduces the same issues that might come up with existing software builds. Another option for an administrator is bringing a clone of the virtual environment online, isolating it from the production environment, and then running the upgrades and testing the cloned environment. This way, if anything negative occurs, the production environment is never impacted.

Clones can also be used to create identical environments for testing and development. When combined with the PS Series Thin Clone feature, clones result in significant space savings. By taking a clone of several virtual machines residing on a datastore volume, the volume can be converted into a Thin Clone template and several space-efficient copies of these VMs can be spun off and given to various developers and test environments. For information on leveraging Thin Clones in your environment, refer to the document, *Dell EqualLogic PS Series Template Volumes and Thin Clones: How and When to Use Them*.

It is very important to note that no matter what the reason is for utilizing clones, they are an exact match of the existing virtual machine. This means that the hostname, IP address, and application namespace are identical. Therefore, it is vital that whenever a cloned virtual environment is brought online, that it is segmented from the production environment to avoid conflict. This can be done with isolated virtual switches, networking changes, or other methods.

Note: Virtual machines that have data drives on multiple datastore volumes are supported by snapshot creation. However, during a clone operation, these virtual machines that have the additional data drives still point to the original volumes. This causes conflicts without some manual configuration steps. As a best practice, keep all of the data of the VMs that need to be cloned on a single datastore volume to avoid any potential issues.

To create a clone from a snapshot:

1. Open the VSM GUI and go to the object snapshots.

   You can also search in the Data Recovery area of VSM and double click the object. Select a snapshot to clone, right click it, and select **Clone from Snapshot**.

   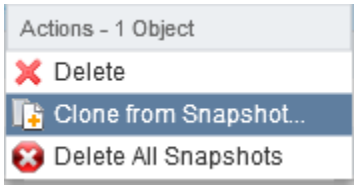

   Hardware Resource: VSM needs to create a name suffix for every volume that it clones for the restore. This is needed so that no two datastores can have the same name and provides a way to differentiate between the original datastore and the cloned datastore.

2. By default, VSM appends *-clone- date/time of Snapshot*. A description can be entered and the deployment hardware resource selected.

   This enables cloning from a production cluster and mounting to a test and development cluster in the same vSphere farm.

3. Select the hardware resource and click **Next**.

Note: This clones the entire Snapshot and consumes space for the datastore volumes that are part of the snapshot.

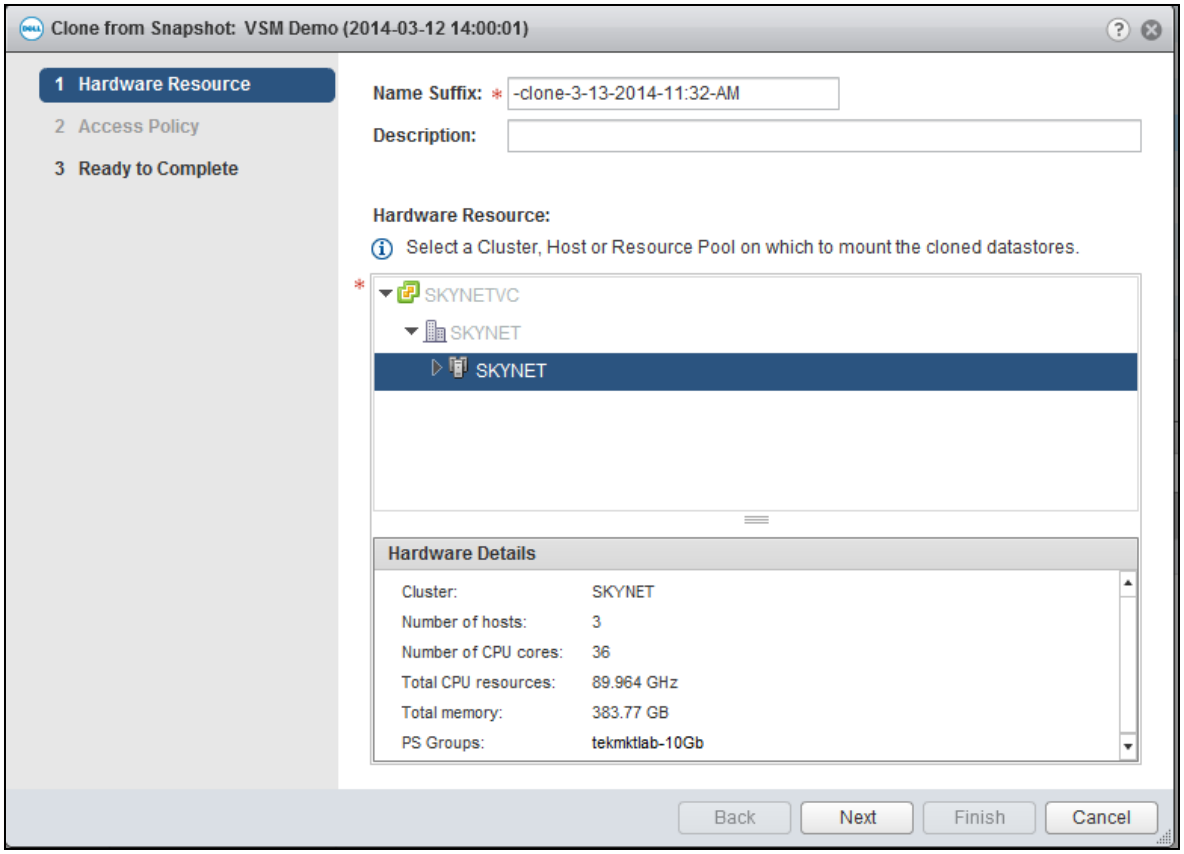

Access Policy: The next step in the process is to choose the access policy for the new cloned volumes.

4. Select **Auto-generate** to create ACLs from the hosts that are in the selected cluster. You can choose to specify a new or existing ACL policy.

   For information about ACL policies in VSM see section 2.5.2.

5. Complete your selections and click **Next**.

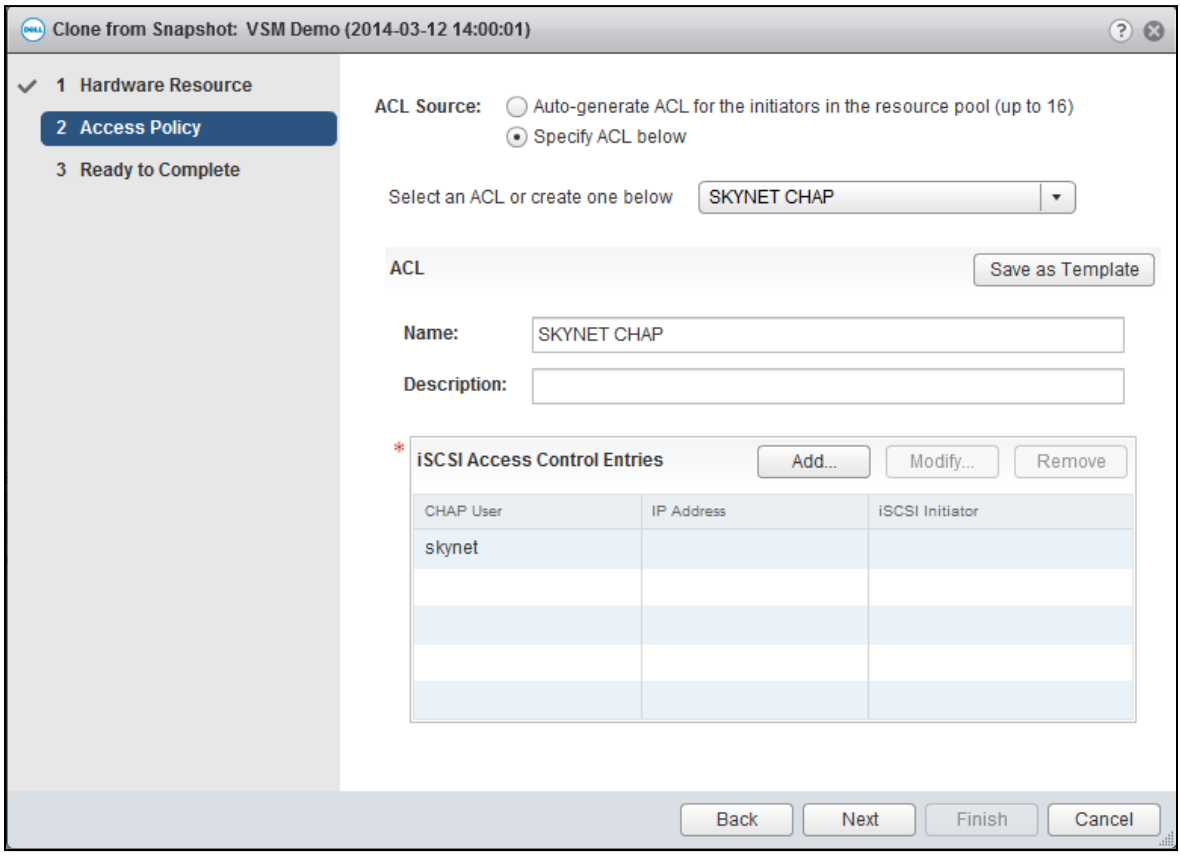

6. **Ready to Complete**: Verify the options are correct and click **Finish**.

During this time, VSM coordinates with the PS Series SAN and creates volume clones of all the snapshots that are part of the VSM snapshot. Once the volumes are cloned, VSM tells vCenter to rescan and bring these new cloned volumes into the environment in a datastore folder.

The new cloned volumes are displayed inside the PS Series Group Manager GUI as well as the new datastore volumes inside vCenter.

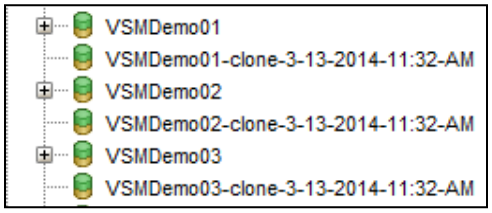

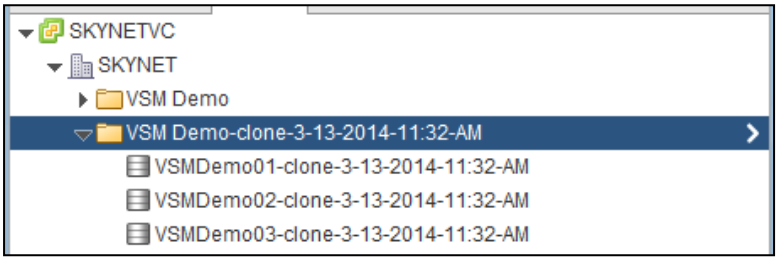

VSM will not register or power on the VMs so that the original VM environment is protected. Once the cloning is complete, you can browse the datastore, register the VMs, isolate them and power them on.

Note: VMs contain VMware snapshots and possibly memory state from the snapshot process. During the process, VSM does not revert or delete these snapshots to protect the original VM. These snapshots need to be managed manually once the VMs are isolated.

6.7 Advanced cloning in selective data recovery

There are times when data restoration needs to be more granular than at the level of the individual datastore or individual virtual machine. The idea behind selective data recovery is creating clones, bringing the information online and then attaching the data drive of the VM from the snapshot back to the original VM. Using clones for this is preferred over bringing a snapshot online. By taking a clone of the snapshot, the original PS Series snapshot data is not modified. If the original PS Series snapshot is online, the data integrity for recovery could be changed (or deleted) by accident.

There are multiple options for selective data recovery, but they all revolve around mounting a point-in-time version of the data disk to a VM (usually the original). In order for this to work, the VM OS must support the ability to hot add data disks. The other option is to have a temporary or standby recovery VM available that can have data drives mounted to it and then used to find the files to recover and copy them back to the original location.

The process is similar to creating a clone but with some additional steps.

1. Find the snapshot containing the data that needs to be recovered. Right click and choose **Clone from Snapshot**.

Note: Even if you are trying to restore just one file, the Smart Clone operation could cause multiple datastores to be cloned and mounted. These will need to be deleted using Datastore Manager when the process is finished.

2. Follow the same process as creating a clone from a snapshot. Once the cloned datastores are scanned and found, the data drive can be attached to the VM or recover the VM for file restoration.
3. Right click the VM receiving the recovery data disk and click **Edit Settings**.
4. Click the new device drop-down menu, choose **New Existing Hard Disk**, and click **Add**.

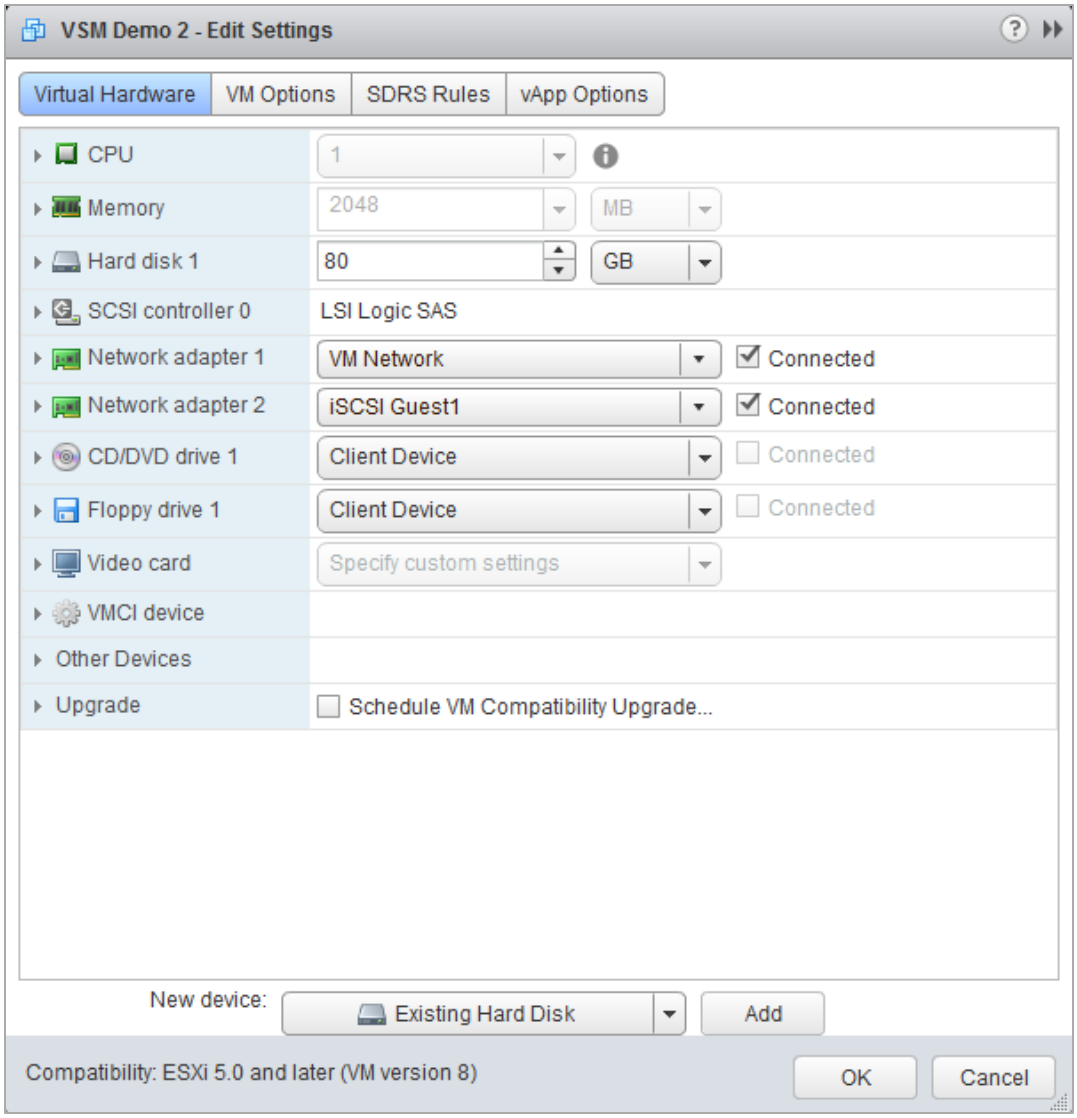
VSM Demo 2 - Edit Settings
Virtual Hardware
VM Options
SDRS Rules
vApp Options
CPU
1
Memory
2048
MB
Hard disk 1
80
GB
SCSI controller 0
LSI Logic SAS
Network adapter 1
VM Network
Connected
Network adapter 2
iSCSI Guest1
Connected
CD/DVD drive 1
Client Device
Connected
Floppy drive 1
Client Device
Connected
Video card
Specify custom settings
VMCI device
Other Devices
Upgrade
Schedule VM Compatibility Upgrade...
New device:
Existing Hard Disk
Add
Compatibility: ESXi 5.0 and later (VM version 8)
OK
Cancel

The next step is to add an existing hard disk to the VM that points to the cloned datastore.

5. Find the VMDK file with the data in the folder to recover and click **OK**. Click **OK** again to commit the change.

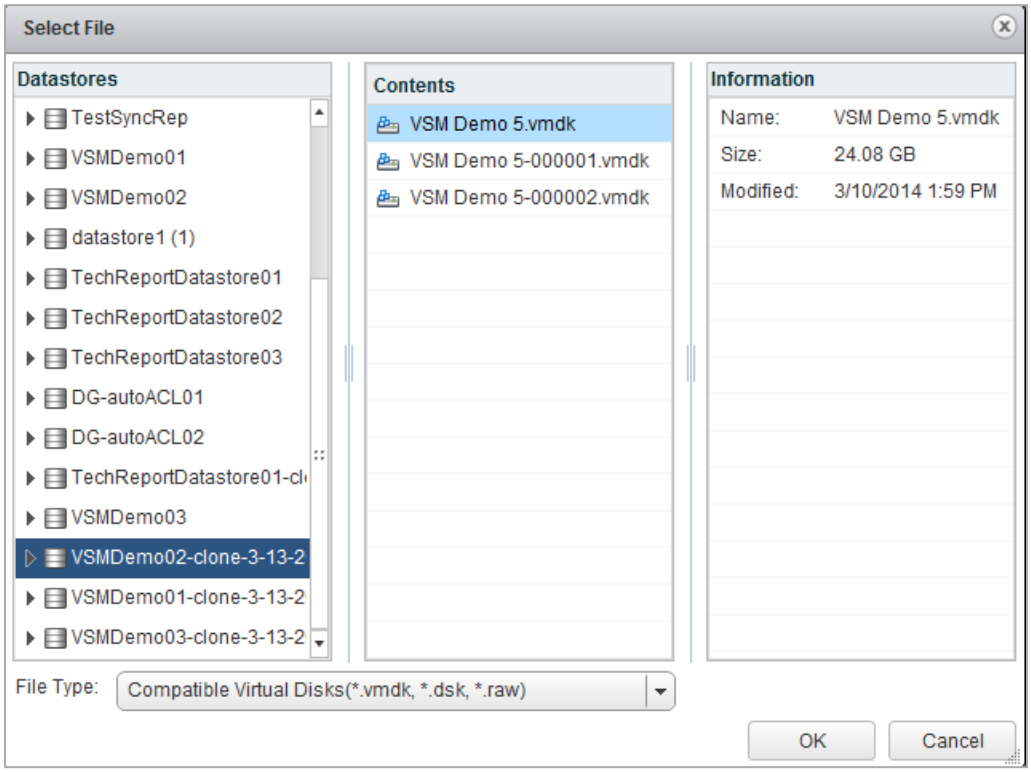

6. Once the recovery data disk has been mounted to the VM, utilize the native OS tools to recover the data.

   For example, in Microsoft Windows Server® 2008 R2, open disk management first.

7. Place the disk online to assign a drive letter to it.
8. Browse the assigned drive letter to see that it is the version of the original data drive from the point in time that the snapshot was taken.
9. Copy or move the files or data that need to be recovered inside the VM.

   If using a recovery VM, move the files back to the original VM.

Clean up the environment once the files have been recovered.

1. Remove the added hard disk by editing the VM settings and removing it.

   If the VM or recovery VM does not support hot add/remove, the VM needs to be powered off to remove the hard disk.

2. Use the VSM Datastore Manager to delete all of the clone volumes that were created during the recovery process.

These clones will be listed as a completed task for the clone. This will ensure proper removal of the iSCSI targets and deletion of PS Series cloned volumes.

6.8 Multilayered data protection approach and data placement

PS Series SANs are integrated with vSphere through the Virtual Storage Manager. With features such as snapshots, they provide an additional layer of protection by offering hypervisor-aware snapshots for virtual machines. These tools and techniques are designed to enhance existing data protection or business-continuance strategies and work in conjunction with other solutions. Traditional backup techniques, as well as the PS Series Auto-Snapshot Manager/Microsoft Edition inside Windows VMs, can be used to protect Microsoft SharePoint®, SQL, and Exchange data or Auto-Snapshot Manager/Linux Edition used to protect Linux data drives.

Leveraging all of these tools together requires a new approach to data protection and data placement. The snapshot within the PS Series SAN is done at the volume level even if the object in vCenter is a folder or a subset of VMs. This means that to meet the SLA and RTO of a particular set of VMs, they should all reside together in the same protection scheme. During VM deployment, a location for the VM so that it can have the required protection options and service level for recovery needs to be identified. As data protection scenarios are built, it will be easier to decide where the VM belongs. Once the tiers are set up and configured by either folders or datastores, VM placement will be easier. In addition to meeting SLAs, VM placement also has an effect on local PS Series snapshot space. Whenever a VM is moved using migrate or storage vMotion, the SAN keeps track of the movement because it is identified as new writes. Leveraging Storage DRS (sDRS) or constantly moving VMs from one volume to another, could dramatically increase the amount of snapshot space consumed on the SAN to keep track of this movement.

VSM Snapshots can also be used in conjunction with a variety of the other PS Series host integration tools for more granular protection of the application data within the virtual machine.

7 Local data protection strategies with vVol datastores

Virtual Volumes significantly changes taking snapshots. Traditional VMware snapshots, its delta VMDKs, and its associated limitations have been replaced by SAN-based traditional volume snapshots. This provides a number of benefits:

- Individual virtual machine snapshots
- Instant creation of virtual machine snapshots
- Snapshots that can be kept indefinitely
- Instant restoration of a virtual machine from a snapshot

7.1 Comparing VMFS and vVol data protection

With traditional VMFS datastores, there are two snapshot options:

- Traditional VMware snapshots:
  - Use familiar vSphere GUI
  - Perform per-virtual-machine snapshots
  - Use best practices that limit the number of snapshots
  - Use best practices that limit how long snapshots should be kept
  - Can impact performance
  - May take a long time to restore
  - Offer application consistency

Note SAN-based snapshots by themselves are unaware of the target snapshot data on the volume, and therefore create crash-consistent snapshots. However, when combined with host-side tools from the PS Series Host Integration Tools (HIT) kits, like VSM, they become application aware, and can create application consistent snapshots.

- PS Series SAN snapshots with VSM:
  - Use the Dell VSM plugin to vSphere
  - Take snapshots of an entire datastore of virtual machines
  - Take snapshots which can be kept indefinitely
  - Cause no performance impact
  - Perform rollback restores quickly, but roll back all virtual machines on the volume
  - May take a long time to perform selective restores
  - Offer application consistency

Traditional VMware snapshots and VSM snapshots both have their particular use cases in addition to pros and cons. Virtual Volumes combines the best features of both into one technology:

- Use native vSphere GUI
- Provide granular, per-virtual-machine snapshots
- Instantly create and restore

- Cause no performance impact
- Offer application consistency
- Allow snapshots to be kept indefinitely

Note: A virtual machine can span a VMFS datastore and a vVol datastore, however, this configuration is not supported for snapshots.

7.2 Protection with VMware vVol snapshots

The workflow for creating a snapshot of a virtual machine stored on a vVol datastore is the same as when it is stored on a VMFS datastore. However, the snapshot creation method has changed:

1. A snapshot is initiated from vSphere
2. vSphere will optionally quiesce the guest filesystem, requiring VMware Tools to be installed within the virtual machine.
3. At this point, things diverge from the previous method:

    a. vSphere communicates with the PS Series array using the VASA Provider, prepares for snapshot, and takes snapshot API calls to the array
    b. The PS Series array creates a volume snapshot of each data virtual volume (also referred to as a VMDK virtual volume) in the virtual machine. These volume snapshots use the same technology that has previously been used for snapshotting traditional volumes.
    c. If the option to snapshot the virtual machine memory was selected, vSphere will request an additional memory virtual volume, where it will write the contents of the virtual machine RAM.

With VMware virtual volume snapshots off-loaded to the PS Series array, the creation and restoration of snapshots is almost instantaneous. As an array-based snapshot, they can be kept indefinitely, as long as sufficient space is available, and the VMware limit of the number of snapshots in a chain limit of 32 has not been exceeded.

7.2.1 Creating a snapshot

Follow this step-by-step process to create a snapshot:

1. From the virtual machine context sensitive menu, or its **Actions** dropdown menu, select **Snapshots** > **Take Snapshot**.

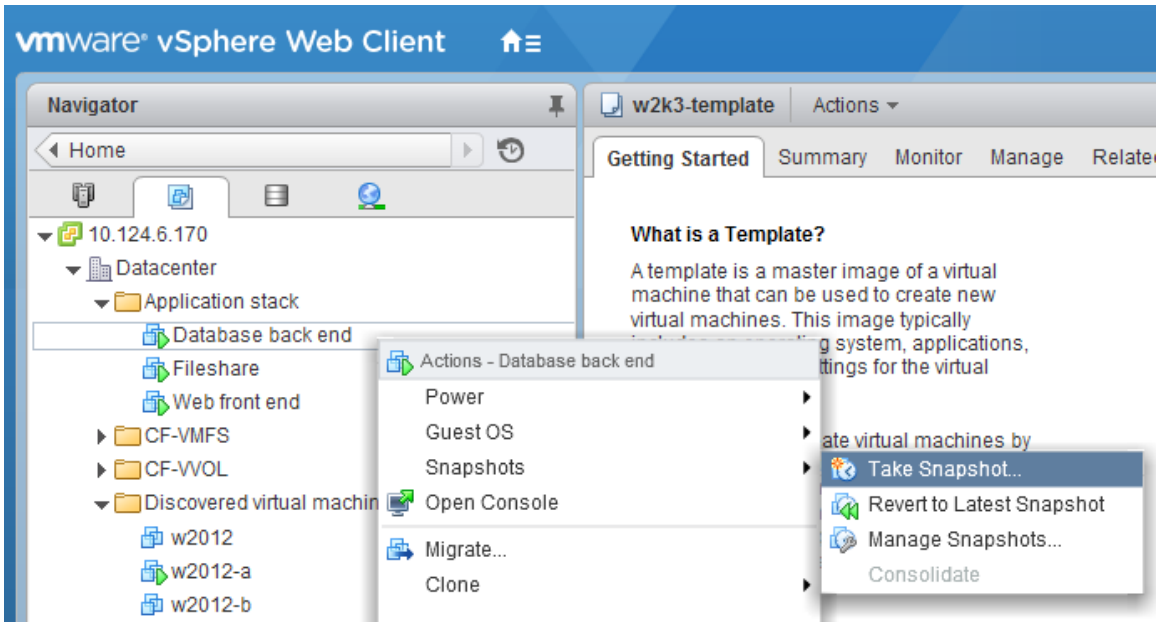

2. Provide a **Name** and **Description** for the snapshot to enable easy identification of what a particular snapshot contains.
3. Optionally, elect to include **Snapshot the virtual machine's memory** (selected by default).

   Including this is a business decision, and will be dictated by the need to have the virtual machines contents at restore time. This option includes the contents of the memory of the virtual machine with the snapshot.

Note: This can extend the time it takes to create the snapshot while the contents of the virtual machines memory are written to disk.

4. Optionally, elect to **Quiesce guest file system**. This causes the VMware Tools to invoke the Microsoft Volume Shadow Copy Service to quiesce I/O on the disk and hold them in memory while the snapshot is created.

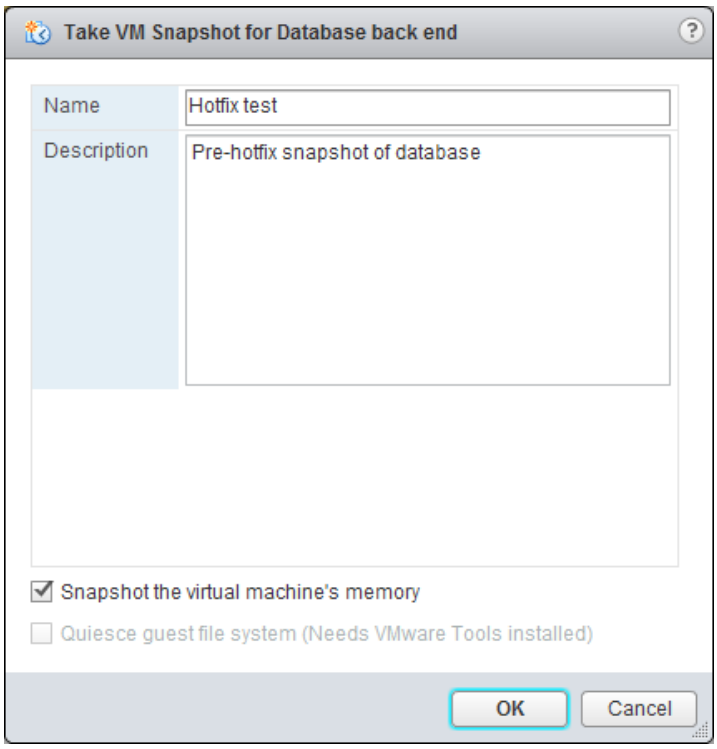

Note: A choice must be made between selecting **Snapshot the virtual machine's memory** and **Quiesce guest file system**. Generally the longer a snapshot is retained, the less valuable the virtual machine memory is.

5. Click **OK** to create the snapshot.

Remember that with Virtual Volumes, VMware will offload the creation of virtual machine snapshots to an array, where the array native snapshot capabilities can be used. This means that snapshots do not need to be deleted after a short period of time, nor does the best practice of only going two or three snapshots deep apply. Array-based snapshots have become a new tool in the vSphere admins toolbox, and can be used to add granularity of protection when used in conjunction with backup programs, or to provide a mean of rapidly restoring from high risk software changes.

7.2.2 Restoring from a snapshot

Restoring a virtual-volume-based virtual machine from an array-based snapshot is a rapid process. However, if a memory dump was included at the time of taking the snapshot, the process will take longer since this will need to be read into memory on the ESXi host. This rapid restoration process is due to the nature of array snapshots. At restore time, the process on the array is:

1. vSphere shuts down the virtual machine, and sends the restore request.
2. The restore request is received from vSphere through the VASA Provider.
3. The currently attached data virtual volumes for the virtual machine are detached from the virtual machine.
4. The virtual machine is reconfigured to point to a specified snapshot of virtual volumes.

5. vSphere is informed of the change.

Because data movement does not occur, the restoration process is just an update to the array internal database and a reconfiguration of the virtual machine. This allows the process to complete in a matter of seconds.

Two methods can be used to restore from a snapshot: the **Manage Snapshot** dialog where a particular restore can be chosen, or the **Revert to Latest Snapshot**, which reverts the virtual machine to the most recent snapshot.

1. From the virtual machine context sensitive menu, or its **Actions** dropdown menu, select **Snapshots**, and then **Manage Snapshot**.
2. Select the snapshot to restore the virtual machine from, and then click **Revert to.**

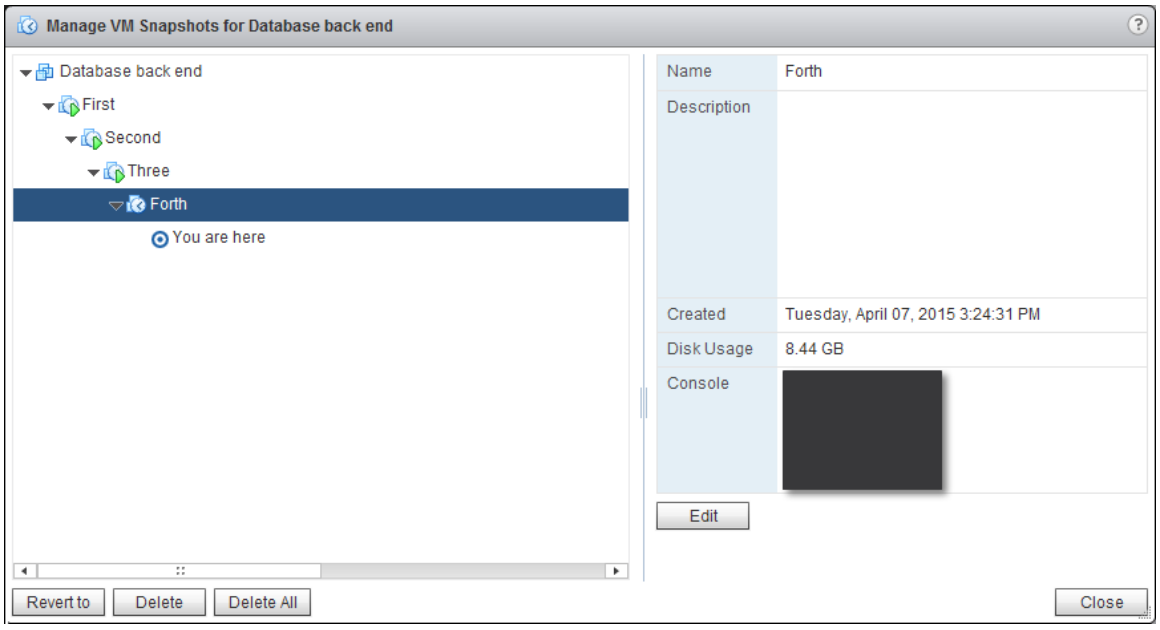

3. Click **Close** to exit the dialogue.

The restore process is initiated and because it is handled by the array, the process completes almost instantly.

Note: Reverting to the latest Snapshot follows the same process, but rather than specifying the particular snapshot to restore from, the most recent snapshot is used for the restore.

7.2.3 Automating protection with snapshot schedules

With VMware Virtual Volumes, VMware snapshots are nearly instant and no longer have the limitation of VMware snapshots on VMFS; they can be used with more frequency and in different ways. When coupled with scheduled tasks, snapshots can be used to complement existing backup strategies. Regular backups are performed as normal (for example, weekly fulls and daily incremental) but complemented with regular snapshots during the business day (such as every four hours or every hour). This provides a shorter RPO with a very short RTO.

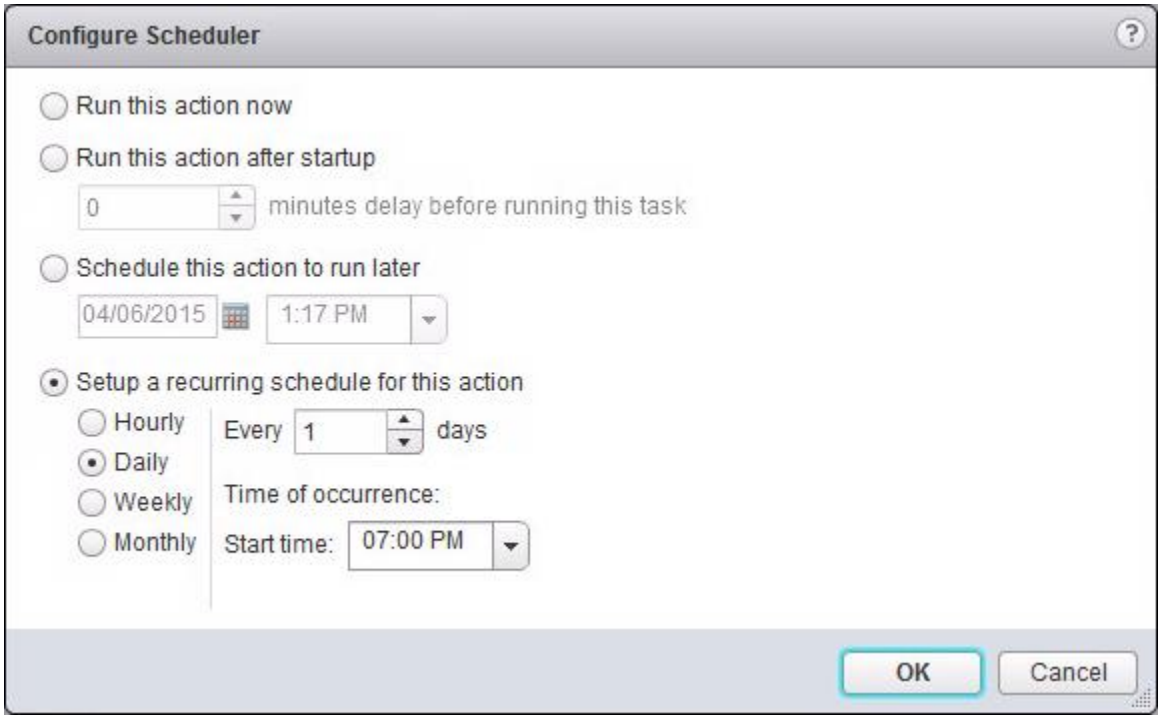

Note: While VMware snapshots can be non-disruptive, selecting **Quiesce guest file system** invokes a process to bring the virtual machine disks into a state suitable for backups. Depending on how active the data is and the application in use, this may cause disruption.

7.3 Additional vVol snapshot functionality with VSM

In a somewhat confusing circular fashion, VSM builds on the VMware snapshot functionality that also enables vSphere to perform. This allows vSphere administrators to extend snapshots to groups of virtual machines, have more granular scheduling options, and use additional restore options.

7.3.1 Protecting groups of virtual machines

Often multiple virtual machines are involved in supporting an application. For example, a web-based application may have one or more web servers, a database backend, and a file server. There are times when it will be appropriate to have a snapshot of the entire application stack at the same time. VSM builds on the VMware hardware-assisted snapshots and provides this functionality. To enable a snapshot of a group of virtual machines at the same time, they must be placed within a virtual machine folder.

Note: The VSM interface can be used to snapshot folders of virtual machines, or individual virtual machines.

1. From the virtual machine folder context sensitive menu, or its **Actions** dropdown menu, select **All Dell VSM Actions** > **Create Snapshot**.

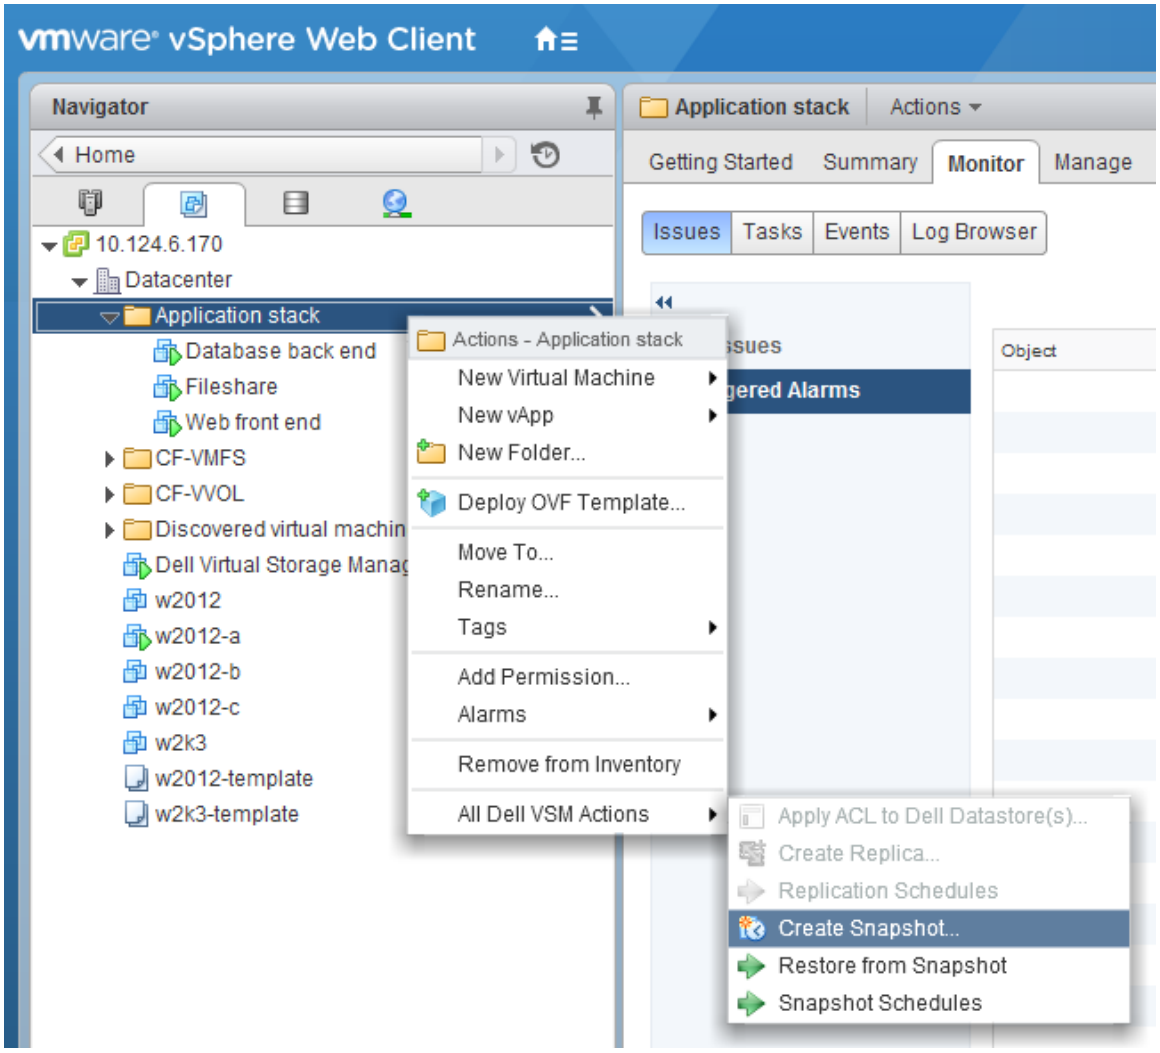

2. Similar to the native VMware snapshot option from vSphere, there are options to quiesce the filesystem and include a copy of the memory.

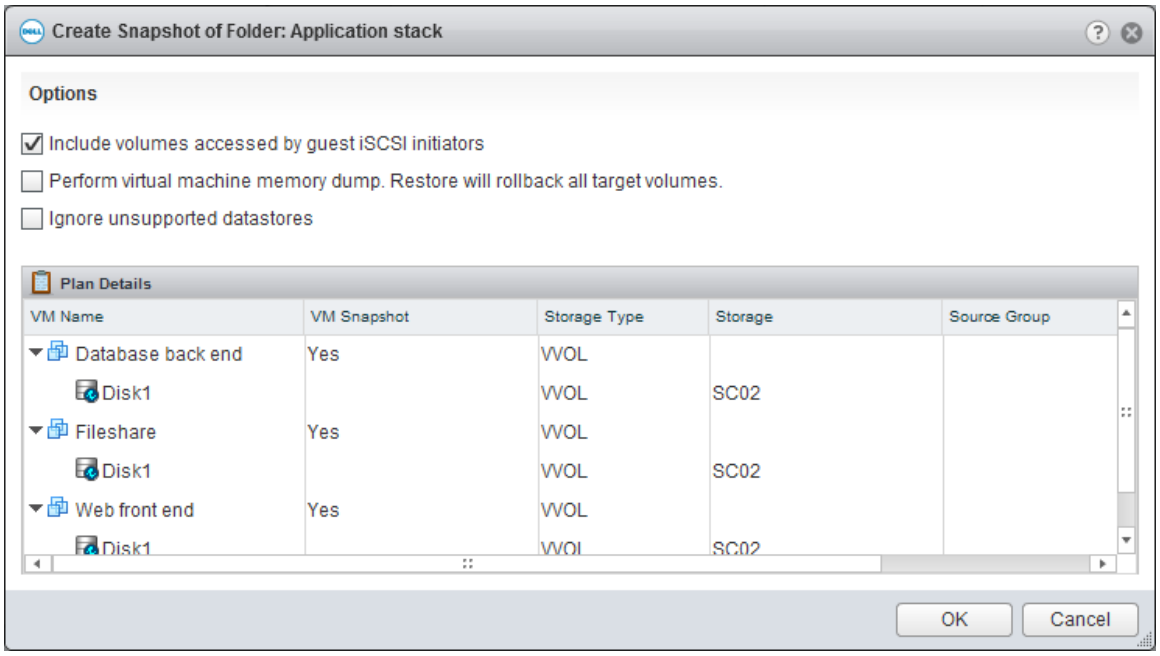

Additionally, there is an option to include iSCSI within the guest mounted iSCSI volumes. This requires the virtual machine to be powered on and the VMware Tools installed. VSM will analyze the virtual machines for iSCSI attached volumes, and snapshot them at the same time as the virtual machine.

Note: A virtual machine folder may consist of VMFS based virtual machines, and vVol-based virtual machines. However, a snapshot of a virtual machine that is spread across a traditional VMFS datastore and a vVol datastore cannot be completed using either the VMware or VMS tools and is not recommended.

3. Click **OK** to start the snapshot creation.
4. The status of the snapshot creation can be observed by clicking on the **VSM Job Status** popup in the bottom right corner of the vSphere Web Client.

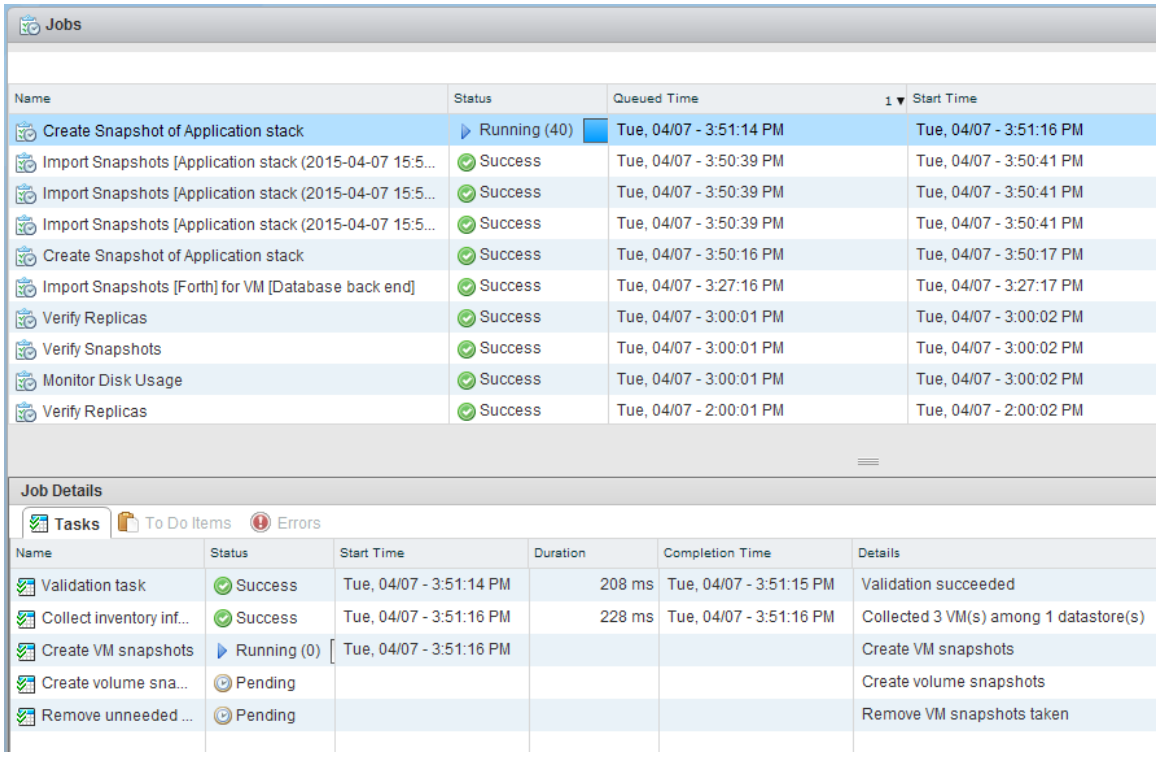

The upper half of the screenshot shown above displays the history of **Jobs** performed, with the current job (**Create Snapshot of Application stack**) highlighted. In the **Jobs Details** section of the screen, the individual tasks that make up the job and their status can be observed.

7.3.2 Automating protection with VSM snapshot schedules

Similar to scheduling snapshots of virtual machines (detailed in section 6.3), VSM extends flexible schedules to the snapshot protection of virtual volume based virtual machines and virtual machine folders of virtual machines, but not to entire datastores of virtual-volume-backed virtual machines, as can be done for VMFS datastores.

VSM snapshot schedules serve the same purpose as VMware snapshot schedules, enabling snapshots to be used to provide additional options to existing backup strategies, resulting in shorter RTO and RPO periods. The advantage that VSM schedules have over VMware schedulers are:

- Additional scheduling options; for example: hourly protection may be desired, but only during the business day, and not outside these hours or on weekends.
- Schedule templates can be created, enabling pre-defined scheduling options to be applied to other virtual machines.
- iSCSI within the guest mounted volumes can be included

Groups of virtual machines can be protected though the use of virtual machine folders. To create a snapshot schedule for a VSM:

1. From the virtual machine or virtual machine folder context sensitive menu, or its **Actions** dropdown menu, select **All Dell VSM Actions** > **Snapshot Schedules**.

   Or with the object selected, click **Dell VSM** from the **Manage** tab, click **Snapshot Schedules**, and then the **Add Schedule** icon  .

2. On the **Add Snapshot Schedule** dialogue, edit the provided name if needed (the default VSM provides a name based off the snapshot target).
3. Select the desired scheduling frequency.
4. Under the **Advanced** section, the following additional options are available:

   - **Include volumes accessed by guest iSCSI initiators**: This option requires the virtual machine to be powered on and VMware Tools installed. At the time of the snapshot, VSM will detect if an iSCSI within the guest volume is attached to the virtual machine and snapshot that volume with the virtual machine.
   - **Preform virtual machine dump:** This option will include the contents of the virtual machine memory with the snapshot.
   - **Ignore unsupported datastores:** If a virtual machine includes a virtual disk located on an unsupported datastore, such as a local datastore, VSM will not attempt to include it in the snapshot.

5. Click **OK** to assign the schedule to the virtual machine or virtual machine folder.

Note: A virtual machine or virtual machine folder can have multiple snapshot schedules associated with it. Should one or more snapshot schedules occur at the same time, they will be taken sequentially.

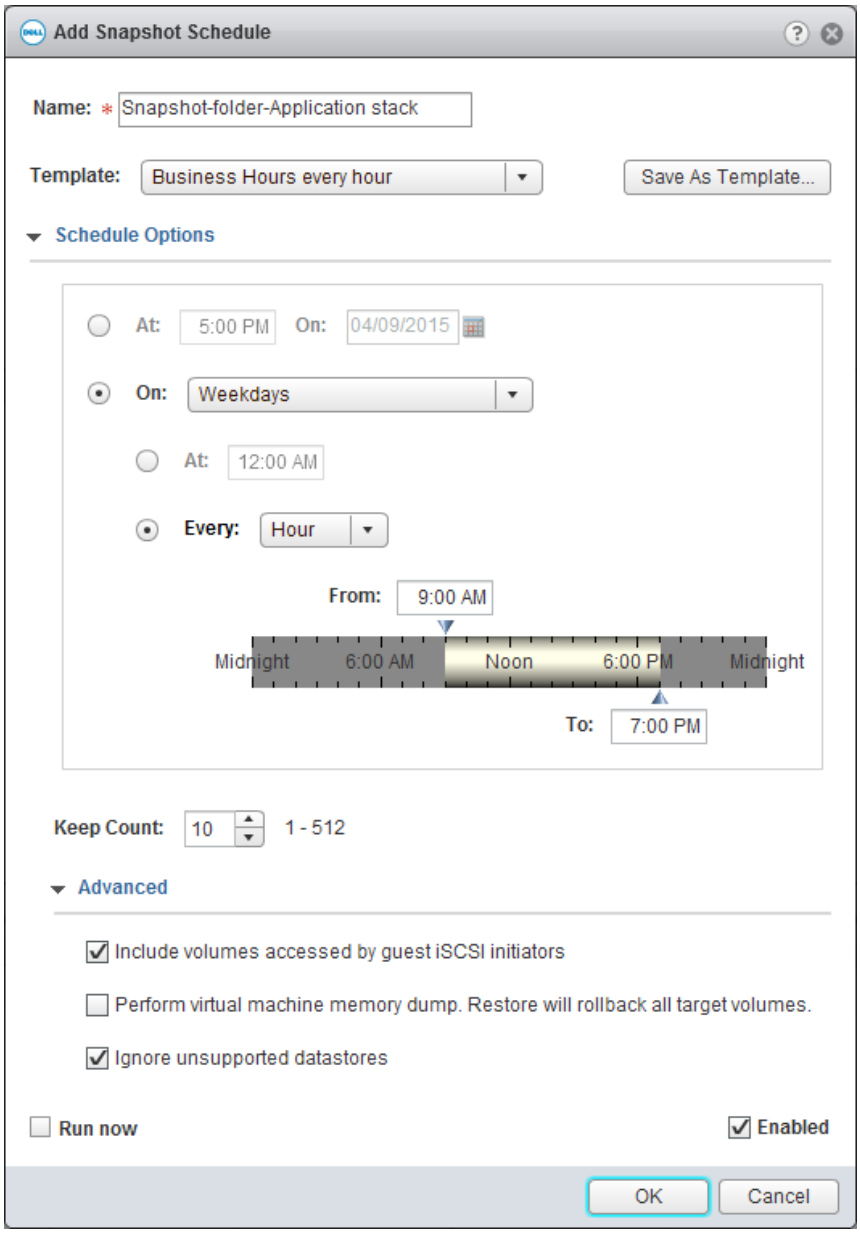

6. Alternatively, provide a **Name** to the schedule and select a previously created schedule from the template dropdown menu, and then click **OK** to continue.

7.4 Recovering virtual machines using VSM

Recovering a virtual machine may be necessary for a variety of reasons such as a bad patch or software build, file corruption, or user error. However, recovering from backups often incurs a long RTO and RPO, resulting in significant downtime. While snapshots are not a replacement for traditional backups, they can be used to complement them, and provide a shorter RTO and RPO.

With Virtual Volumes, the recovery process is significantly faster than traditional backup recovery process, and faster than recovering a virtual machine from a VMFS datastore snapshot. The short recovery time is because the traditional, lengthy part of the recovery (the movement of significant amounts of data), is replaced by unmapping the data virtual volume and then remapping it to the selected snapshot. This recovery process is handled by the array in a few minutes.

7.4.1 Recovering complete virtual machine

Regardless as to how the VSM snapshot of the virtual machine was initialized (manually, as part of a scheduled, or included within a virtual machine folder object), the recovery process is the same.

1. Initializing a recovery from a snapshot using one of the following methods:
   - From the **Dell VSM** button under a virtual machine or virtual machine folder **Manage** tab, and then select **Snapshots.**
   - From the **Dell VSM** page, select **Data Recovery** under **VSM Inventory**. Select **Snapshots** from the **Object Type** dropdown, and click the **Search** button. If there are a large number of virtual machines being protected with snapshots, filter the results by putting a suitable search term in the **Managed Object Name** textbox. Select the desired object to be recovered.

Note: An individual virtual machine can be recovered when the protected object is a virtual machine folder without impacting or altering the other virtual machines protected within that virtual machine folder.

2. Select the particular snapshot to restore from, and click the **Selective Restore** icon  . When a snapshot is selected, the **Restore Details** pane displays information about the contents of the snapshot. For example, if the object that was protected was a virtual machine folder, then details of the virtual machines and their virtual disks will be displayed.
3. From the **Selective Restore Inventory** dialogue, select the virtual machine to be restored, and click **Next** to continue.
4. Review the summary of the task, and click **Finish.**

7.4.2 Recovering individual files or virtual disks

Occasionally, it is preferable to not restore an entire virtual machine, but to restore individual files or an individual virtual disk. In such instances, a clone vVol VM from snapshot recovery action can be used to assist in this recovery. This process will create a new virtual machine from the selected snapshot, from which the lost data can be recovered.

1. Initialize a clone vVol VM from snapshot recovery from the **Dell VSM** button under a virtual machine **Manage** tab, and then select **Snapshots.** From the VSM Actions drop-down (), select **Clone VVOL VM from Snapshot**.

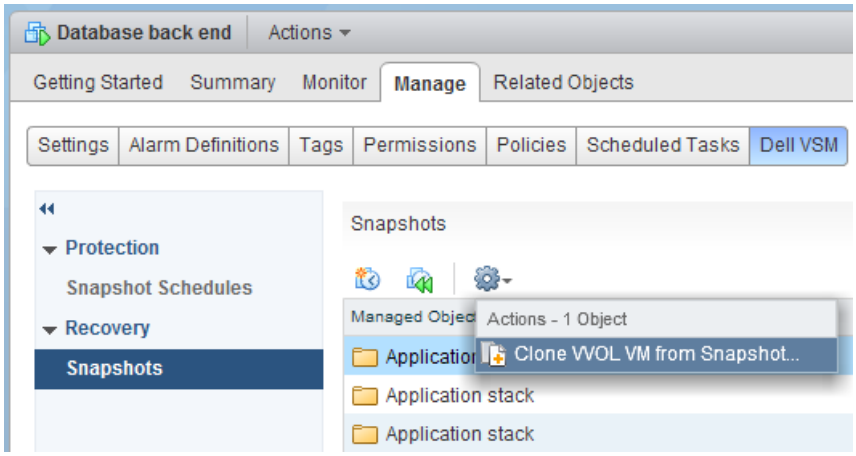

2. Provide a **Name Suffix** for the clone. By default, VSM will append –clone-<*created time_stamp*>. Click **OK** to begin the task.

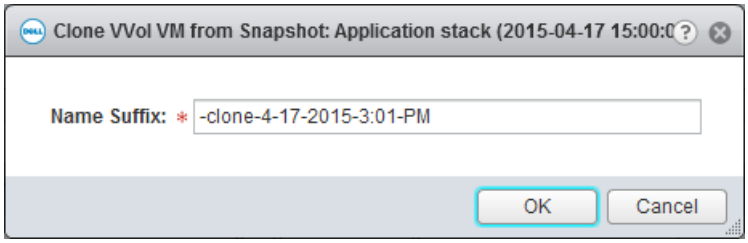

Once the clone vVol VM from snapshot task is completed, the recovery of the data can be performed.

3. Select the virtual machine and edit the setting that will add in the virtual disk from the recovered clone.
4. Once the virtual disk is mounted, a file-level copy can be performed at the guest OS level.
5. When the recovery is complete, remove the virtual disk and delete the clone.

If the guest OS does not support hot adding and removing virtual disks, it will be necessary to power off the virtual machine. This is covered in detail in section 6.7, “Advanced cloning in selective data recovery.”

Note: The cloned virtual machine will be identical to its snapshot source. Therefore, additional steps should be taken prior to powering on the virtual machine to prevent network conflicts.

7.5 Creating clones from snapshots

The ability to clone a virtual machine has been a long-time feature of vSphere, enabling virtual machines (including a completely installed and configured operating system) to be deployed in minutes. Recent versions of vSphere have enabled this even while the virtual machine is running. However, this clone functionality has not enabled the ability to turn back the clock and get a clone of the state a virtual machine was previously in.

There are several reasons why this may be desired:

- Ability to test the impact of changes on a known identical copy of the production environment
- Concerns of impacting production environment by cloning live production virtual machine
- Providing a copy of a production environment for testing, or data-mining

Now that VMware snapshot functionality has been offloaded to the array, the ability to leverage the array and VSM snapshot capabilities come into play. Using the same process described in section 7.4.2, a clone of a virtual machine can be created from a snapshot in a matter of seconds without impacting the production environment or enabling the previous state of the virtual machine to be accessed.

Note: Regardless of the reason for creating a clone of a virtual machine, it will be an exact match for the existing virtual machine. This means that the hostname, IP address, and application name space is identical, and should be isolated or altered so it does not conflict with the production environment.

8 Summary

The Dell Virtual Storage Manager is a vCenter plugin that provides a whole suite of tools for managing and protecting virtualized environments. By leveraging VSM snapshots for local data protection, environments can augment their existing backup strategies to provide a much finer window of recovery. Also included with VSM is the PS Series VASA Provider which enables Storage Policy Based Management and Virtual Volumes. Installation guidance is provided to enable customers to hit the ground running with as many of the PS Series storage features as possible. As businesses are growing, their virtual infrastructures and tools (like VSM) are needed to keep up with the growth and provide manageable recovery points and data protection.

A Additional resources

A.1 Technical support and customer service

Offering online and telephone-based support and service options, Dell support service can answer your questions about PS Series arrays, groups, volumes, array software, and host software. Availability varies by country and product, and some services might not be available in your area.

Visit Dell.com/support or call 800-945-3355 (United States and Canada).

For international support of Dell PS Series products, visit http://www.dell.com/support/contents/us/en/555/article/Product-Support/Dell-Subsidiaries/equallogic

Note: If you do not have access to an Internet connection, contact information is printed on your invoice, packing slip, bill, or Dell product catalog.

For PS Series software and documentation, visit eqlsupport.dell.com (login required).

A.2 Dell online services

Learn more about Dell products and services using this procedure:

1. Visit Dell.com or the URL specified in any Dell product information.
2. Use the locale menu or click on the link that specifies your country or region.

A.3 Dell PS Series storage solutions

To learn more about current and upcoming Dell PS Series solutions, visit the EqualLogic Dell TechCenter page. Here you can find articles, demos, online discussions, technical documentation, and more details about the PS Series product family.

For PS Series technical content, visit the EqualLogic Technical Content page on Dell TechCenter.

Dell Storage technical content can be found on the Storage Applications Engineering page.

A.4 Related documentation

See the following referenced or recommended resources related to this document.

Vendor	Document Title
Dell	*Dell PS Series Arrays: Advanced Storage Features in VMware vSphere*
Dell	*Dell EqualLogic PS Series Template Volumes and Thin Clones: How and When to Use them*
Dell	*EqualLogic PS Series Architecture: Snapshot Space Borrowing Overview*
VMware	*vSphere Security Guide*
VMware	KB 1015180: Understanding virtual machine snapshots in VMware ESXi and ESX

B Configuration details

The following table shows the software and firmware used for the preparation of this paper.

Table 1 Software and firmware versions

Vendor	Model	Software revision
Dell	PS Series SAN	6.0, 7.0, and 8.0
Dell	Virtual Storage Manager	4.0 and 4.5
VMware	vCenter™	5.5 and 6.0
VMware	ESX®/ESXi™	5.1, 5.5, and 6.0

C Virtual Volumes terminology

While reading this paper, it is important to have an understanding of the following vVols-relevant terminology.

VASA Provider: The VASA Provider plays an important role in enabling a vVol environment. The VASA Provider offers out-of-band management access to the SAN from vCenter. It enables vCenter to communicate with the SAN in ways that the current SCSI protocol does not. Through this communication channel, vCenter sends operational requests for interacting with the Virtual Volumes that back Virtual-Volume-based virtual machines.

The Dell VASA Provider ships as part of the Virtual Storage Manager plugin for vCenter, which also provides enhanced storage management functionality to vCenter.

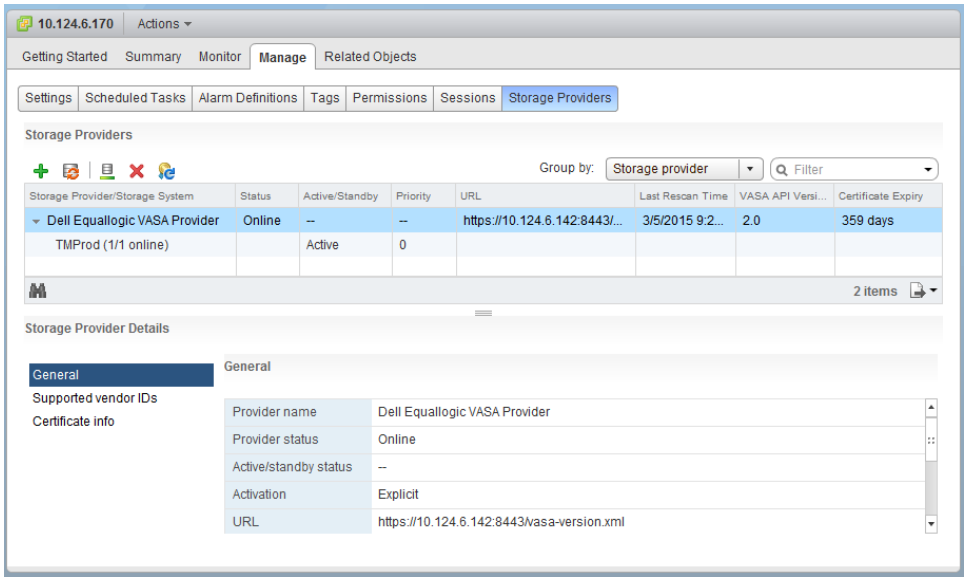

Figure 13 VASA Provider status as shown in vSphere Web Client

Protocol endpoint: The protocol endpoint is a unique volume on the SAN, it has a size of zero megabytes and a LUN ID of 256.

It is the SAN endpoint of the communication between the ESXi host and the Virtual Volumes on the SAN. The transport protocol (iSCSI in the case of PS Series) endpoint is where the communication is turned over to internal SAN protocols. This unique volume can be thought of as a multiplexer LUN that acts as both the target and the initiator, and enables ESXi hosts to see a single volume while multiple independent volumes fan out behind it in the SAN. These multiple independent volumes are included in a virtual machine.

The protocol endpoint is also where access controls are placed and initiators are queried to insure that they are permitted access to the storage container and Virtual Volumes. VSM manages these access controls directly from vCenter.

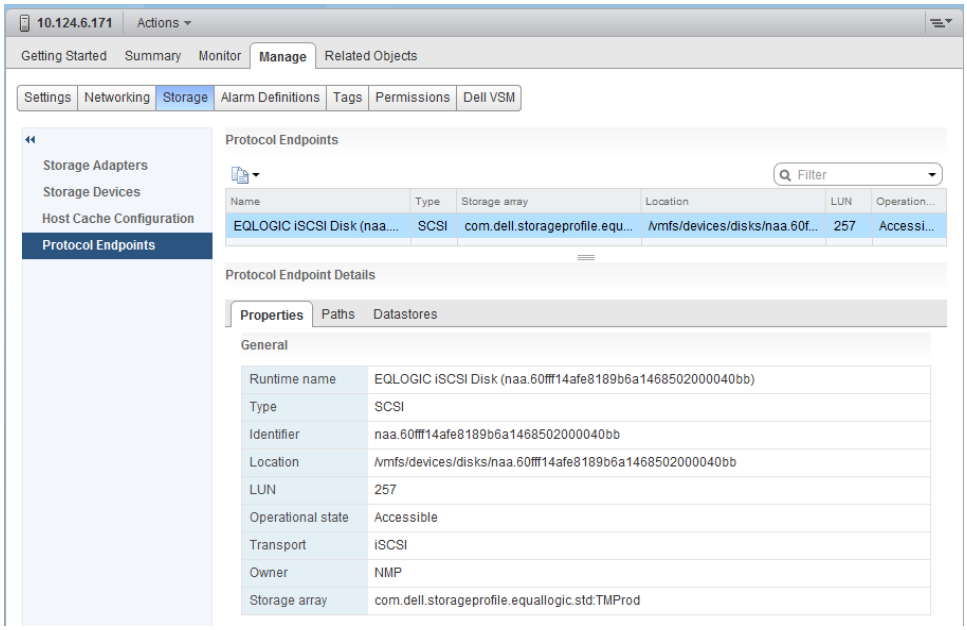

Figure 14 Protocol endpoint as shown in vSphere Web Client

Storage container: A storage container is reserved space on the SAN that can be increased and decreased as needs change (PS Series storage requires thick, or 100 percent, space reservation). It can also be conceptualized as a type of folder object on the SAN for organizing multiple volumes together.

Storage containers are seen and treated as regular datastores by vSphere, and are referred to as a vVol-type datastore. They can be browsed as typically done when seeking virtual machine log files. They enable many vSphere workflows to remain unchanged even though significant changes have occurred.

Multiple storage containers can exist within a PS Series group, up to 32, but they cannot span pool or be migrated from one pool to another pool.

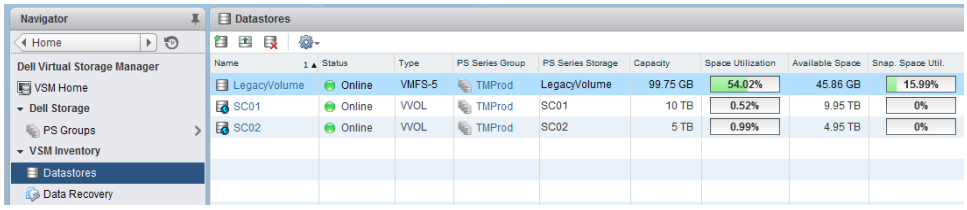

Figure 15 Datastores of type VMFS-5 and of type VVOL as shown in VSM plugin for vSphere Web Client

Virtual Volumes: At a high level, an individual Virtual Volume is a regular volume to the SAN and it can be manipulated by various SAN functions. However, from the vSphere perspective, an individual virtual volume is part of a virtual machine; a complete virtual machine consists of several Virtual Volumes of different types. The SAN is aware of the Virtual Volumes that belong to other Virtual Volumes because of information communicated from vSphere through the VASA Provider.