



VMware vSphere® 5.0 Evaluation Guide

Volume Two - Advanced Storage Features

TECHNICAL WHITE PAPER

Table of Contents

About This Guide	4
System Requirements	4
Hardware Requirements	4
Servers	4
Storage	4
Networking	4
Software and Licensing Requirements	5
VMware vSphere	5
Guest operating systems	5
Evaluation Guide Environment Setup	5
Server Configuration	6
Logical Network Setup	6
Storage Setup	8
Virtual Machine Setup	8
VMware vSphere 5.0 Evaluation Guide, Volume Two – Worksheet	9
Help and Support During the Evaluation	10
VMware Contact Information	10
Providing Feedback	10
vSphere Advanced Storage Features	11
Enabling Storage I/O Control	11
Performing a VMware vSphere® VMFS Live Upgrade	13
Create a VMFS-3 Datastore	13
Move Virtual Machines to the VMFS-3 Datastore	18
Initiate the Live Upgrade from VMFS-3 to VMFS-5	18
Testing vSphere Storage DRS	20
Introduction	20
Creating a Datastore Cluster	20
Exploring Your Datastore Cluster	24
Provisioning a New Virtual Machine	27
Space Balancing	32

Using Profile-Driven Storage	36
Introduction	36
Create a Virtual Machine Storage Profile with a User-Defined Storage Capability	36
Assign a Storage Capability to a Datastore	38
Create a Virtual Machine Storage Profile Containing the User-Defined Storage Profile	39
Assign a VM Storage Profile to a Virtual Machine	41
Check Whether the Virtual Machine Is Running on “Compliant” Storage	42
Bring a Virtual Machine into Compliance	42
Evaluating Storage I/O Control	44
Introduction	44
Create a Virtual Machine on a Datastore	44
Generate I/O to the Datastore from the Virtual Machine	45
Enable Storage I/O Control	47
Monitoring the Effect of Limiting IOPS	48
Monitoring the Effects of Shares	49
Summary	51

About This Guide

The purpose of the *VMware vSphere 5.0 Evaluation Guide, Volume Two – Advanced Storage Features* is to support a self-guided, hands-on evaluation of VMware vSphere® 5.0 (“vSphere”) advanced storage features such as vSphere Storage I/O Control (SIOC) and vSphere Storage DRS.

This guide covers evaluation cases that are suitable for IT professionals who have an existing VMware virtualization environment and want to evaluate features in vSphere that enable greater storage automation and consolidation while maintaining service levels.

System Requirements

To ensure the best experience when using this guide, the user must configure hardware and software as detailed in the following section.

Hardware Requirements

This guide makes the following assumptions about your existing physical infrastructure:

Servers

You must have at least three dedicated servers capable of running VMware ESXi™ 5.0 to provide resources for this evaluation.¹

Storage

You have shared storage with enough space available to allow creating three 100GB dedicated datastores. Shared storage can be SAN or NAS. This document assumes SAN-based storage.

Networking

You will need at least three virtual networks configured to separate virtual machine, VMware vSphere® vMotion® and vSphere management. These networks can be set up on a single virtual switch with multiple port groups or across multiple virtual switches. For the purposes of this evaluation guide, the configuration includes a single vSphere standard switch with three port groups.

For more detailed requirements, see the following table.

HARDWARE	MINIMUM	WHAT'S USED IN THIS GUIDE
ESXi	Three ESXi hosts CPU – Two processors of 2GHz Memory – 6GB Network – 2x 1Gb network adaptor	Three ESXi hosts (Cisco CS 1.3.1) CPU – Two quad-core Intel Xeon “Nehalem” processors of 2.6GHz Memory – 48GB Network – 4x 10GB network adaptor
Storage	Two datastores (100GB each)	Three datastores (Fibre Channel – 100GB each)
Network	One VLAN for carrying virtual machine traffic; one VLAN for carrying management traffic	Separate VLANs for ESXi management, vMotion, and virtual machine traffic

1. These servers must be on the *VMware vSphere 5.0 Hardware Compatibility List (HCL)*.

Software and Licensing Requirements

This guide makes the following assumptions about your existing software infrastructure:

VMware vSphere

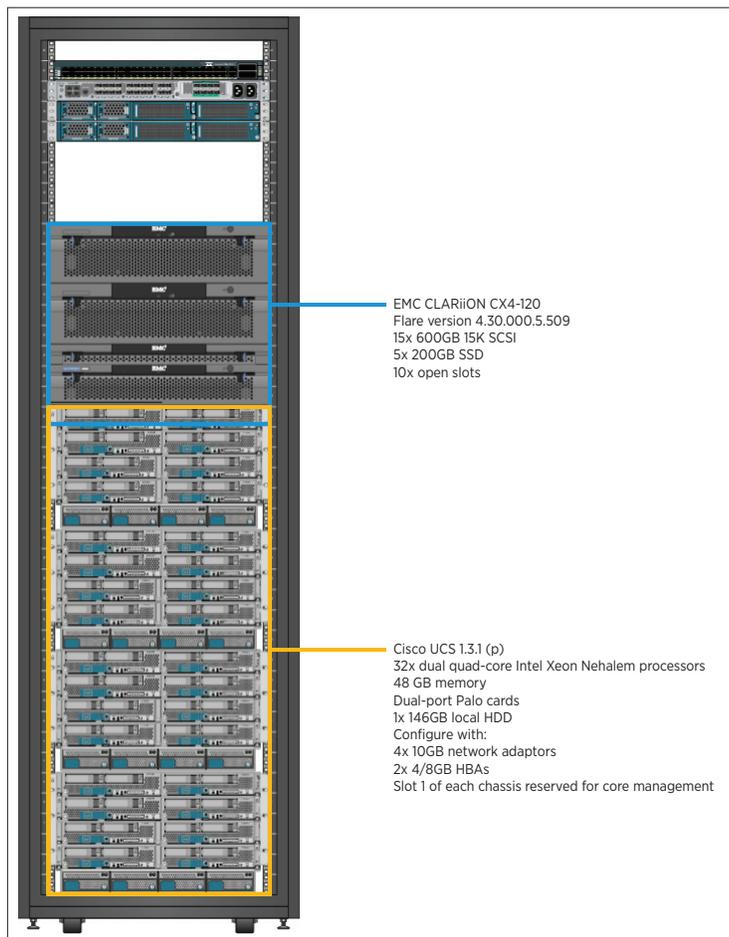
This volume of the *VMware vSphere 5.0 Evaluation Guide* requires vSphere 5.0 and licensing for vSphere Enterprise Plus. The vSphere 5.0 evaluation license available from the VMware evaluation portal provides Enterprise Plus functionality for 60 days and is the best choice for performing the vSphere 5.0 evaluations.

Guest operating systems

This volume of the *VMware vSphere 5.0 Evaluation Guide* will require five or six virtual machines running Windows 2003 or Windows 2008.

Evaluation Guide Environment Setup

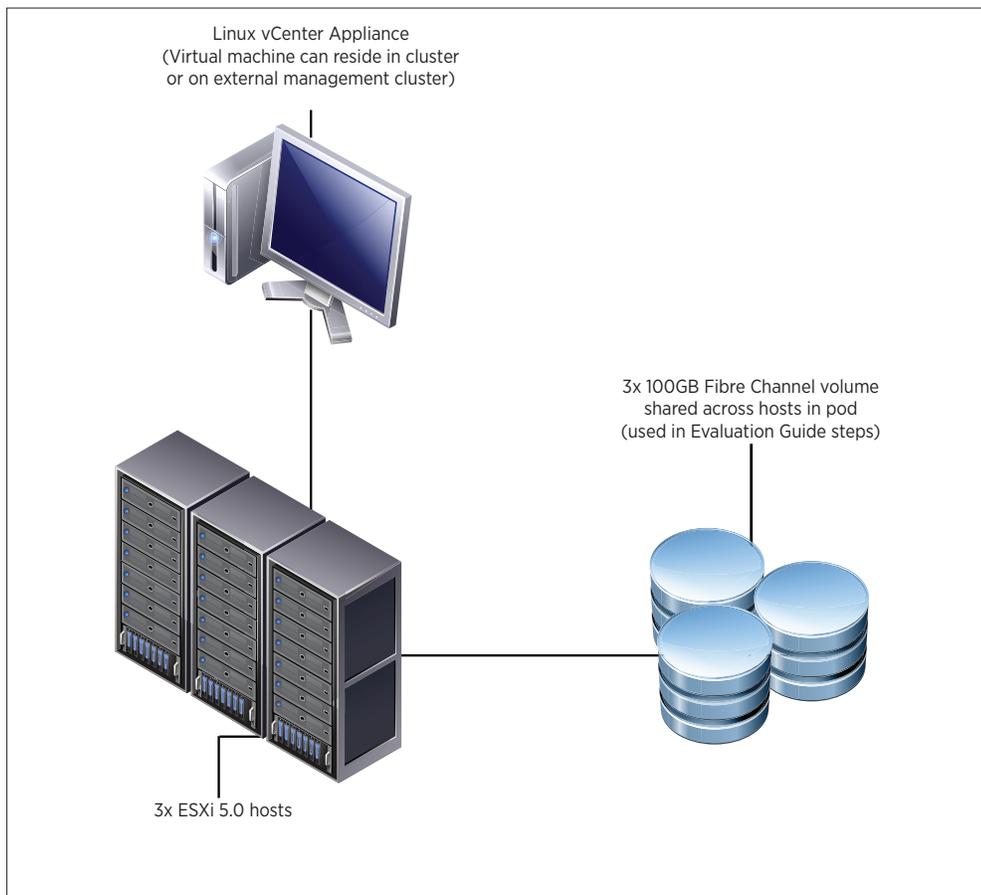
The VMware Technical Marketing lab was built using a combination of Cisco UCS server hardware and EMC CLARiiON CX-4 Fibre Channel (FC) storage. The environment consisted of eight identical four-node “pods,” with most pods configured as a three-node ESXi cluster and a fourth node for management. In many cases, additional resources have been configured in the Technical Marketing test-bed configuration to support other evaluation projects and are present in the diagrams. The user can configure only what is called for in the following and can safely ignore additional resources in screenshots and topology diagrams. The following picture shows the Technical Marketing test rack.



Server Configuration

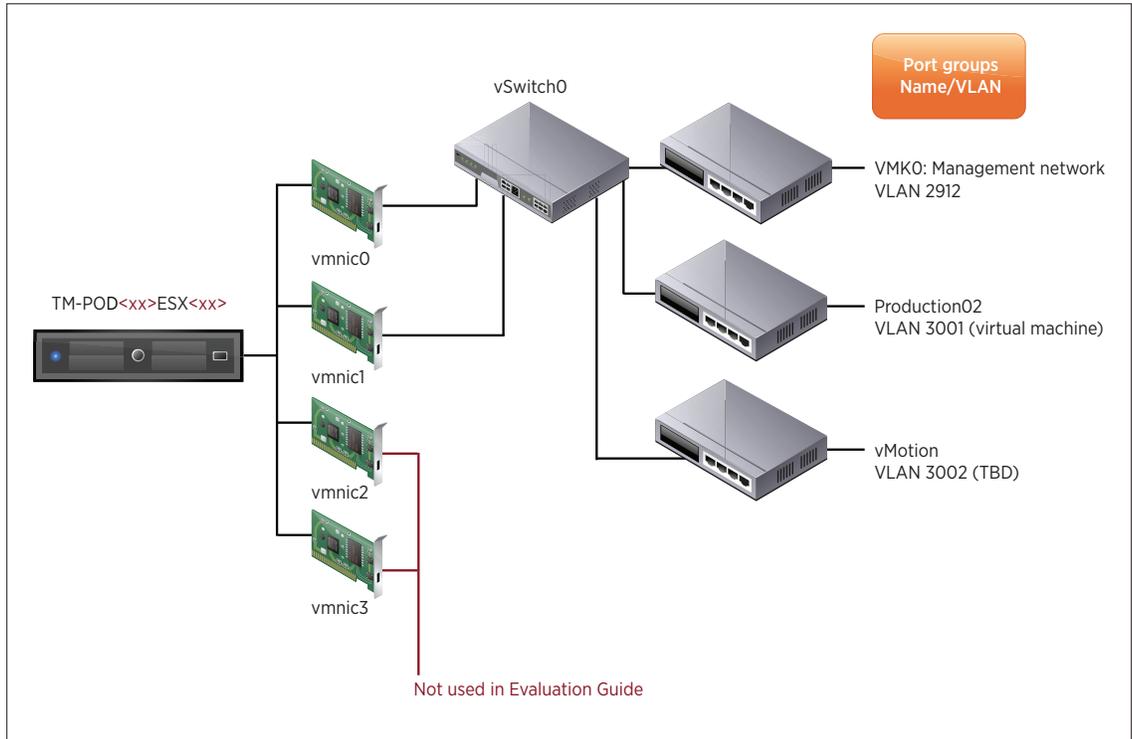
The *VMware vSphere Evaluation Guide* is based on three modern “server class” systems with adequate processor and memory to host six to eight minimally configured virtual machines used for testing. The servers used for this evaluation do not need to be extremely powerful, just reliable and on the vSphere 5.0 HCL.

Each server must have at least 2x 1GB or 2x 10GB network adaptor and proper connection to shared storage. The following diagram summarizes the evaluation guide test-bed configuration.

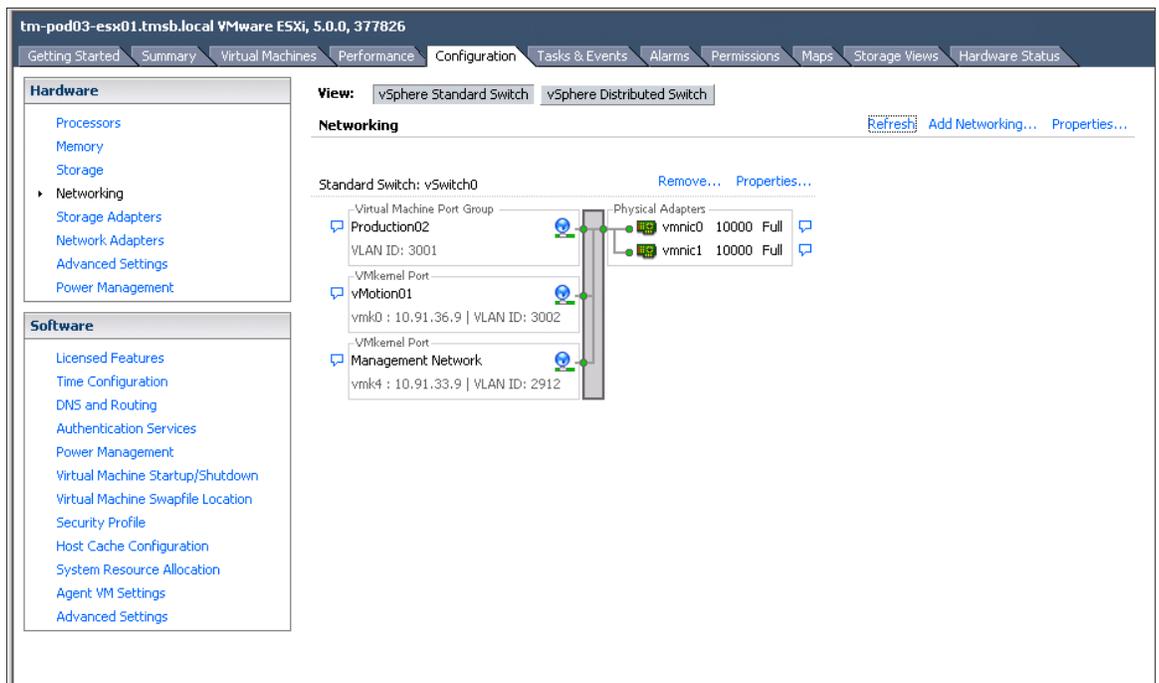


Logical Network Setup

VMware vSphere 5.0 Evaluation Guide, Volume Two, uses a simple network configuration consisting of three logical networks. The first is for vSphere management traffic, including vSphere High Availability. The second is for vMotion and the third is for virtual machine traffic. Each logical network is configured as a port group on a standard switch, with a corresponding VLAN configured to provide physical isolation of the network traffic.

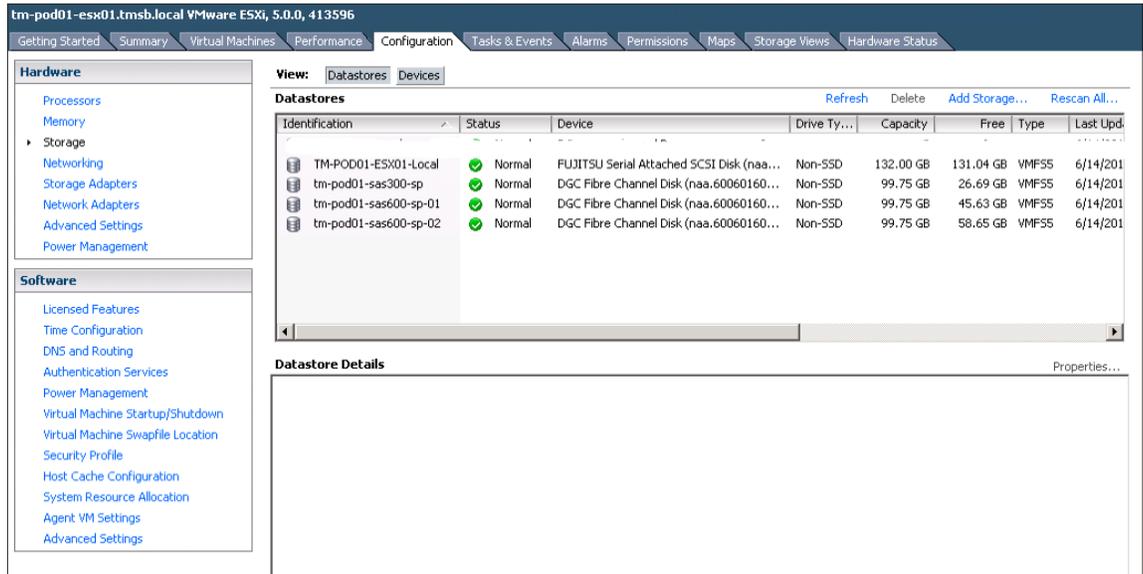


On the vSphere side, the network configuration looks like the following:



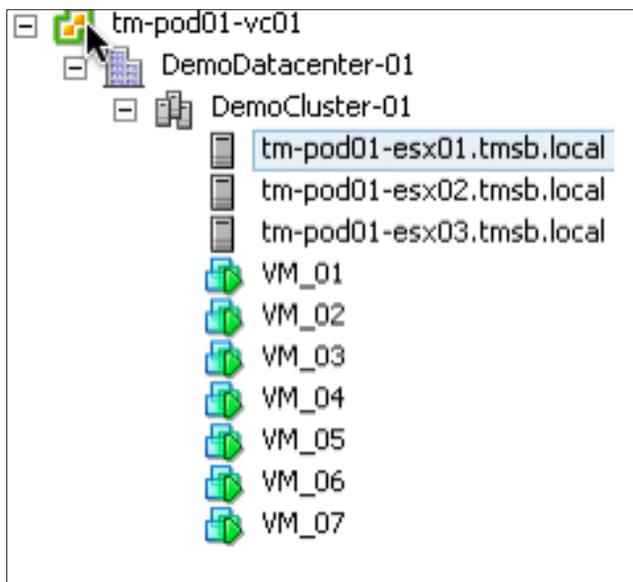
Storage Setup

VMware vSphere 5.0 Evaluation Guide, Volume Two, uses a storage configuration consisting of three 100GB FC LUNs presented to each host, enabling creation of three datastores.



Virtual Machine Setup

VMware vSphere 5.0 Evaluation Guide, Volume Two, uses a total of seven virtual machines for testing. This volume will require Windows 2003 or Windows 2008 guest operating systems. It is up to the user to configure virtual machines that can be brought up to a running state for testing. The following diagram shows VM_01 through VM_07 configured in the Technical Marketing test lab:



VMware vSphere 5.0 Evaluation Guide, Volume Two – Worksheet

You can use the following worksheet to organize your evaluation process.

HARDWARE CHECKLIST:	
All hardware has been validated against the <i>VMware Hardware Compatibility List</i> (HCL).	
Each host has 2x 1GB or 2x 10GB network cards connected to a common switch (will be configured as a network adaptor team).	
Each host has required HBA/network adaptor to access shared storage.	

SOFTWARE CHECKLIST:	
VMware vSphere/VMware ESXi installation media is available.	
VMware vCenter™ Server appliance is downloaded.	
VMware vSphere® Client™ is installed.	
ESXi host 1 host name.	
ESXi host 2 host name.	
ESXi host 3 host name.	
Subnet, netmask and default gateway for management network.	
Subnet, netmask and default gateway for virtual machine network.	
Subnet, netmask and default gateway for vMotion network.	

STORAGE CHECKLIST:	
All servers can see at least three common 100GB LUNs (or NFS exports).	
Datastore 1 name.	
Datastore 2 name.	
Datastore 3 name.	

Help and Support During the Evaluation

This guide provides an overview of the steps required to ensure a successful evaluation of VMware vSphere. It is not meant to substitute product documentation. Refer to online vSphere product documentation for more detailed information (see the following links). You can also consult the online VMware knowledge base if you have any additional questions. If you require further assistance, contact a VMware sales representative or channel partner.

VMware vSphere and vCenter resources:

- Product documentation:
<http://www.vmware.com/support/pubs/>
- Online support:
<http://www.vmware.com/support/>
- Support offerings:
<http://www.vmware.com/support/services>
- Education services:
<http://mylearn1.vmware.com/mgrreg/index.cfm>
- Support knowledge base:
<http://kb.vmware.com>
- VMware vSphere® PowerCLI Toolkit Community:
http://communities.vmware.com/community/developer/windows_toolkit
(or type Get-VIToolkitCommunity within PowerCLI)
- PowerCLI Blogs:
<http://blogs.vmware.com/vipowershell>

VMware Contact Information

For additional information or to purchase VMware vSphere, the VMware global network of solutions providers is ready to assist. If you would like to contact VMware directly, you can reach a sales representative at 1-877-4VMWARE (650-475-5000 outside North America) or email sales@vmware.com. When emailing, include the state, country and company name from which you are inquiring. You can also visit <http://www.vmware.com/vmwarestore/>.

Providing Feedback

We appreciate your feedback on the material included in this guide. In particular, we would be grateful for any guidance on the following topics:

- How useful was the information in this guide?
- What other specific topics would you like to see covered?
- Overall, how would you rate this guide?

Send your feedback to the following address: tmdocfeedback@vmware.com, with “VMware vSphere 5.0 Evaluation Guide” in the subject line. Thank you for your help in making this guide a valuable resource.

vSphere Advanced Storage Features

Enabling Storage I/O Control

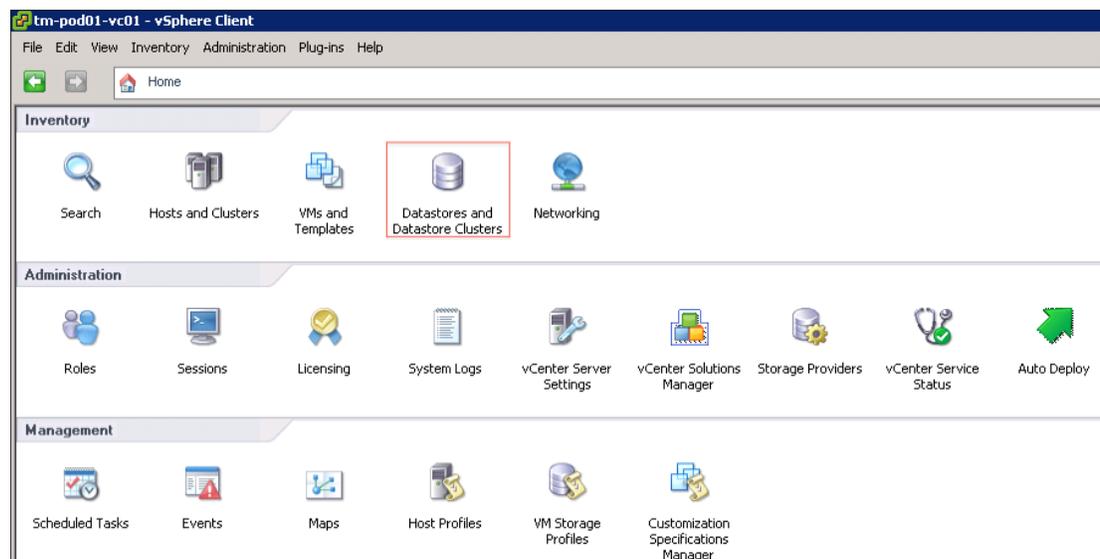
vSphere Storage I/O Control (SIOC) was initially introduced in vSphere 4.1 to provide I/O prioritization of virtual machines running on a cluster of VMware ESX® servers that had access to a shared, iSCSI or FC, storage pool. It extended the familiar constructs of shares and limits, which existed for CPU and memory, to address storage utilization through a dynamic allocation of I/O queue slots across a cluster of ESX servers.

Enabling Storage I/O Control to Avoid Denial of Service on Shared Storage

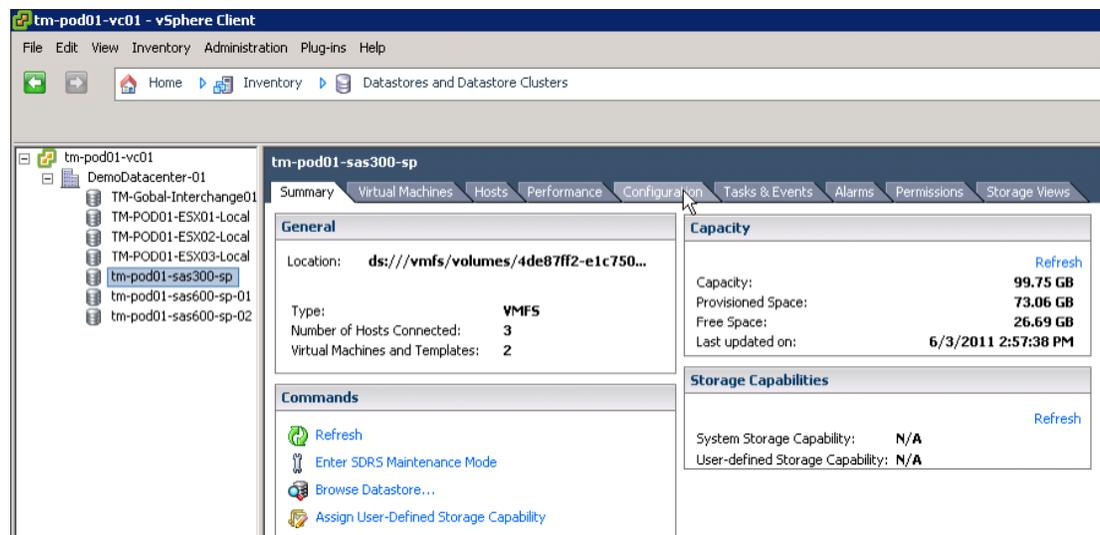
This next section will display how easy it is to enable SIOC:

Step 1:

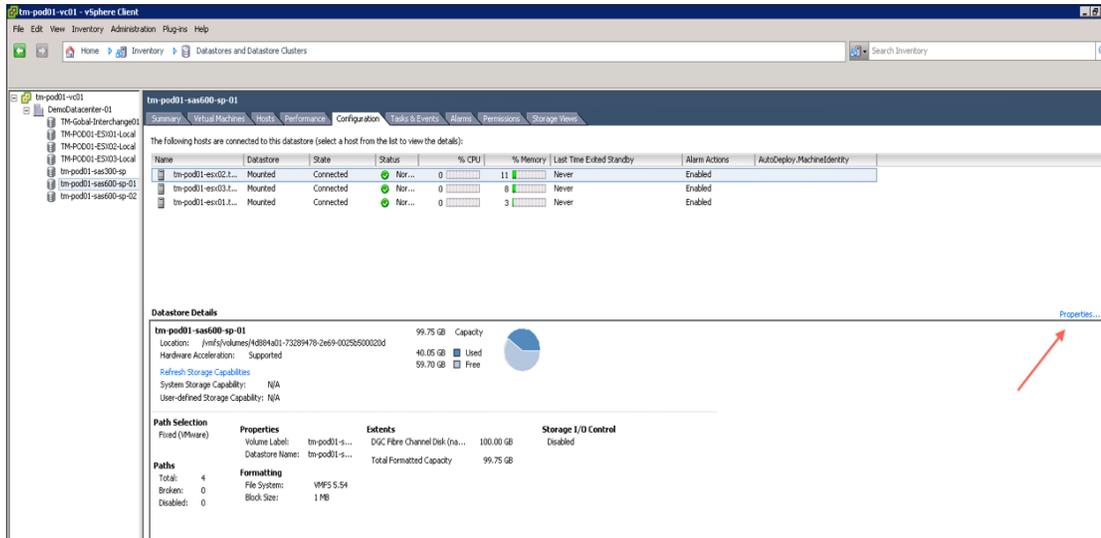
1. Go to the **Datstores and Datastore Clusters** view.



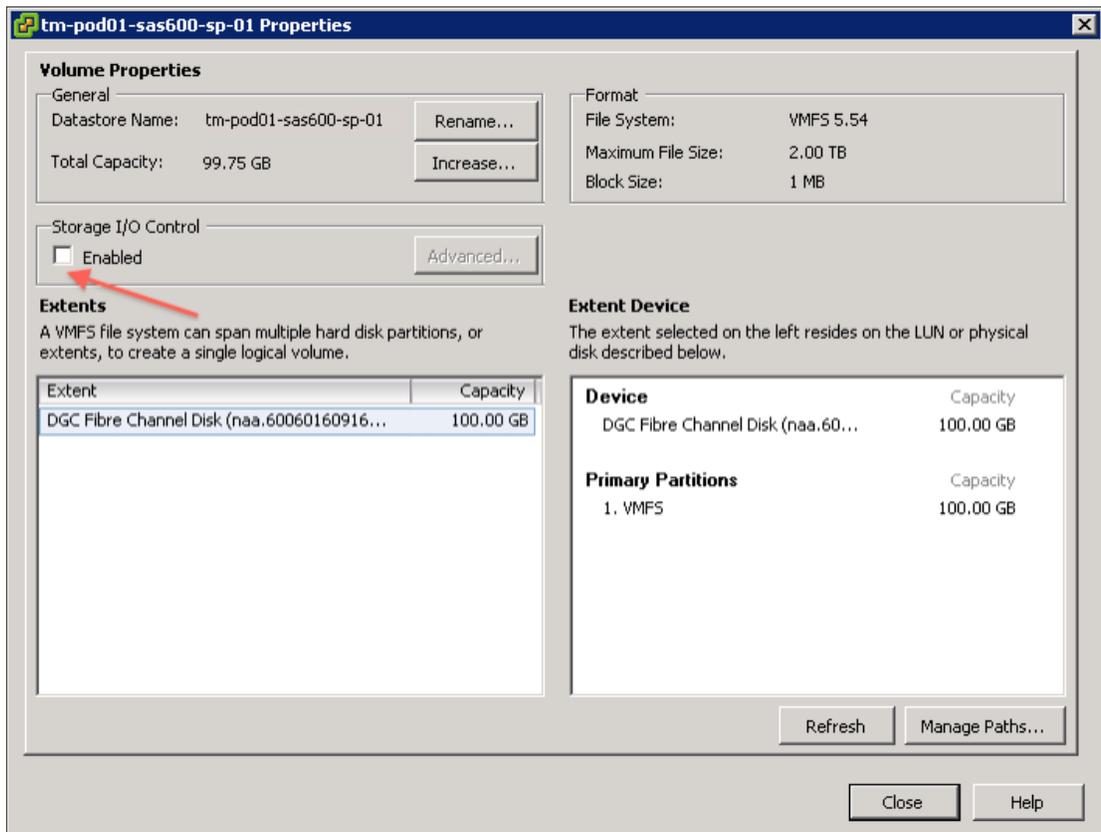
2. Select a datastore.



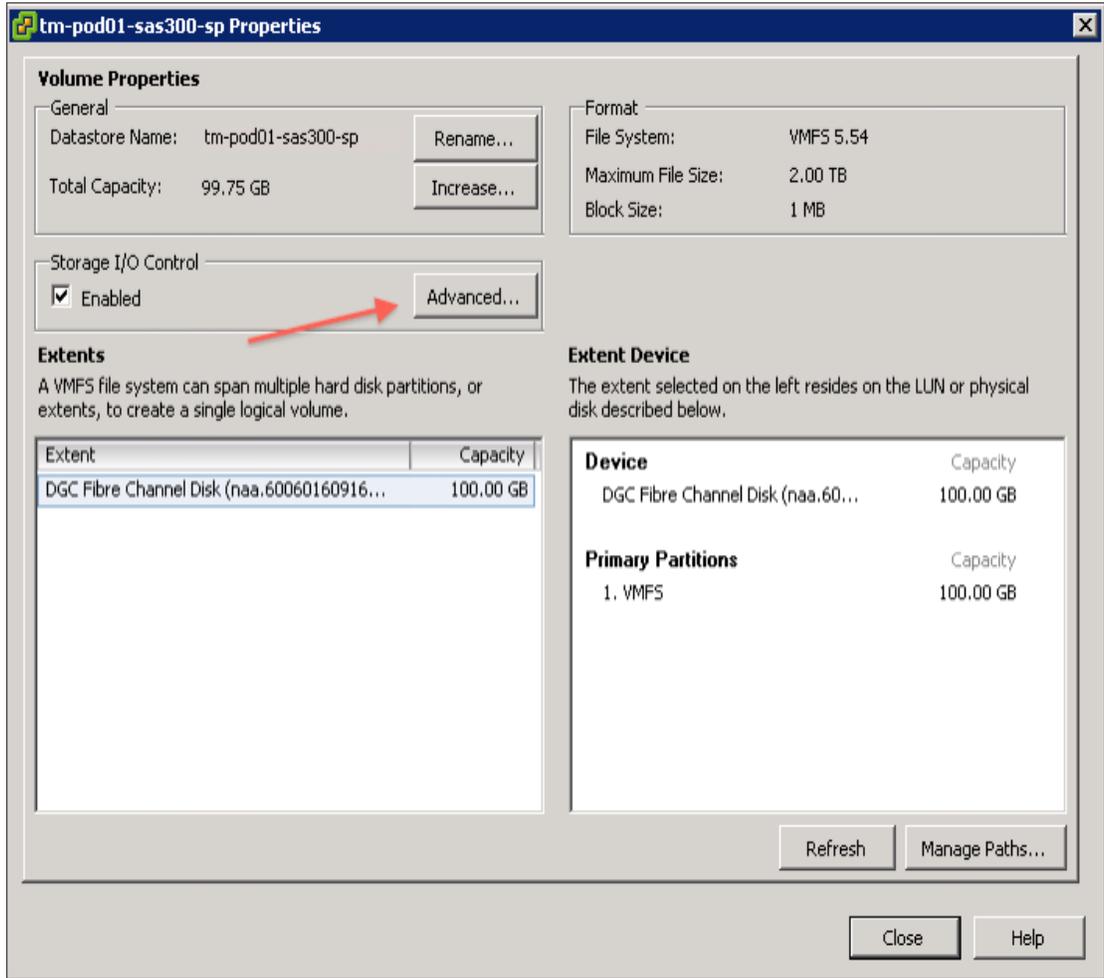
3. Click the **Configuration** tab and click **Properties**.



4. Click **Enabled** in the **Storage I/O Control** section.



- The latency threshold can be configured separately when you click **Advanced**. We will leave it set to the default.



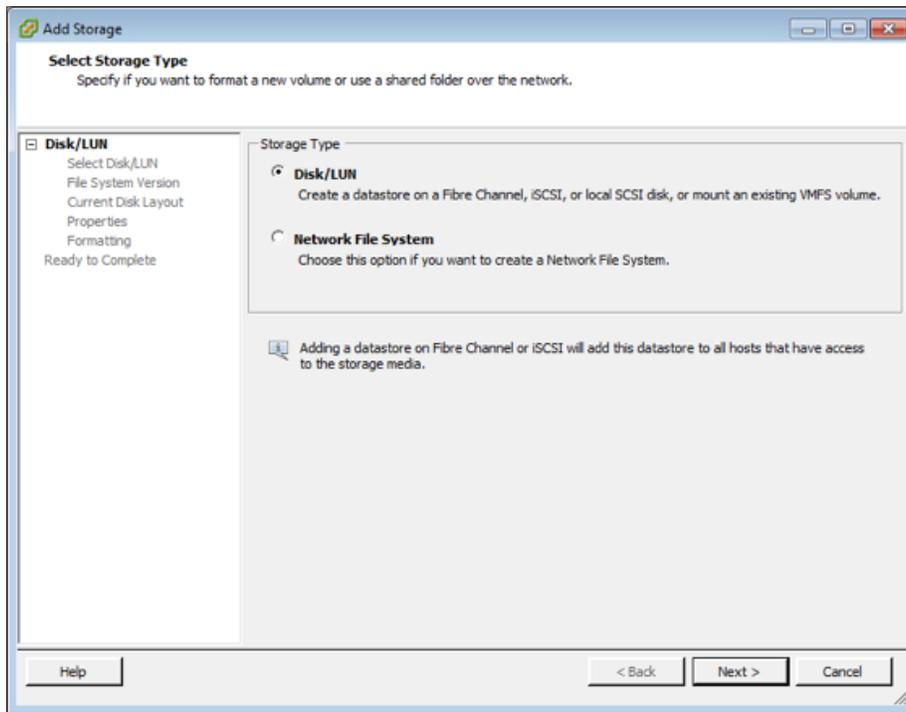
Performing a VMware vSphere® VMFS Live Upgrade

Create a VMFS-3 Datastore

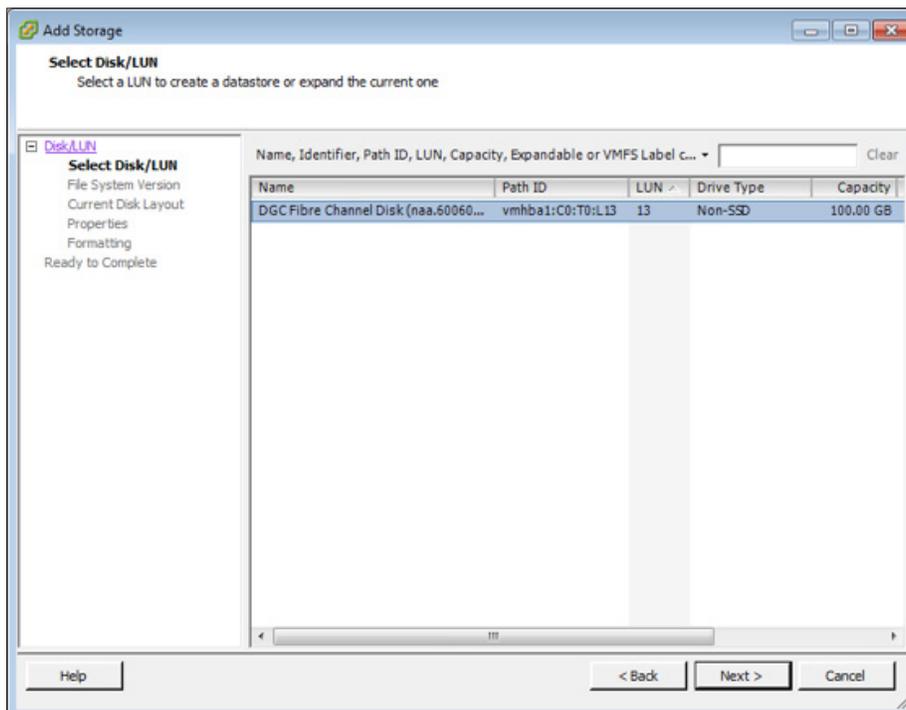
This first step is necessary only if you do not already have a VMFS-3 datastore. If you already have a VMFS-3 datastore, proceed to step 2. If you do not have a VMFS-3 datastore, first select an ESXi host from the vCenter inventory, click the **Configuration** tab, and in the **Hardware** window, choose **Storage**. This will display the current list of datastores. Click on the link **Add Storage...** :



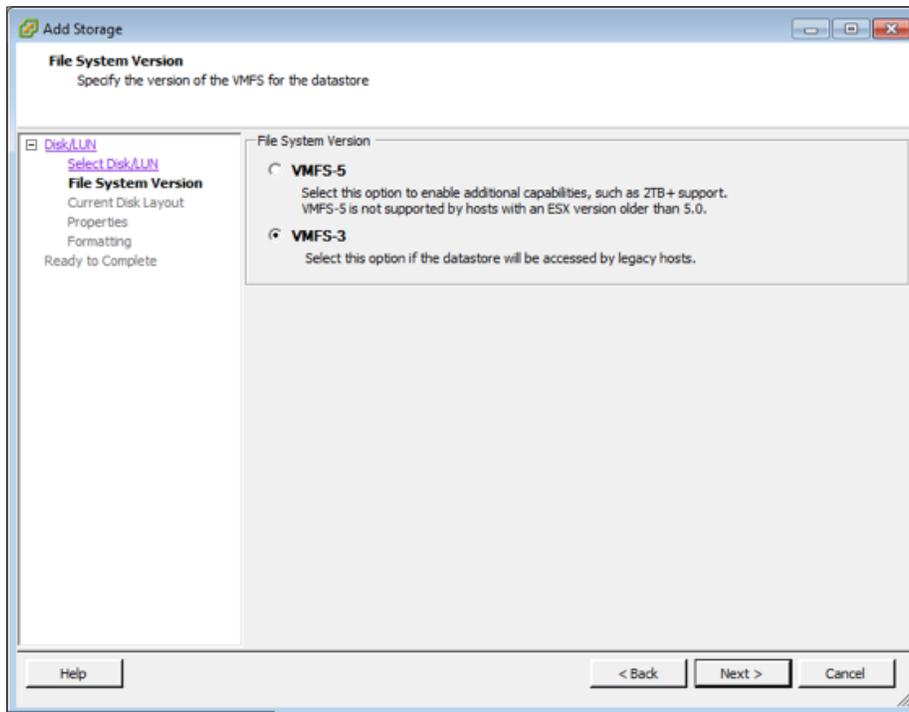
This will launch the **Add Storage** wizard. From the first screen, select **Disk/LUN**:



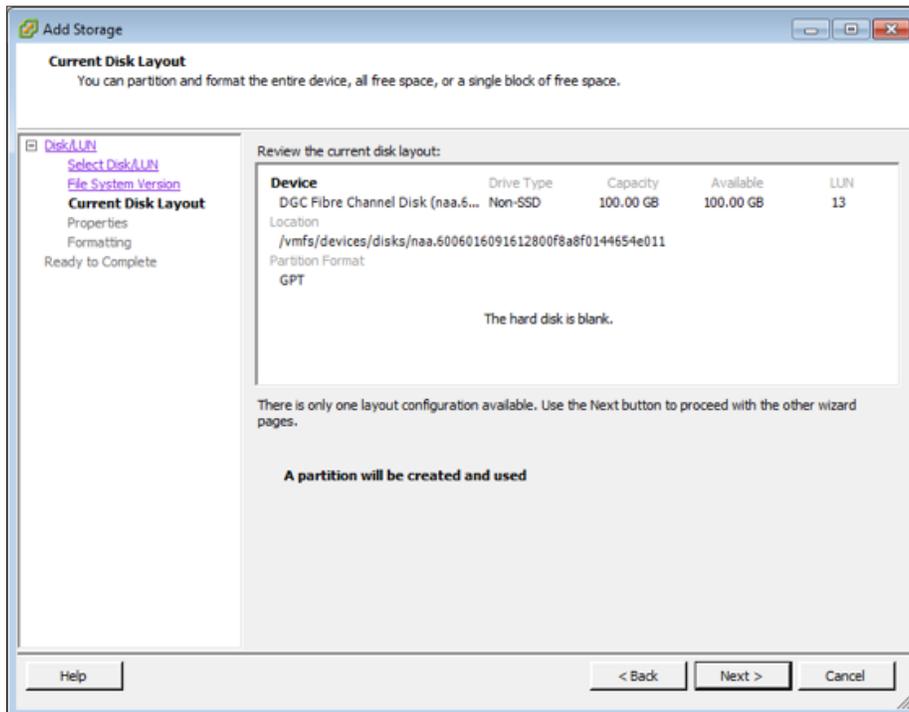
Choose a spare LUN for the VMFS-3 datastore. In this example, a spare 100GB LUN is selected.



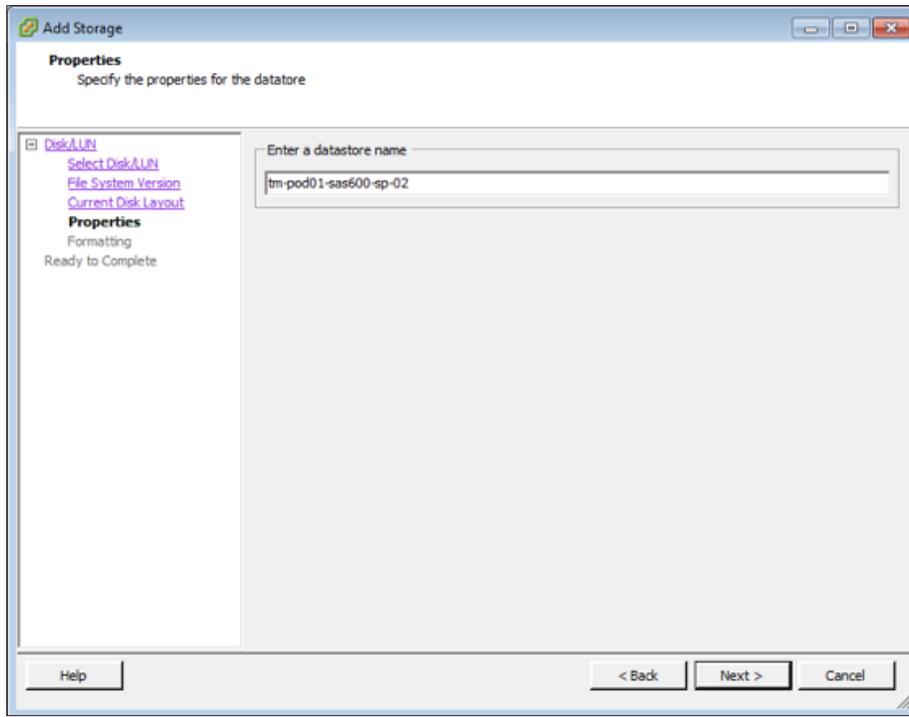
Select the filesystem version. For the purposes of this exercise, you should choose VMFS-3.



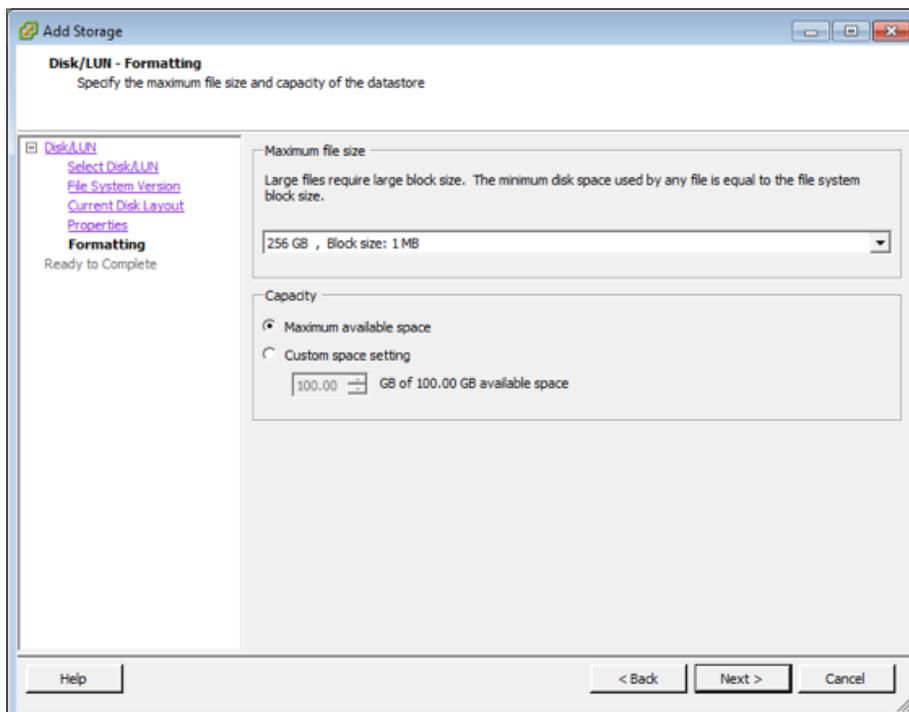
This displays the disk layout. The partition format used for VMFS in vSphere 5.0 has changed from master boot record (MBR) to GUID partition table (GPT).



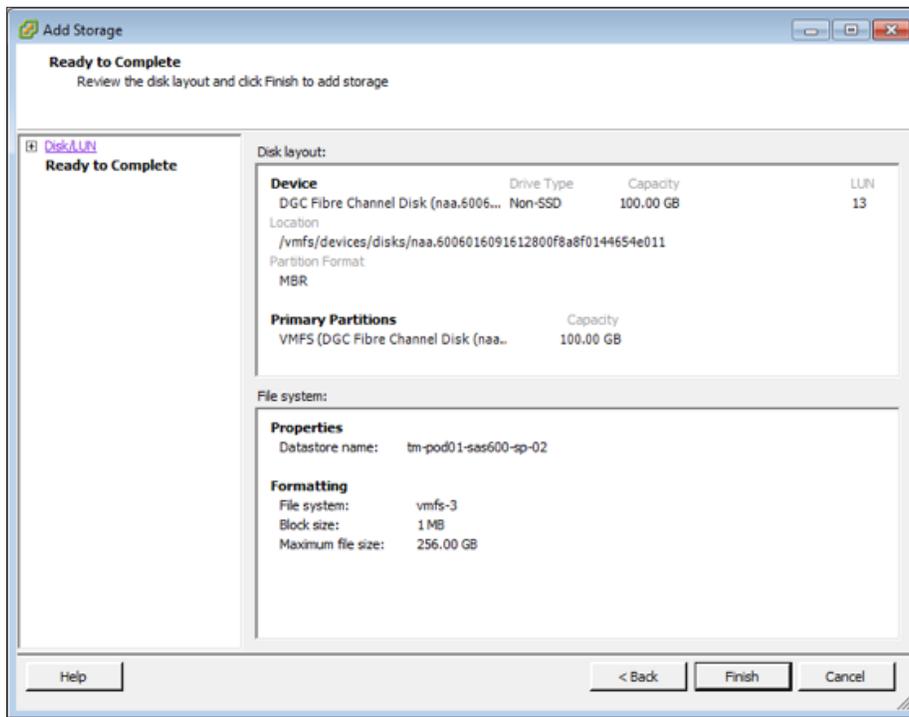
Give the datastore a name.



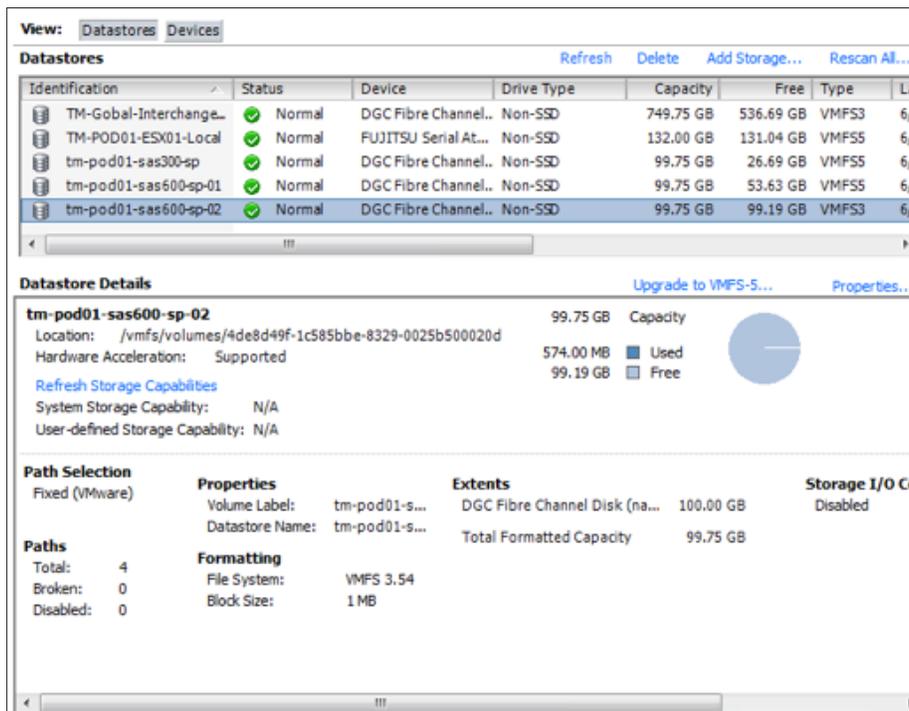
Select a block size. This impacts the largest file size in VMFS-3. You can also choose to use only part of the disk for VMFS-3. In this example, we will leave the block size at 1MB, giving us a maximum file size of 256GB. We will also use all available space on the disk for this datastore.



Click **Finish** to initiate the creation of the VMFS-3 datastore.



When the datastore is created, select it from the **Storage** view. In the **Datastore Details**, notice that it is VMFS-3. There is also an option in the **Datastore Details** to perform an **Upgrade to VMFS-5**. We will return to this in a while.

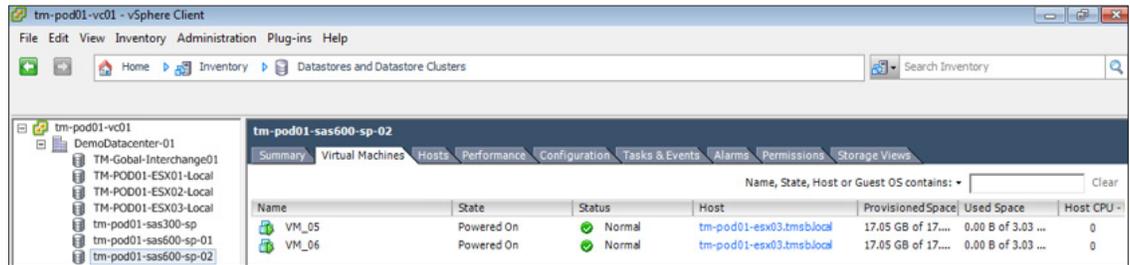


This completes step 1 of the **VMFS Live Upgrade** evaluation.

Move Virtual Machines to the VMFS-3 Datastore

If you have virtual machines on other datastores, and your environment contains a license for vSphere Storage vMotion, you can hot-migrate a number of virtual machines to this VMFS-3 datastore. If you do not have running virtual machines, create a new one, or deploy one from a template, to the VMFS-3 filesystem. At the end of this step, you should have at least one virtual machine running on this datastore.

To verify that the virtual machines are running on your VMFS-3 datastore, navigate to **Datastore** and **Datastore Clusters**, choose your VMFS-3 datastore and then select the **Virtual Machines** tab. In this example, there are two virtual machines running on the VMFS-3 datastore.

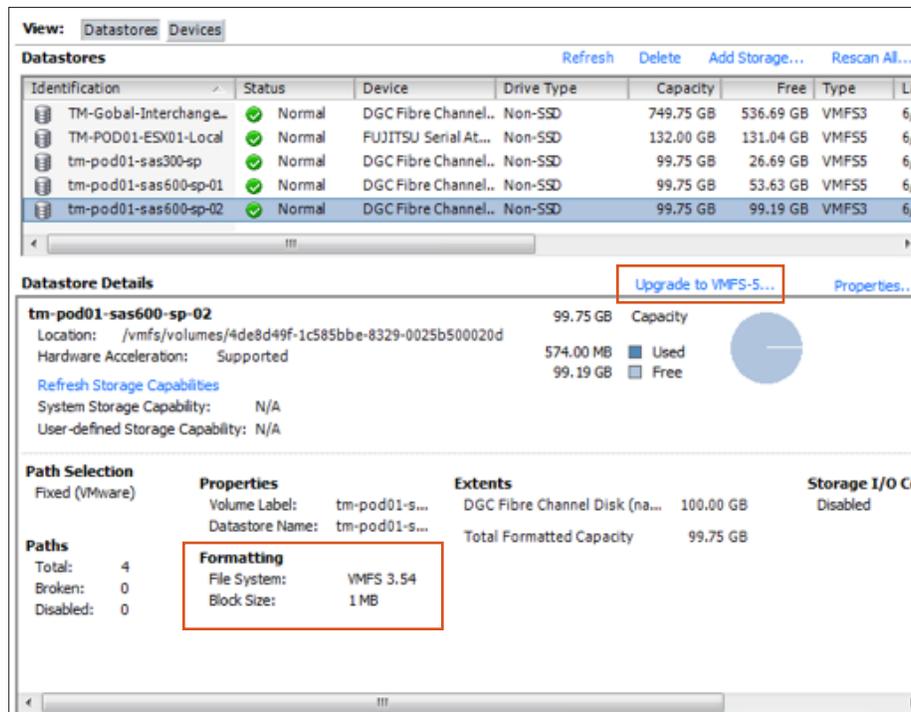


We have running virtual machines to demonstrate that the VMFS can be upgraded without impacting the running virtual machines using that datastore.

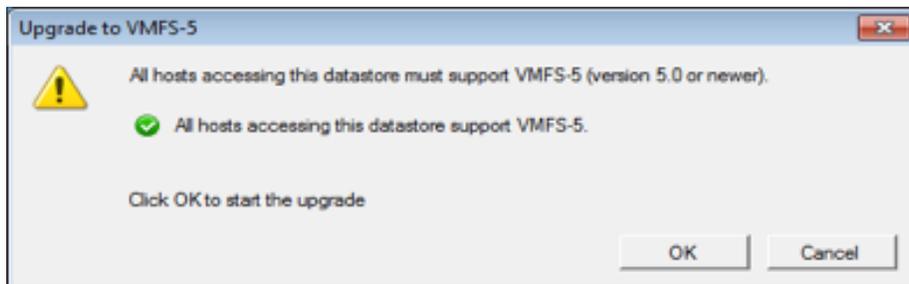
This completes step 2 of the **VMFS Live Upgrade** evaluation.

Initiate the Live Upgrade from VMFS-3 to VMFS-5

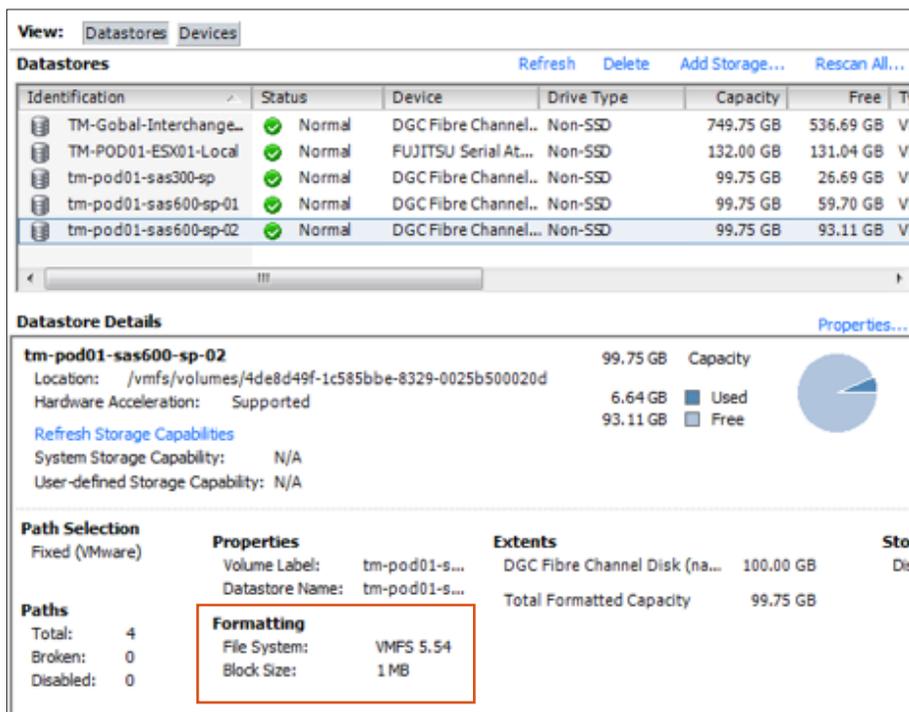
Return to the **Storage** view that we saw previously, where there was a link to **Upgrade to VMFS-5**.



Click the **Upgrade to VMFS-5** link. The first thing that vCenter does is verify that all hosts accessing the datastore are running ESXi 5.0. If any hosts accessing this datastore are not running ESXi 5.0, the upgrade is not allowed. In this example, all hosts are ESXi 5.0:



Click **OK** to proceed with the upgrade. You should see an **Upgrade VMFS** task commence in the task bar. After a moment, the task completes and your VMFS-3 filesystem is now a VMFS-5 filesystem:



While this upgrade of VMFS-3 to VMFS-5 was taking place, the virtual machines continued to run on the datastore. There was no need to move them to other datastores during the upgrade, which is something that was necessary in previous upgrades on VMFS.

This concludes the **VMFS Live Upgrade** evaluation.

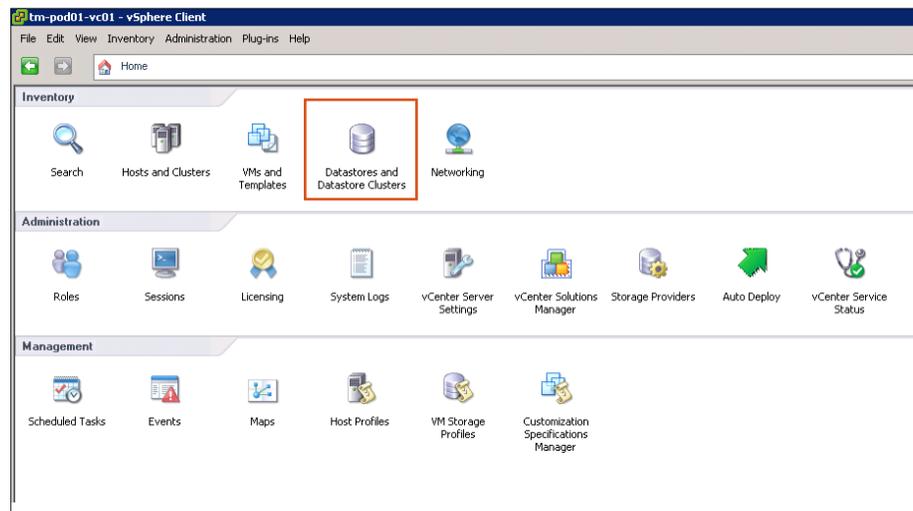
Testing vSphere Storage DRS

Introduction

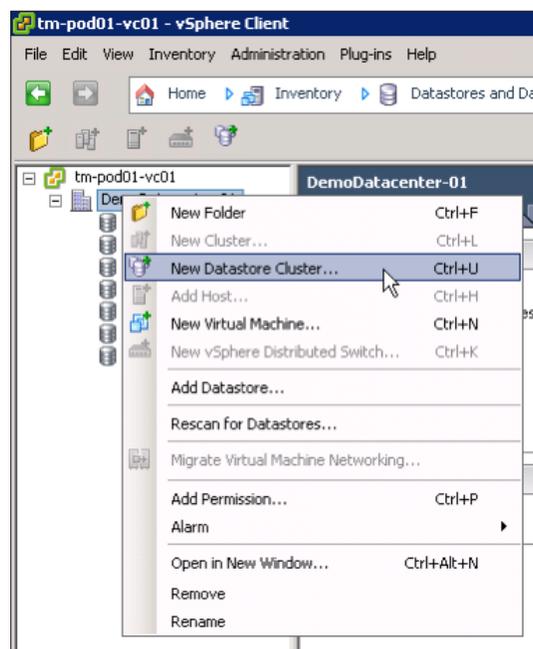
Virtual machine provisioning has always imposed operational challenges. Monitoring datastore capacity and I/O load has proven to be very difficult and as a result is often neglected. This can lead to hot spots and over- or under-utilized datastores over time. vSphere Storage DRS is a new feature introduced in vSphere 5.0 that helps prevent these problems. It provides smart virtual machine placement and load balancing mechanisms based on I/O and space capacity. Storage DRS will help decrease operational effort associated with the provisioning of virtual machines and the monitoring of the storage environment.

Creating a Datastore Cluster

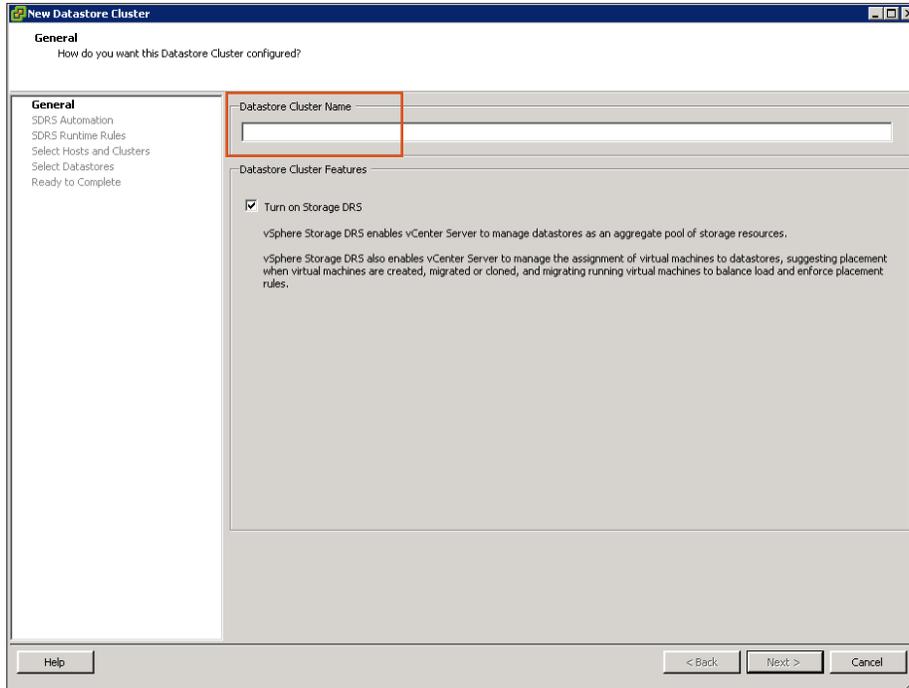
1. From the **vCenter Home** view, select **Datastores and Datastore Clusters**.



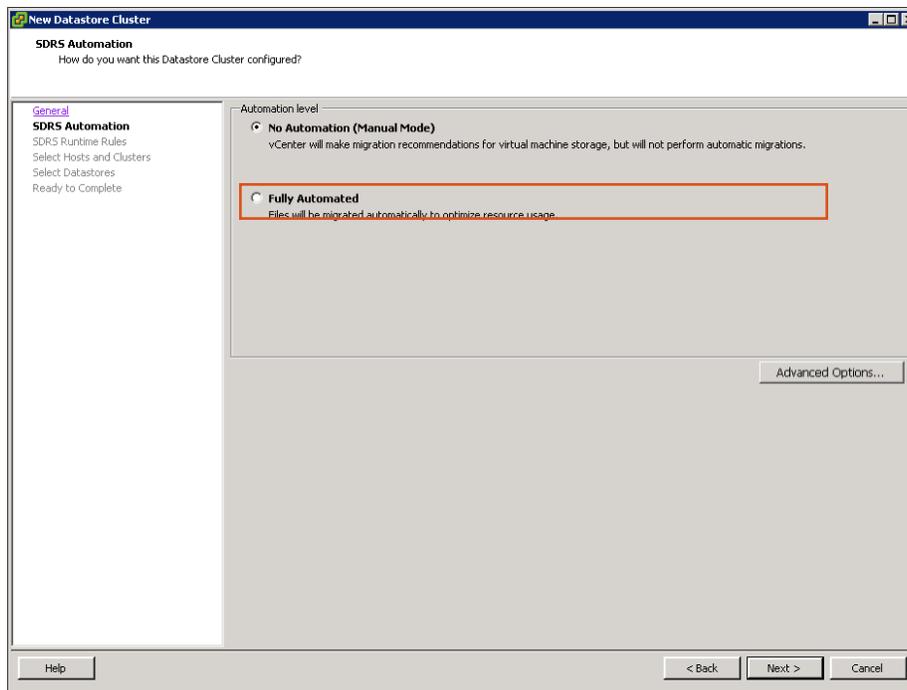
2. Select your **Datacenter** object. Right-click it and select **New Datastore Cluster**.



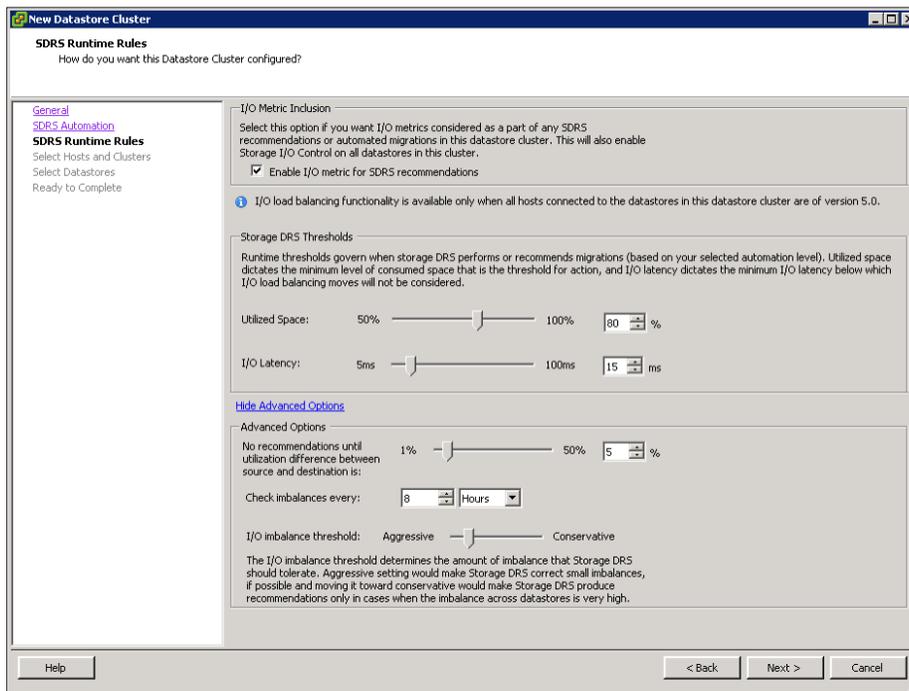
3. Give the **New Datastore Cluster** a name and click **Next**.



4. Select **No Automation (Manual Mode)** and click **Next**. (**Manual Mode** means that Storage DRS will only make recommendations and that the user must apply these. **Fully Automated** means that Storage DRS will make recommendations and apply these directly by migration virtual machines or virtual disks to the proposed destination datastore.)

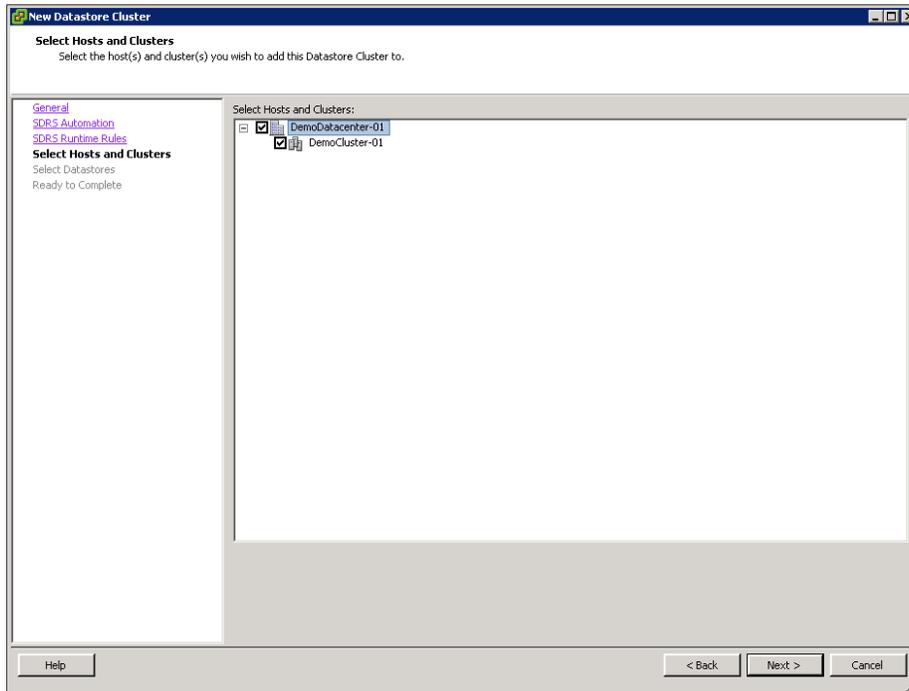


- Click **Show Advanced Options**. In the top part of the screen, the threshold for both **Utilized Space** and **I/O Latency** are shown. Storage DRS will make recommendations only when either of the two is exceeded. At the bottom of the screen, you will see the utilization difference, invocation period and the imbalance threshold. The utilization difference is the minimal difference between the source and the destination. Based on this value, Storage DRS will filter out those datastores whose utilization difference is below the given threshold during the selection of a destination. The default is set to 5%. The aggressiveness factor determines the amount of I/O imbalance Storage DRS should tolerate. The invocation period, 8 hours by default, determines how often Storage DRS will evaluate the environment and possibly generate recommendations.

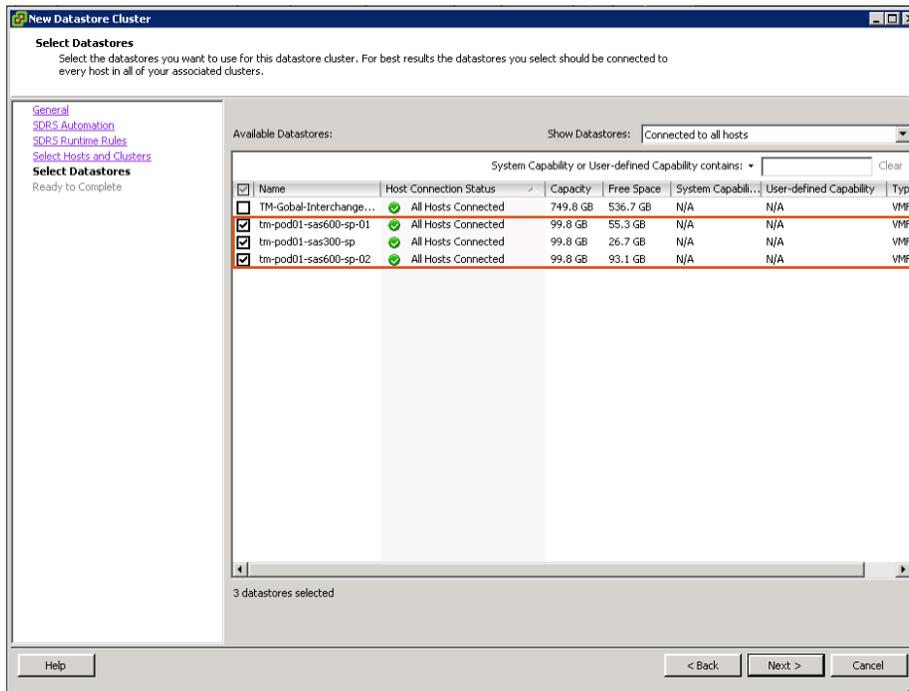


- Leave all the settings to default and click **Next**. Storage DRS enables Storage I/O Control automatically when I/O metric is enabled.

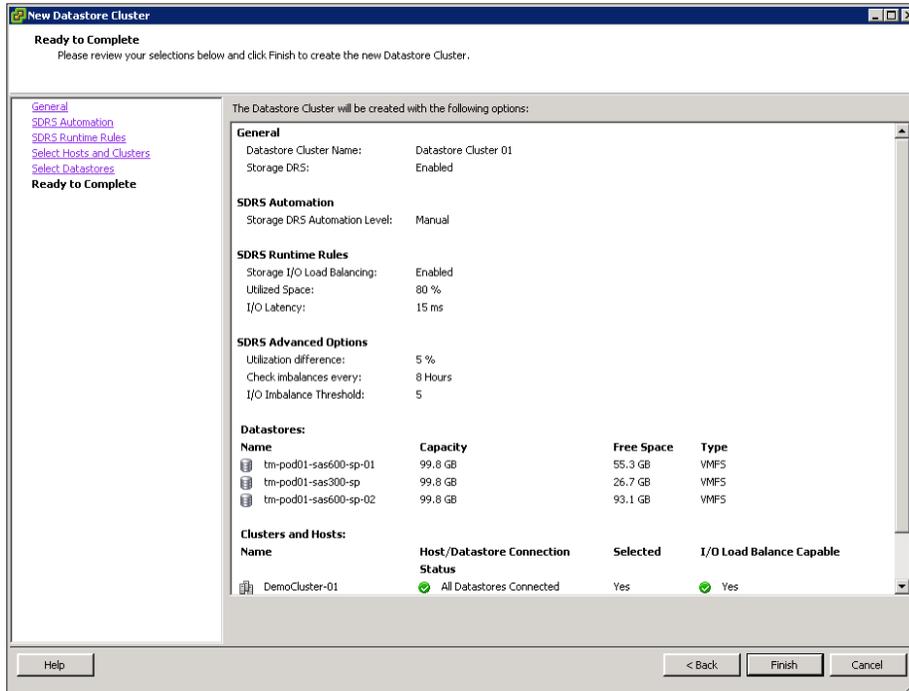
7. Select the cluster to which you want to add this datastore cluster.



8. Select the datastores that should be part of this datastore cluster. We are using datastores that already contain virtual machines. Creating a datastore cluster is a nondisruptive task and can be done if needed during production hours.



- Review your selections. Ensure that all hosts are connected to the datastores in the datastore cluster. Click **Finish**.



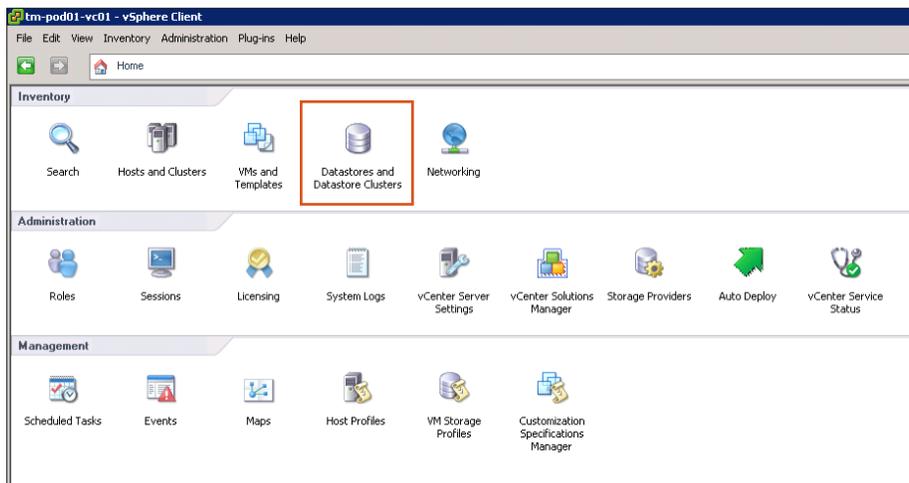
- The datastore cluster will now be created and a new object should appear on the **Datastores and Datastore Clusters** view. This object should contain the selected datastores.

This completes step 1 of **Testing vSphere Storage DRS**.

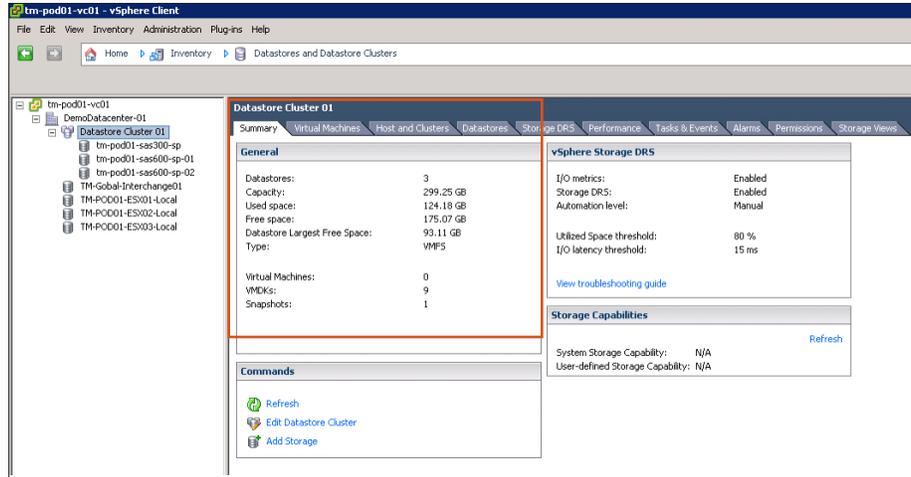
Exploring Your Datastore Cluster

In this section, we will show some of the important tabs of the **Datastore Cluster** object. These can be used for monitoring and managing your datastore cluster.

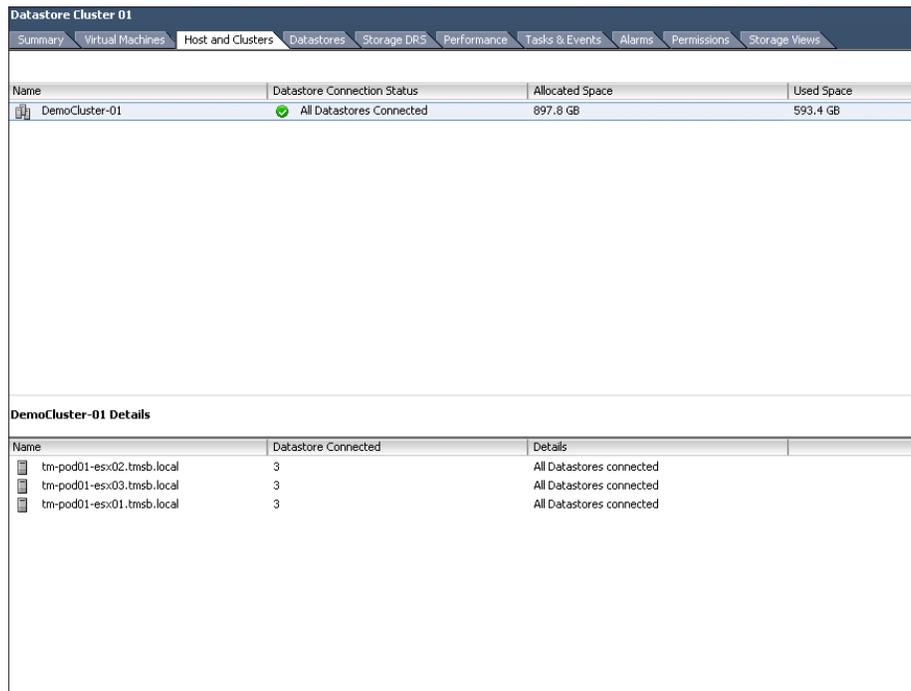
- From the **vCenter Home** view, select **Datastores and Datastore Clusters**.



2. Select your newly created **Datastore Cluster**. The **Summary** tab will show the number of datastores that are part of the datastore cluster and details regarding capacity and how Storage DRS was configured.



3. The **Hosts and Clusters** tab shows the datastore cluster and the health of the datastore cluster. It also shows, when the datastore cluster is selected, which hosts are connected and how many datastores are connected from this datastore cluster to the host.



- The **Datstores** tab shows all connected datstores and their characteristics. Storage I/O Control is enabled on these datstores.

Identification	Status	Host Connection Status	Device	Drive Type	Capacity	Free	Type	Last Update
tm-pod01-sas300...	Normal	All Hosts Connected	naa.6006016091...	Non-SSD	99.75 GB	26.69 GB	N/A	6/9/2011 11:16:56 PM
tm-pod01-sas600...	Normal	All Hosts Connected	naa.6006016091...	Non-SSD	99.75 GB	55.27 GB	N/A	6/9/2011 12:36:42 PM
tm-pod01-sas600...	Normal	All Hosts Connected	naa.6006016091...	Non-SSD	99.75 GB	93.11 GB	N/A	6/9/2011 12:36:42 PM

- The **Storage DRS** tab is one of the main tabs in this view. If there are any recommendations, they will be displayed in this tab. It is also possible to manually run Storage DRS by clicking **Run Storage DRS**.

Datstore Cluster Properties

Migration Automation Level: **Manual**

Utilized Space Threshold: **80%**

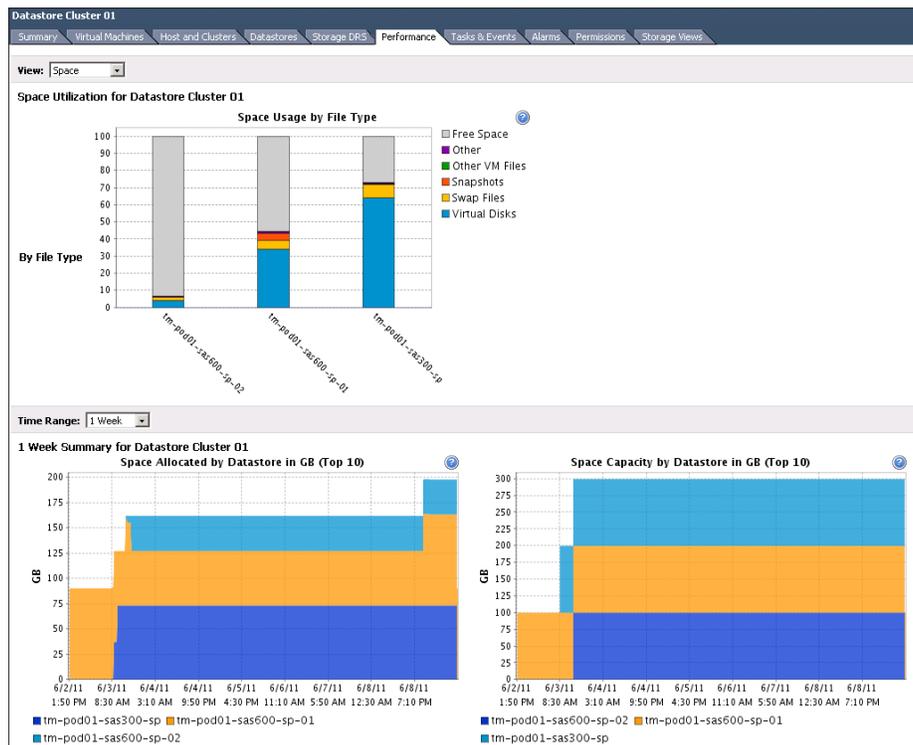
I/O Latency Threshold: **15 ms**

Storage DRS Recommendations

Apply	Recommendation	Reason	Space Util.	Space Util.	Space Util.	Space Util.	I/O Late...	I/O Late...
[Empty table body]								

Override Storage DRS recommendations Apply Recommendations

- The **Performance** tab shows the current and trending space utilization or performance statistics when selected from the pull-down list.



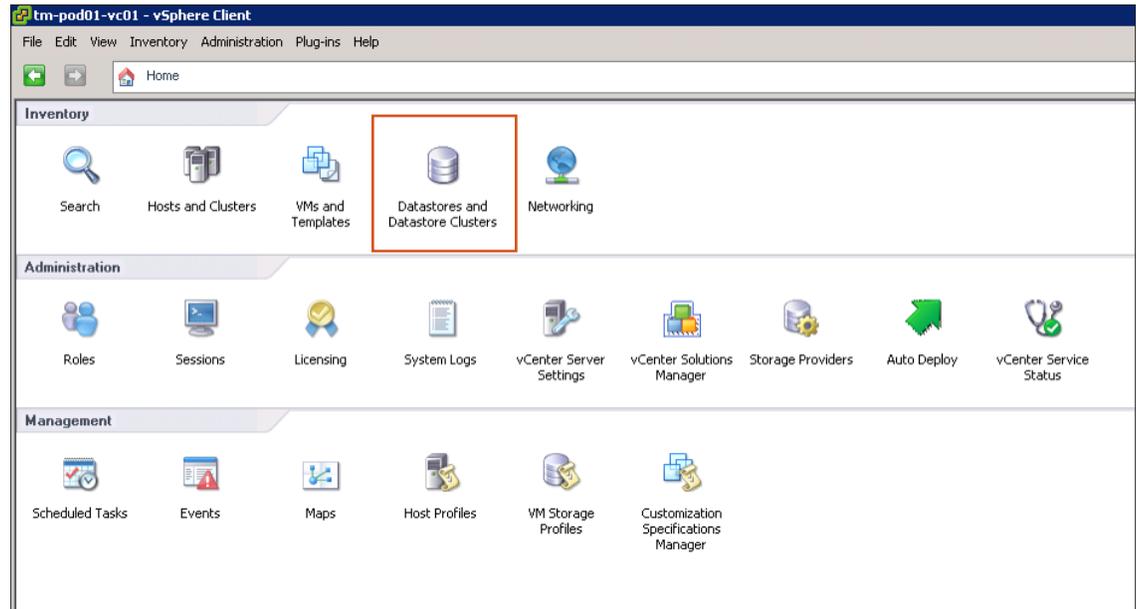
- The **Tasks & Events** tab shows all recent tasks and events. This tab is very useful for troubleshooting purposes and to validate the successful completion of tasks.

This completes step 2 of **Testing Storage DRS**.

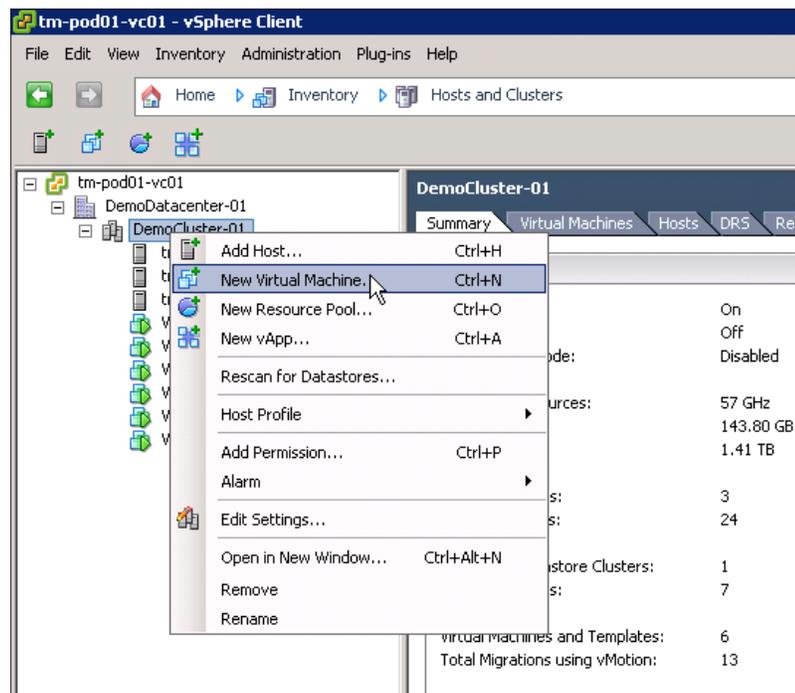
Provisioning a New Virtual Machine

In this section, we will create a new virtual machine and provision it to the newly created datastore cluster. Storage DRS will place the virtual machine, based on the current disk space utilization and I/O latency.

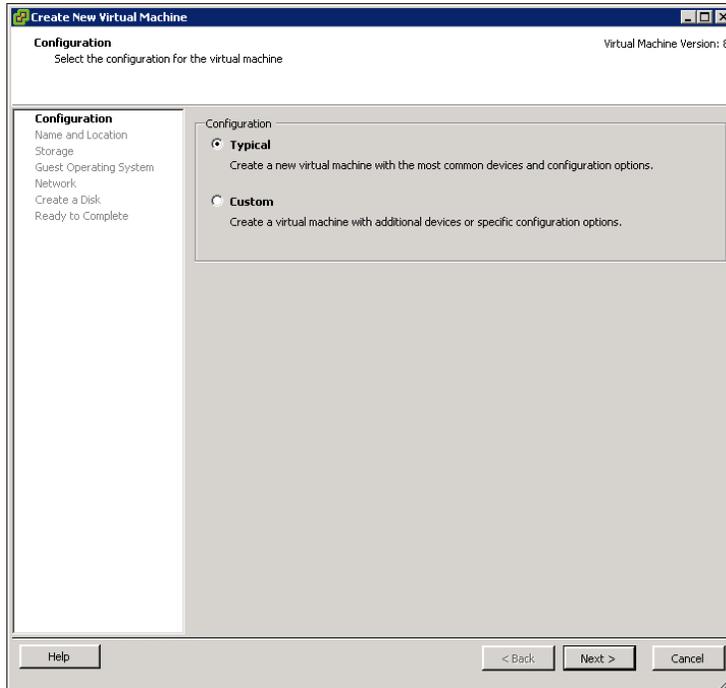
1. From the **vCenter Home** view, select **Hosts and Clusters**.



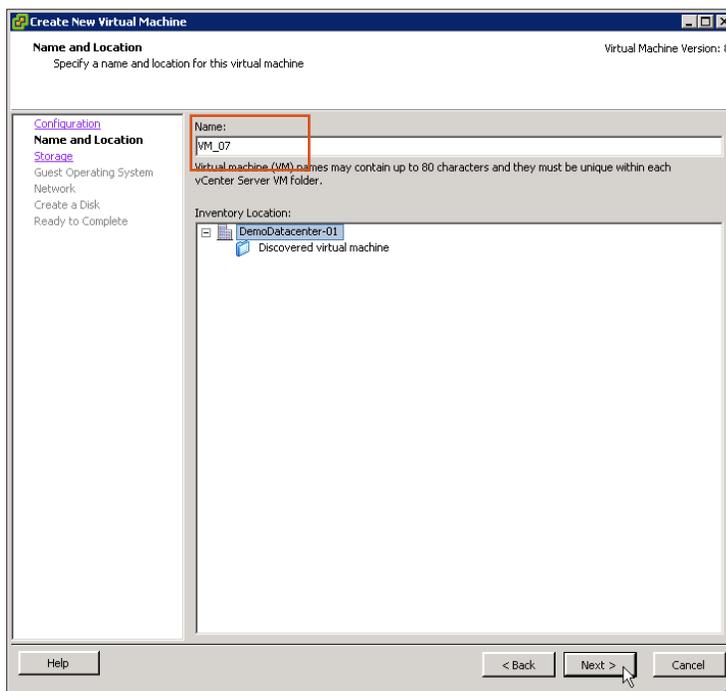
2. Right-click your **Cluster** object and click **New Virtual Machine**.



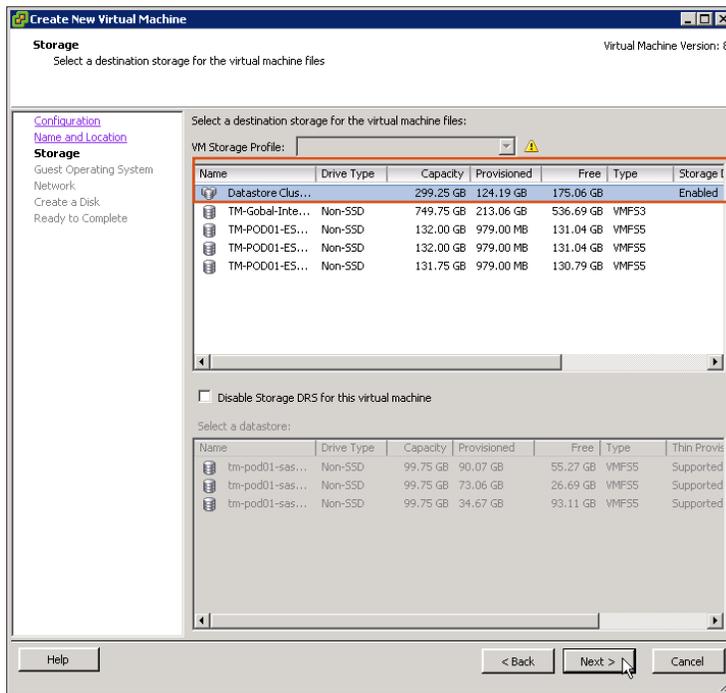
3. Select **Typical** and click **Next**.



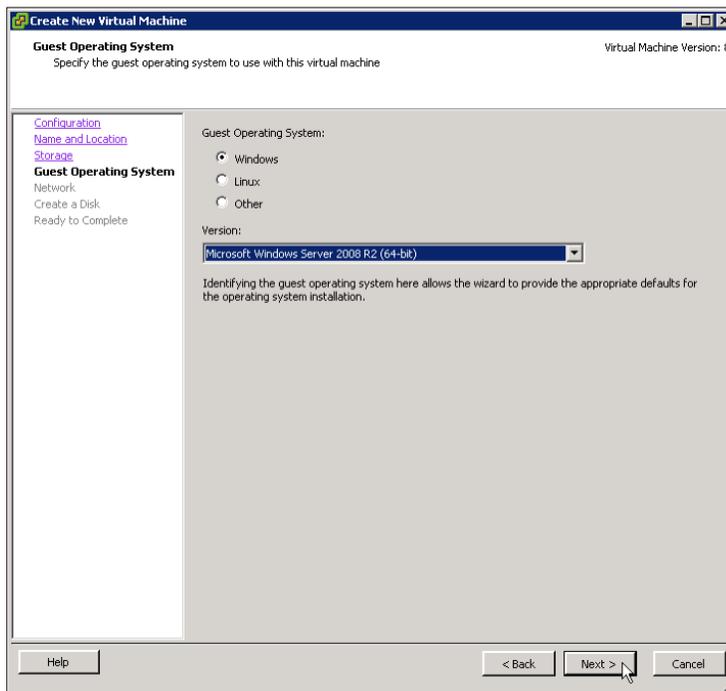
4. Give the virtual machine a unique name and click **Next**.



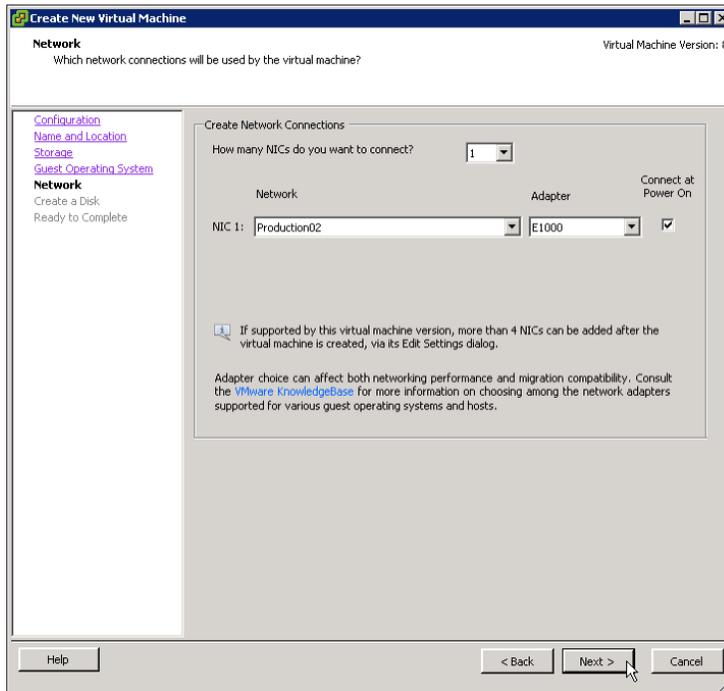
5. Select the datastore cluster where this virtual machine must be stored and click **Next**.



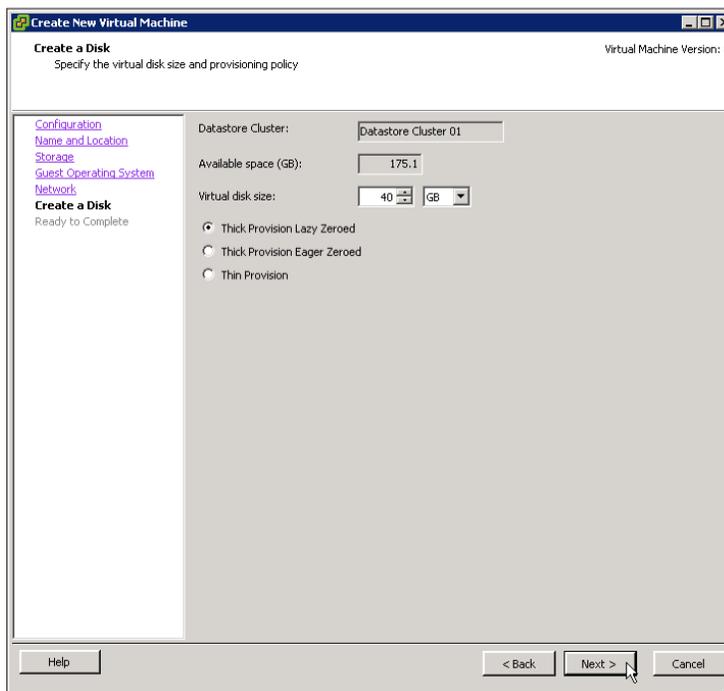
6. Select an operating system. In our example, we use Windows 2008, 64-bit. Click **Next**.



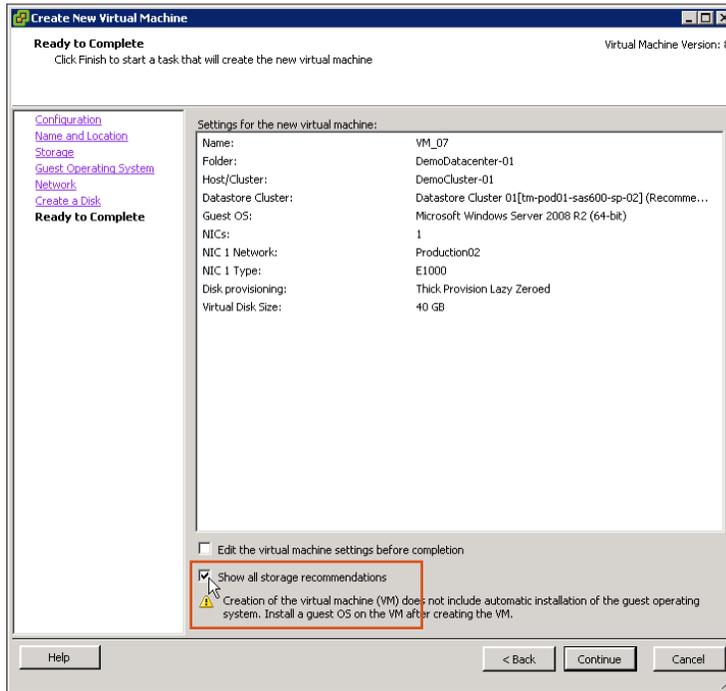
7. Select the correct port group and click **Next**.



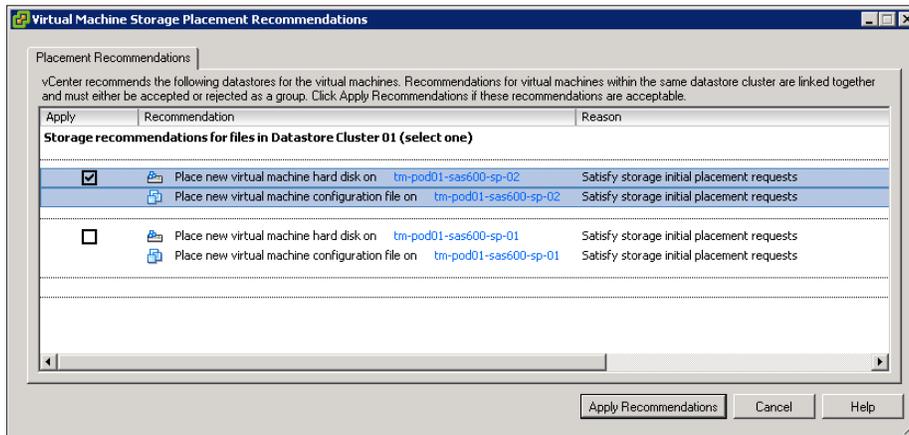
8. Depending on the available disk space, it might be necessary to decrease the size. In most cases, the default setting should be fine. Click **Next**.



9. Select **Show all storage recommendations** to see which datastore is recommended by Storage DRS as the destination for this virtual machine. Click **Continue**.



10. Storage DRS, when possible, makes several recommendations, enabling you to manually select a different datastore. We will use the recommended datastore by clicking **Apply Recommendations**.



11. The virtual machine will now be created.

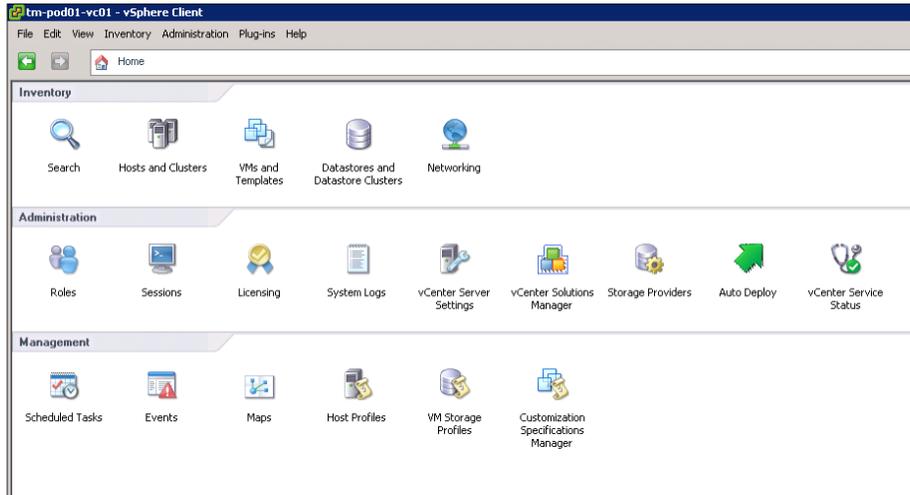
Recent Tasks								
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Apply Storage DRS recommend...		Completed		root	tm-pod01-vc01	6/9/2011 2:30:25 PM	6/9/2011 2:30:25 PM	6/9/2011 2:30:32 PM

This completes step 3 of **Testing Storage DRS**. If the used datastores were newly created, we recommend going through step 3 multiple times to create multiple virtual machines. In addition, we recommended installing an operating system to enable the possibility of creating load and also to ensure that disk space is allocated to the VMDK.

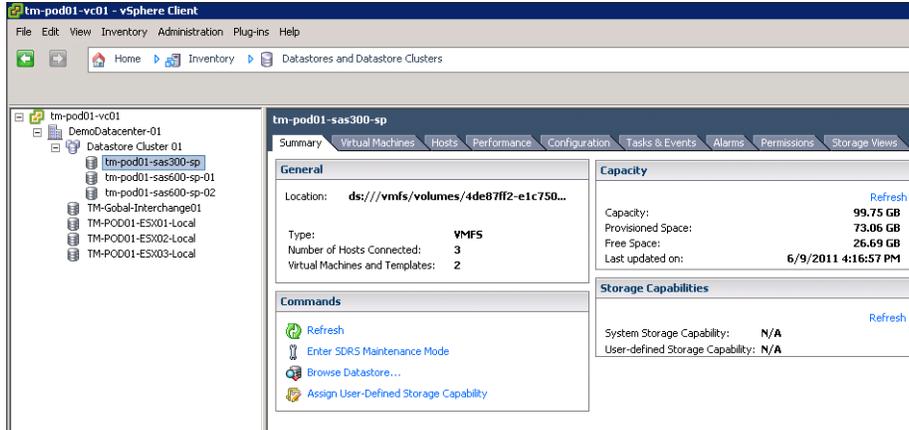
Space Balancing

In this section, we will create an imbalance from a disk space perspective, to see what recommendations Storage DRS will make. To complete this step, it is required to have multiple virtual machines stored on your datastore cluster.

1. From the **vCenter Home** view, select **Datastores and Datastore Clusters**.

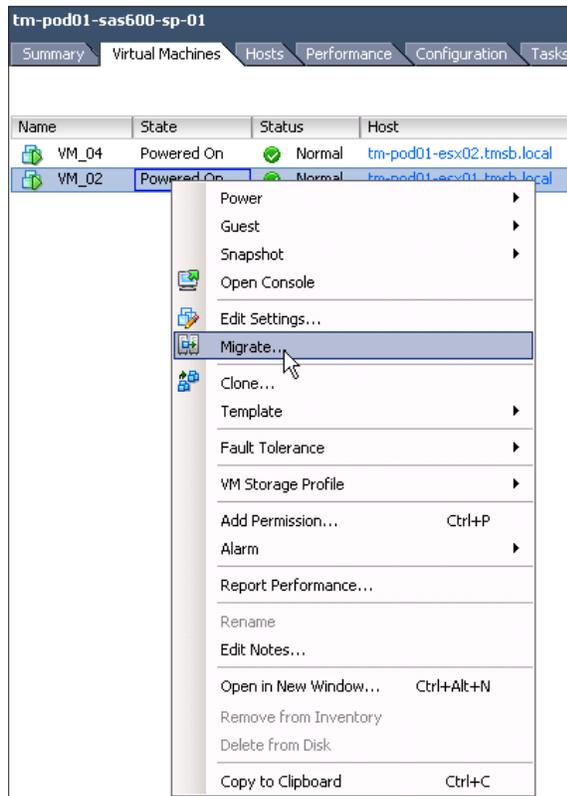


2. Select one of the datastores in your datastore cluster. Click the virtual machine tab and find the virtual machine with the most **Used Space**.

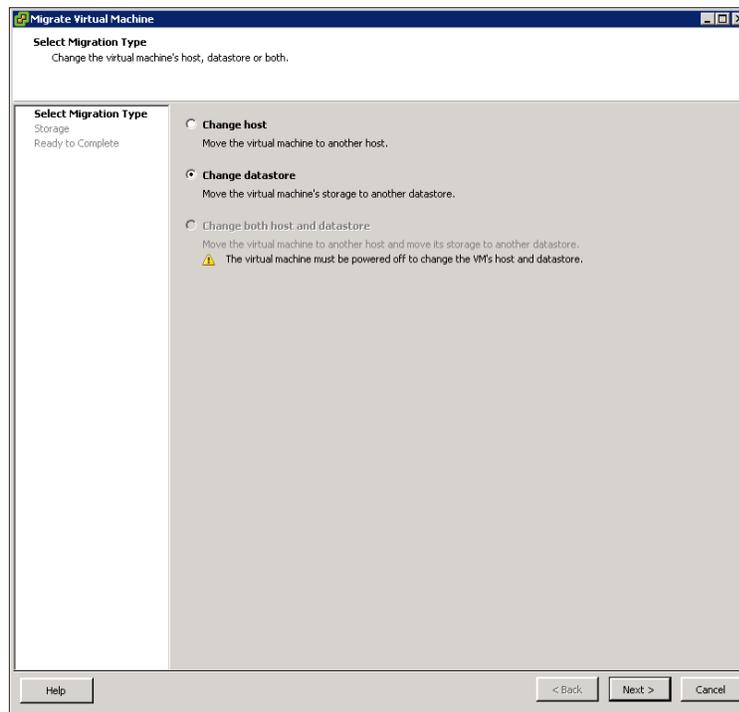


3. Migrate a virtual machine to the datastore that has the least amount of free space. Find a combination that will exceed the configured 80% space utilization.

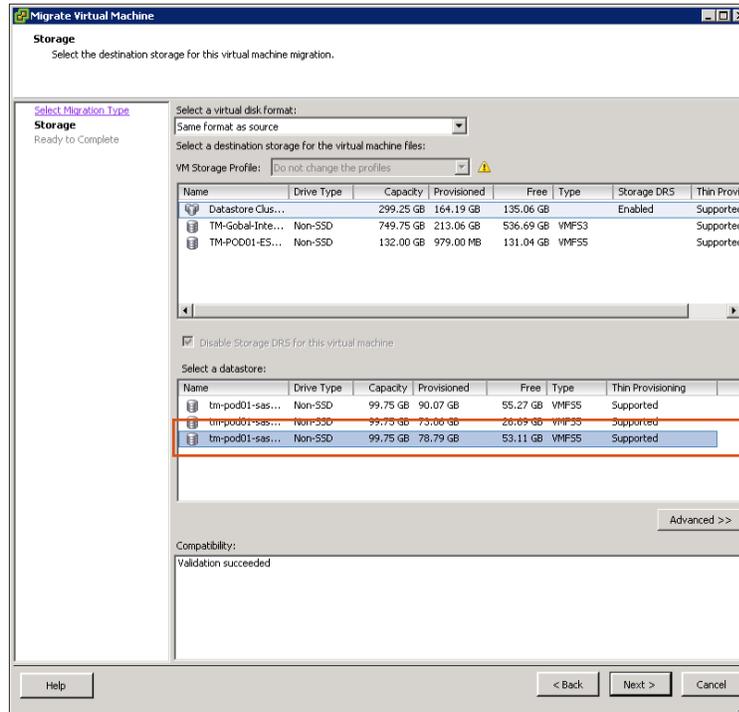
- a. Right-click the virtual machine and click **Migrate**.



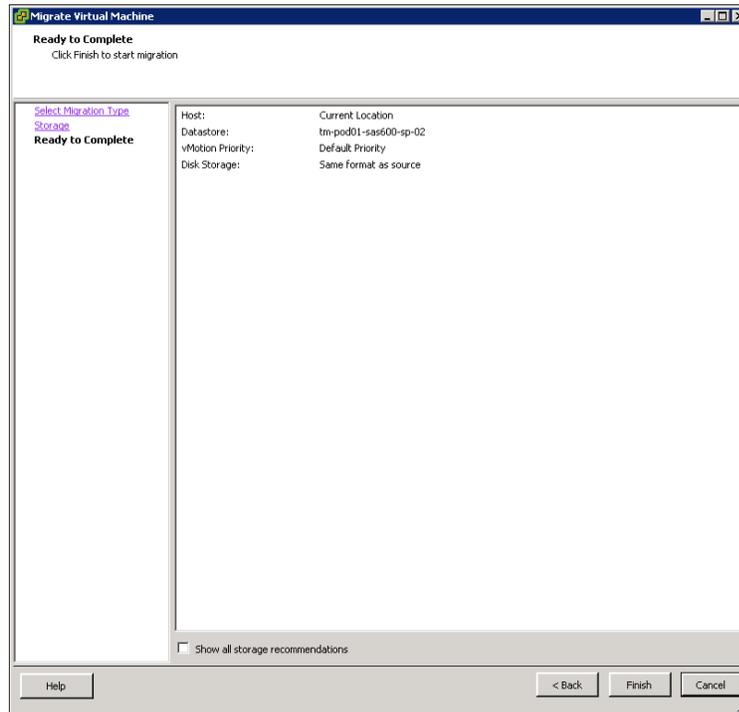
- b. Select **Change Datastore** and click **Next**.



- c. In the bottom section of the window, find the datastore that will exceed the utilization threshold after the migration. Select it and click **Next**.



- d. Review the selections and click **Finish**.



- e. The virtual machine is now migrated live to the selected datastore. Validate that the migration has been successfully completed.

Recent Tasks								
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Apply Storage DRS recommend...		Completed		root	tm-pod01-vc01	6/9/2011 5:03:25 PM	6/9/2011 5:03:25 PM	6/9/2011 5:09:50 PM

4. If the migration has been successfully completed, select your datastore cluster and click the **Storage DRS** tab.

The screenshot shows the 'Storage DRS' configuration page for 'Datastore Cluster 01'. The 'View' section has tabs for 'Recommendations', 'Faults', and 'History'. The 'Storage DRS Recommendations' table is empty. A red box highlights the 'Run Storage DRS' button in the top right corner.

5. Click **Run Storage DRS** to manually start the process. Storage DRS will now check whether any of the thresholds have been exceeded (space and I/O) and will make a recommendation when it is possible to solve the imbalance. When Storage DRS is set to **Fully Automated**, it will automatically solve the imbalance.

The screenshot shows the 'Storage DRS Recommendations' table populated with three entries. Each entry includes a recommendation icon, a description of the migration, the reason (Balance datastore space usage), and various utilization metrics. A red box highlights the 'Apply Recommendations' button in the bottom right corner.

6. Click **Apply Recommendations** to solve the imbalance. Using Storage vMotion, Storage DRS will now migrate the virtual machines that were recommended to be migrated. As can be seen in the following screenshot, Storage DRS will make multiple recommendations to solve the imbalance, if required.

Recent Tasks								
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Execute Storage vMotion for St...	VM_07	In Progress		System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	
Execute Storage vMotion for St...	VM_06	40%	Migrati...	System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM
Execute Storage vMotion for St...	VM_05	29%	Migrati...	System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM

7. Validate that the migration has been successfully completed.

Recent Tasks								
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Execute Storage vMotion for St...	VM_07	Completed		System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:25:33 PM
Execute Storage vMotion for St...	VM_06	Completed		System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:25:28 PM
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM
Execute Storage vMotion for St...	VM_05	Completed		System	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:25:25 PM
Apply a Storage DRS recomme...		Completed		root	tm-pod01-vc01	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM	6/9/2011 5:24:13 PM

You have now successfully completed step 4 and the exercise **Testing Storage DRS**.

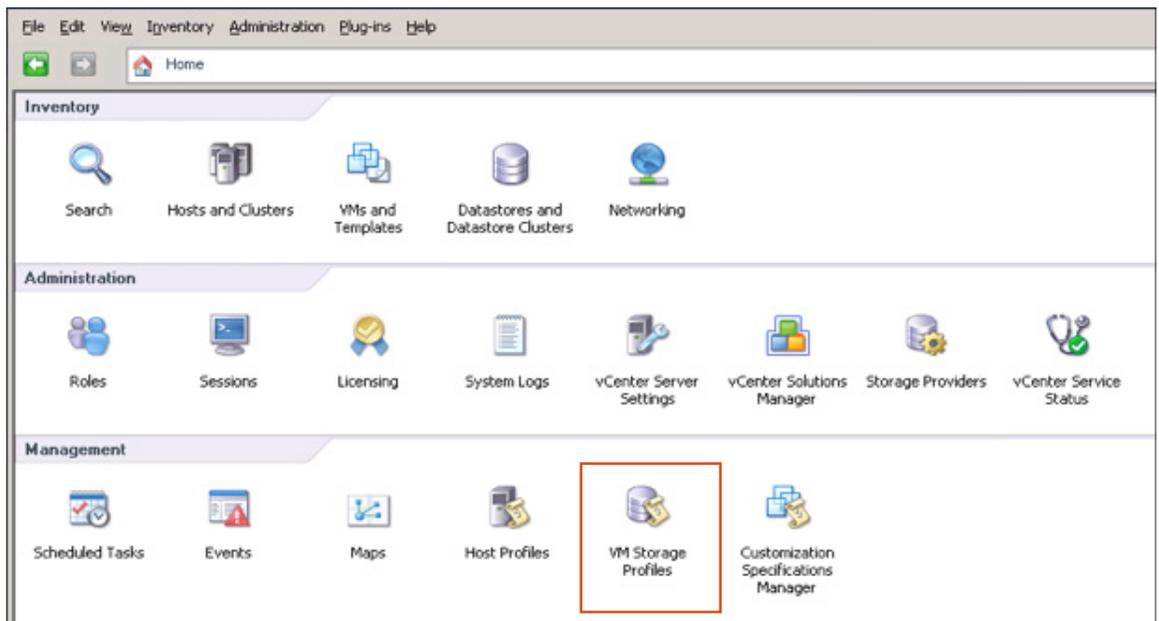
Using Profile-Driven Storage

Introduction

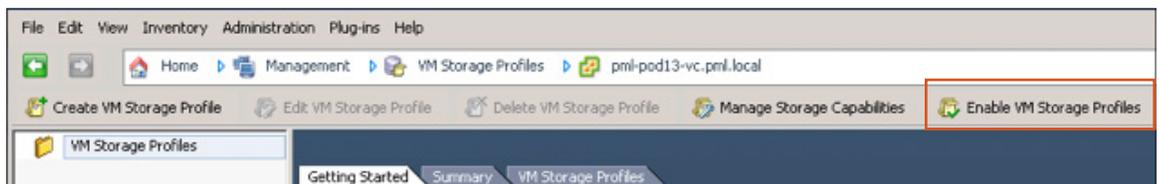
Large-scale storage configurations are difficult to manage. It is difficult for administrators to correctly identify storage characteristics in vSphere, so it is difficult for them to know whether virtual machines are being deployed or migrated to the correct datastore. Virtual machine storage profiles predefine classes of virtual machine storage. This reduces placement errors during provisioning, migration and cloning by monitoring virtual machine storage placement against predefined virtual machine storage profiles.

Create a Virtual Machine Storage Profile with a User-Defined Storage Capability

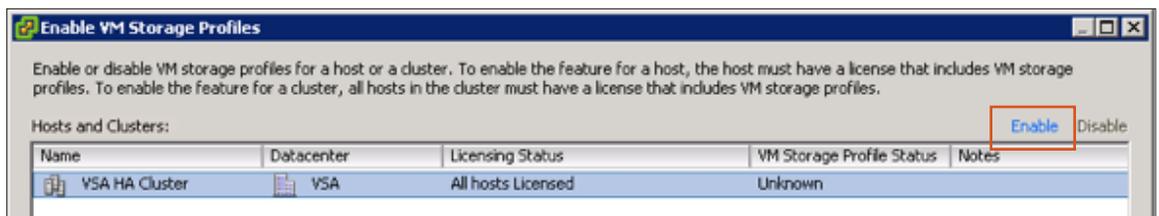
From the **vCenter Home** view, select **VM Storage Profiles**.



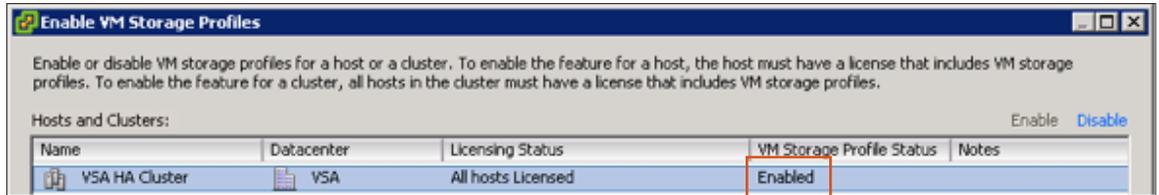
You must now click the **Enable VM Storage Profiles** button, located in the toolbar underneath the navigation bar at the top right of the window:



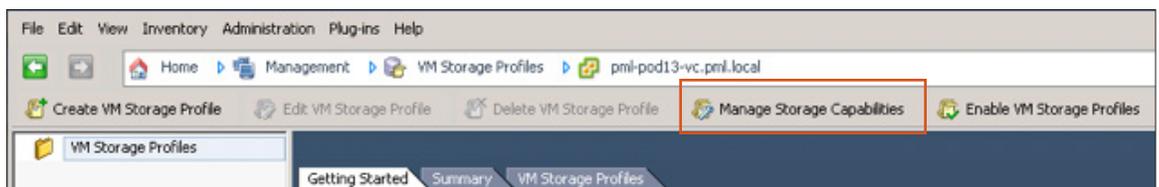
If your hosts are in a cluster, you can enable the **VM Storage Profiles** cluster-wide. If your hosts are not in a cluster, you must enable them individually. In this example, the hosts are in a cluster, so they can be enabled cluster-wide. Single-click the **Enable** link.



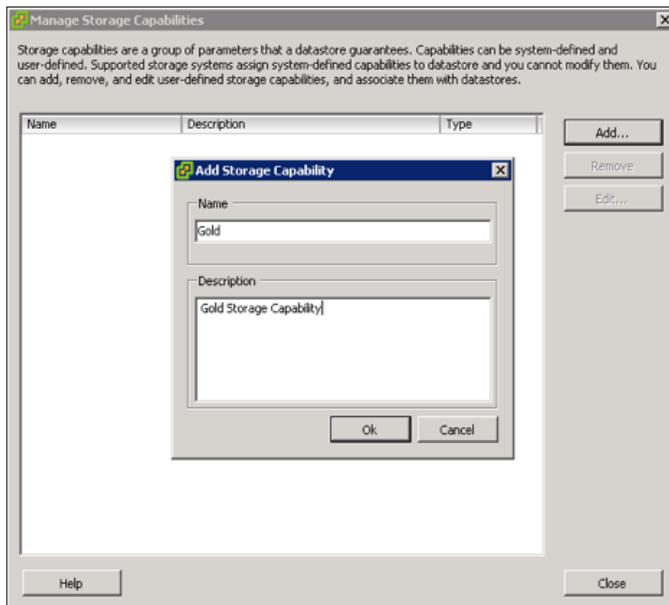
After **VM Storage Profiles** has been enabled, the **VM Storage Profile** will change to **Enabled**:



Close the **Enable VM Storage Profiles** window by clicking the **Close** button, located in the lower right-hand corner. You must now click the **Manage Storage Capabilities** button, located in the toolbar underneath the navigation bar at the top of the window:



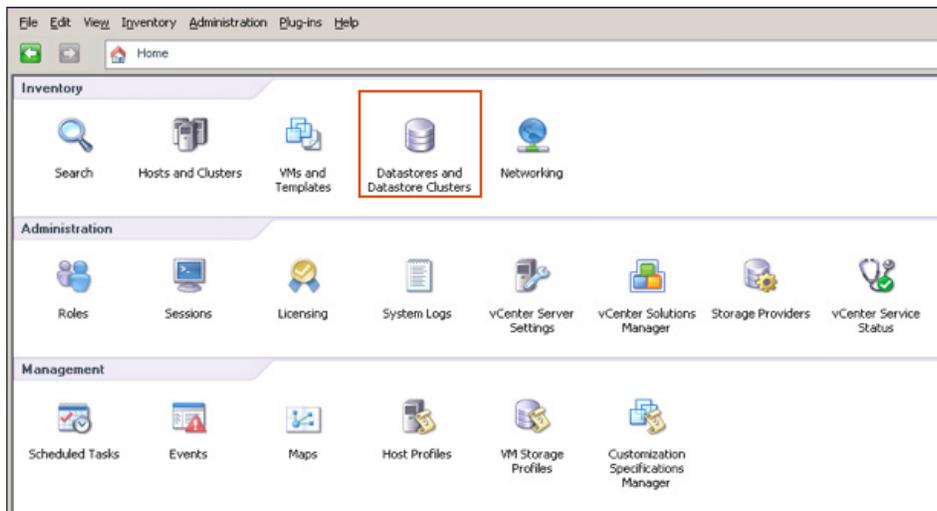
The **Manage Storage Capabilities** window will appear. Click the **Add** button. Give your storage capability a name; in this example, **Gold**. You can provide an optional description.



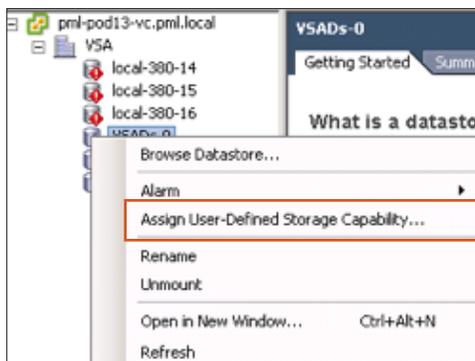
Click **OK** and then **Close**. This completes the first step of the **Storage-Driven Profiles Evaluation**.

Assign a Storage Capability to a Datastore

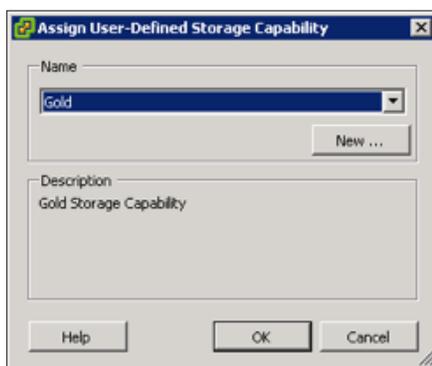
This datastore should not contain any virtual machines. This is not a requirement. From the **vCenter Home** view, select **Datstores and Datastore Clusters**.



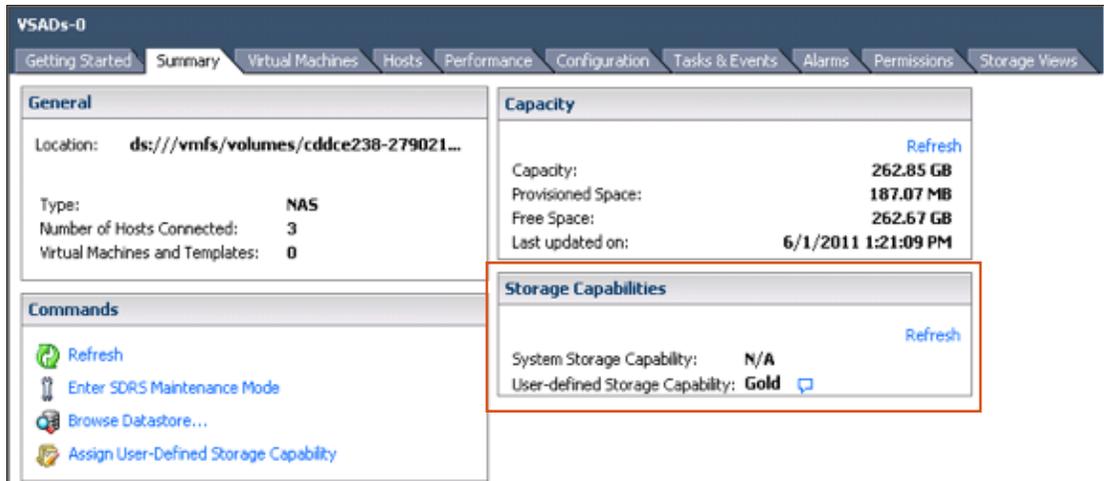
From the vCenter **Inventory** on the left-hand side, right-click the datastore to which you want to assign the storage capability. Select **Assign User-Defined Storage Capability...** :



From the drop-down menu, choose the **User-Defined Storage Capability** defined in the first step; in this example, **Gold**:

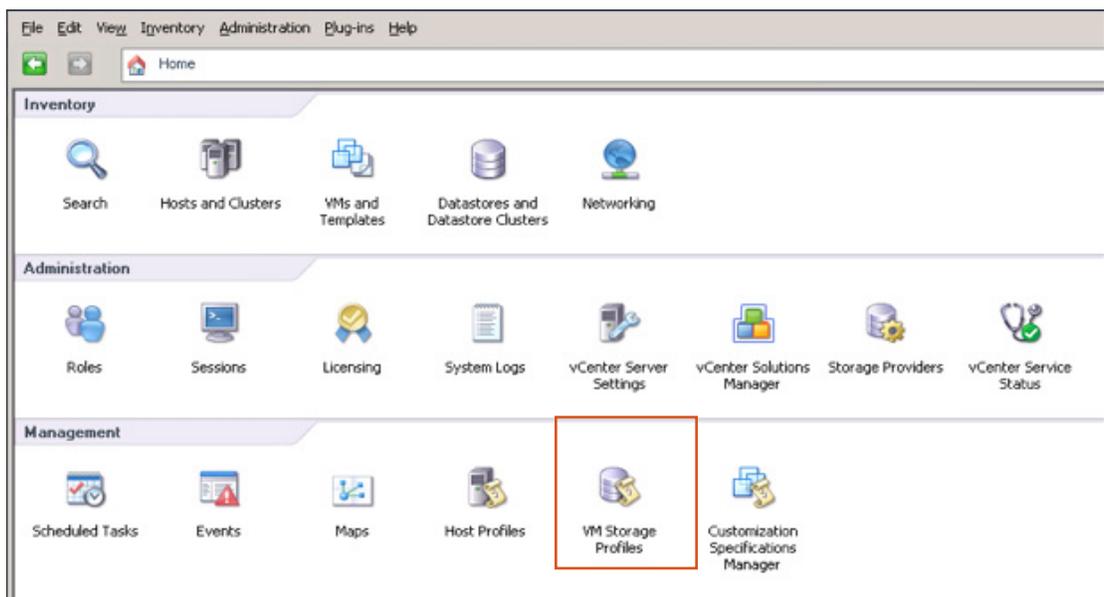


Click **OK**. Stay in the **Datastores and Datastore Clusters** view and select the **Summary** tab for the datastore that was assigned the storage capability. You should observe that the **Gold** capability is now visible:

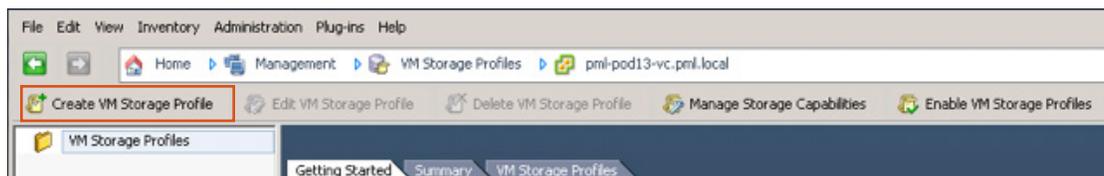


Create a Virtual Machine Storage Profile Containing the User-Defined Storage Profile

From the **vCenter Home** view, once again select **VM Storage Profiles**.



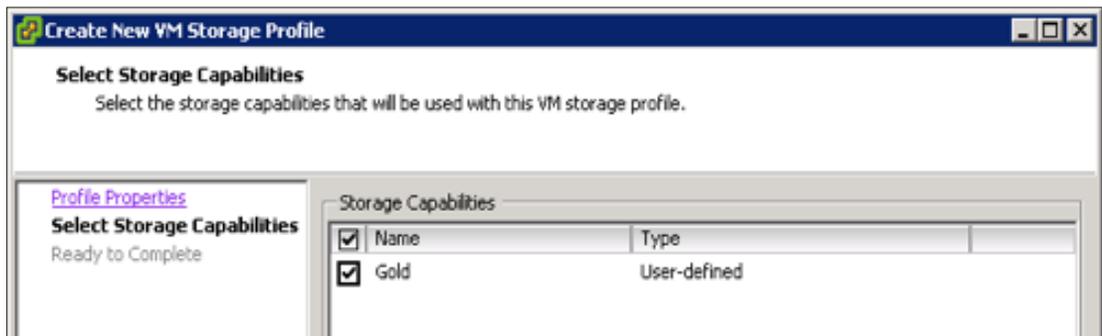
Click the **Create VM Storage Profile** button, located in the toolbar underneath the navigation bar at the top left of the window:



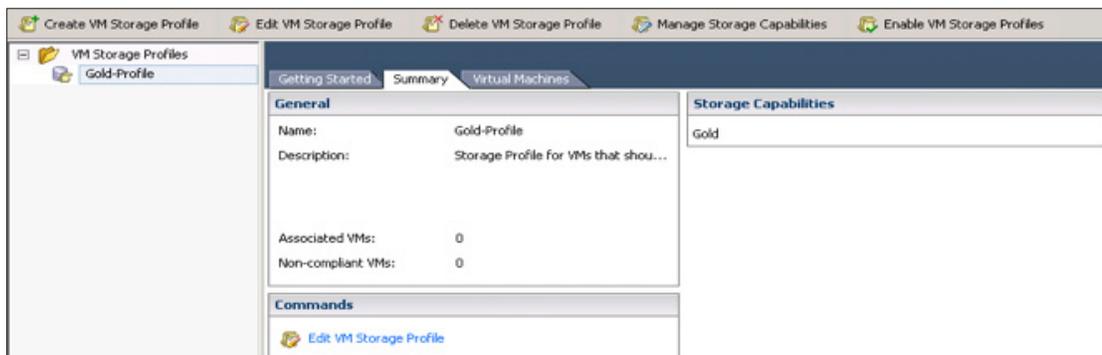
This launches the **Create New VM Storage Profile** wizard. The first step is to provide a name and an optional description to the profile; in this example, **Gold-Profile**:



Click **Next** to proceed to the **Select Storage Capabilities** window. At this point, there is only a single user-defined storage capability, **Gold**. Check the adjacent box to select it:



Click **Next**. Then click **Finish** to complete the creation of the VM Storage Profile. In the **Inventory** panel, select the newly created **Gold-Profile** and **Summary** tab to view the details of the VM Storage Profile.

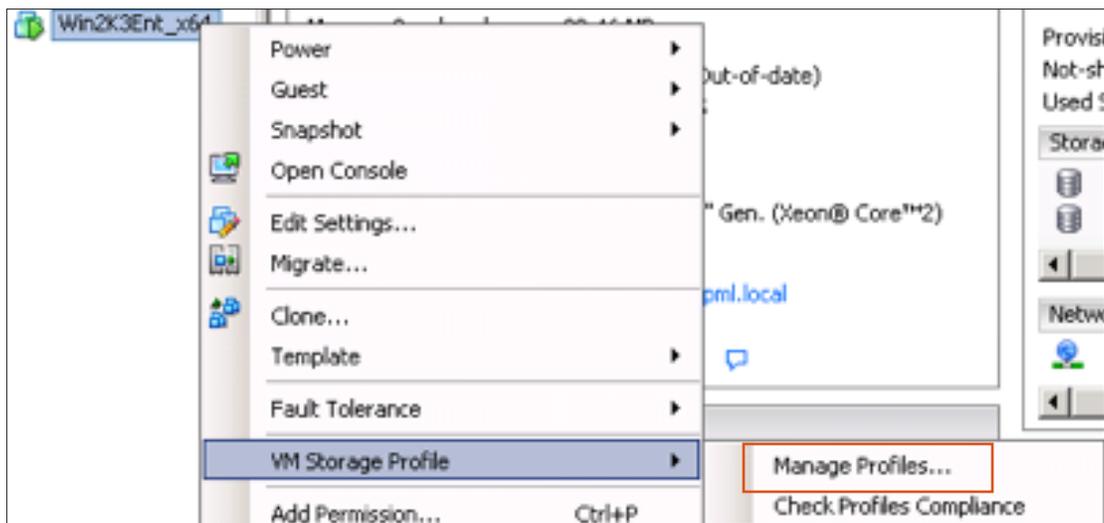


Assign a VM Storage Profile to a Virtual Machine

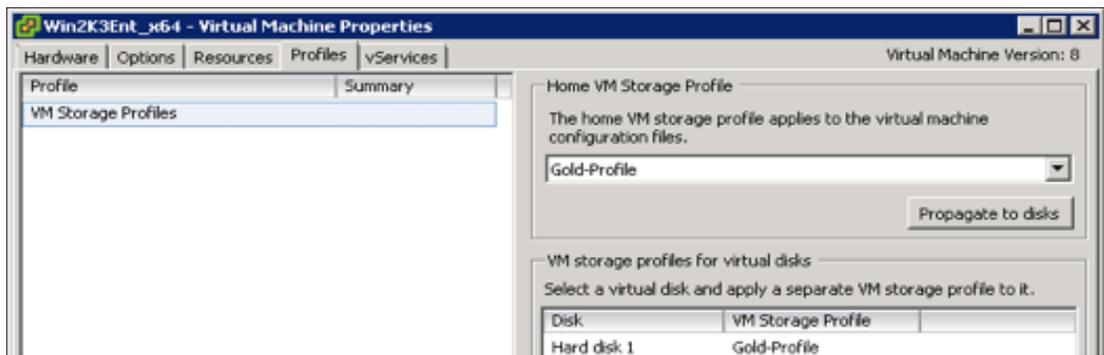
There is an assumption made here that there is a virtual machine available in your environment. In this evaluation guide, we will be associating a VM Storage Profile with an already existing virtual machine. However, VM Storage Profiles can also be associated with a virtual machine during its initial creation, meaning that the virtual machine can be placed on “compatible” storage from the outset.

First, go to the **Hosts and Clusters** view. Identify a virtual machine that you want to use as part of the VM Storage Profiles evaluation. This virtual machine requires a disk but can have any guest OS installed. In this example, I have chosen a virtual machine running Windows 2008 R2 (64-bit).

Right-click the virtual machine and select **VM Storage Profile**. Then select **Manage Profiles**.



This opens a new **Profiles** tab in **Virtual Machine Properties**. From the drop-down list, select the profile that you created earlier. In this example, it is called Gold-Profile. You will also see a button called **Propagate to disks**. By default, the VM Storage Profile applies only to the virtual machine’s configuration files. You must use the **Propagate to disks** button to also include the virtual machine’s disks (VMDKs) in the profile. Click the **Propagate to disks** button. The hard disk(s) of the virtual machine are included in the profile.



Click **OK**.

Check Whether the Virtual Machine Is Running on “Compliant” Storage

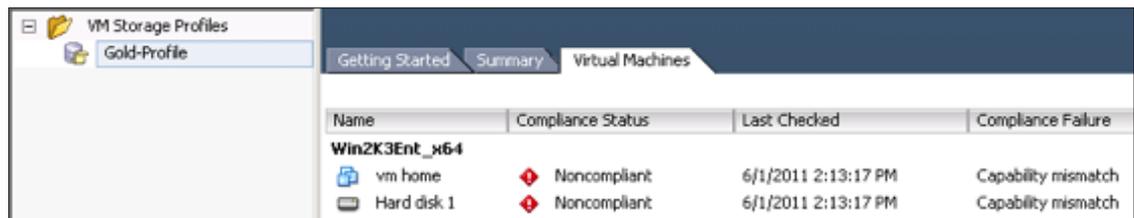
In the **VM Summary** tab, click the **Refresh** button in the **VM Storage Profiles** window.



Now we have a datastore with the user-defined storage capability called **Gold**. We also have a VM Storage Profile called **Gold-Profile** with the same capability, and now we have a virtual machine with that profile attached. However, because the virtual machine currently resides on a datastore *without* that storage capability, the virtual machine is deemed noncompliant; that is, it is not on a datastore with the necessary storage capabilities. The following is what is reported when we refresh the VM Storage Profiles window:



Further details about the reason for noncompliance can be found back in the **VM Storage Profiles** view. Select the **Gold-Profile** and then the **Virtual Machines** tab. Because we have only one virtual machine, the display will be short:

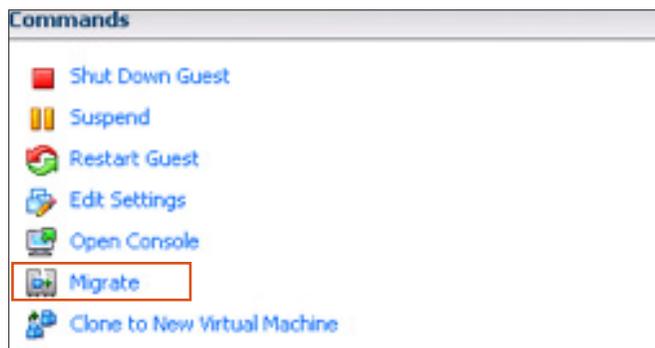


There are two entries displayed here. One entry is for the virtual machine’s configuration files and the other one is for the virtual machine’s hard disk.

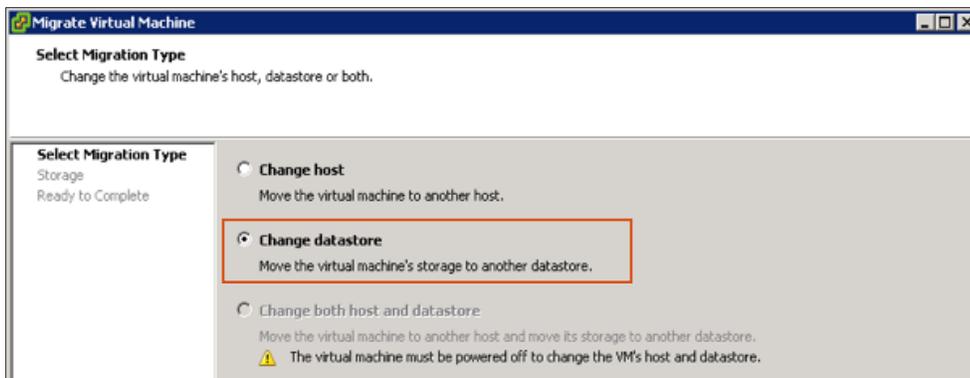
Bring a Virtual Machine into Compliance

To bring this virtual machine into VM Storage Profile compliance, you must migrate it to a datastore that has the correct storage capabilities. The easiest way to do this is via Storage vMotion, which will enable you to migrate a running virtual machine from one datastore to another.

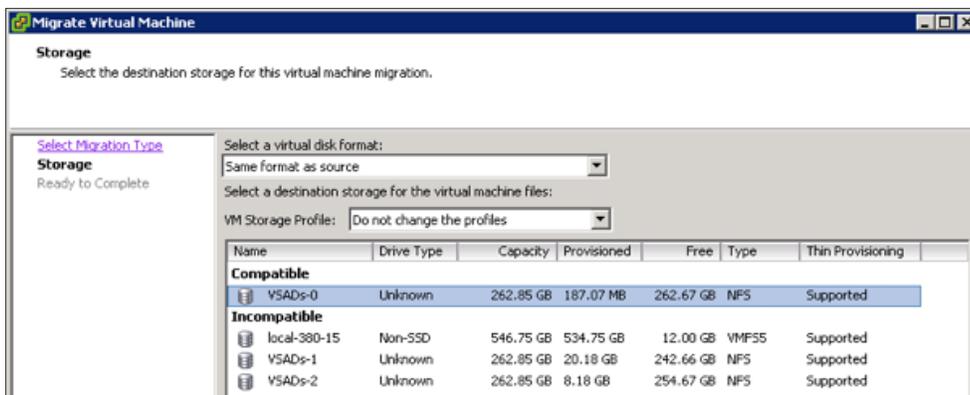
To initiate a Storage vMotion instance, select the virtual machine from the **Hosts and Clusters** view. In the **Summary** tab, select the **Migrate** option in the **Commands** window:



When the **Select Migration Type** window appears, choose the option to **Change datastore**:



VM Storage Profiles are integrated into the migration wizard. On the next screen, select the destination storage. VM Storage Profiles are used to ensure that only those datastores that contain the storage capabilities as defined in the storage are presented as **Compatible**. In this example, only datastore VSADs-0 has the **Gold-Profile**:



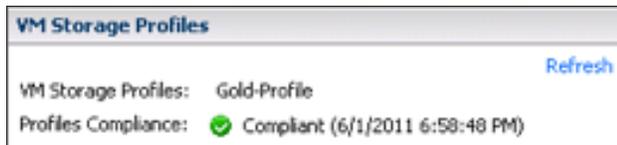
Because this virtual machine already has a VM Storage Profile called **Gold-Profile** associated with it, the **Do not change the profiles** option is identical to selecting Gold-Profile from the pull-down menu. This is also identical to the **Storage** window one would see during the initial creation of a virtual machine, so the correct **Compatible** storage can be chosen for the virtual machine right from the start.

By choosing the compatible datastore from the list, when our migration completes we will know that the virtual machine will reside on a datastore that has the same storage capabilities as those defined in the VM Storage Profile, that is, **Gold**.

Click **Next** and **Finish** to start the migration. Observe the status of the Storage vMotion via the **Recent Tasks** view:

Recent Tasks								
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Relocate virtual machi...	Win2K3Ent_x64	27%	Migrating t...	PML\Administ...	pn1-pod13-vc...	6/1/2011 2:50:03 PM	6/1/2011 2:50:03 PM	

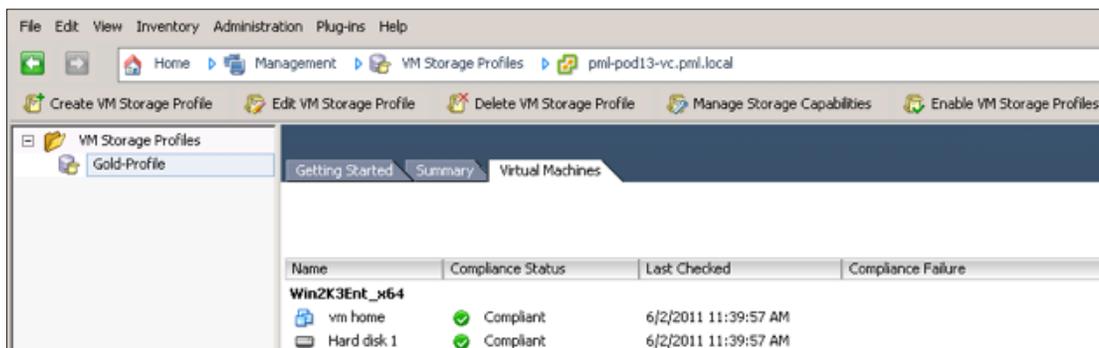
After the virtual machine has completed migrating to the new **Compliant** datastore, **Refresh** the **Storage Profiles** window of the virtual machine again to check whether it is in compliance. It should be compliant this time.



As a final step, again use the **VM Storage Profiles** view. As before, select the **Gold-Profile** and then the **Virtual Machines** tab to see the compliance state of your virtual machine. You will probably need to click the **Check Compliance Now** button, located in the upper right-hand corner of the screen:



After you have run the compliance check, the virtual machine's configuration files and disk should both be in compliance:



This completes the **Profile-Driven Storage** evaluation steps.

Evaluating Storage I/O Control

Introduction

Storage I/O Control enables cluster-wide control of disk resources, which prevents a single virtual machine from monopolizing all the I/O to a particular datastore. In this part of the evaluation guide, we will see how one can tune the IOPS that a particular virtual machine can generate to a shared datastore.

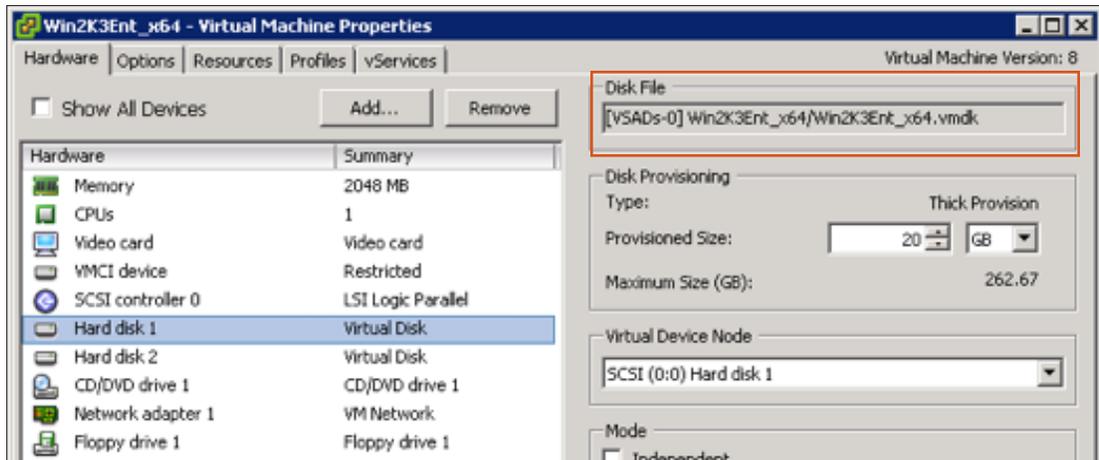
Priority is established using shares, although specific limits based on IOPS can also be implemented. In this part of the storage evaluation guide, we will examine the features of SIOC and how they can assist you in ensuring "fairness" across all your virtual machines from an I/O perspective.

Create a Virtual Machine on a Datastore

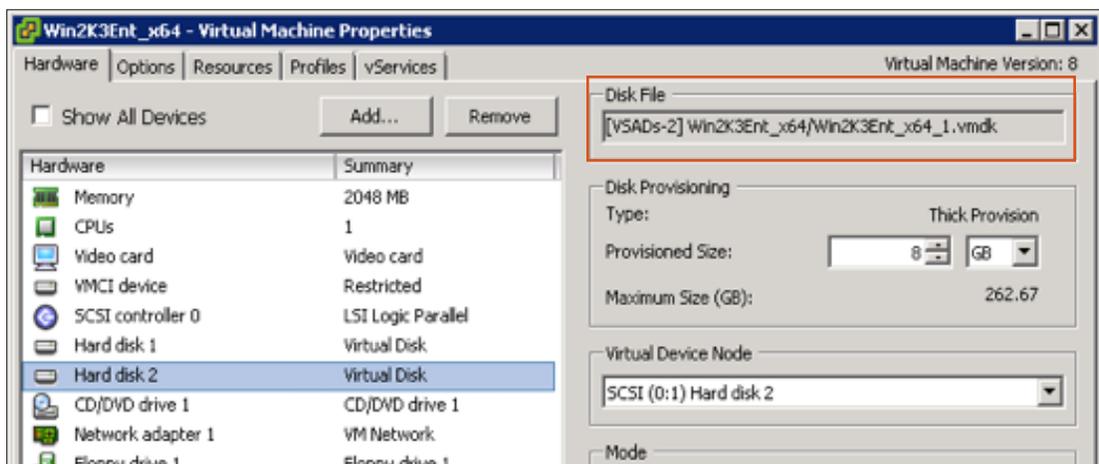
To look at the performance of the SIOC, we will deploy two virtual machines to the datastores in your environment. We must deploy the virtual machines on different ESXi 5.0 hosts. In this example, virtual machines running Microsoft Windows 2003 x64 as the guest OS are deployed. These virtual machines have two disks. One (the boot disk) is on one NFS datastore; the other (the data disk) is on another NFS datastore. The data disk of each virtual machine is placed on the same datastore.

The ability to use SIOC on NFS datastores is a new feature of vSphere 5.0. **iometer** (<http://www.iometer.org>) has also been installed onto the guest OS of each virtual machine, so that a certain amount of I/O load can be driven to the virtual machine's data disk.

Disk 1: Boot Disk – resides on first NFS datastore.

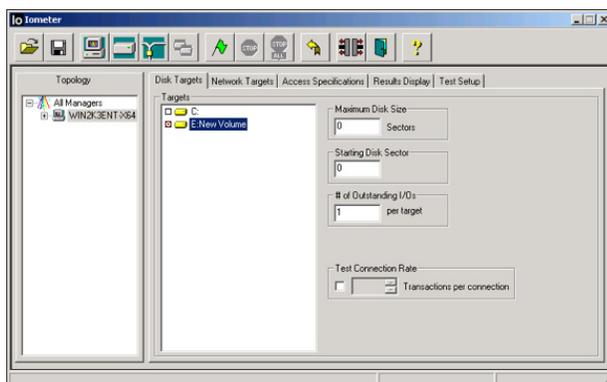


Disk 2: Data Disk – resides on second NFS datastore. This must be the same for both virtual machines.



Generate I/O to the Datastore from the Virtual Machine

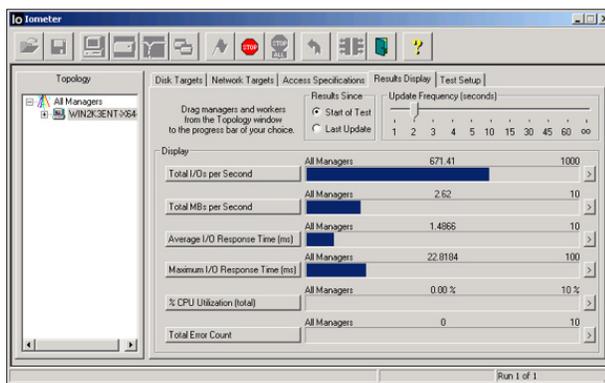
Power up both virtual machines, and open a console to each of them. On the desktop of each one, there is an **lometer** icon. We will use lometer to generate I/O to the shared datastore. Launch lometer. Select the virtual machine from the **Topology** view. Then select the **Disk Targets** tab. Finally, select the disk on the second NFS datastore. In this example, it is the **E: New Volume** drive.



Next, select the **Access Specifications** tab. This will show that we are doing **4K I/Os, 75% Read, 0% random** (sequential) operation.



Finally, select the **Results Display** tab. You should see that the **Update Frequency** is set to 2 seconds. Click the **green flag icon** to start the I/O. Save the results file to the default location. In the **Results Display** tab, you can now begin to see IOPS and latency information being updated every 2 seconds.



A better view can be seen by clicking the arrows [**>**] at the end of the **Iometer** display screen. Tune the display using the **Range** value in the lower left-hand corner.



Repeat these steps using the same configuration setup (4K I/Os, 75% Read, 0% Random) on the other virtual machine on the other ESXi server. Leave Iometer running on both virtual machines. I/O is now being generated to the same shared datastore from two virtual machines on two different ESXi hosts.

The I/O on the first virtual machine should start to gradually decrease because there is now additional contention on the shared datastore. This is normal.

Now we can start to look at the Storage I/O Control feature.

Enable Storage I/O Control

Now that both virtual machines, from different ESXi 5.0 hosts, are generating I/O to the same shared datastore, we can enable SIOC on that datastore. This will allow us to manage which virtual machine's I/O gets priority on the datastore.

In the VMware vSphere® Client™, select **Home**. In the **Inventory**, select **Datastores and Datastore Clusters**. Next, select the shared datastore to which the virtual machines are issuing I/O. In the **Configuration** tab, in the upper part of the display, you will see details on which ESXi 5.0 hosts are using the datastore. In this case, there are three ESXi hosts that have this NFS datastore mounted:

Name	Datastore	State	Status	% CPU	% Memory	Memory Size	CPU Count	NIC Count	Uptime
ts03-h380-15.pml.local	Mounted	Connected	Warning	0	10	24565.73 MB	2	4	12 days
ts03-h380-14.pml.local	Mounted	Connected	Normal	1	9	24565.73 MB	2	4	12 days
ts03-h380-16.pml.local	Mounted	Connected	Warning	0	10	24565.73 MB	2	4	12 days

In the lower half of the display, details regarding the actual configuration of the datastore are shown:

Datastore Details

VSADs-2
 Server: 10.20.196.119
 Folder: /exports/31bb7049-56d6-4a31-a788-f3d203b4b360

262.65 GB Capacity
 8.18 GB Used
 254.67 GB Free

Refresh Storage Capabilities
 System Storage Capability: N/A
 User-defined Storage Capability: N/A

Storage I/O Control
 Disabled

Storage I/O Control is currently disabled. To enable it, click on the **Properties** link, located to the right of the **Datastore Details** box. This will open the datastore **Properties** window. On the left side of the window is a checkbox for **Storage I/O Control**. Click **Enabled** and then **Close** the **Properties** box.

VSADs-2 Properties

Storage I/O Control
 Enabled

Edit Congestion Threshold
 Congestion Threshold: 30 ms [Reset]

Setting improper congestion threshold values might be detrimental to the performance of the datastore.

[Close] [Help]

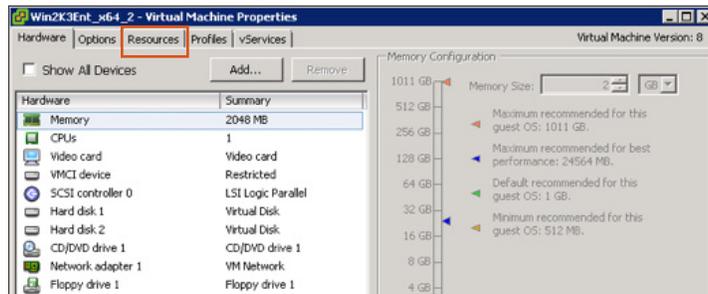
Immediately, a new task is launched to enable Storage I/O Control. If you want to modify the congestion threshold (that is, the latency value at which SIOC is activated), click the **Advanced** button. We will not be modifying this value during this exercise, but you can verify that the current threshold is **30ms**. This means that if cluster-wide I/O latency to this datastore exceeds 30ms, SIOC will commence.

However, our latency value for I/O driven by Iometer is very low, typically **2-3ms**, well below the 30ms threshold that exists by default to trigger SIOC.

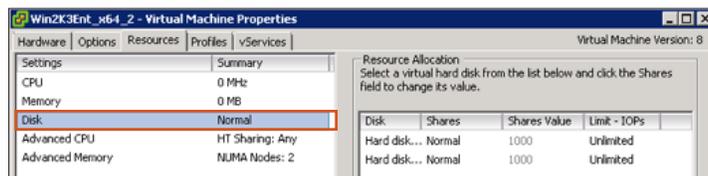
We will next modify the Iometer configuration to generate a much larger number of I/Os. This will also cause the I/O latency value to rise above the SIOC trigger value of 30ms.

Monitoring the Effect of Limiting IOPS

To see the effect of SIOC's being used to limit IOPS, we must first do some configuration steps on the virtual machine resources. In the vSphere Client, go to **Home**. From **Inventory**, select **Hosts and Clusters**. Click the first Windows 2003 virtual machine. In the **Summary** tab, click **Edit Settings**.

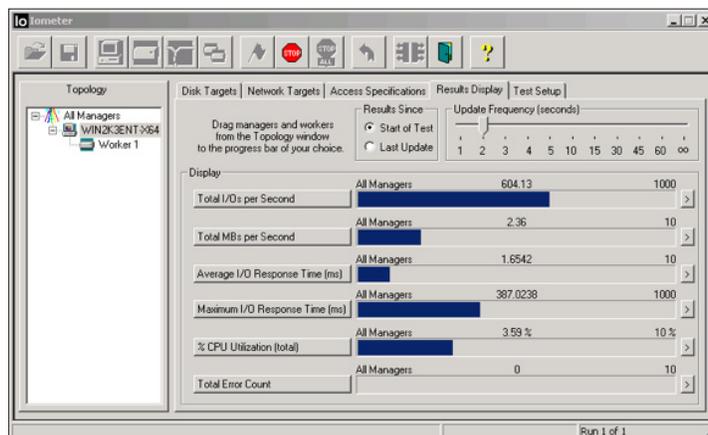


Next, select the **Resources** tab. In the **Settings** list, select the **Disk** entry.

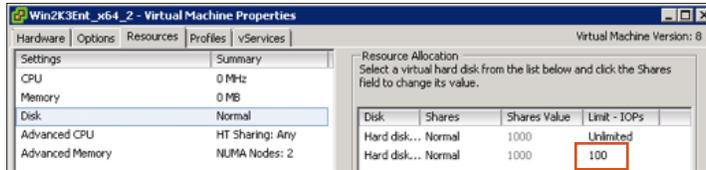


The **Hard disk 2** entry is the disk on the shared datastore that has SIOC enabled. The **Shares** value is set to **Normal** and the **Limit - IOPS** value is set to **Unlimited**. The settings are identical on the other virtual machine. This means that even if SIOC did trigger on this shared datastore, the two virtual machines would get equal priority when it came to I/O to the datastore.

We will now look at SIOC's enforcing of the IOPS limit. Check back to the **vSphere Client Performance** tab or the virtual machine's iometer results to see the number of IOPS currently being generated. The value in this exercise is approximately 500–600 IOPS.



For the first virtual machine, modify the **Resource Allocation** entry for **Disk** and set the **Limit – IOPS** to **100**. Simply select the **Unlimited** value and type **100** into the **Limit – IOPS** field. Click **OK**.



Now click the first virtual machine's console and monitor iometer's **Total I/Os per Second**. You will see a very gradual decrease in the IOPS being generated by this virtual machine.

Conversely, if you monitor the second virtual machine's console and watch the iometer display for IOPS, you should observe a gradual increase in the IOPS.

These are very gradual decreases and increases in IOPS in each of the virtual machines over a long period of time. They do not immediately limit the IOPS, because this could have an adverse effect on any running applications.

Take a few moments to observe this operation before continuing with the next part of the evaluation.

Monitoring the Effects of Shares

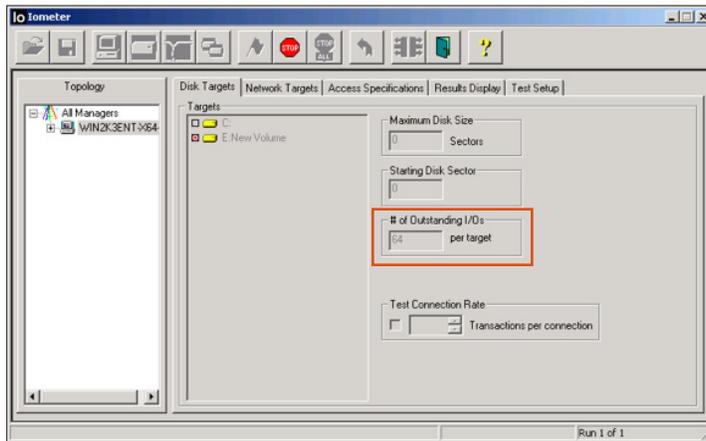
After monitoring the IOPS' reducing on the first virtual machine and increasing on the second virtual machine, stop iometer and change the **Limit – IOPS** value from **100** back to **Unlimited**. Restart iometer and enable the IOPS' returning to a very similar value on both virtual machines.

Next, to see the effect of shares on the I/O, modify the shares value on one of the virtual machines to be **High (2000)** rather than the default **Shares** value of **Normal (1000)**. Because the I/O value is not causing any latency issues, this won't have any effect on the I/O of your virtual machines. (If you are observing changes in the IOPS, this might be due to the last exercise.)

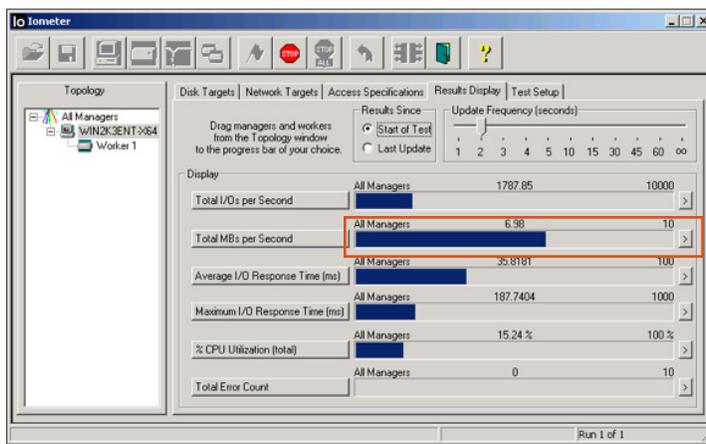


Next, we will modify the iometer settings. We will set the value for outstanding I/O to be **64**. This should mean that the latency value for I/Os becomes high enough (greater than 30ms) to trigger congestion. This should also mean that SIOC will consider the virtual machine with 1000 shares to have a lower priority than the virtual machine with 2000 shares regarding the scheduling of I/O.

Open the console to the first virtual machine. Launch Iometer. Open Iometer configuration file **iometer.icf**. On the main **Disk Targets** tab, change the **# of Outstanding I/Os per target** from **1** to **64**. Restart the I/O by clicking the **green flag icon**.



Repeat this on the other virtual machine. Go to the **Results Display** tab and click the **green flag icon** to start I/O. The latency value will now increase with both virtual machines on two separate ESXi hosts driving I/O. Latency in this example is now in the **30-40ms** range, but this might be different in your case. It doesn't matter. The point is to go above the latency threshold defined in SIOC.



Again, there will be a gradual movement towards the prioritizing of shares. You should observe a gradual increase in the IOPS for the virtual machine with 2000 shares and a gradual decrease in IOPS for the virtual machine with 1000 shares. Take a few moments to observe these changes. This completes the evaluation of **Storage I/O Control**.

Summary

VMware vSphere 5.0 adds many new storage features to an already rich set of capabilities supported in vSphere 4.1. It reduces complexity while providing greater scalability. Virtual machine provisioning historically has imposed operational challenges. Monitoring and manually balancing workloads, or provisioning virtual machines based on I/O capacity and space utilization, have proven to be very difficult and have often been neglected, leading to hot spots and over- or underutilized datastores. vSphere Storage DRS provides smart virtual machine placement and load balancing mechanisms based on I/O and space capacity. VM Storage Profiles can be used during the provisioning of virtual machines and disks, enabling placement based on the requirements of the virtual machines and the offered storage tiers. These features result in a decrease in the operational effort associated with the provisioning and monitoring of virtual machines and storage environments.

