5.1 Provisioning storage from IBM Storwize V5000 and making it available to the host

This section describes the setup process and shows how to create volumes and make them accessible from the host. The following basic process is used to setup your environment:

1. Create volumes.
2. Map volumes to the host.
3. Discover the volumes from the host and specify multipath settings

Complete the following steps to create the volumes:

1. Open the All Volumes window of the IBM Storwize V5000 GUI to start the process of creating volumes, as shown in Figure 5-1.

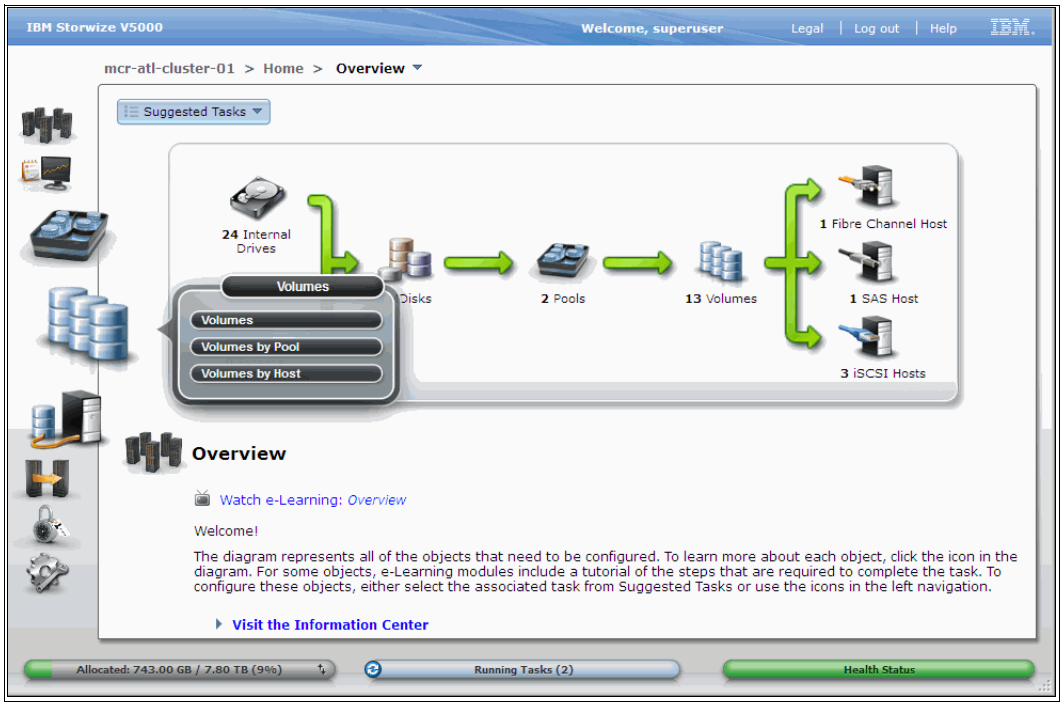

Figure 5-1 GUI Volumes option

2. Highlight and click **Volumes** and the window in which all current volumes are listed opens, as shown in Figure 5-2.

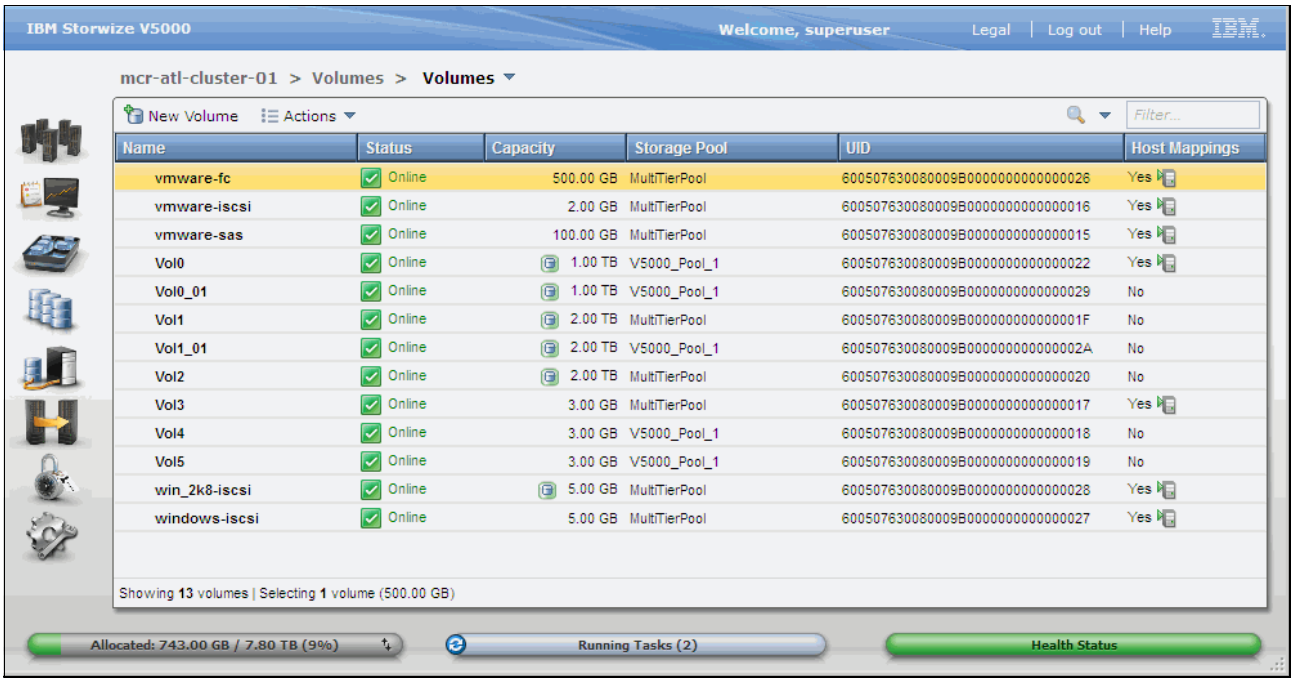

Figure 5-2 Volume listings

3. If this is a first-time setup, there are no volumes listed. Click **New Volume** in the upper left of the window. The New Volume window opens, as shown in Figure 5-3.

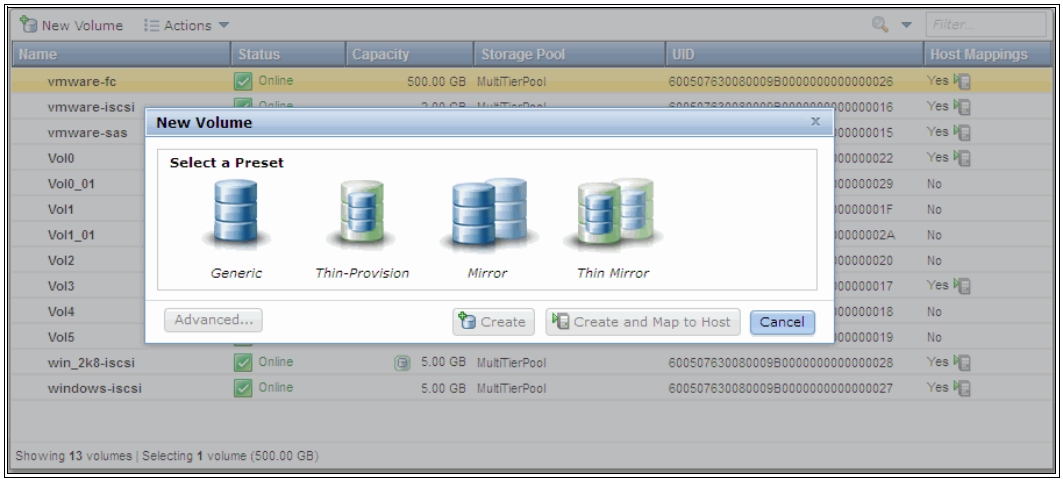

Figure 5-3 New Volume window

By default, all volumes that you create are striped across all available MDisks in that storage pool. The GUI for the IBM Storwize V5000 provides the following preset selections for the user:

– Generic: A striped volume that is fully provisioned, as described in 5.1.1, “Creating a generic volume” on page 164. Fully provisioned means that the volume capacity reflects the same size physical disk capacity.

– Thin-provisioned: A striped volume that is space-efficient. This means that the volume capacity does not reflect the physical capacity that is available to the volume. There are choices available in the Advanced menu to help determine how much space is initially fully allocated and how large the volume can grow, as described in 5.1.2, “Creating a thin-provisioned volume” on page 167.

– Mirror: A striped volume that consists of two striped copies and is synchronized to protect against loss of data if the underlying storage pool of one copy is lost, as described in 5.1.3, “Creating a mirrored volume” on page 169.

– Thin-mirror: Two synchronized copies, which are thin provisioned, as described in 5.1.4, “Creating a thin-mirror volume” on page 174.

4. Select the type of volume that you want to create. For more information, see the following sections:

   – 5.1.1, “Creating a generic volume” on page 164
   – 5.1.2, “Creating a thin-provisioned volume” on page 167
   – 5.1.3, “Creating a mirrored volume” on page 169
   – 5.1.4, “Creating a thin-mirror volume” on page 174

5.1.1 Creating a generic volume

The most commonly used type of volume is the generic volume. This type of volume is fully provisioned, that is, the volume size reflects the physical disk capacity that is allocated to the volume. The host and the IBM Storwize V5000 see the fully allocated space without a mirror.

Complete the following steps to create a generic volume:

1. Choose a generic volume, as shown in Figure 5-4.

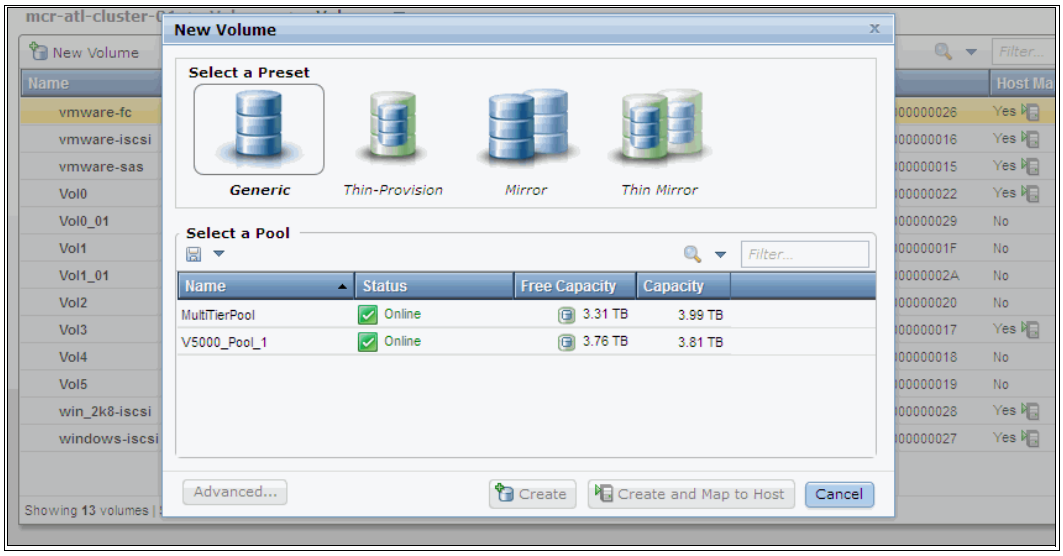

Figure 5-4 Provisioning a Generic volume

2. Select the pool in which the volume is to be created. Select the pool by clicking it. In our example, click the pool that is called **V5000_Pool_1**. The result is shown in Figure 5-5.

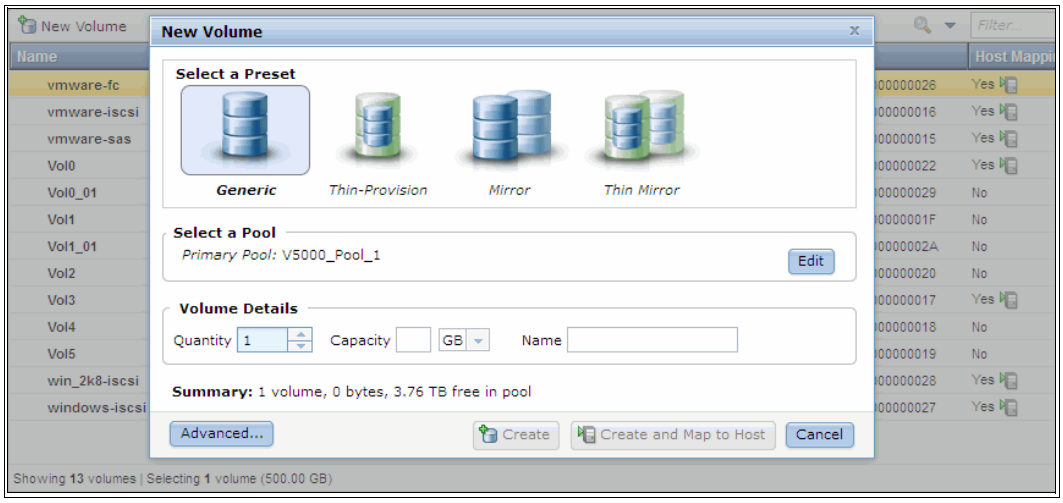

Figure 5-5 Pool selection

Important: The Create and Map to Host option is disabled if no host is configured on the IBM Storwize V5000. For more information about configuring the host, see Chapter 4, “Host configuration” on page 153.

There are advanced options available, as shown in Figure 5-6.

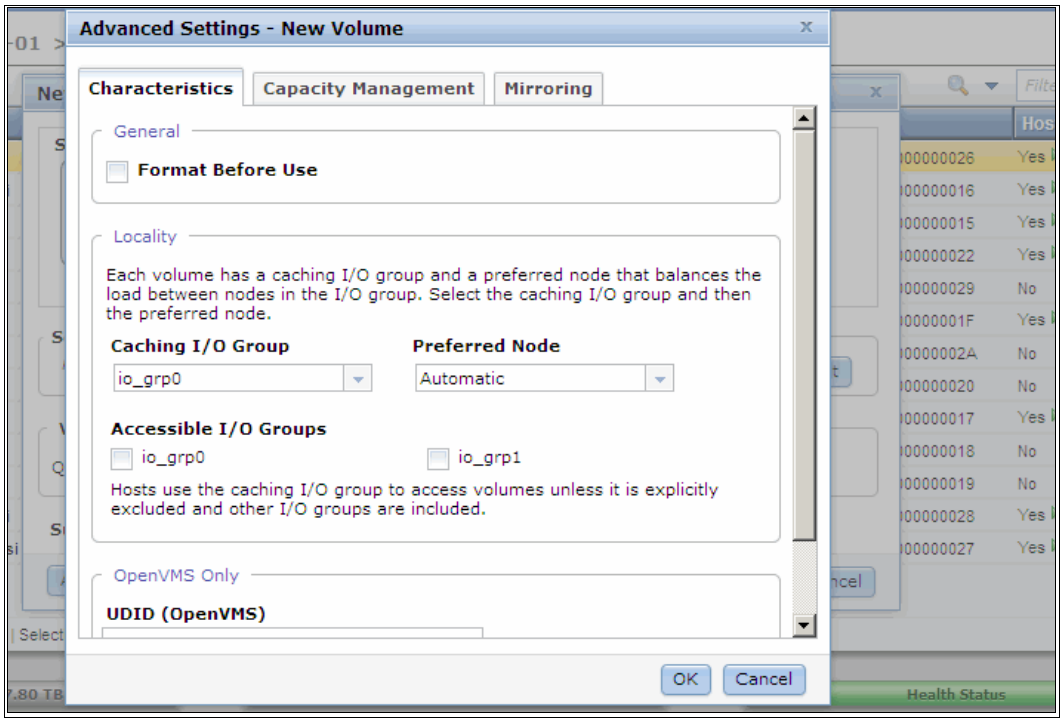

Figure 5-6 Generic advanced options

For Generic volumes, capacity management and mirroring do not apply. If you have two I/O Groups within your IBM Storwize V5000 (two control enclosures that are configured as one cluster), you can specify which I/O Group is used to access the volume and provide the volume caching. Similarly, there is an option to set the preferred node within the I/O Group that owns the volume initially. The recommendation is to set Preferred Node to automatic and allow the IBM Storwize V5000 to balance the volume I/O across the two node canisters in the I/O Group.

Where the Caching I/O Group is concerned, caution must be exercised. Hosts might be able to communicate with both I/O Groups depending on the method of connectivity that is used in any zoning that is employed. Setting the Caching I/O Group option to automatic or to the wrong I/O Group results in the volume being unavailable to the correct host. Ensure that the host you want the volume to be accessed from is correctly zoned and attached to the I/O Group that you define as the caching I/O Group for that volume or that your host is connected and zoned to all node canisters in both I/O Groups. This feature is useful if you have a host that has limited connectivity; for example, an SAS direct attach host. It might be connected only to the node canisters of I/O Group0, but you might want to provision a volume from I/O Group1.

Important: Ensure that the Caching I/O Group is set correctly.

3. Enter a volume name and size. Click **Create and Map to Host** to create and map the volume to a host or click **Create** to complete the task and leave mapping the volume to a later stage. The generic volume is created, as shown in Figure 5-7.

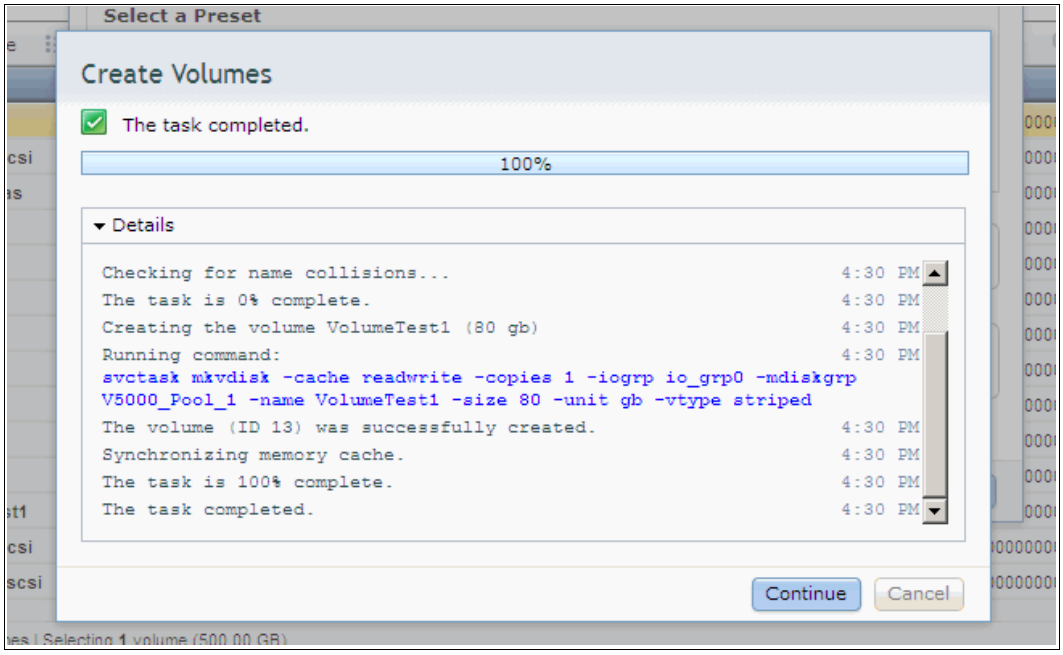

Figure 5-7 Volume creation complete

4. Click **Continue**. For more information, see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 177.

Volumes can also be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 181.

5.1.2 Creating a thin-provisioned volume

Volumes can be configured to be thin-provisioned. A thin-provisioned volume behaves the same as a fully provisioned volume regarding application reads and writes. However, when a thin-provisioned volume is created, it is possible to specify two capacities: the real physical capacity that is allocated to the volume from the storage pool, and its virtual capacity that is available to the host. The real capacity determines the quantity of extents that are initially allocated to the volume. The virtual capacity is the capacity of the volume that is reported to all other components (for example, FlashCopy and cache) and to the host servers.

To create a thin-provisioned volume, complete the following steps:

1. Select **Thin-Provision**, as shown in Figure 5-8.

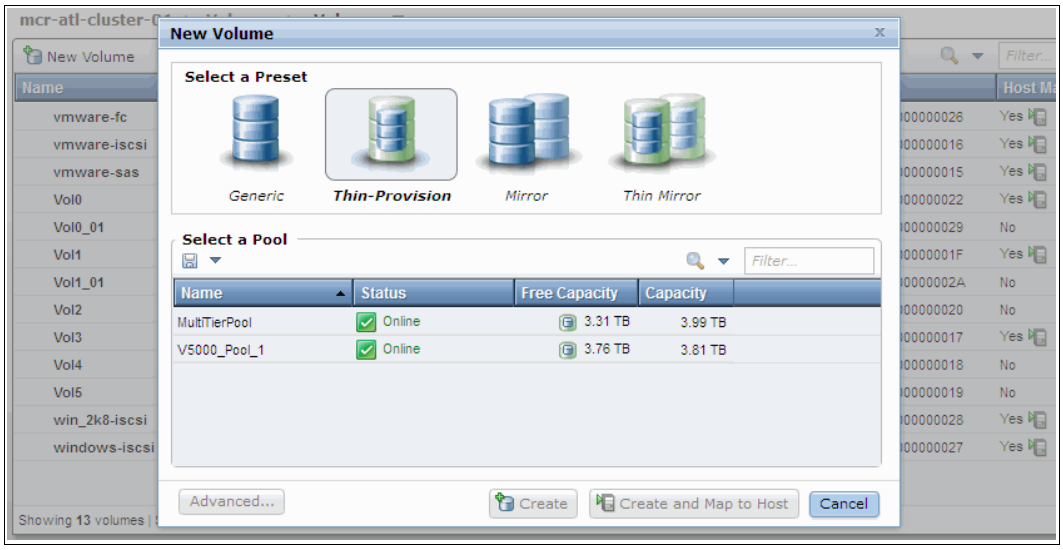

Figure 5-8 Creating a thin-provisioned volume

2. Select the pool in which the thin-provisioned volume should be created by clicking it and entering the volume name and size. In our example, we clicked the pool that is called **V5000_Pool_1**. The result is shown in Figure 5-9.

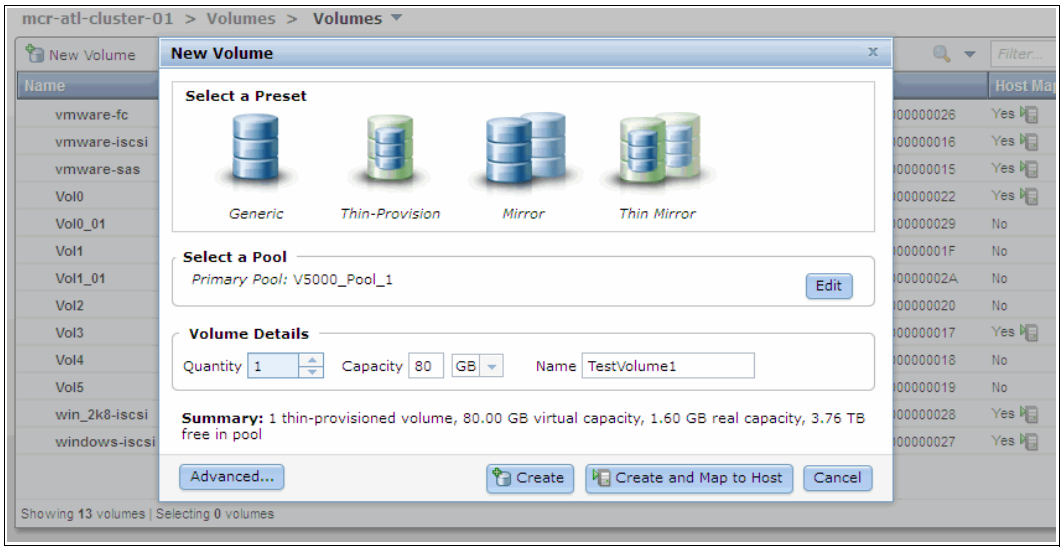

Figure 5-9 Enter the volume name and size

Under the Volume Name field is a summary that shows that you are about to create a thin-provisioned volume, the virtual capacity is to be configured (the volume size you specified), the space that is physically allocated (real capacity), and the available physical capacity of the pool. By default, the real capacity is 2% of the virtual capacity, but you can change this setting in the Advanced options. By selecting this option, the window defaults to Capacity Management, as shown in Figure 5-10.

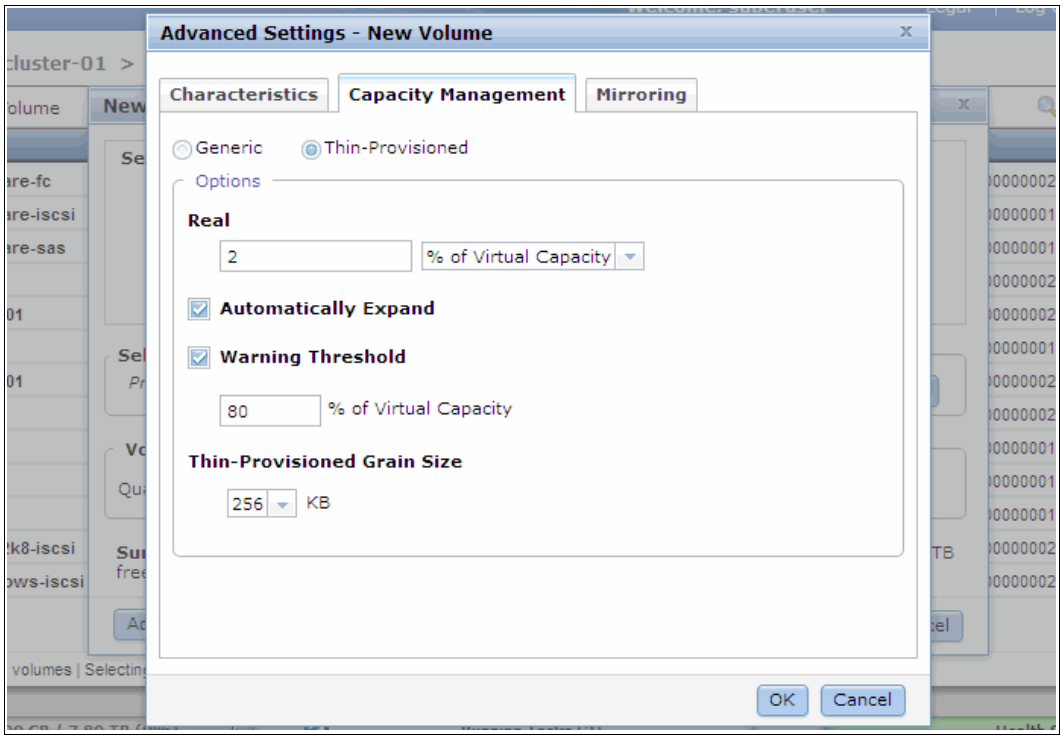

Figure 5-10 Advanced Settings: New Volume

The following advanced options are available:

- Real: Specify the size of the physical capacity space that is used during creation.
- Automatically Extend: This option enables the automatic expansion of real capacity as the physical data size of the volume grows.
- Warning Threshold: Enter a threshold for receiving capacity alerts. The IBM Storwize V5000 sends an alert when the physically allocated capacity reaches 80% of the virtual capacity in this case (which is the default setting).
- Thin-Provisioned Grain Size: Specify the grain size for real capacity.

3. Make your choices, if required, and click **OK** to return to New Volume window, as shown in Figure 5-9 on page 167.

4. Click **Create and Map to Host** to create and map the volume to a host, or click **Create** to complete the task and leave mapping the volume to a later stage. The volume is created, as shown in Figure 5-11.

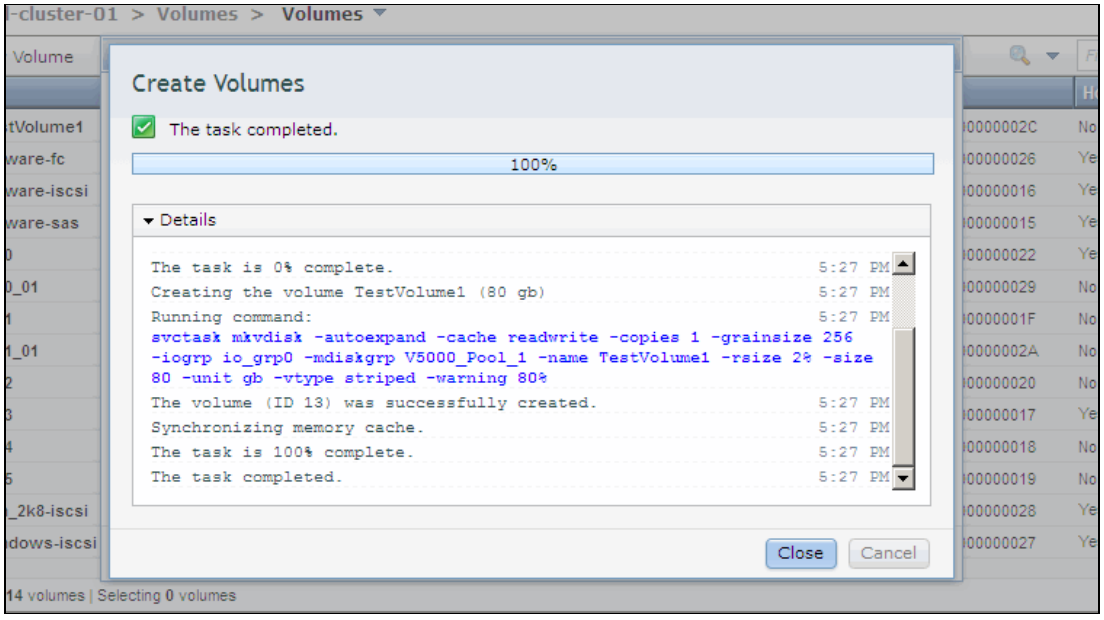

Figure 5-11 Thin volume creation complete

If you decided to map the host, click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 177.

The volumes can be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 181.

5.1.3 Creating a mirrored volume

IBM Storwize V5000 offers the capability to mirror volumes, which means a single volume is presented to the host, but two copies exist in the storage back end, usually in different storage pools (all reads are handled by the primary copy). This feature is similar to host-based software mirroring, but it provides a single point of management for all operating systems and provides storage high availability to operating systems that do not support software mirroring.

By using this setup with the mirror copies in different storage pools, you can protect against array failures (for example, multiple disk failures). More advanced features also are available to you, as described in Chapter 8, “Advanced host and volume administration” on page 349.

The mirroring feature improves availability, but it is not a disaster recovery solution because both copies are accessed by the same node pair and are addressable only by a single cluster.

For more information about a disaster recovery solution with mirrored copies that are spanning I/O Groups in different locations, see Chapter 10, “Copy services” on page 449.

To create a mirrored volume, complete the following steps:

1. Select **Mirror**, as shown in Figure 5-12.

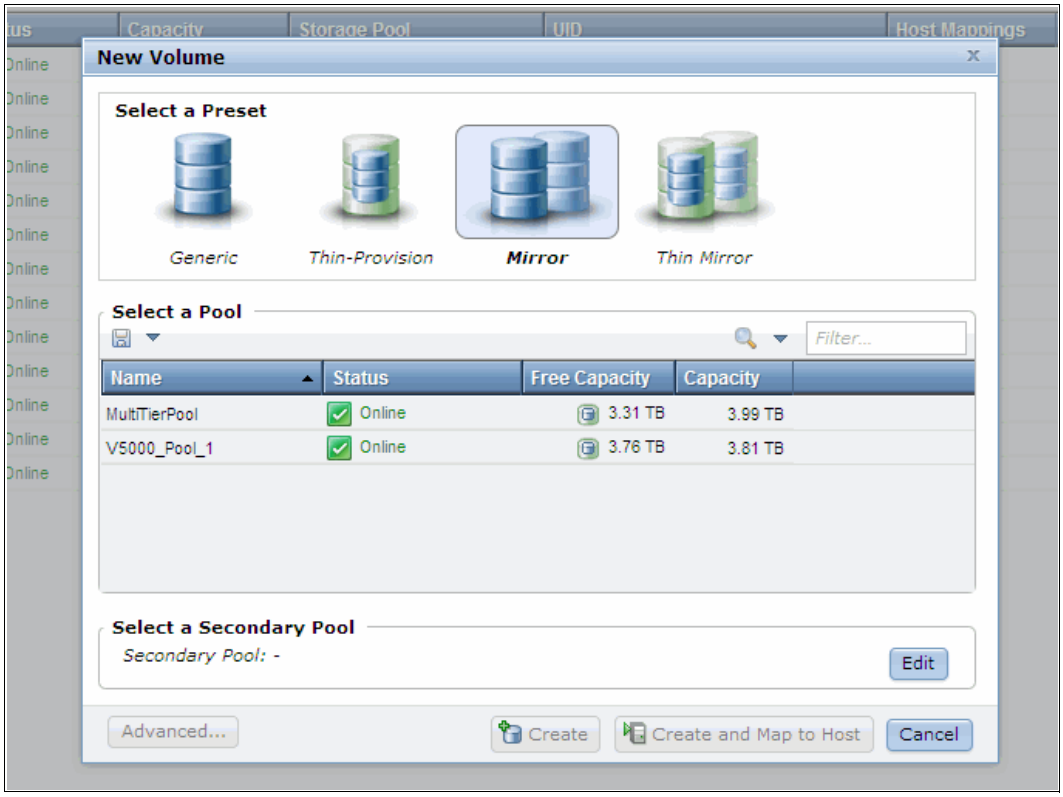

Figure 5-12 Create a mirrored volume

2. Select the primary pool by clicking it and the view changes to the secondary pool, as shown in Figure 5-13 on page 171.

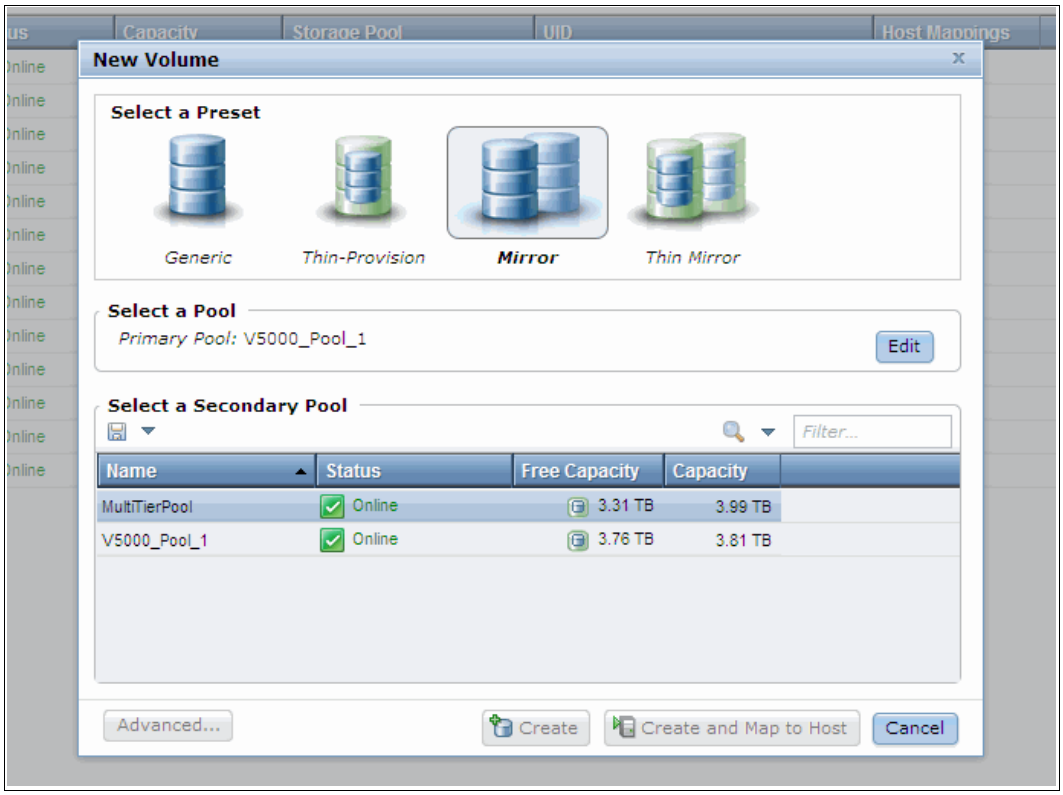

Figure 5-13 Selecting primary storage pool

3. Select the secondary pool by clicking it. Enter a volume name and the required size, as shown in Figure 5-14.

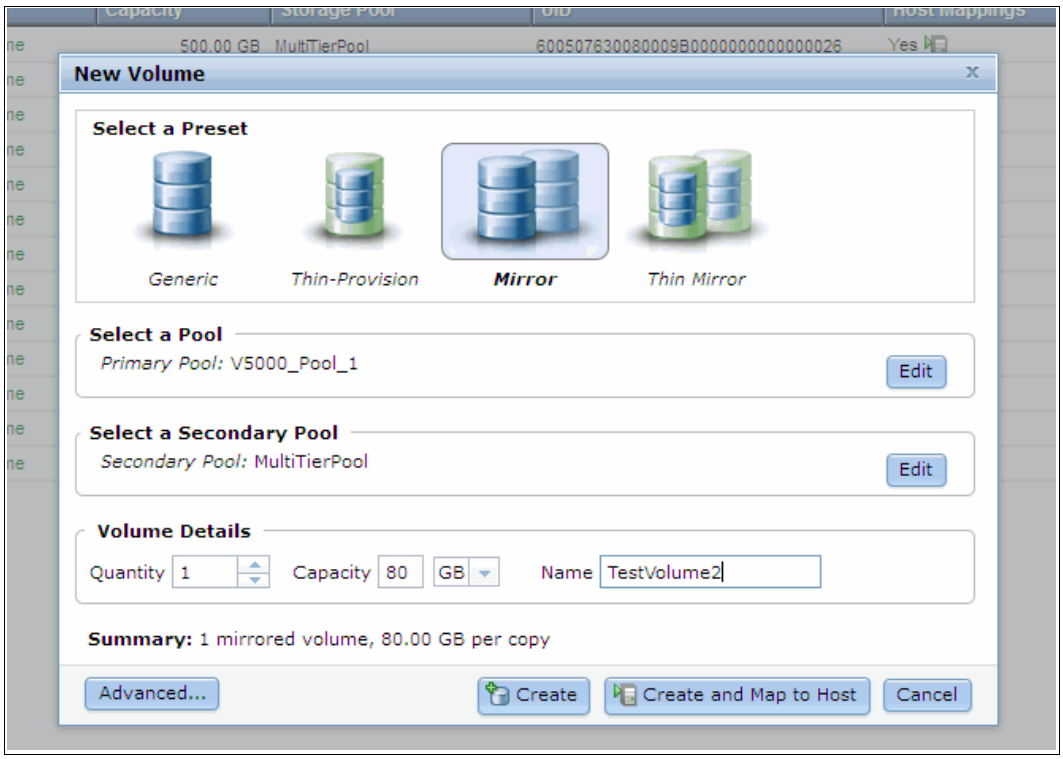

Figure 5-14 Select a secondary pool, volume name, and size

Storage pools: Before a mirrored volume is created, it is best to create at least two separate storage pools and use different pools for the primary and secondary pool when you are entering the information in the GUI to create the volume. In this way, the two mirror copies are created on different MDisks (and, therefore, different physical drives) and protect against a full MDisk failure in a storage pool. For more information about storage pools, see Chapter 7, “Storage pools” on page 295.

4. The summary shows you the capacity information about the pool. If you want to select advanced settings, click **Advanced** and then click the **Mirroring** tab, as shown in Figure 5-15.

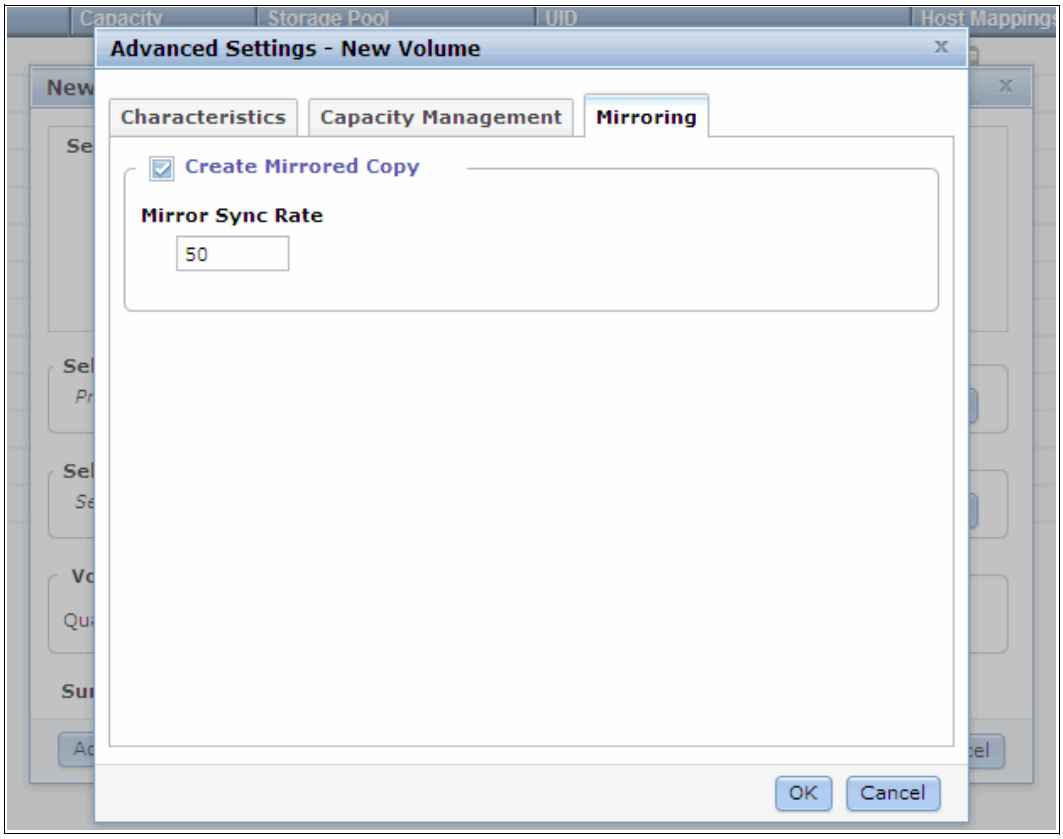

Figure 5-15 Advanced mirroring features

5. In the advanced mirroring settings, you can specify a synchronization rate. Enter a Mirror Sync Rate 1 - 100%. With this option, you can set the importance of the copy synchronization progress. This sets the preference to synchronize more important volumes faster than other mirrored volumes. By default, the rate is set to 50% for all volumes. If for any reason the mirrors loose synchronization, this parameter governs the rate at which the various mirrored volumes resynchronize.

   Click **OK** to return to the New Volume window, as shown in Figure 5-14 on page 171.

6. Click **Create and Map to Host** and the mirrored volume is created, as shown in Figure 5-16. If you do not want to map the hosts, click **Create** to complete the task and exit to the GUI.

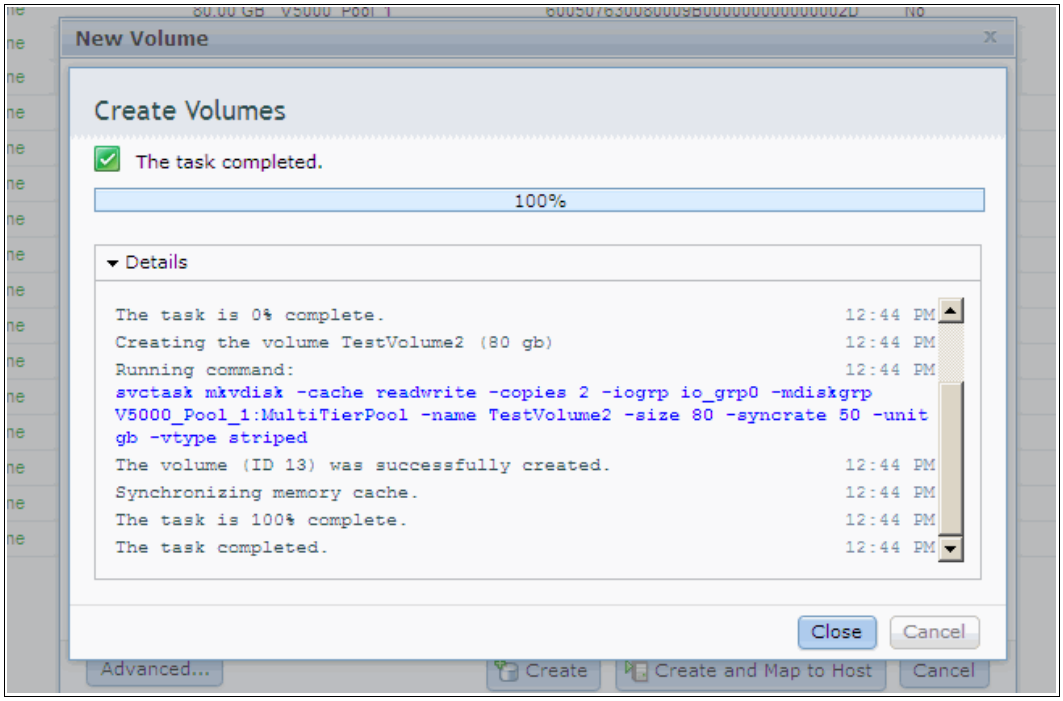

Figure 5-16 Mirrored volume task complete

7. If you decided to map the host, click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 177.

   The volumes can be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 181.

5.1.4 Creating a thin-mirror volume

By using a thin-mirror volume, you can allocate the required physical space on demand (as described in 5.1.2, “Creating a thin-provisioned volume” on page 167) and have several copies of a volume available (as described in 5.1.3, “Creating a mirrored volume” on page 169).

To create a thin-mirror volume, complete the following steps:

1. Select **Thin Mirror**, as shown in Figure 5-17.

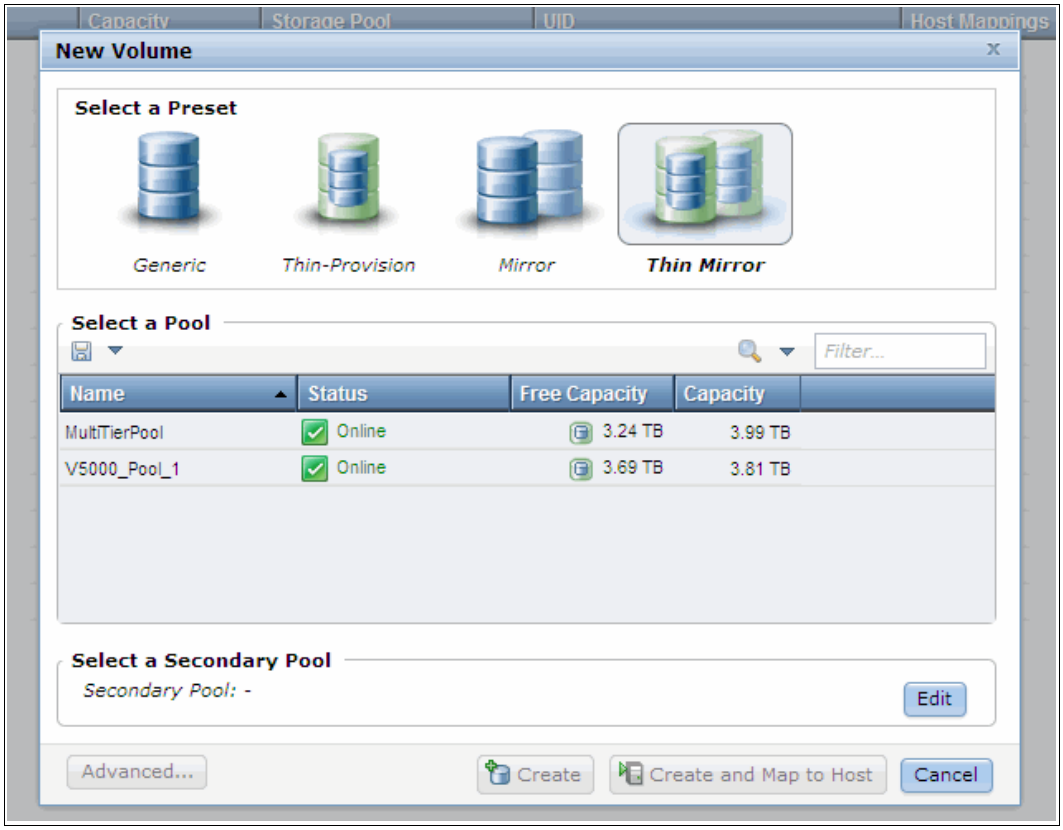

Figure 5-17 Create a Thin Mirror

2. Select the primary pool by clicking it and the view changes to the secondary pool, as shown in Figure 5-18.

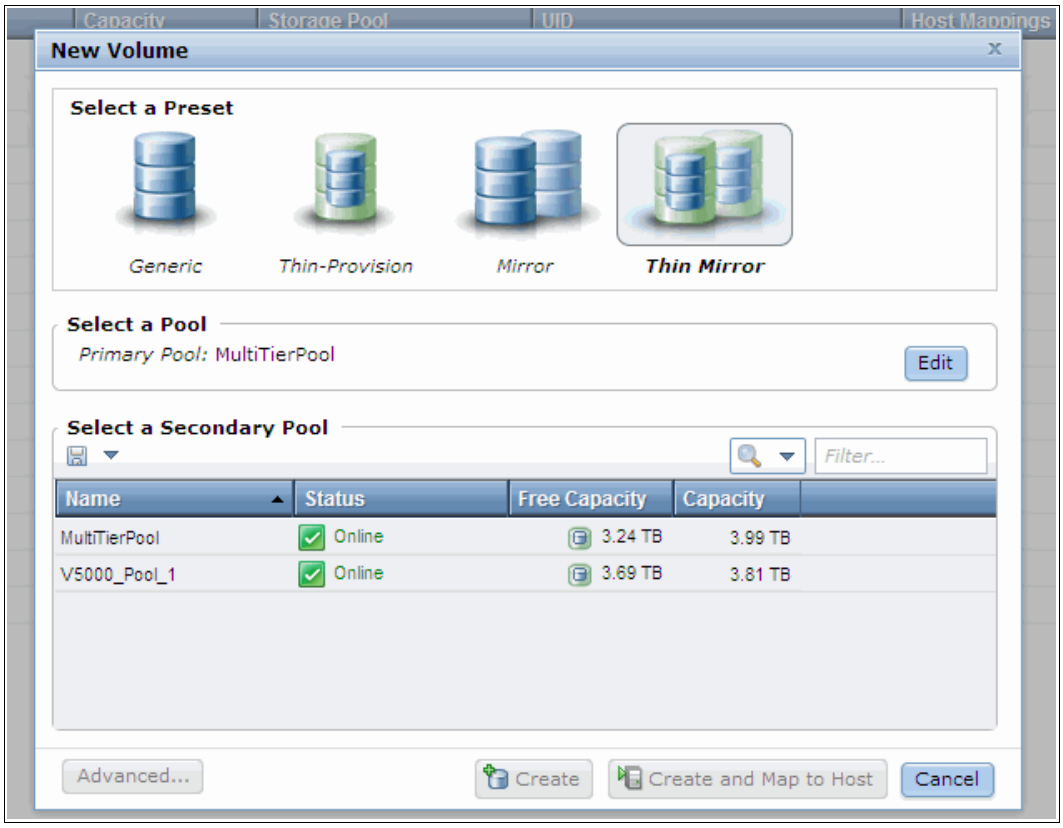

Figure 5-18 Selecting storage pools

3. Select the pool for the secondary copy and enter a name and a size for the new volume, as shown in Figure 5-19.

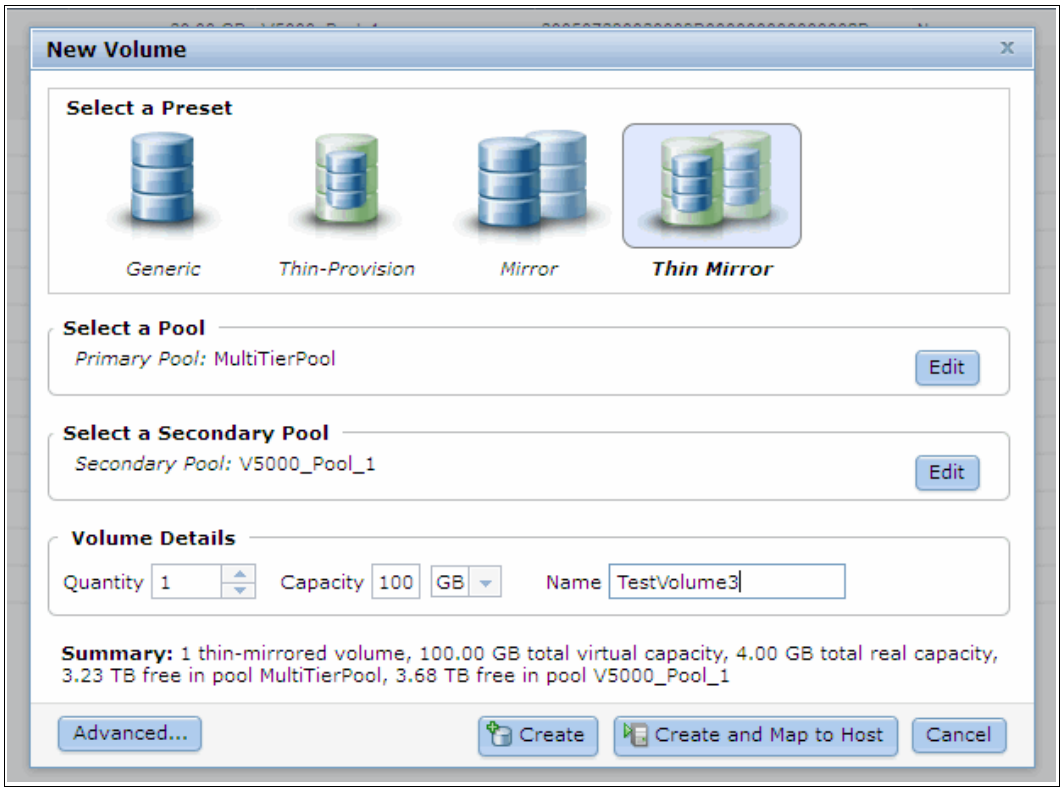

Figure 5-19 Enter a volume name and size

4. The summary shows you the capacity information and the allocated space. You can click **Advanced** and customize the thin-provision settings (as shown in Figure 5-10 on page 168) or the mirror synchronization rate (as shown in Figure 5-15 on page 172). If you opened the advanced settings, click **OK** to return to the New Volume window.

5. Click **Create and Map to Host** and the mirrored volume is created, as shown in Figure 5-20. If you do not want to map the hosts, click **Create** to complete the task and exit to the GUI.

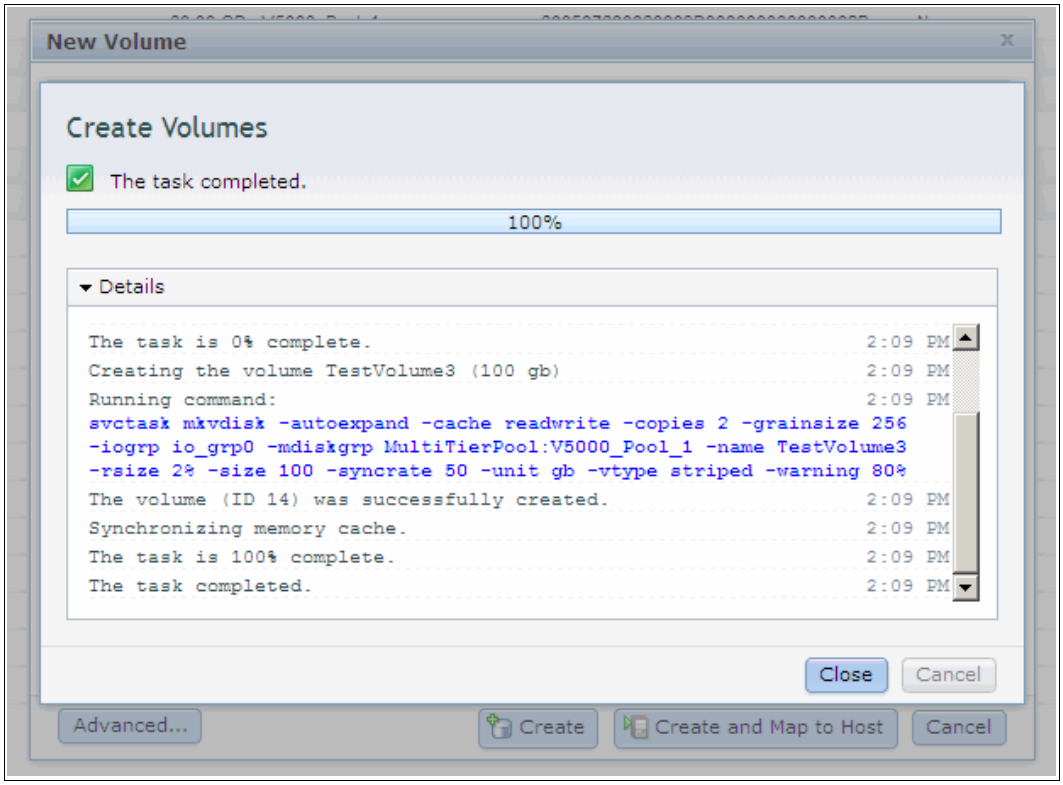

Figure 5-20 Thin Mirror Volume task complete

6. If you decided to map the host, click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 177.

   The volumes can be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 181.

5.2 Mapping a volume to the host

The first part of this section describes how to map a volume to a host if you click **Create and Map to Host**. The second part of this section describes the manual host mapping process that is used to create customized mappings.

5.2.1 Mapping newly created volumes to the host by using the wizard

We continue to map the volume that we created in 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162. We assume that you followed the procedure and clicked **Create and Map to Host** followed by **Continue** when the volume create task completed, as shown in Figure 5-21 on page 178.

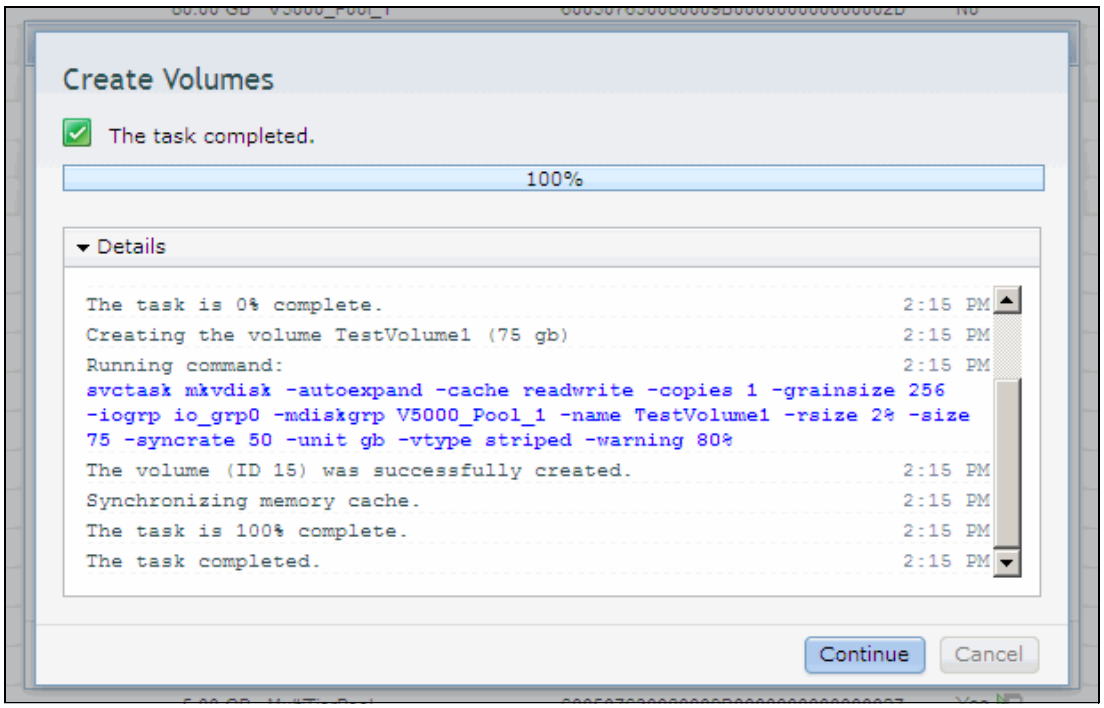

Figure 5-21 Continue to host mapping

To map the volumes, complete the following steps:

1. Select the host I/O Group to which the host is connected. The default setting is All I/O Groups, as shown in Figure 5-22.

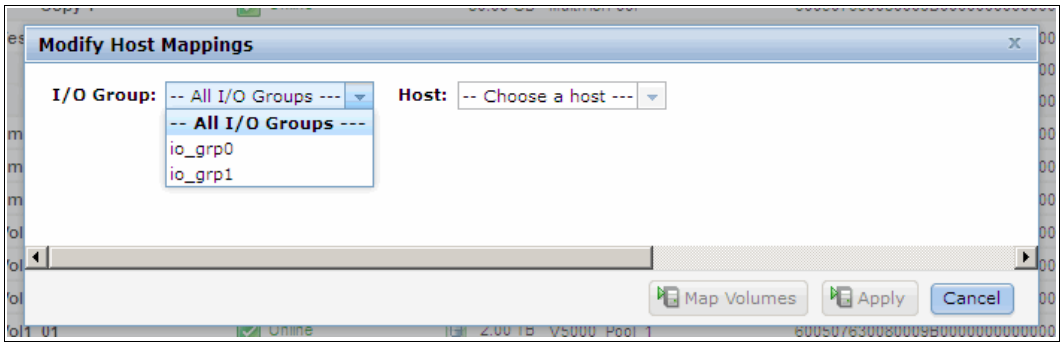

Figure 5-22 Choose an I/O Group

Selecting the correct I/O Group is important if there are more than one group. As we described in 5.1.1, “Creating a generic volume” on page 164, when a volume is created, it is possible to define the caching I/O Group or the I/O Group that owns the volume and be used to access it. Therefore, your host must communicate with the same I/O Group for the mapping to be successful. Additionally, when hosts are defined, they should be masked correctly, as described in Chapter 4, “Host configuration” on page 153. If they are so masked, the filters that are shown in Figure 5-22 show the correct hosts that are available on each I/O Group.

2. Select the host to which the volume is to be available, as shown in Figure 5-23.

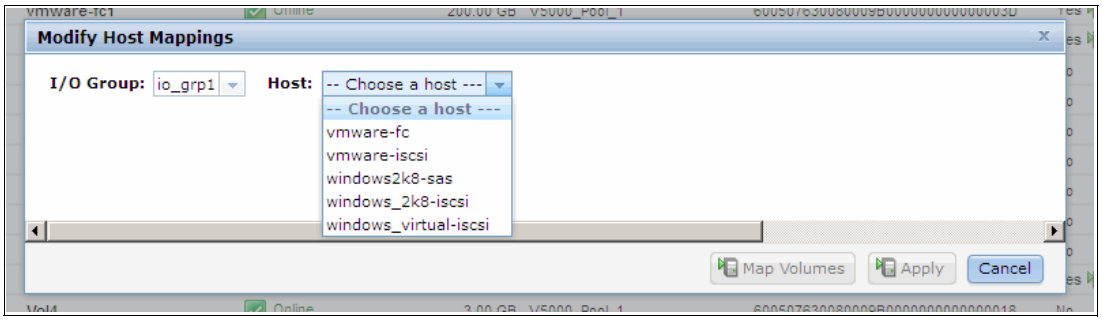

Figure 5-23 Select a host

3. The Modify Host Mappings window opens and your host and the created volume is already selected. Click **Map Volumes** and the volume is mapped to the host, as shown in Figure 5-24.

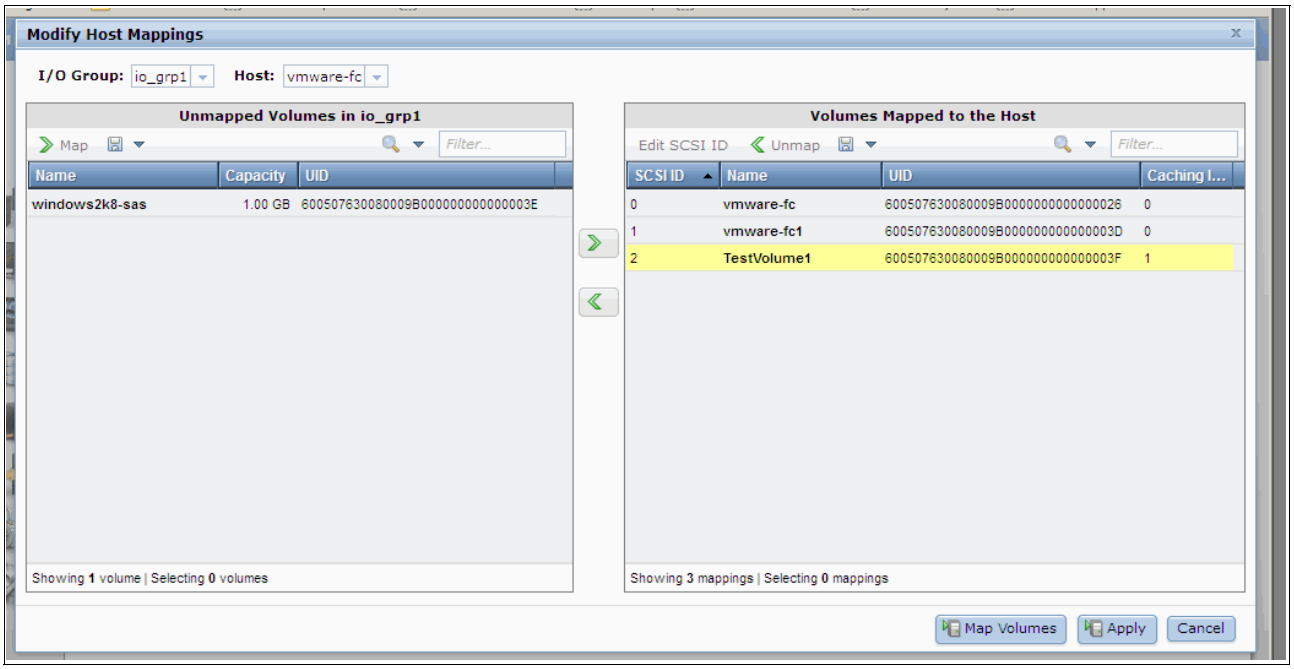

Figure 5-24 Modify Host Mappings window

The new volume to be mapped is highlighted. To continue the process and complete the mapping, you can click **Apply** or **Map Volumes.** The only difference between the two options is that after the mapping task completes (as shown in Figure 5-25 on page 180), the Modify Host Mappings window closes automatically. Clicking **Apply** leaves the Modifying Host Mappings window open.

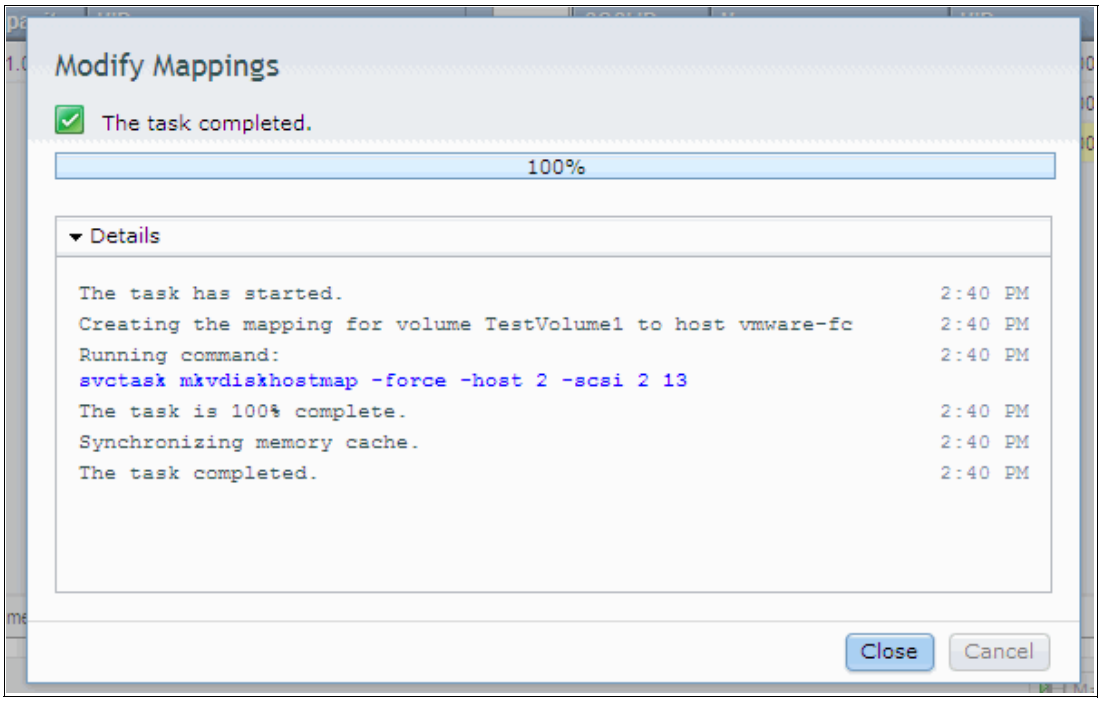

Figure 5-25 Host mapping task complete

4. After the task completes, click **Close**. If you selected the Map Volumes option, the window returns to the Volumes display and the newly created volume is shown. We see that it is already mapped to a host, as shown in Figure 5-26.

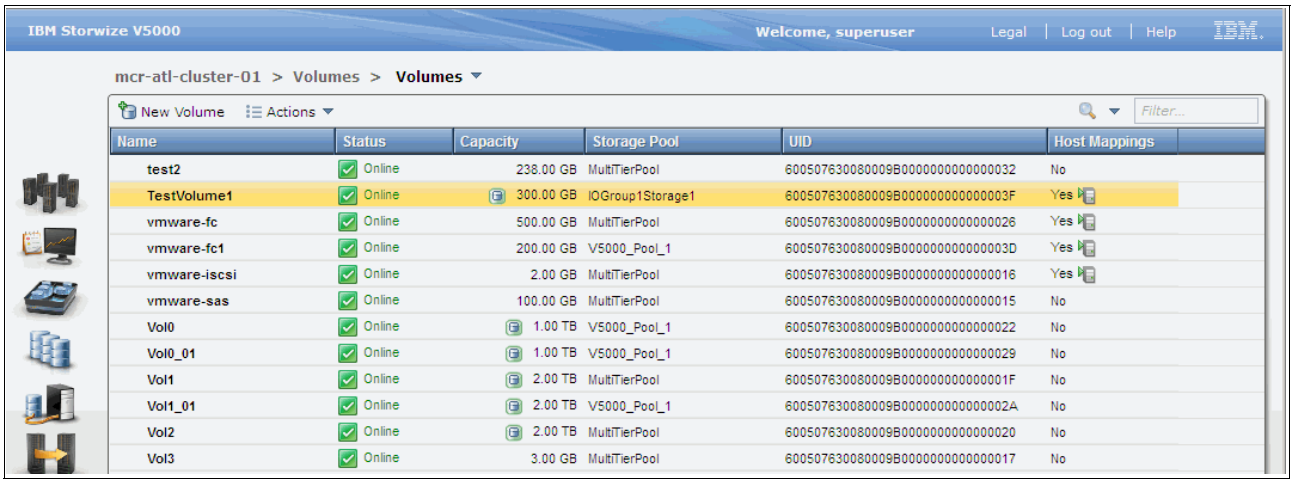

Figure 5-26 New volume that is mapped to host

The host can access the volumes and store data on it. For more information about discovering the volumes on the host and changing host settings (if required), see 5.3, “Discovering the volumes from the host and specifying multipath settings” on page 185.

You also can create multiple volumes in preparation for discovering them later. Mappings also can be customized. For more information about advanced host configuration, see Chapter 8, “Advanced host and volume administration” on page 349.

5.2.2 Manually mapping a volume to the host

We assume that you followed the procedure that is described in 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162 and clicked **Create**.

To manually map a volume to the host, complete the following steps:

1. Open the Hosts window, as shown in Figure 5-27.

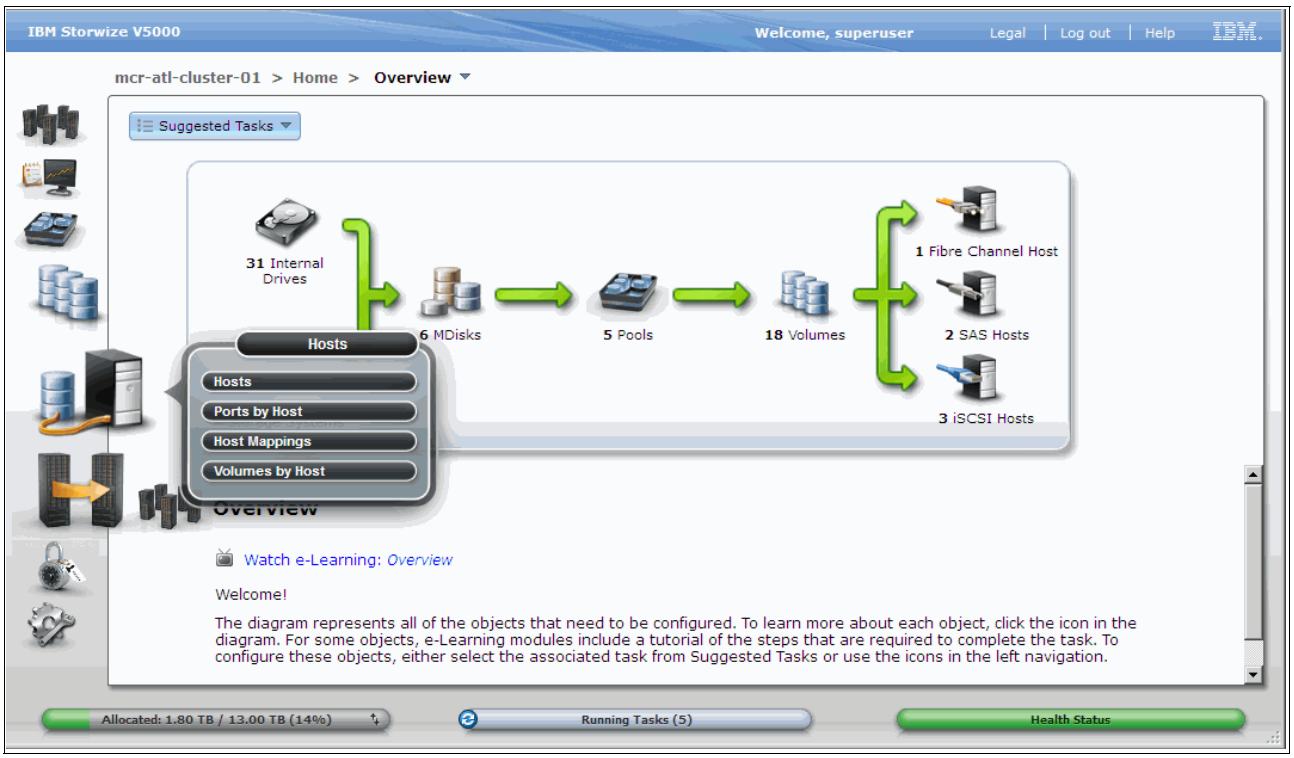

Figure 5-27 Hosts window

2. Right-click the host to which a volume is to be mapped and select **Modify Mappings**, as shown in Figure 5-28.

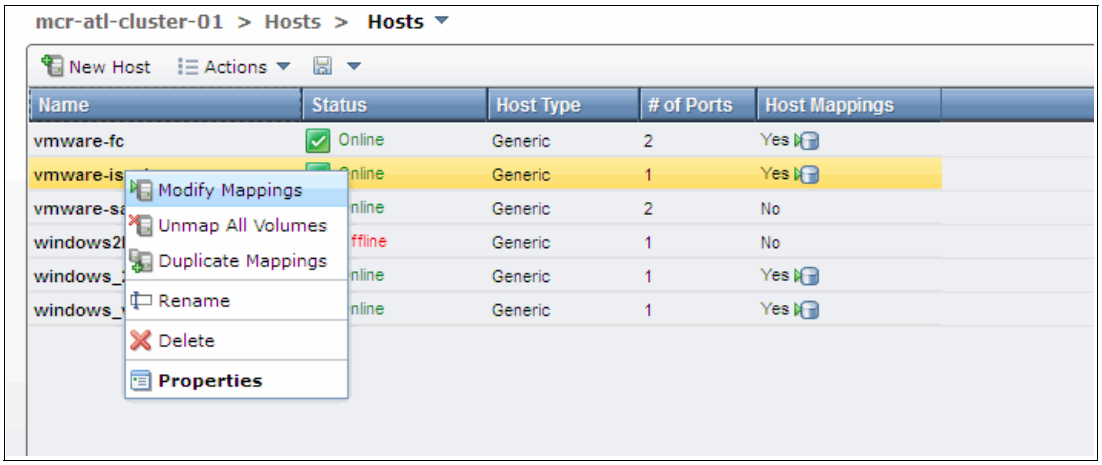

Figure 5-28 Modify mappings selection

3. The Modify Host Mappings window opens. By default, the window shows volumes for all I/O Groups. Selecting the correct I/O Group is important if there is more than one group. As described in 5.1.1, “Creating a generic volume” on page 164, when a volume is created, it is possible to define the caching I/O Group or the I/O Group that owns the volume and be used to access it. Therefore, your host must communicate with the same I/O Group for the mapping to be successful. Additionally, when hosts are defined, they should be masked correctly, as described in Chapter 4, “Host configuration” on page 153.

   Select the volume that you want to map from the Unmapped Volumes pane, as shown in Figure 5-29.

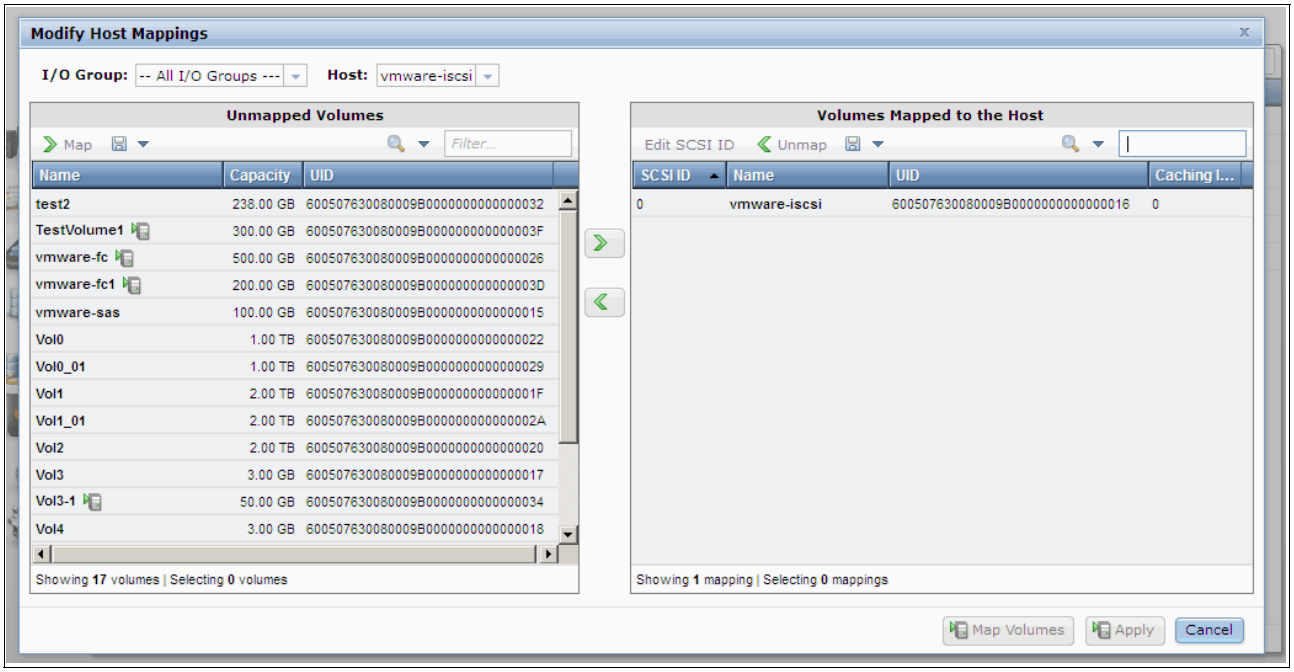

Figure 5-29 Modify host mappings window

The volume is highlighted and the green, right-pointing arrow is active, as shown in Figure 5-30.

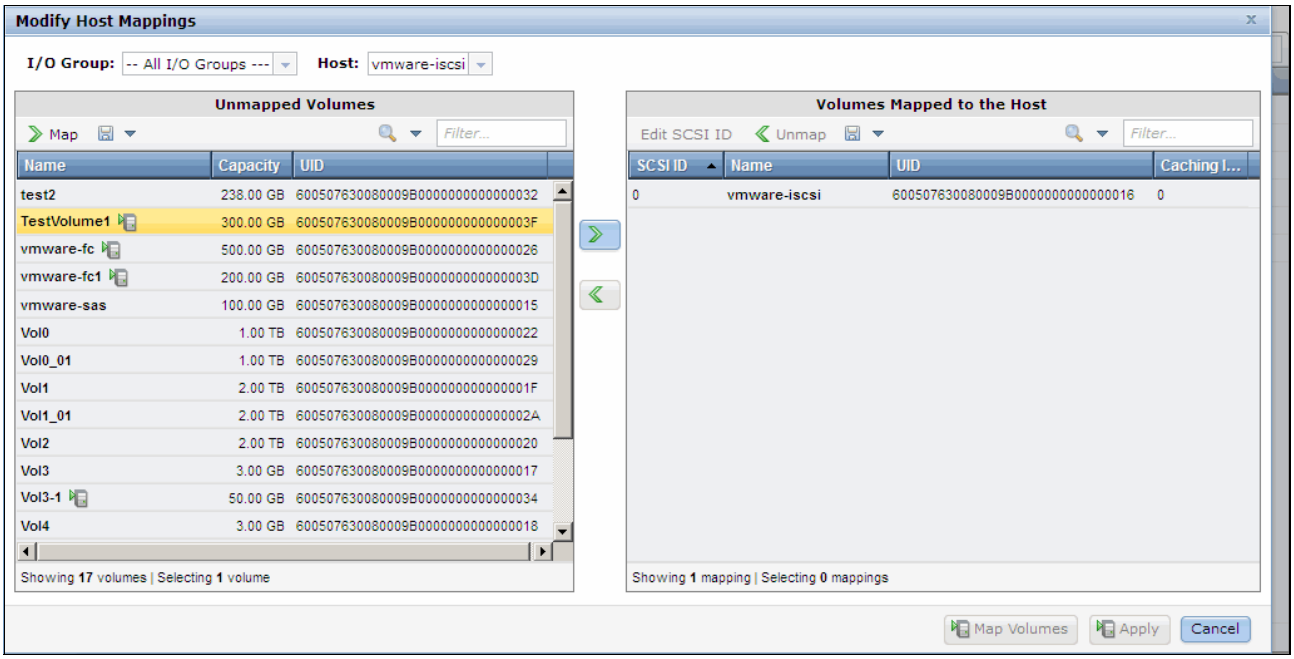

Figure 5-30 Volume mapping selected

Unmapped pane: The Unmapped pane shows all the volumes that are not mapped to the selected host. Some of the volumes might display a mappings icon because they are already mapped to other hosts.

4. Click the right-pointing arrow. The volume is moved to Volumes Mapped to the Host pane, as shown in Figure 5-31 on page 184. Repeat this step for all the volumes that you want to map. To continue and complete the mapping, you can click **Apply** or **Map Volumes**. The only difference between these options is that after the mapping task completes (as shown in Figure 5-31 on page 184), the Modify Host Mappings window closes automatically. Clicking **Apply** but leaves the Modifying Host Mappings window open.

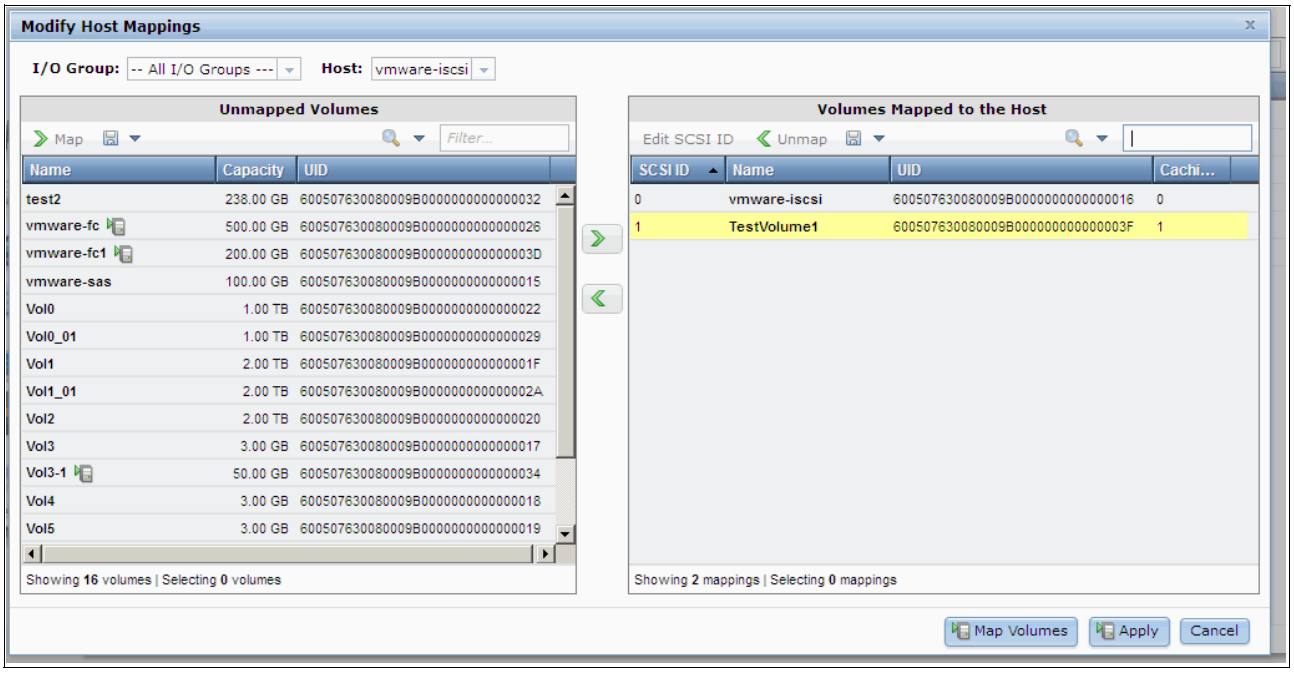

Figure 5-31 Modify host mappings window

5. After the task completes, click **Close**, as shown in Figure 5-32. If you selected the **Map Volumes** option, the window returns to the Hosts display. If you clicked **Apply**, the GUI still displays the Modify Host Mappings window.

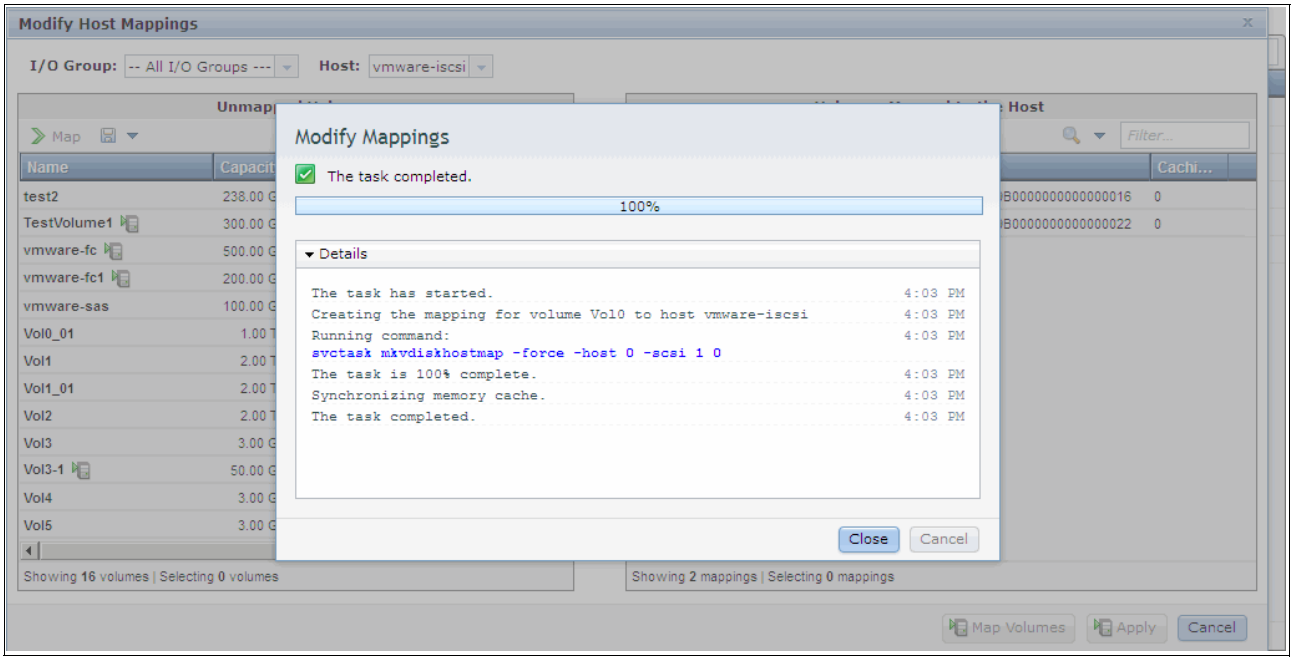

Figure 5-32 Modify mapping complete

The volumes are now mapped and the host can access the volumes and store data on them. For more information about discovering the volumes on the host and changing host settings (if required), see 5.3, “Discovering the volumes from the host and specifying multipath settings” on page 185.

5.3 Discovering the volumes from the host and specifying multipath settings

This section describes how to discover the volumes that were created and mapped in 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162 and 5.2, “Mapping a volume to the host” on page 177, and set more multipath settings, if required.

We assume that you completed all of the following tasks (which are described in this book) so that the hosts and the IBM Storwize V5000 are prepared:

- Prepare your operating systems for attachment, including installing MPIO support. For more information, see Chapter 4, “Host configuration” on page 153.
- Create hosts by using the GUI. For more information, see Chapter 4, “Host configuration” on page 153.
- Perform basic volume configuration and host mapping. For more information, see 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162, and 5.2, “Mapping a volume to the host” on page 177.

This section describes how to discover Fibre Channel, iSCSI, and serial-attached SCSI (SAS) volumes from Windows 2008 and VMware ESX 5.x hosts.

In the IBM Storwize V5000 GUI, click **Hosts**, as shown in Figure 5-33.

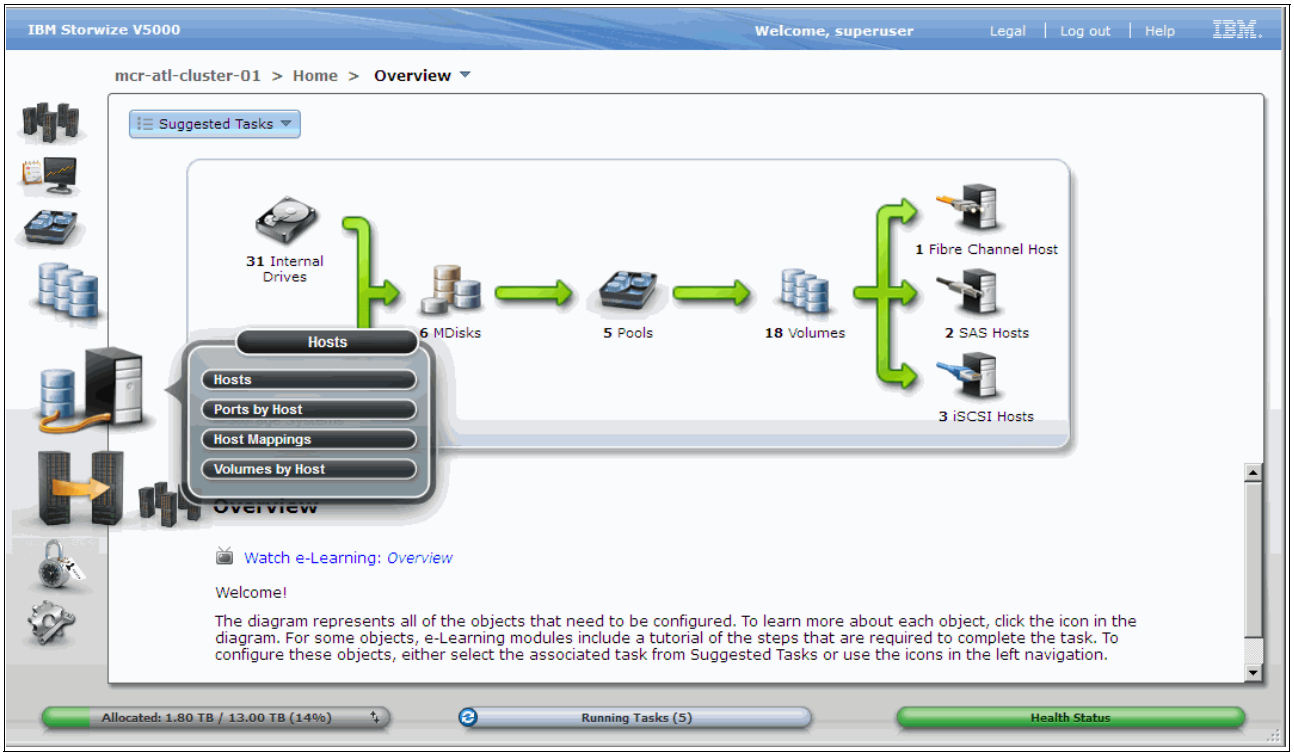

Figure 5-33 Open all hosts

The view that opens gives you an overview of the configured hosts and shows if they are mapped, as shown in Figure 5-34.

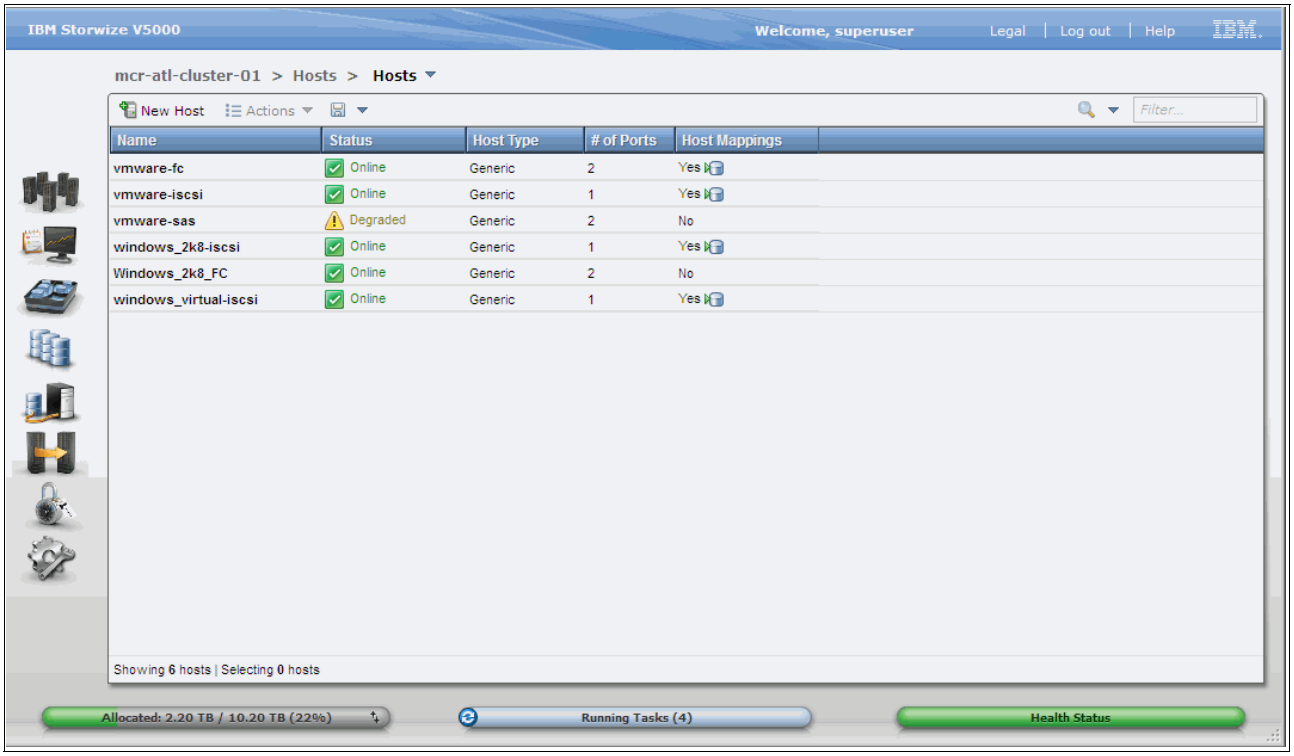

Figure 5-34 All Hosts view

5.3.1 Windows 2008 Fibre Channel volume attachment

To attach the Fibre Channel volume in Windows 2008, complete the following steps:

1. Right-click your Windows 2008 Fibre Channel host in the Hosts view and select **Properties**, as shown in Figure 5-35.

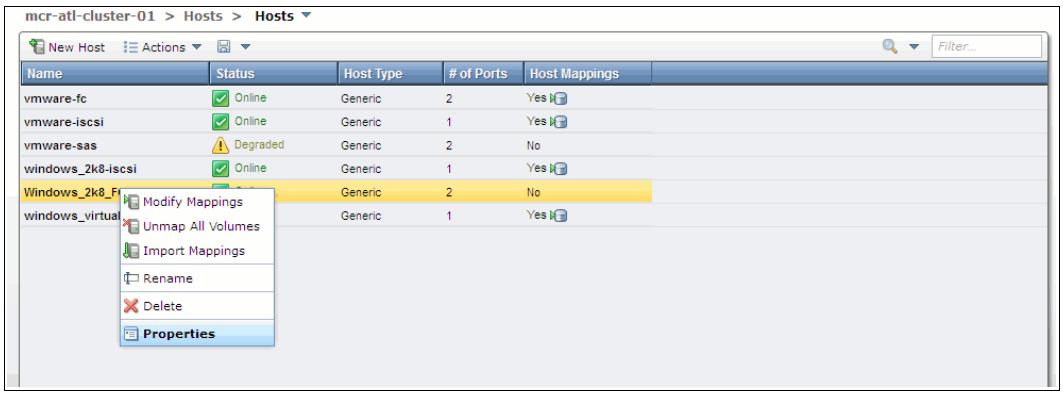

Figure 5-35 Host properties

2. Browse to the Mapped Volumes tab, as shown in Figure 5-36.

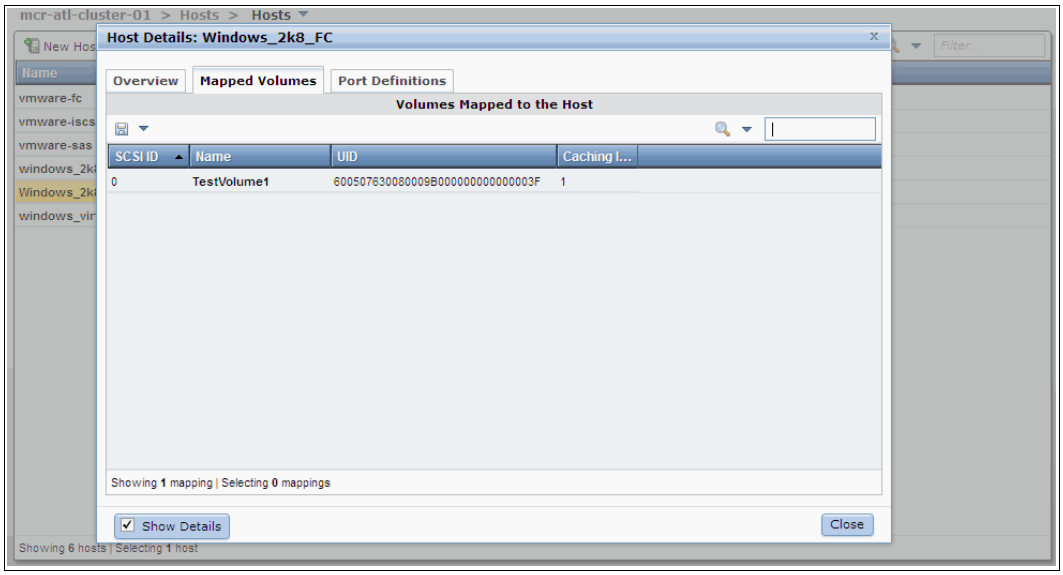

Figure 5-36 Mapped volumes to a host

The host details show you which volumes are mapped to the host. You also see the volume UID and the SCSI ID. In our example, one volume with SCSI ID 0 is mapped to the host.

3. If MPIO is not already installed on your Windows 2008 host and IBM Subsystem Device Driver is not yet installed, follow the procedure that is described in Chapter 4, “Host configuration” on page 153.

4. Log on to your Microsoft host and click **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A command-line interface (CLI) opens. Enter `datapath query device` and press Enter to see whether there are IBM Storwize V5000 disks connected to this host, as shown in Example 5-1.

Example 5-1 Datapath query device

```
C:\Program Files\IBM\SDDDSM>datapath query device

Total Devices : 3


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507630080009B000000000000003F
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL          0          0
    1     Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL         23          0
    2     Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL          0          0
    3     Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL         21          0
```

The output provides information about the connected volumes. In our example, one disk is connected (Disk 1) and four paths to the disk are available (State = Open).

Important: Correct SAN switch zoning must be implemented to allow only eight paths to be visible from the host to any one volume. Volumes with more than eight paths are not supported. For more information, see Chapter 2, “Initial configuration” on page 27.

5. Open the Windows Disk Management window (as shown in Figure 5-37) by clicking **Start** → **Run,** enter `diskmgmt.msc`, and click **OK**.

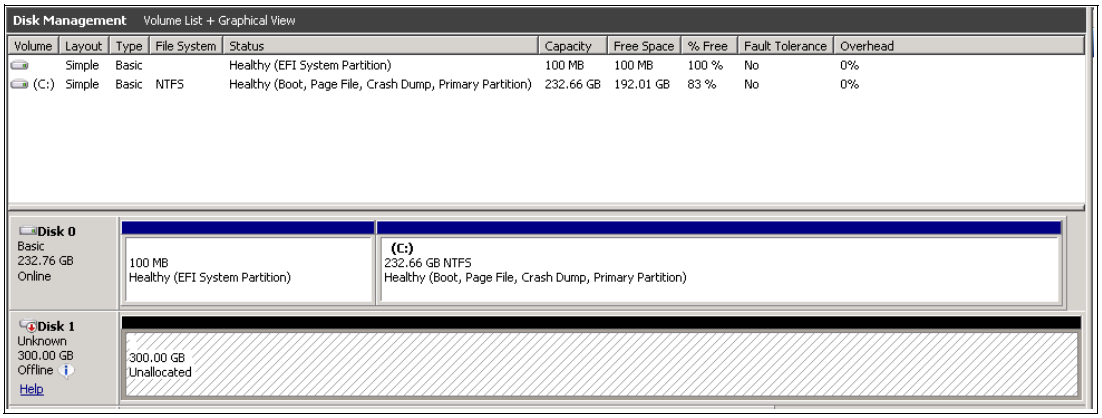

Figure 5-37 Windows Disk Management

6. Right-click the disk in the left pane and select **Online** if the disk is not online already, as shown in Figure 5-38.

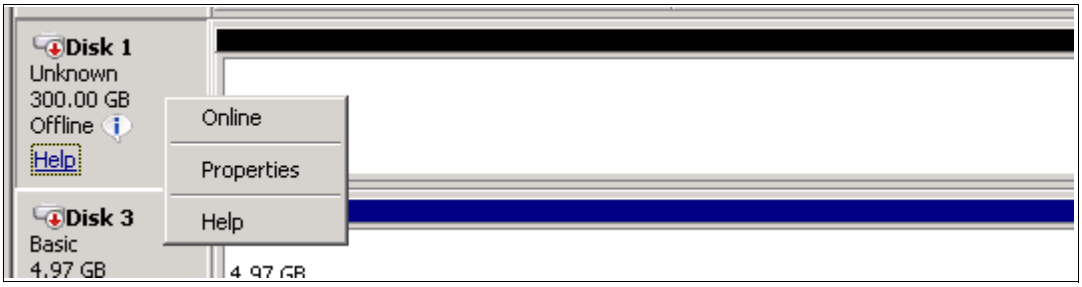

Figure 5-38 Setting a disk online

7. Right-click the disk again and then click **Initialize Disk**, as shown in Figure 5-39.

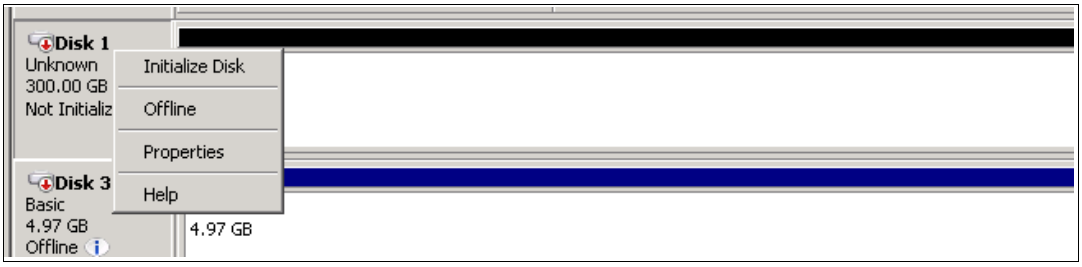

Figure 5-39 Initializing disk

8. Select an initialization option and click **OK**. In our example, we selected MBR, as shown in Figure 5-40.

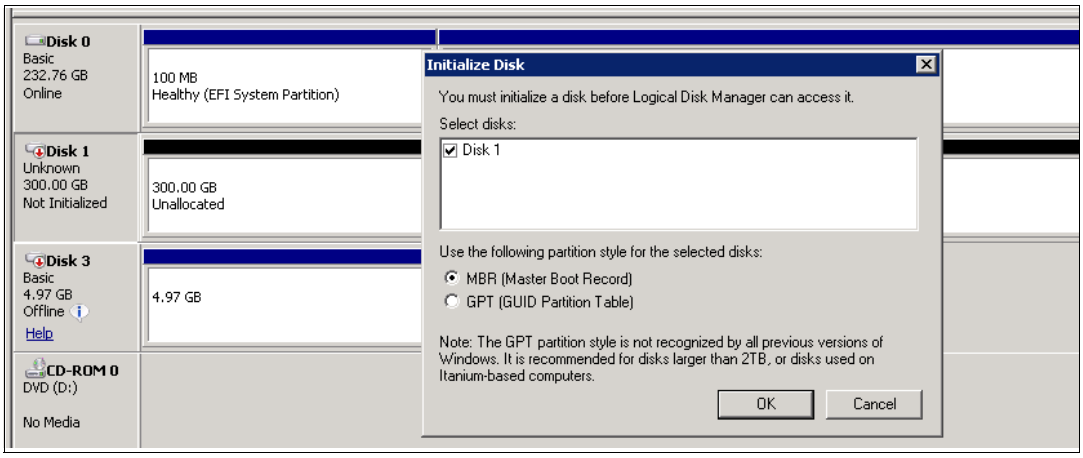

Figure 5-40 Initialize Disk option

9. Right-click the pane on the right side and click **New Simple Volume**, as shown in Figure 5-41.

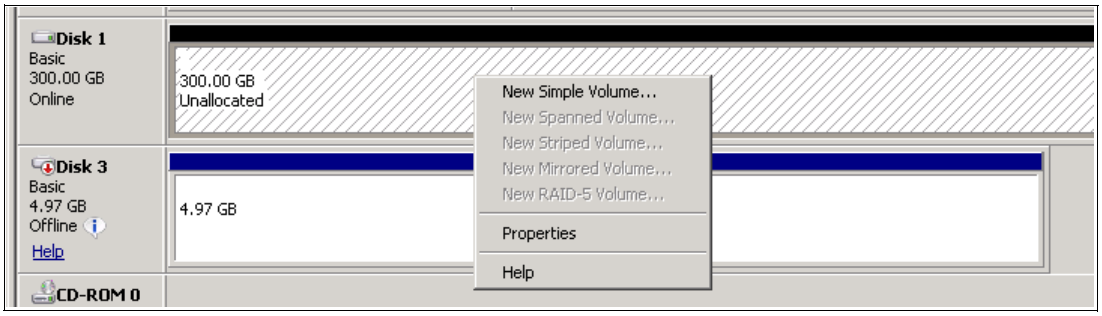

Figure 5-41 New Simple Volume

10. The New Simple Volume wizard starts, as shown in Figure 5-42 on page 190. Follow the wizard and the volume is ready to use from your Windows host, as shown in Figure 5-43 on page 190. In our example, we mapped a 300 GB disk on the IBM Storwize V5000 to a Windows 2008 host by using Fibre Channel connectivity.

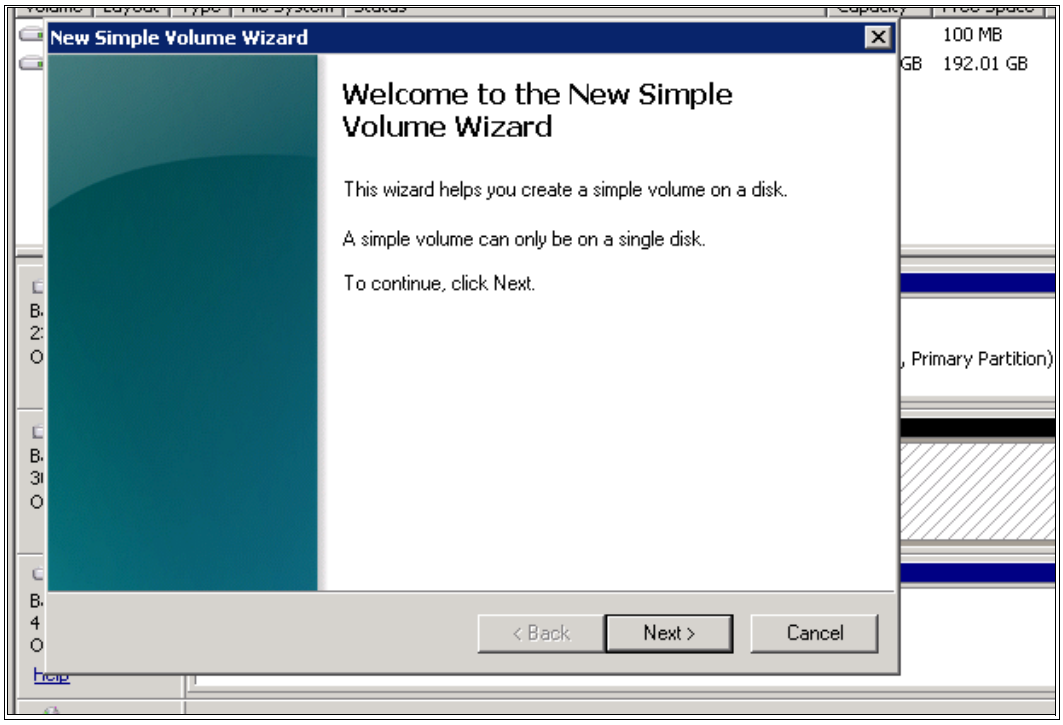

Figure 5-42 New Volume wizard

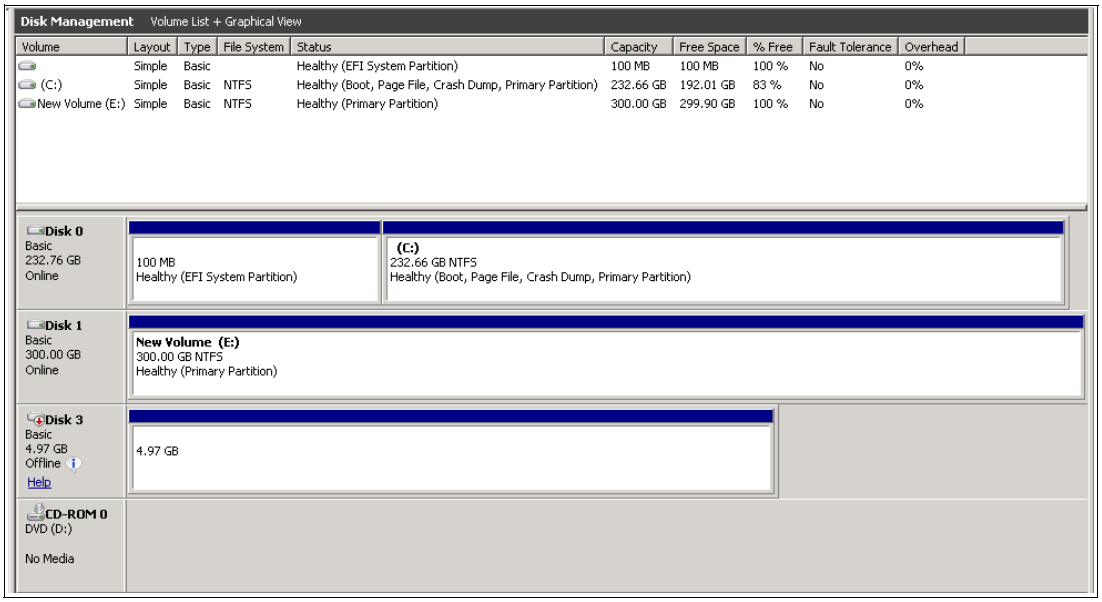

Figure 5-43 Volume is ready to use

Windows device discovery: Windows often automatically discovers new devices, such as, disks. If you completed all the steps that are presented here and do not see any disks, click **Actions** → **Rescan Disk** in Disk Management to discover the new volumes, as shown in Figure 5-44 on page 191.

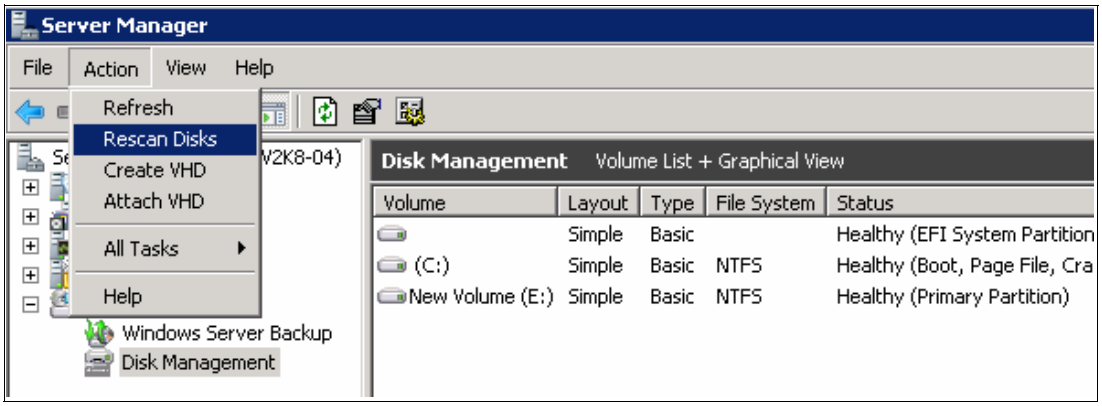

Figure 5-44 Windows disk rescan

The basic setup is now complete and the IBM Storwize V5000 is configured. The host is prepared and can access the volumes over several paths and store data on the storage subsystem.

5.3.2 Windows 2008 iSCSI volume attachment

To perform iSCSI volume attachment in Windows 2008, complete the following steps:

1. Right-click your Windows 2008 iSCSI host in the Hosts view and click **Properties,** as shown in Figure 5-45.

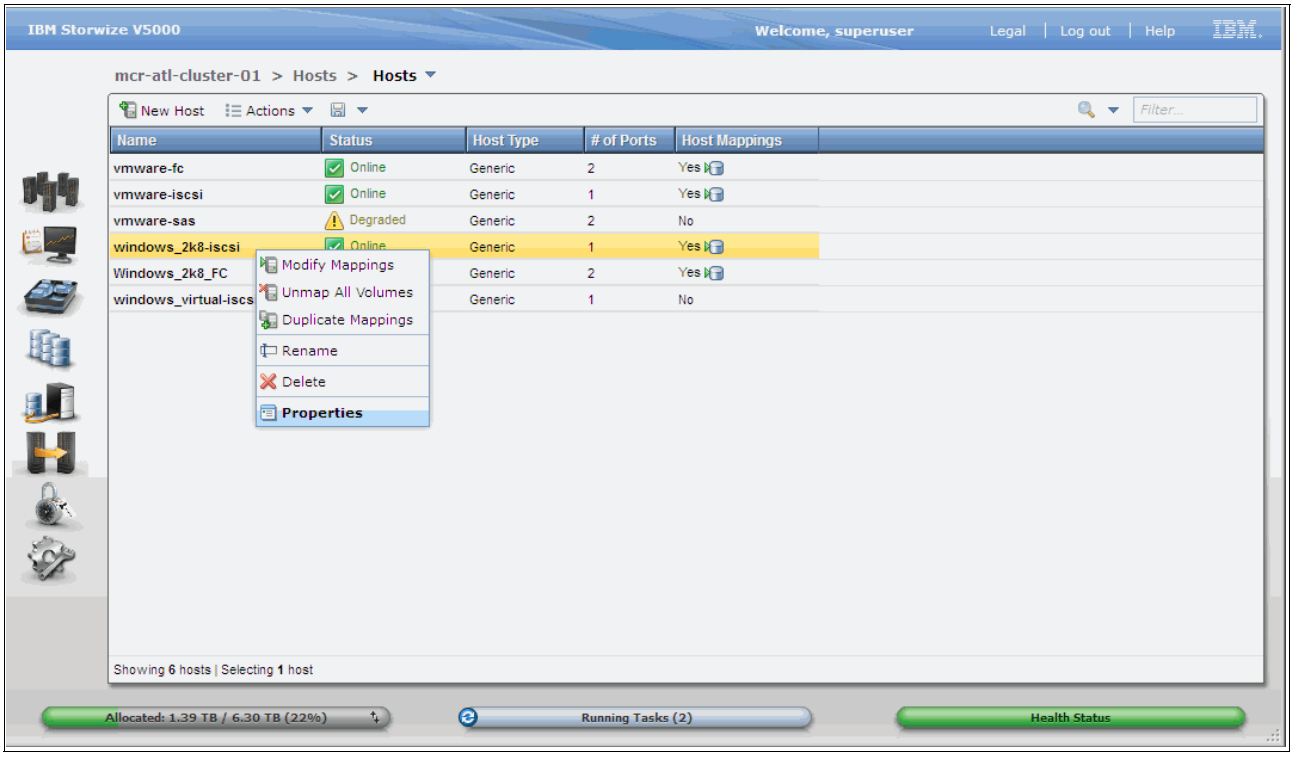

Figure 5-45 All Hosts view

2. Browse to the Mapped Volumes tab, as shown in Figure 5-46.

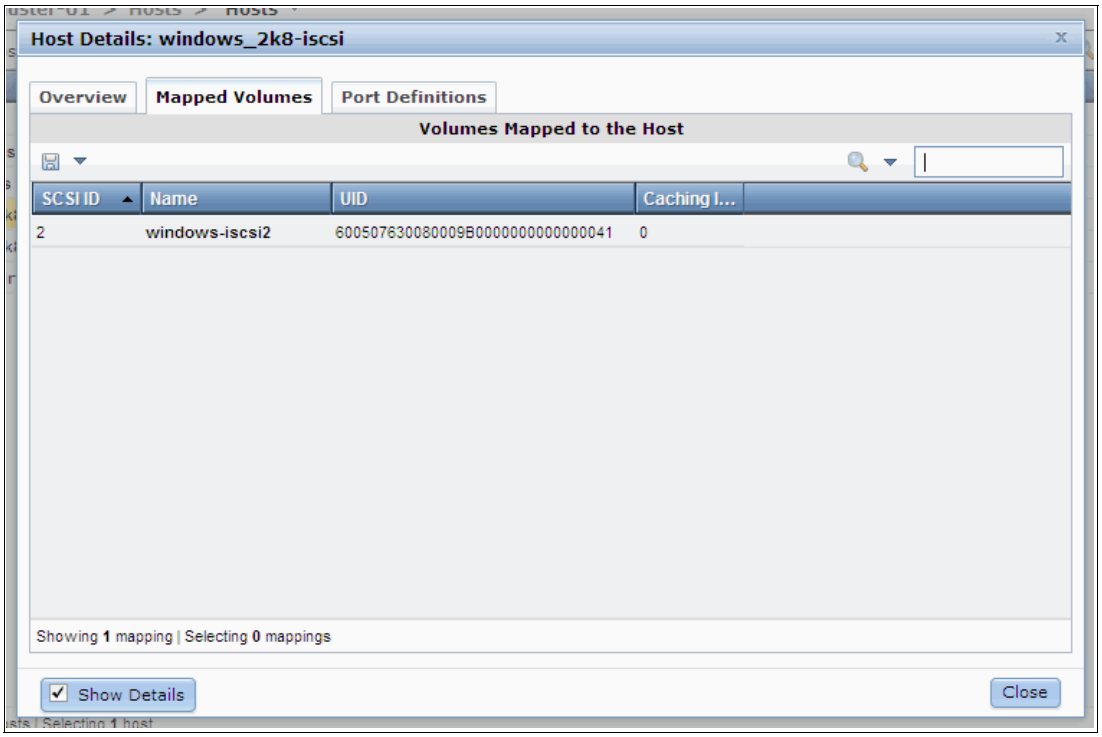

Figure 5-46 Mapped volumes on an iSCSI host

The host details show you which volumes are mapped to the host. You also can see the volume UID and the SCSI ID. In our example, one volume with SCSI ID 2 is mapped to the host.

3. Log on to your Windows 2008 host and click **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Configuration tab, as shown in Figure 5-47.

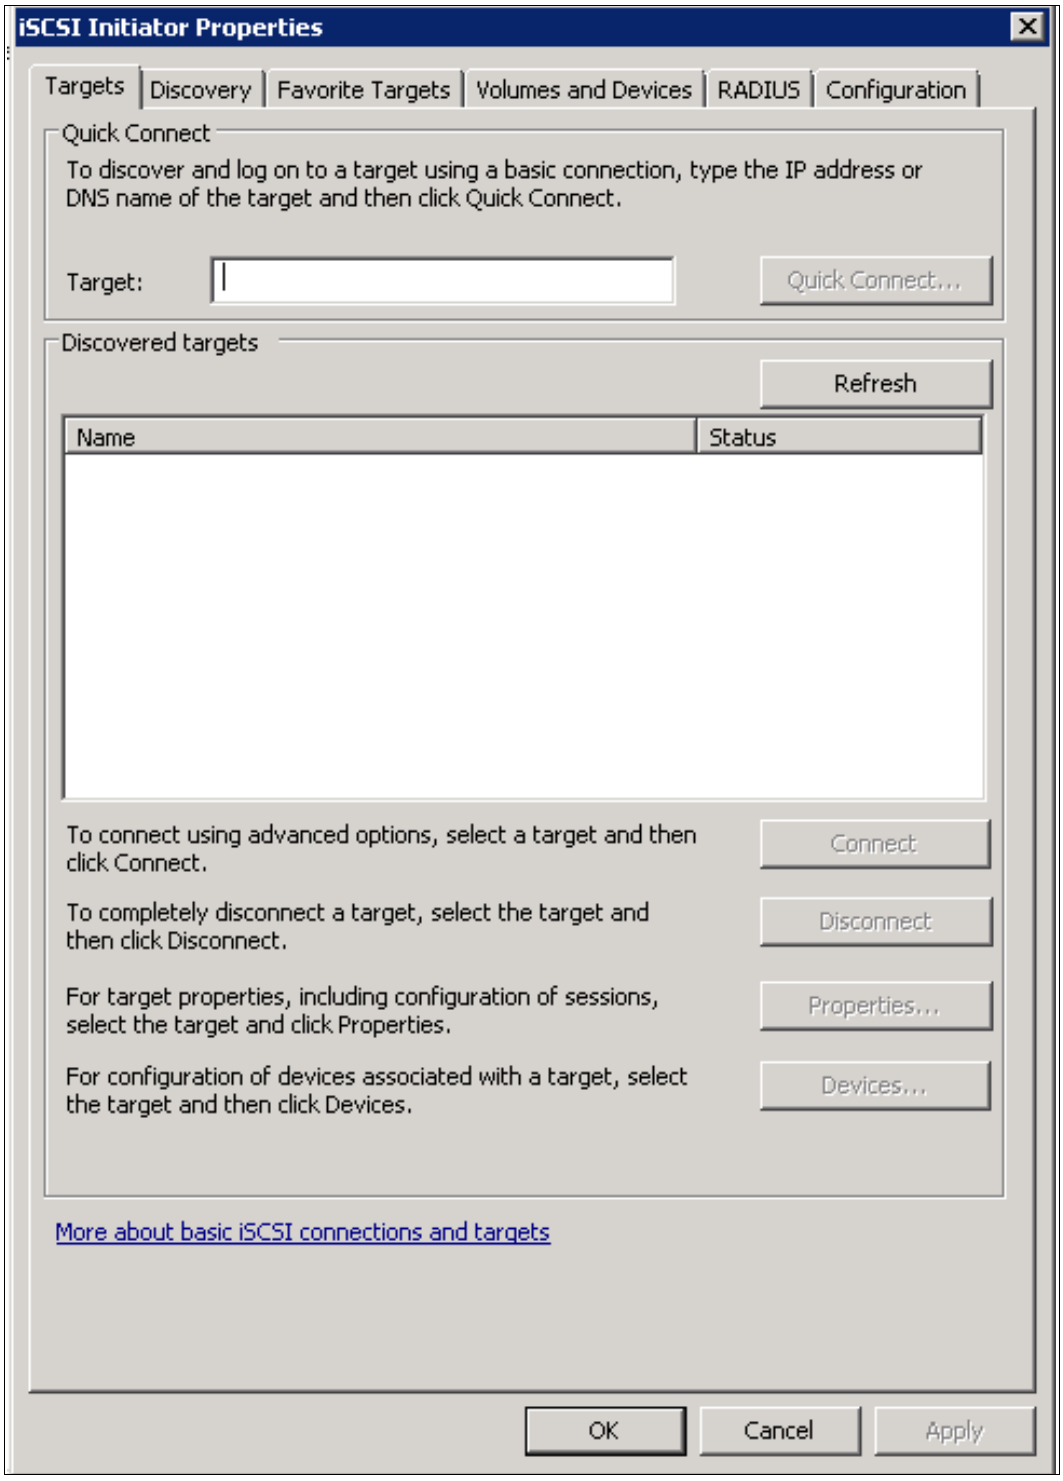

Figure 5-47 Windows iSCSI Configuration tab

4. Enter the IP address of one of the IBM Storwize V5000 iSCSI ports in the Target field at the top of the panel and click **Quick Connect**, as shown in Figure 5-48.

iSCSI IP addresses: The iSCSI IP addresses are different for the cluster and canister IP addresses. They are configured as described in Chapter 4, “Host configuration” on page 153.

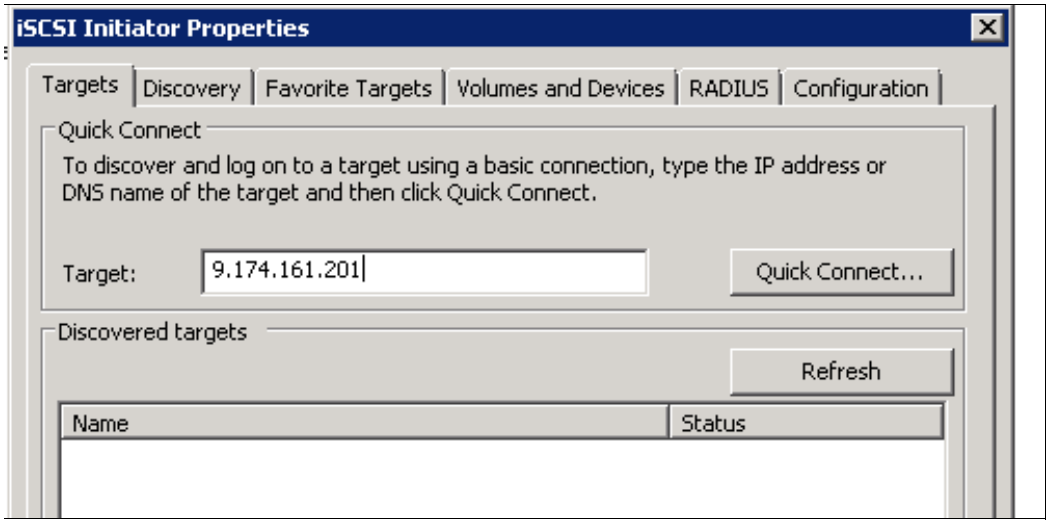

Figure 5-48 iSCSI Quick Connect

The IBM Storwize V5000 initiator is discovered and connected, as shown in Figure 5-49.

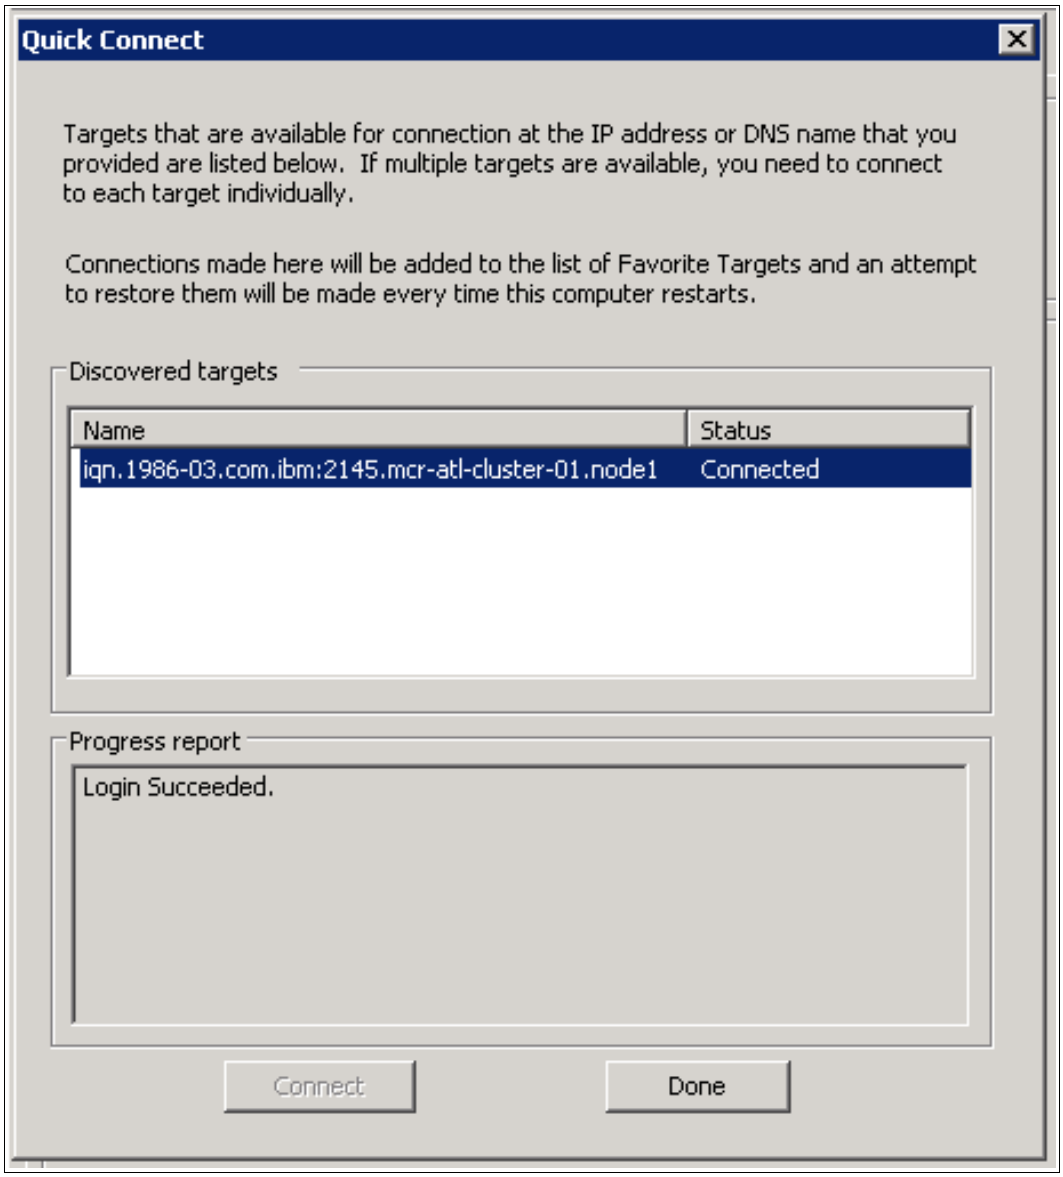

Figure 5-49 iSCSI Initiator target is connected

5. Click **Done** to return to the iSCSI Initiator Properties window.

The storage disk is connected to your iSCSI host, but only a single path is used. To enable multipathing for iSCSI targets, complete the following steps:

1. If MPIO is not already installed on your Windows 2008 host, follow the procedure that is described in 4.2.1, “Windows 2008 R2: Preparing for FC attachment” on page 155. IBM Sub System Device Driver is not required for iSCSI connectivity.

2. Click **Start** → **Administrative Tools** → **MPIO**, click the **Discover Multi-Paths** tab, and select **Add support for iSCSI devices**, as shown in Figure 5-50.

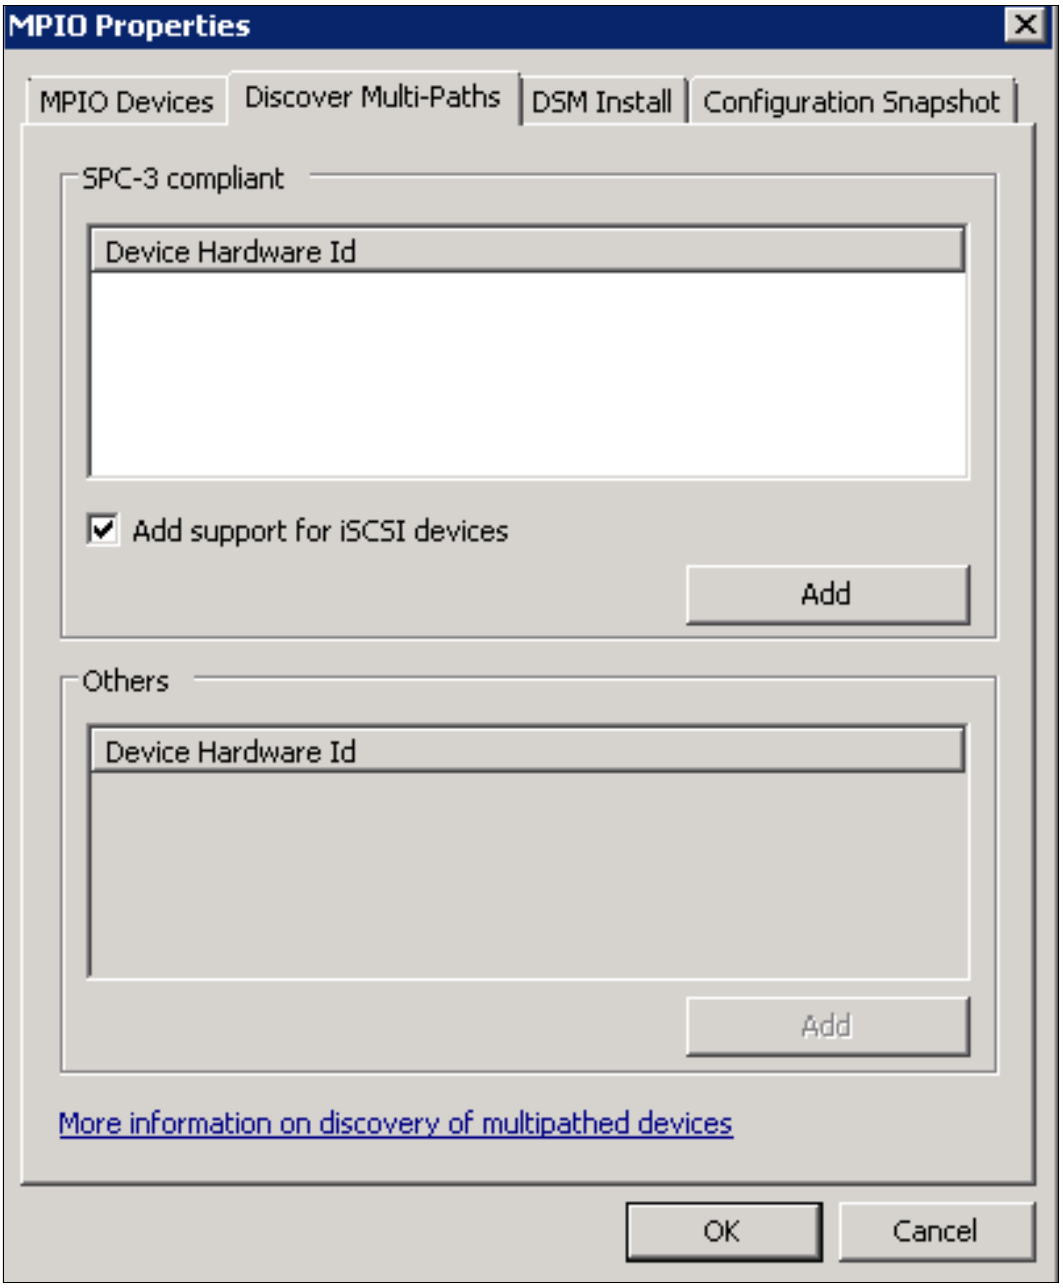

Figure 5-50 Enable iSCSI MPIO

Important: In some cases, the Add support for iSCSI devices option is disabled. To enable this option, you must already have a connection to at least one iSCSI device.

3. Click **Add** and confirm the prompt to reboot your host.

4. After the reboot process is complete, log on again and click **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Configuration tab. Browse to the Discovery tab, as shown in Figure 5-51.

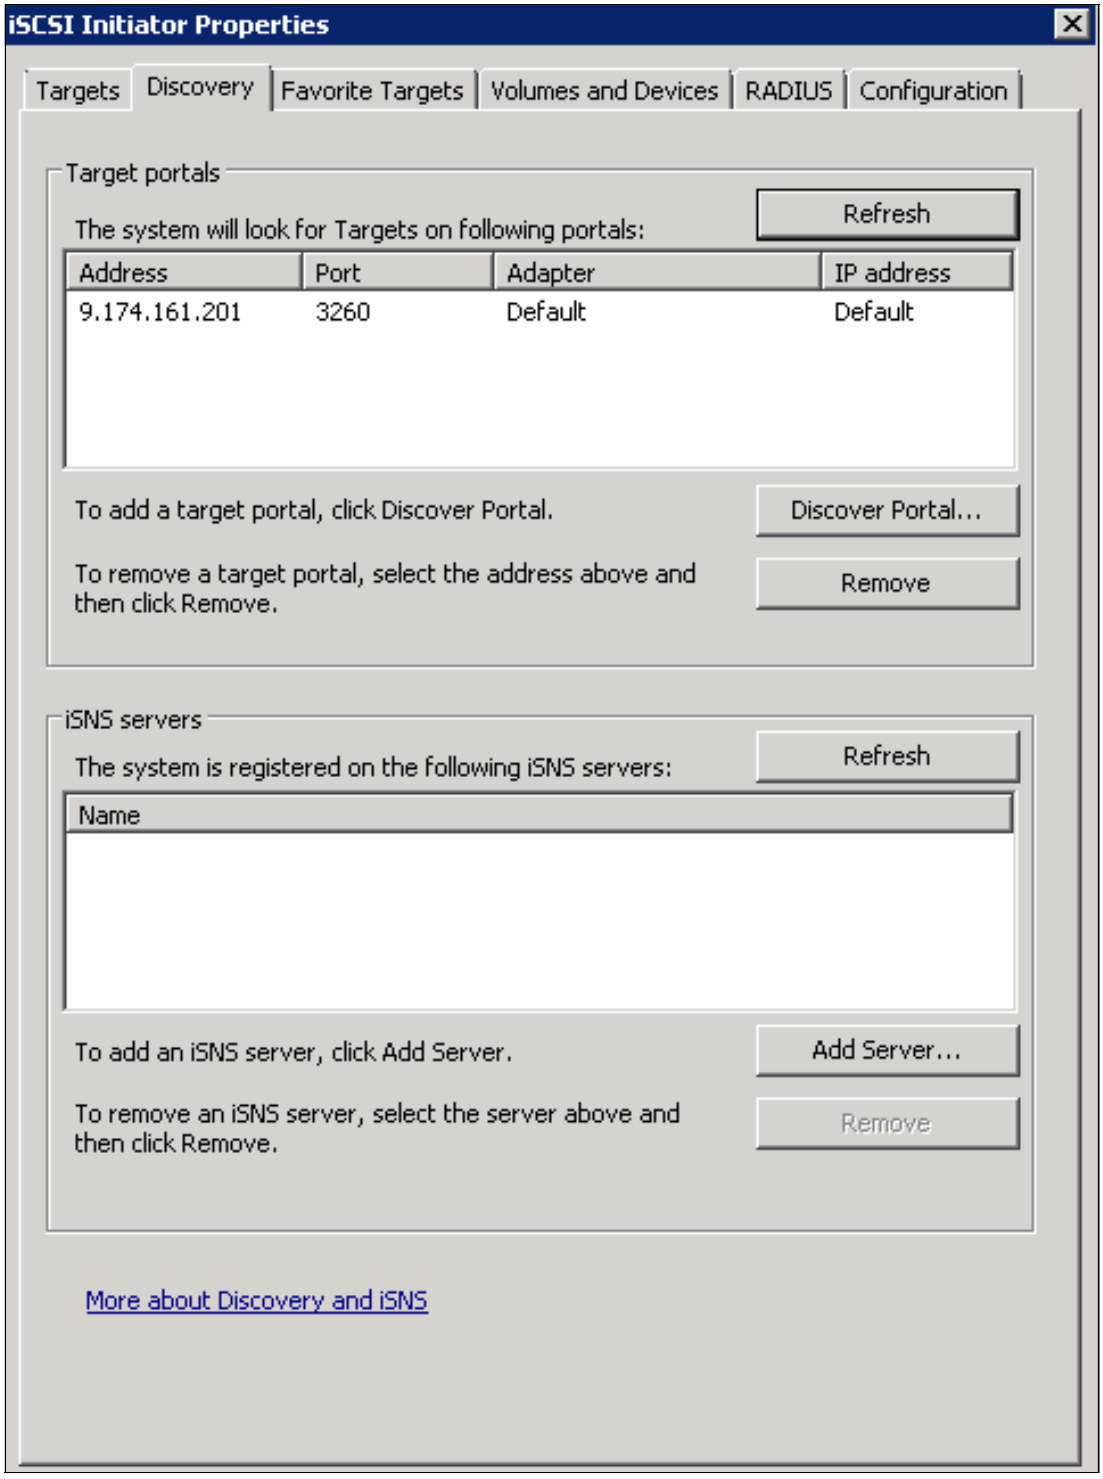

Figure 5-51 iSCSI Properties Discovery tab

5. Click **Discover Portal...**, enter the IP address of another IBM Storwize V5000 iSCSI port (as shown in Figure 5-52), and click **OK**.

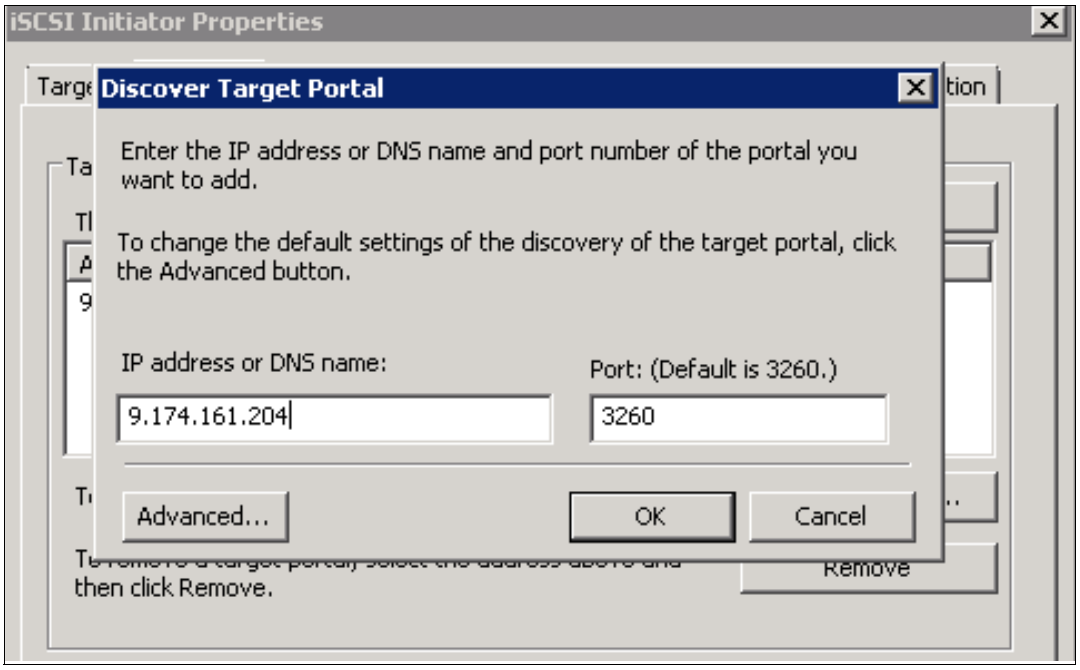

Figure 5-52 Discover Target Portal window

6. Return to the Targets tab (as shown in Figure 5-53) and you see that the new connection there is listed as Inactive.

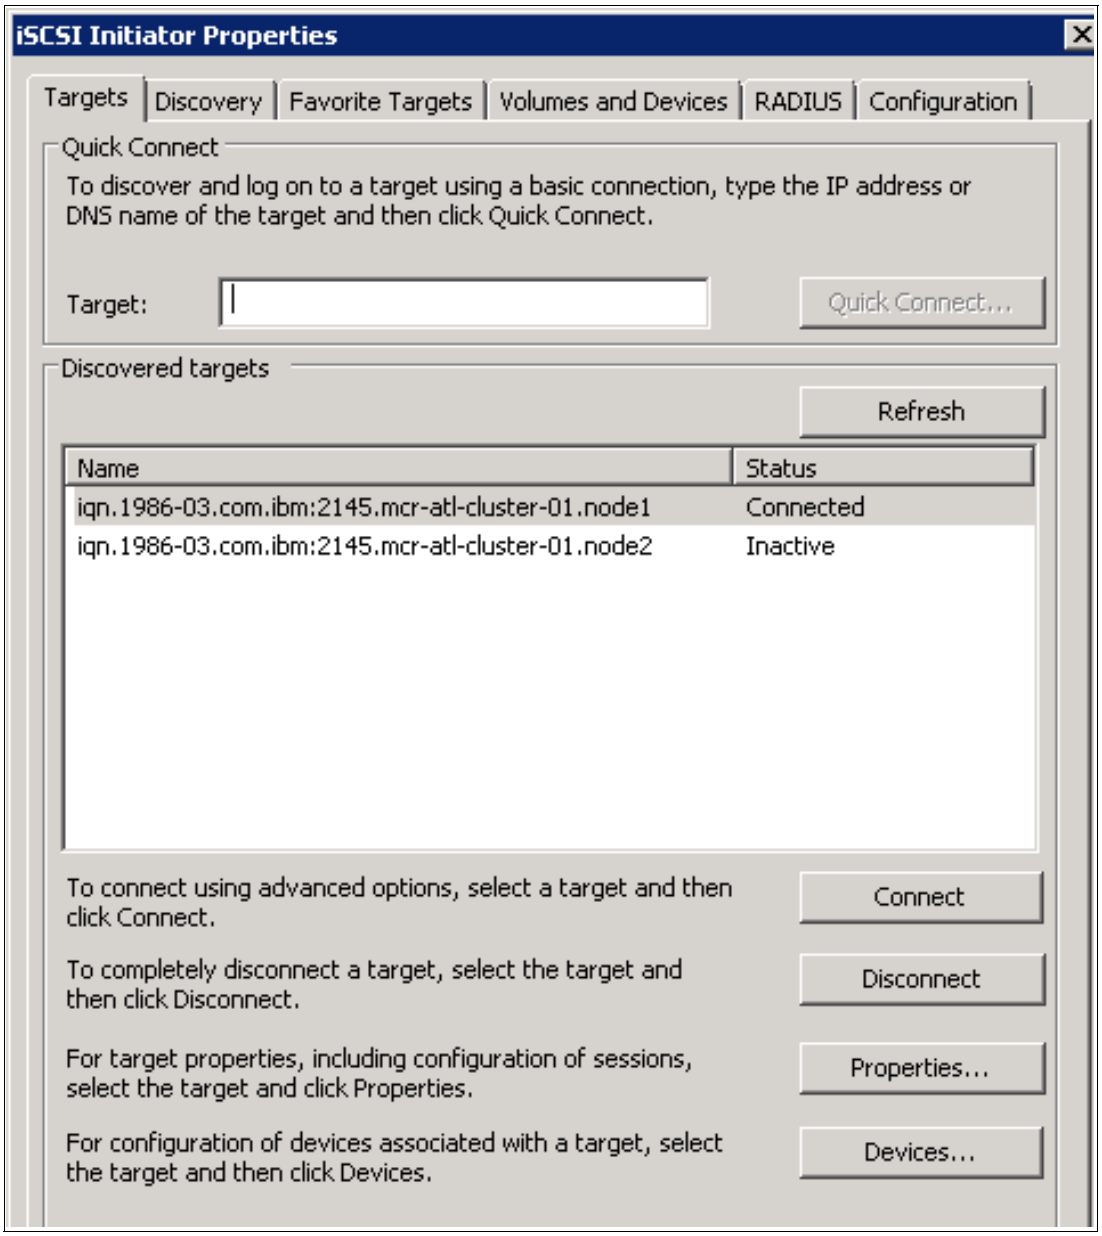

Figure 5-53 Inactive target ports

7. Highlight the inactive port and click **Connect**. The Connect to Target window opens, as shown in Figure 5-54.

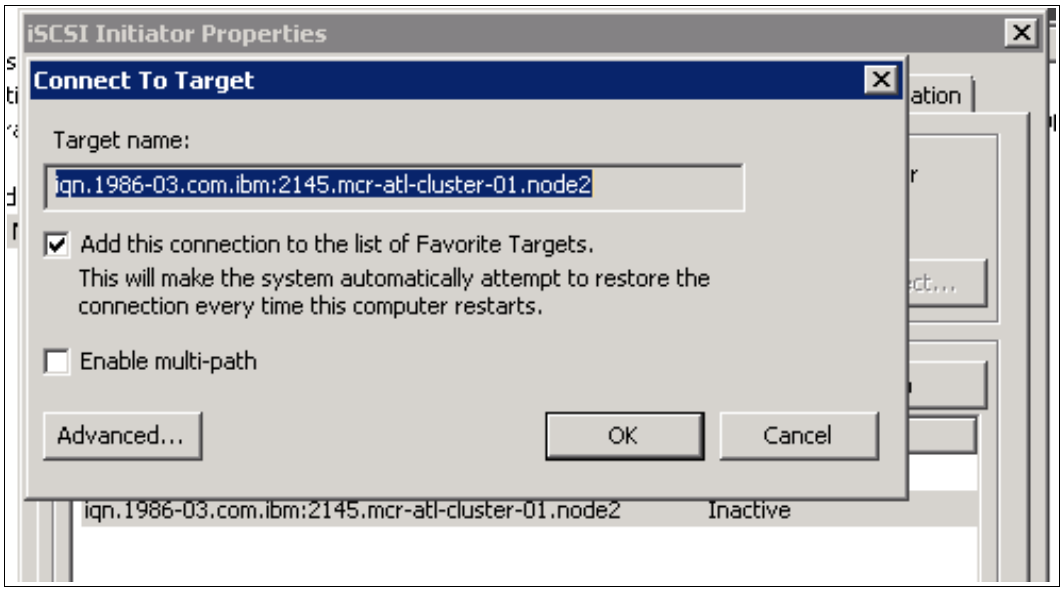

Figure 5-54 Connect to a target

8. Select **Enable Multipath** and click **OK**. The second port is now Connected, as shown in Figure 5-55.

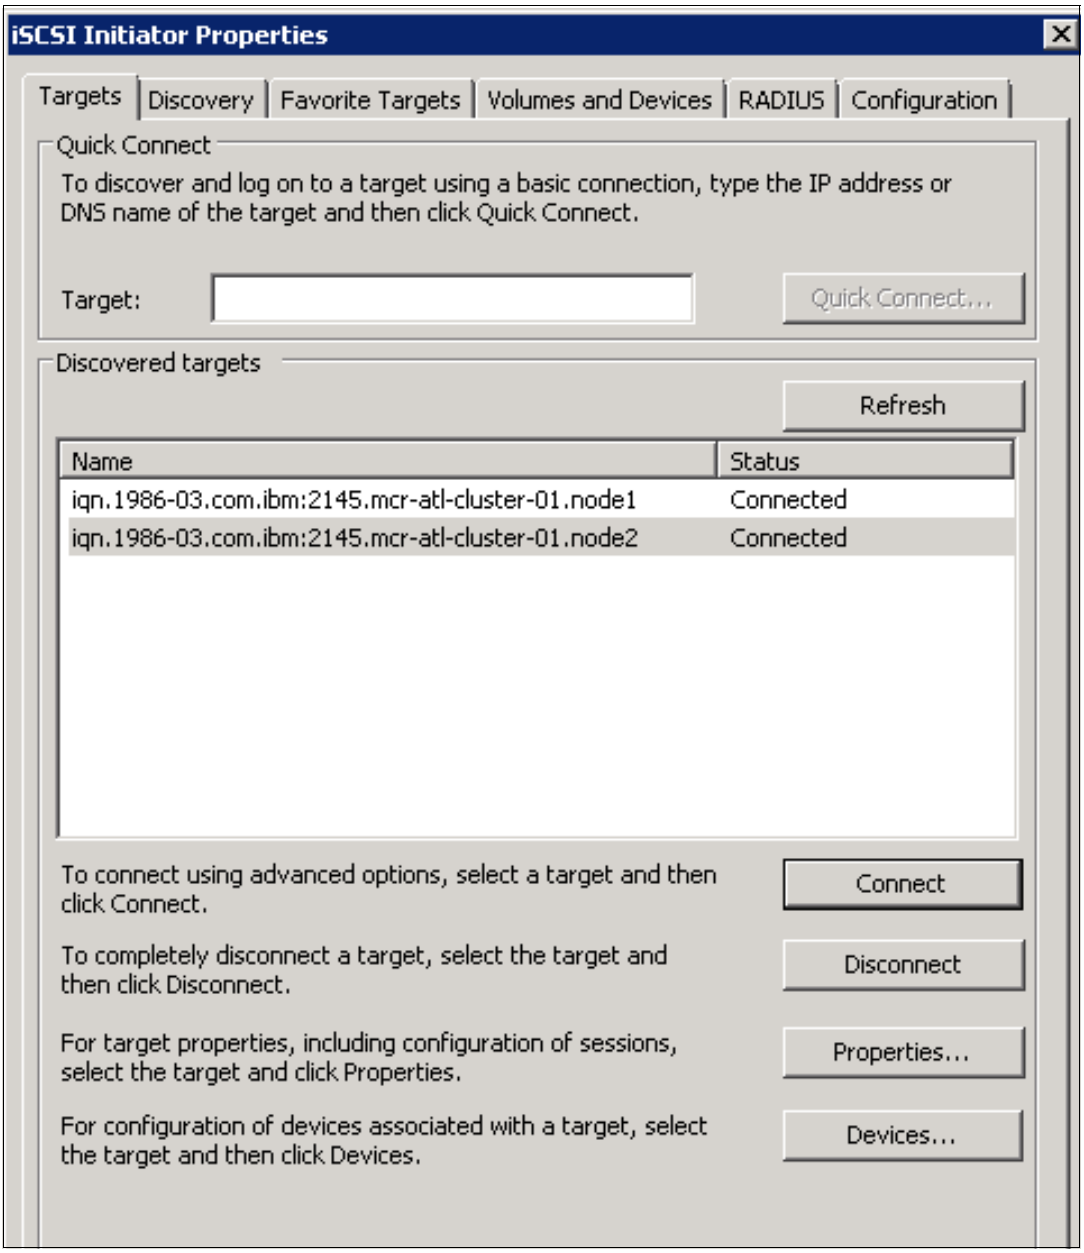

Figure 5-55 Second target port connected

Repeat this step for each IBM Storwize V5000 port you want to use for iSCSI traffic. It is possible to have up to four port paths to the system.

9. Open the Windows Disk Management window (as shown in Figure 5-56) by clicking **Start** → **Run**, entering `diskmgmt.msc`, and then clicking **OK**.

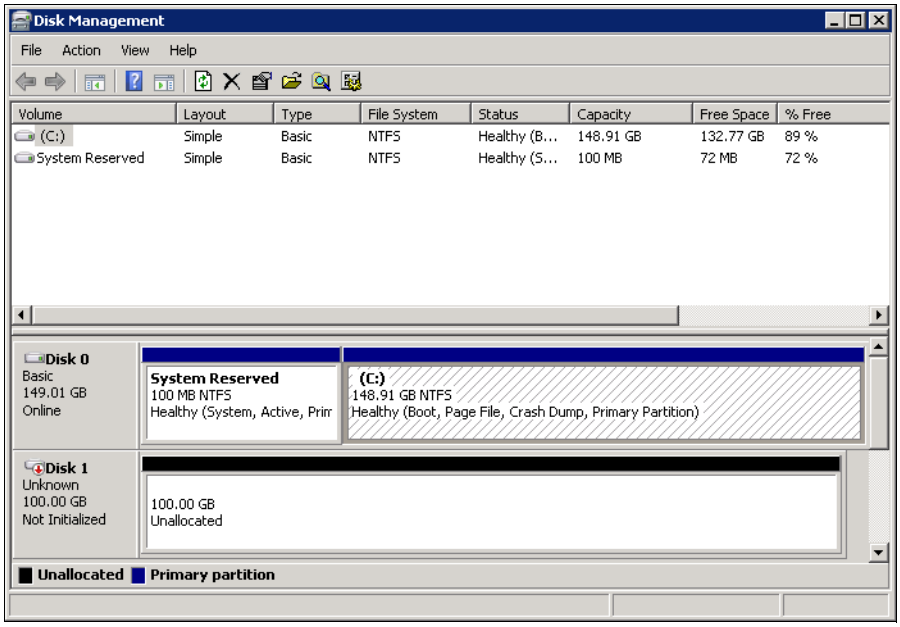

Figure 5-56 Windows Disk Management

10. Set the disk online, initialize it, and then create a file system on it as described in step 6 - 10 of 5.3.1, “Windows 2008 Fibre Channel volume attachment” on page 186. The disk is now ready to use, as shown in Figure 5-57. In our example, we mapped a 5 GB disk to a Windows 2008 host that uses iSCSI connectivity.

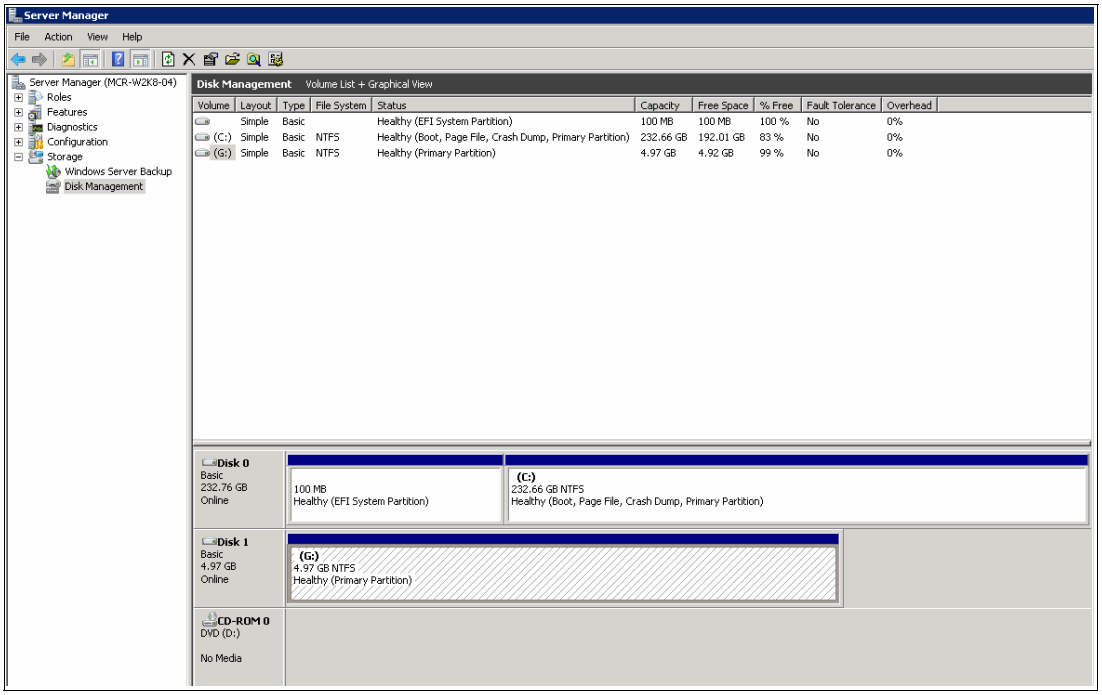

Figure 5-57 Disk is ready to use

5.3.3 Windows 2008 Direct SAS volume attachment

To attach an SAS volume in Windows 2008, complete the following steps:

1. Right-click your Windows 2008 SAS host in the Hosts view and select **Properties**, as shown in Figure 5-58.

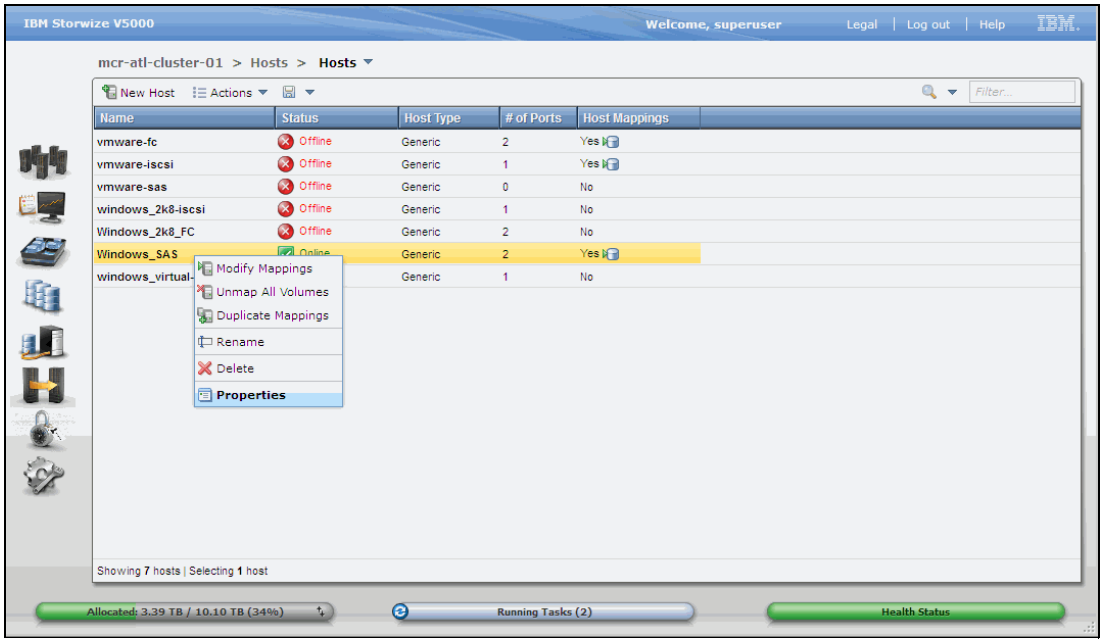

Figure 5-58 Windows SAS host from host view

2. Browse to the **Mapped Volumes** tab, as shown in Figure 5-59.

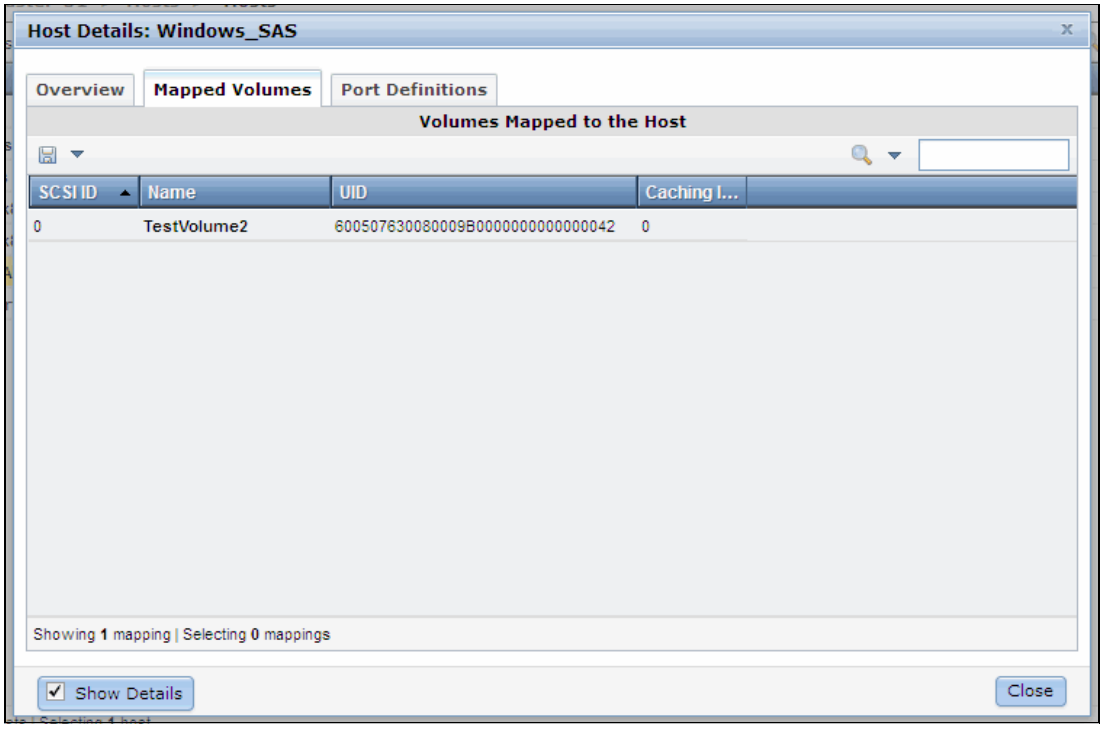

Figure 5-59 SAS host mapped volumes

The Mapped Volumes tab shows you which volumes are mapped to the host. You also see the volume UID and the SCSI ID. In our example, one volume with SCSI ID 0 is mapped to the host.

3. If MPIO is not already installed on your Windows 2008 host and IBM Subsystem Device Driver is not yet installed, follow the procedure that is described in 4.2.1, “Windows 2008 R2: Preparing for FC attachment” on page 155.
4. Log on to your Microsoft host and click **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A CLI opens. Enter `datapath query device` and press Enter to see whether there are IBM Storwize V5000 disks connected to this host, as shown in Example 5-2.

Example 5-2 SDDDSM output SAS attached host

```
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation.  All rights reserved.

C:\Program Files\IBM\SDDDSM>datapath query device

Total Devices : 1


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507630080009B0000000000000042
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL         70          0
    1     Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL          0          0

C:\Program Files\IBM\SDDDSM>
```

The output provides information about the connected volumes. In our example, there is one disk connected (Disk 1) and two paths to the disk are available (State = Open).

5. Open the Windows Disk Management window (as shown in Figure 5-60) by clicking **Start** → **Run**, entering `diskmgmt.msc`, and then clicking **OK**.

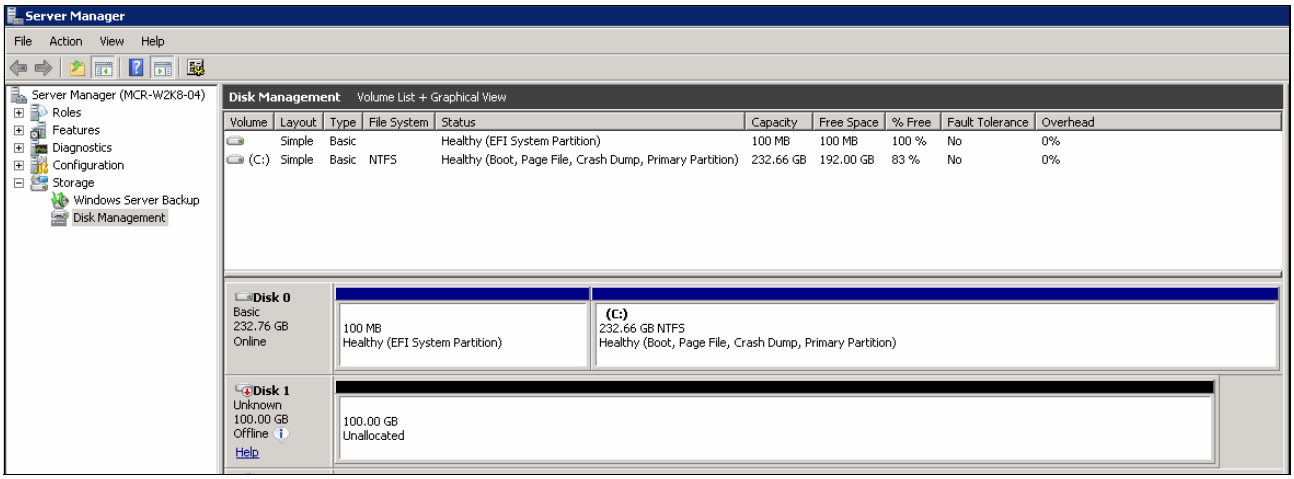

Figure 5-60 Windows Disk Management

6. Right-click the disk in the left pane and select **Online** if the disk is not online already, as shown in Figure 5-61.

Figure 5-61 Setting volume online

7. Right-click the disk again and then click **Initialize Disk**, as shown in Figure 5-62.

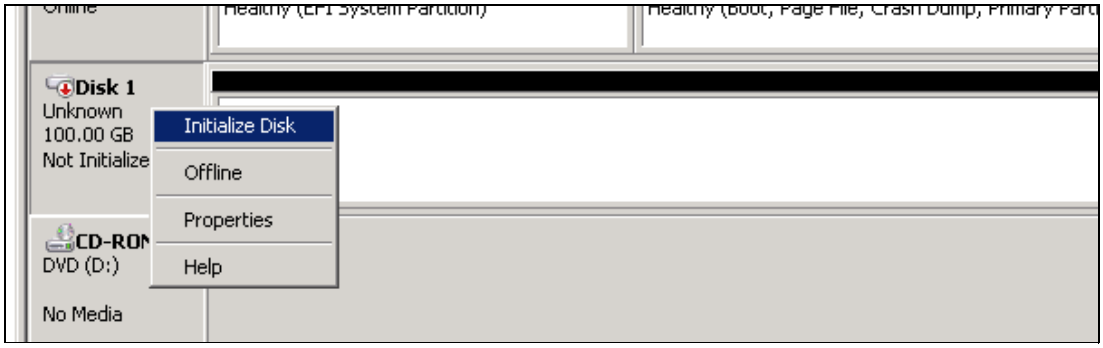

Figure 5-62 Initializing disk

8. Select an initialization option and click **OK**. In our example, we selected MBR, as shown in Figure 5-63.

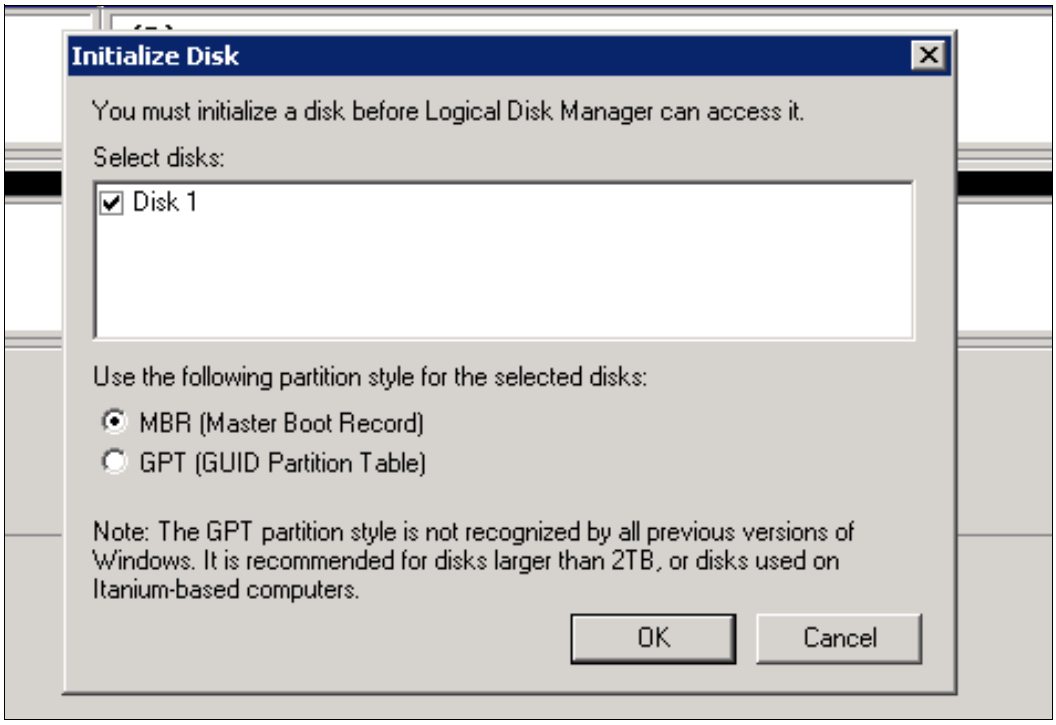

Figure 5-63 Initialize disk option

9. Right-click the pane on the right side and click **New Simple Volume**, as shown in Figure 5-64.

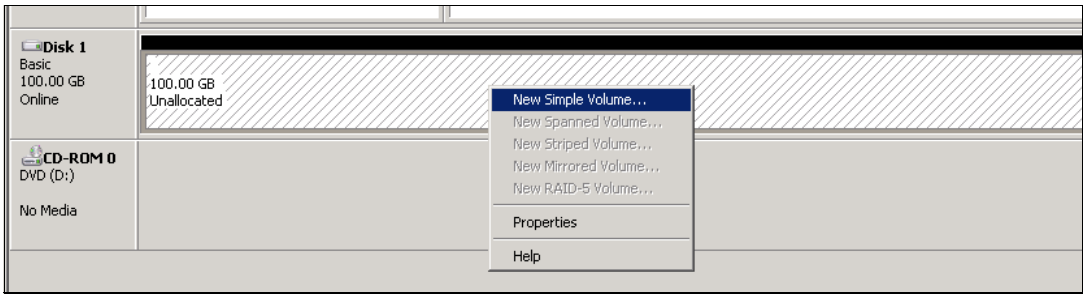

Figure 5-64 New simple volume: SAS attach

10. The New Simple Volume wizard starts, as shown in Figure 5-65. Follow the wizard and the volume is ready to use from your Windows host, as shown in Figure 5-66 on page 207. In our example, we mapped a 100 GB disk on the IBM Storwize V5000 to a Windows 2008 host that uses SAS direct attach connectivity.

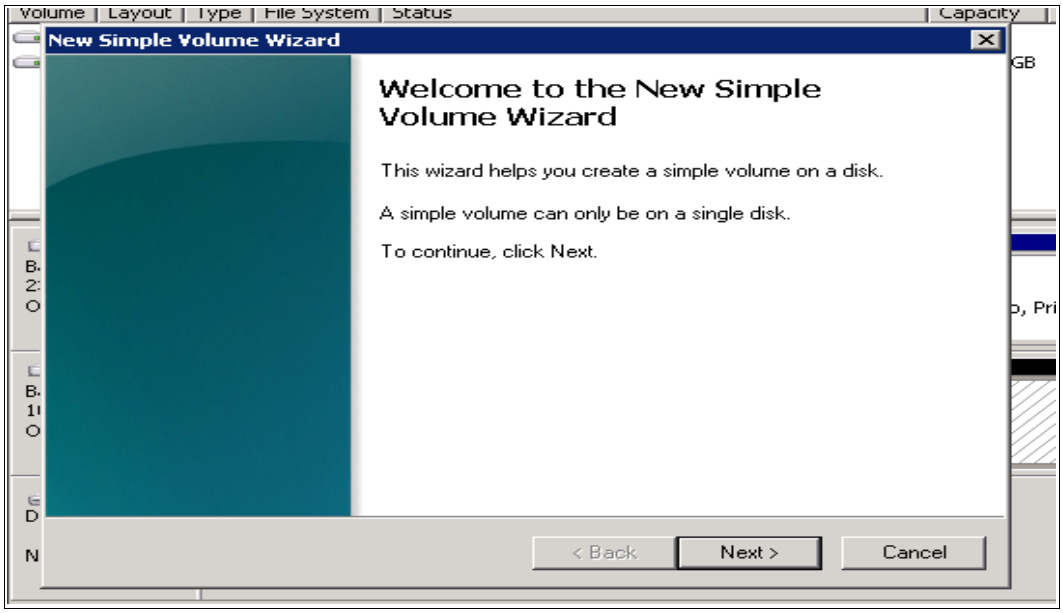

Figure 5-65 Simple volume wizard

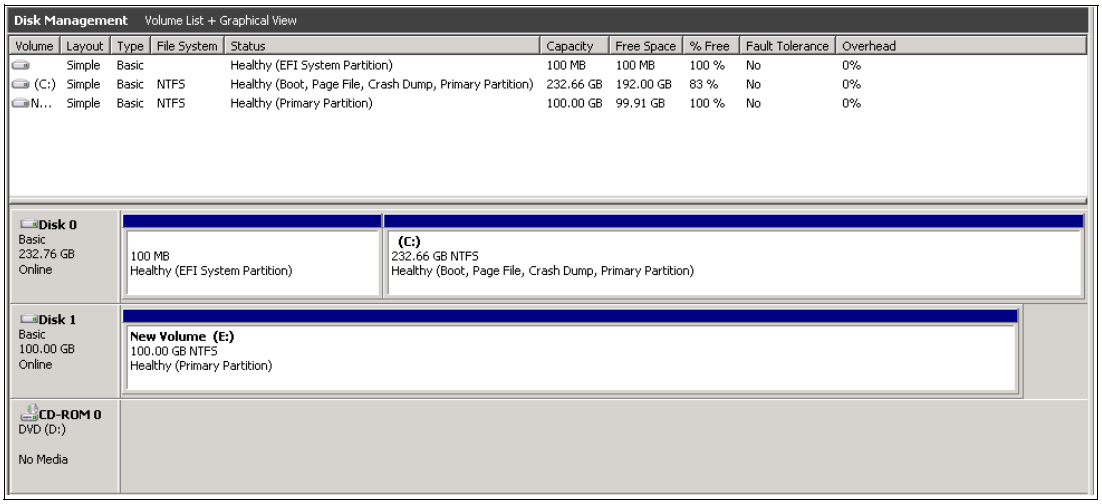

Figure 5-66 SAS attached volume ready to use

5.3.4 VMware ESX Fibre Channel volume attachment

To perform VMware ESX Fibre Channel attachment, complete the following steps:

1. Right-click your VMware ESX Fibre Channel host in the Hosts view and select **Properties**, as shown in Figure 5-67.

Figure 5-67 Example ESX FC host

2. Browse to the **Mapped Volumes** tab, as shown in Figure 5-68.

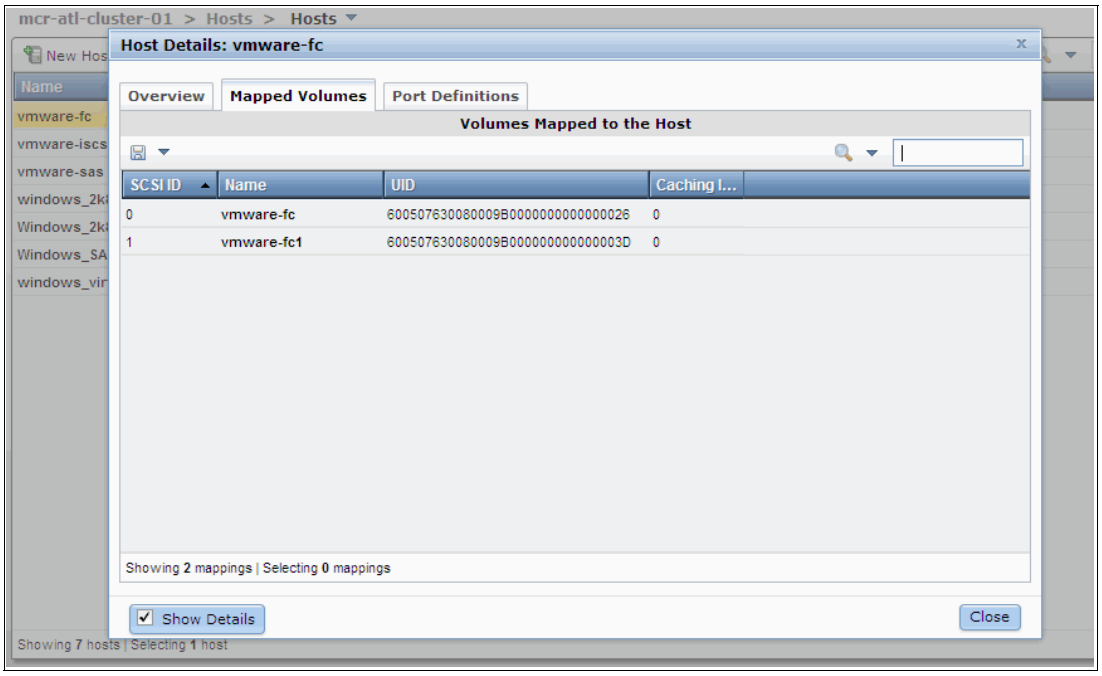

Figure 5-68 Mapped volumes to ESX FC host

In the Host Details window, there are two volumes connected to the ESX FC host that use SCSI ID 0 and SCSI ID 1. The UID of the volumes is also displayed.

3. Connect to your VMware ESX Server by using the vSphere client. Browse to the **Configuration** tab and select **Storage Adapters**, as shown in Figure 5-69.

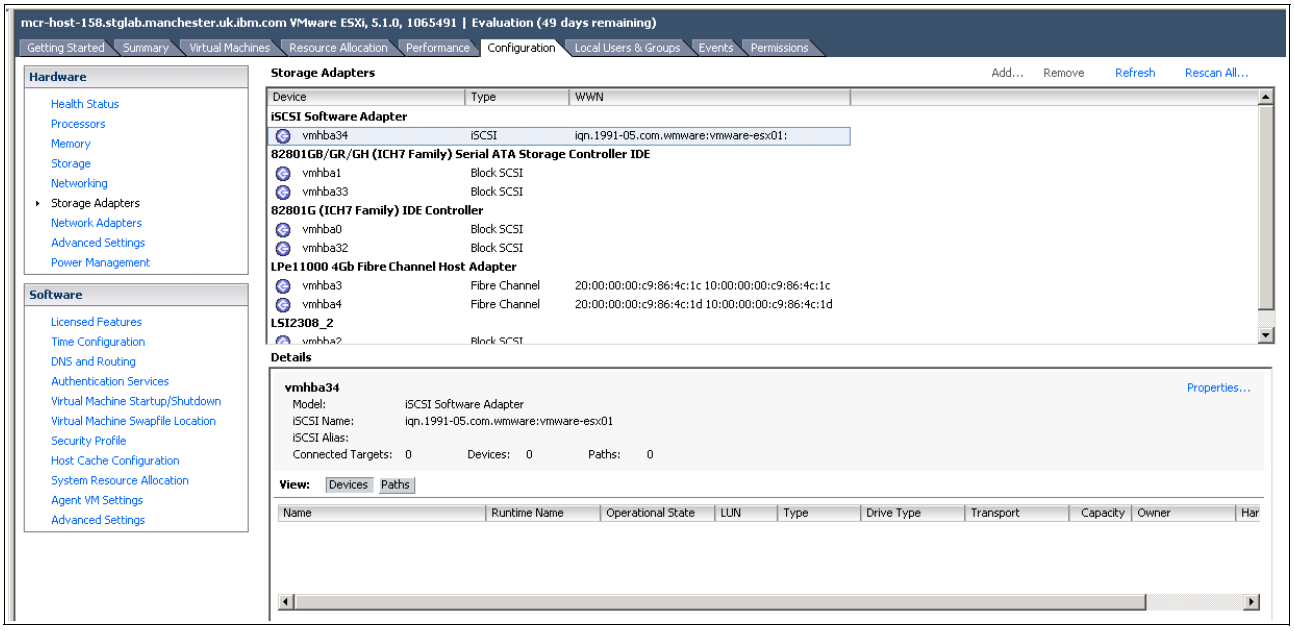

Figure 5-69 vSphere Client: Storage adapters

4. Click **Rescan All...** in the upper right corner and click **OK** in the resulting pop-up window, as shown in Figure 5-70. This scans for new storage devices.

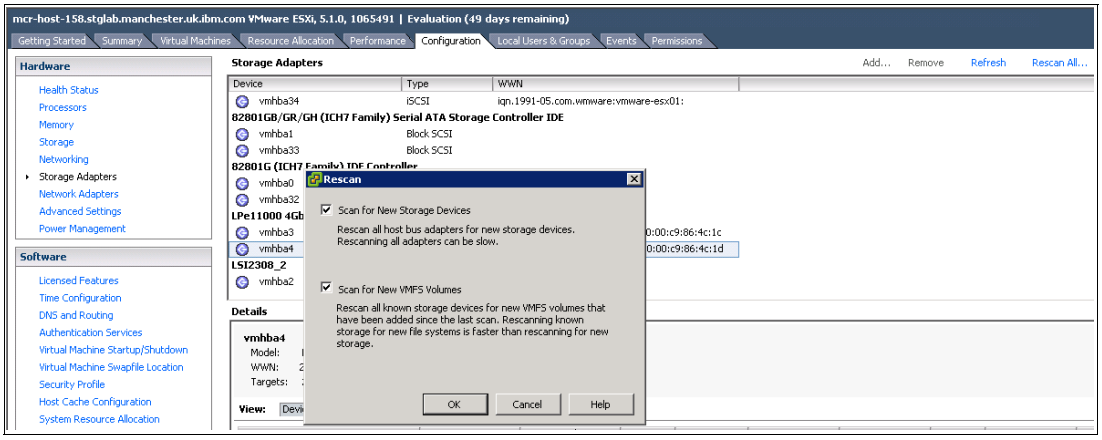

Figure 5-70 Rescan

The mapped volumes on the IBM Storwize V5000 should now appear against the Fibre Channel adapters.

5. Select **Storage** and then click **Add Storage**, as shown in Figure 5-71.

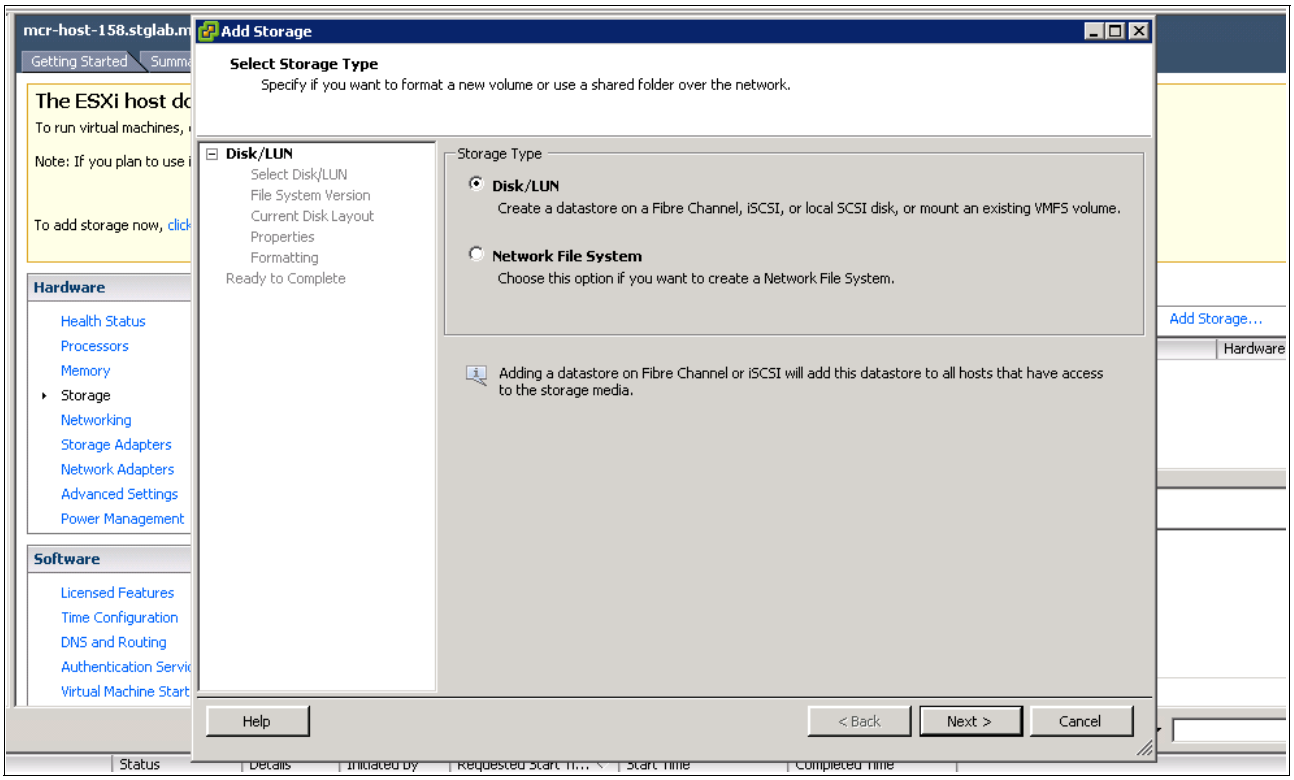

Figure 5-71 vSphere Client: Add Storage

6. The Add Storage wizard opens. Click **Select Disk/LUN** and click **Next**. The IBM Storwize V5000 disks appear, as shown in Figure 5-72. In our example, they are the Fibre Channel Disks. We continue with the 500 GB volume, which we highlight and then click **Next**.

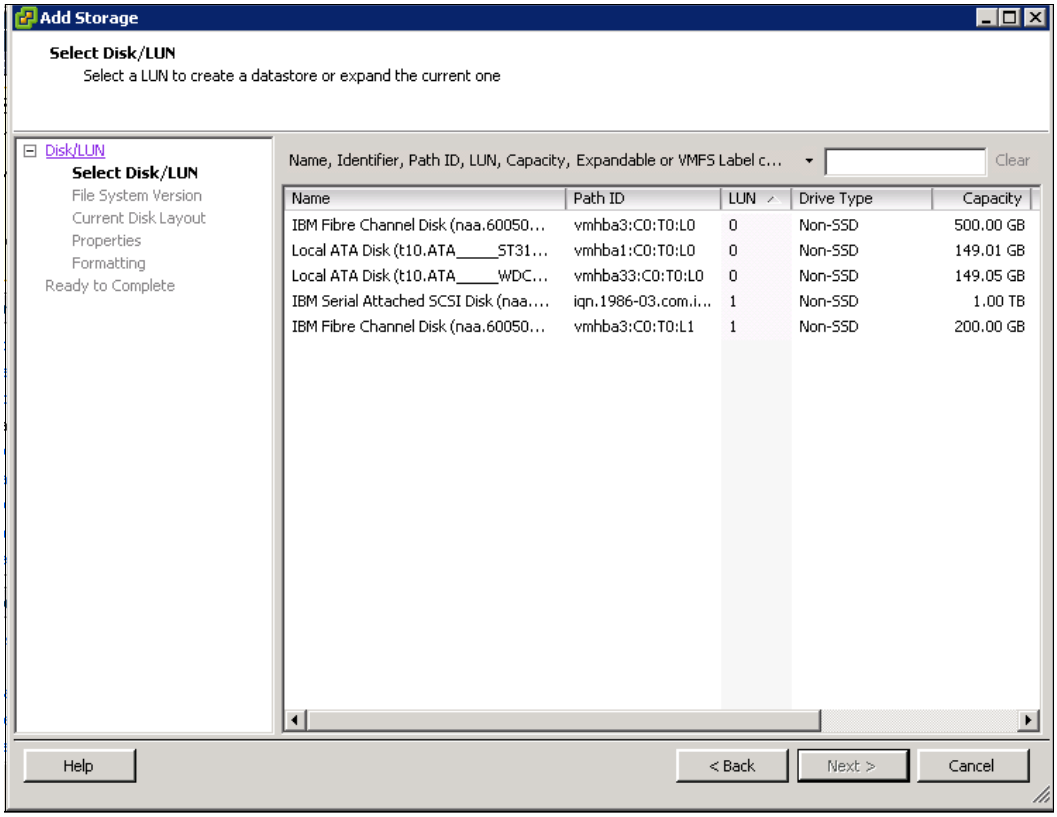

Figure 5-72 Select FC Disk

7. Select a File System Version option. In our example, we selected **VMFS-5**, as shown in Figure 5-73.

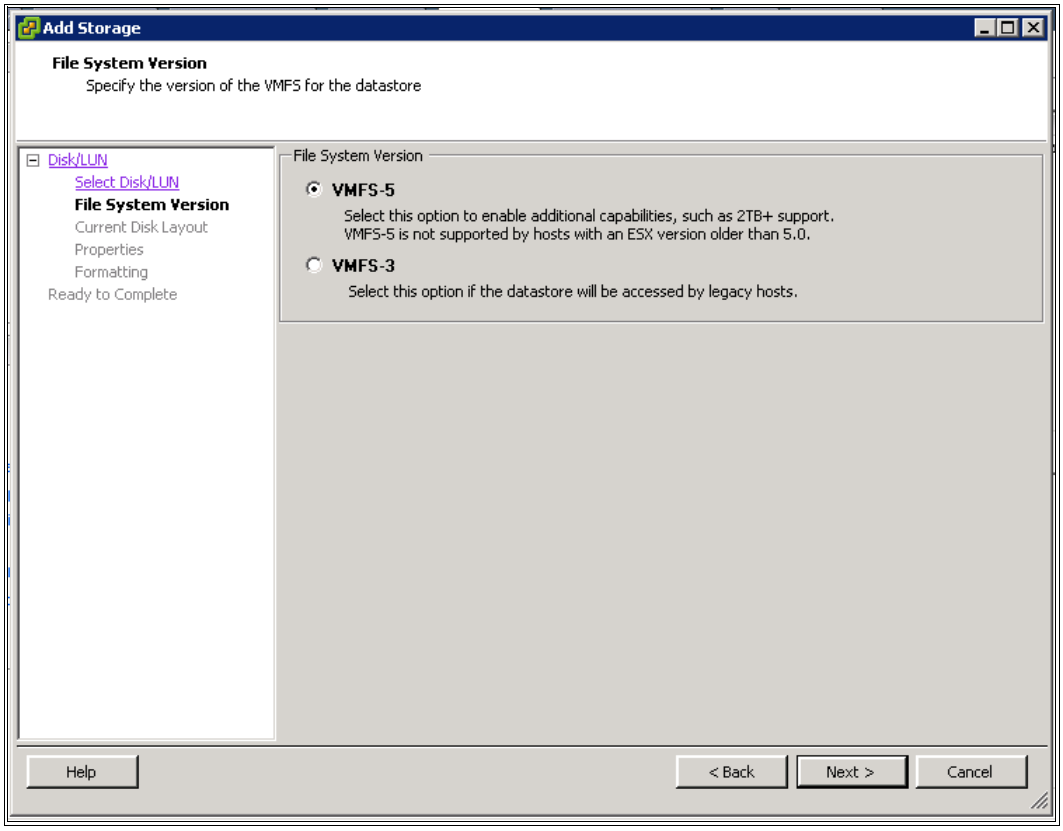

Figure 5-73 Select File System Version

8. Click **Next** to move through the wizard. A summary window of the current disk layout is shown, followed by the option to name the new Datastore. In our example, we chose `RedbookTestOne`, as shown in Figure 5-74.

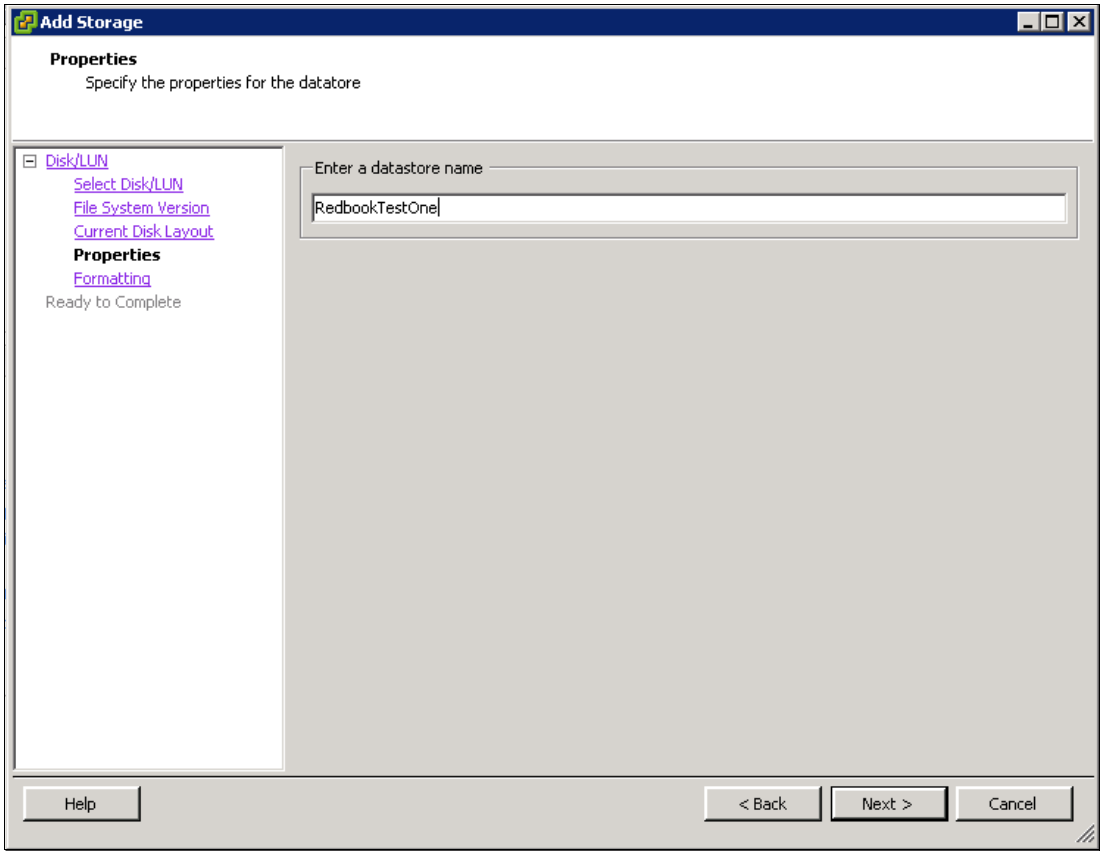

Figure 5-74 Enter a Datastore name

9. Click **Next** and the final window presents the choice of creating the Datastore with the default maximum size of the volume or a proportion of it. After you click **Finish**, the wizard closes and you return to the storage view. In Figure 5-75, you see that the new volume was added to the configuration.

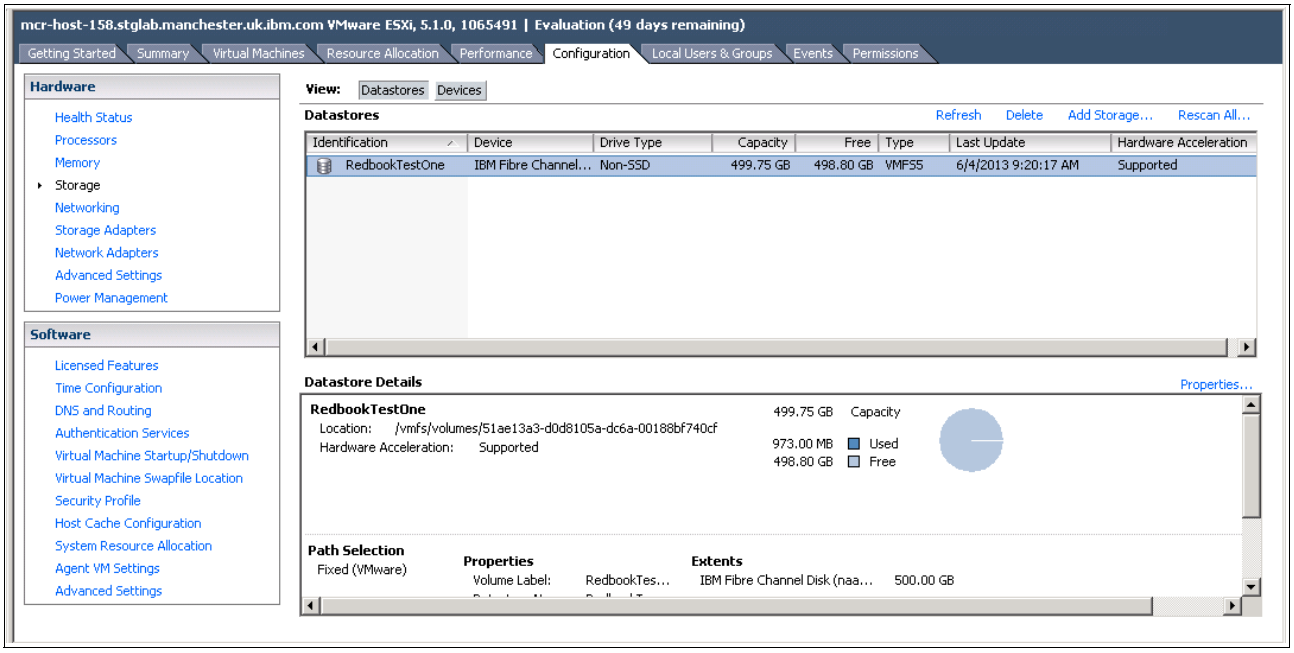

Figure 5-75 Add Storage task complete

10. Highlight the new Datastore and click **Properties** (as shown in Figure 5-76) to see the details of the Datastore, as shown in Figure 5-77 on page 214.

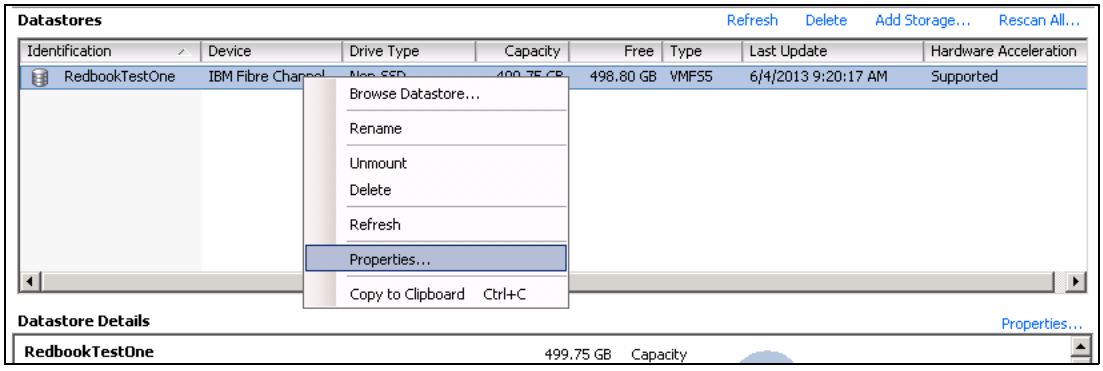

Figure 5-76 Datastore properties

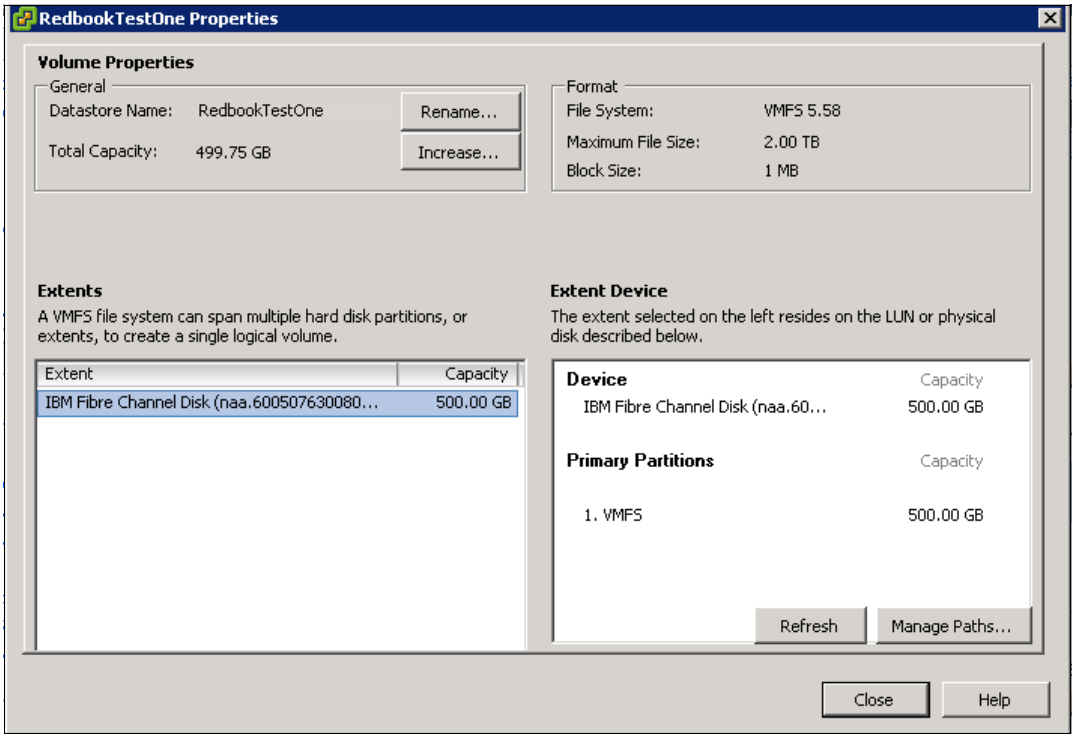

Figure 5-77 Datastore property details

11. Click **Manage Paths** to customize the multipath settings. Select **Round Robin** (as shown in Figure 5-78) and click **Change**.

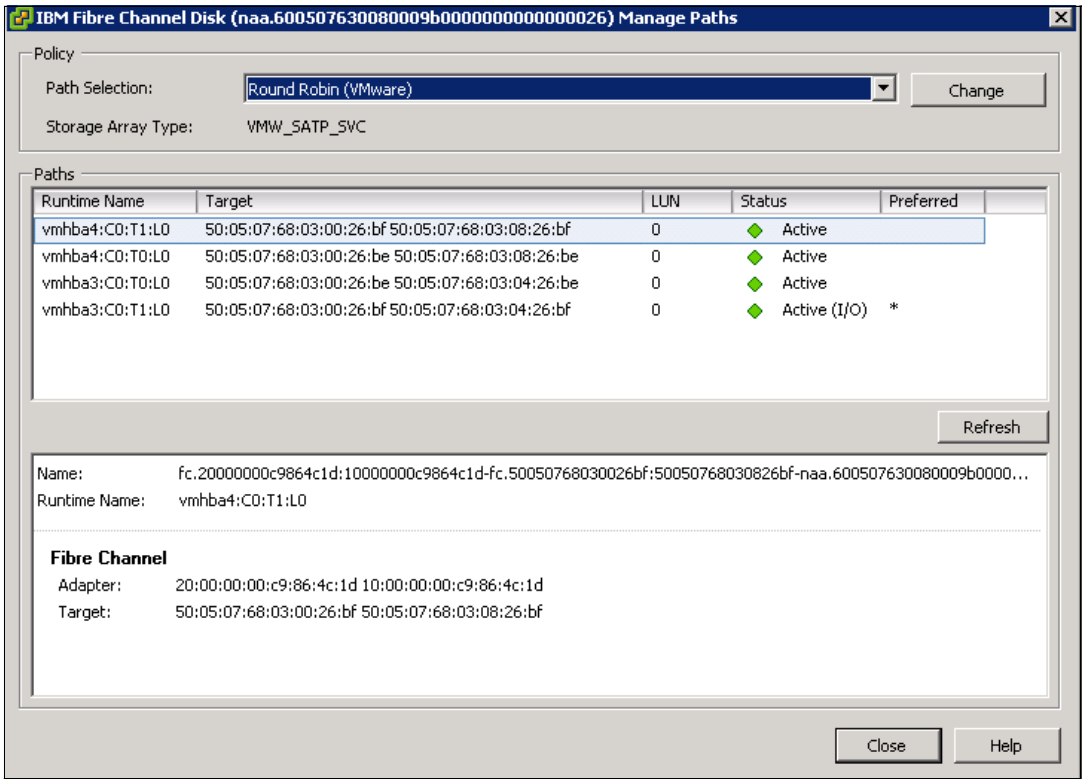

Figure 5-78 Select a Datastore multipath setting

When the change completes, click **Closed** and the storage disk is available and ready to use with your VMware ESX server that uses Fibre Channel attachment.

5.3.5 VMware ESX iSCSI volume attachment

To perform a VMware ESX iSCSI attachment, complete the following steps:

1. Right-click your VMware ESX iSCSI host in the Hosts view and select **Properties,** as shown in Figure 5-79.

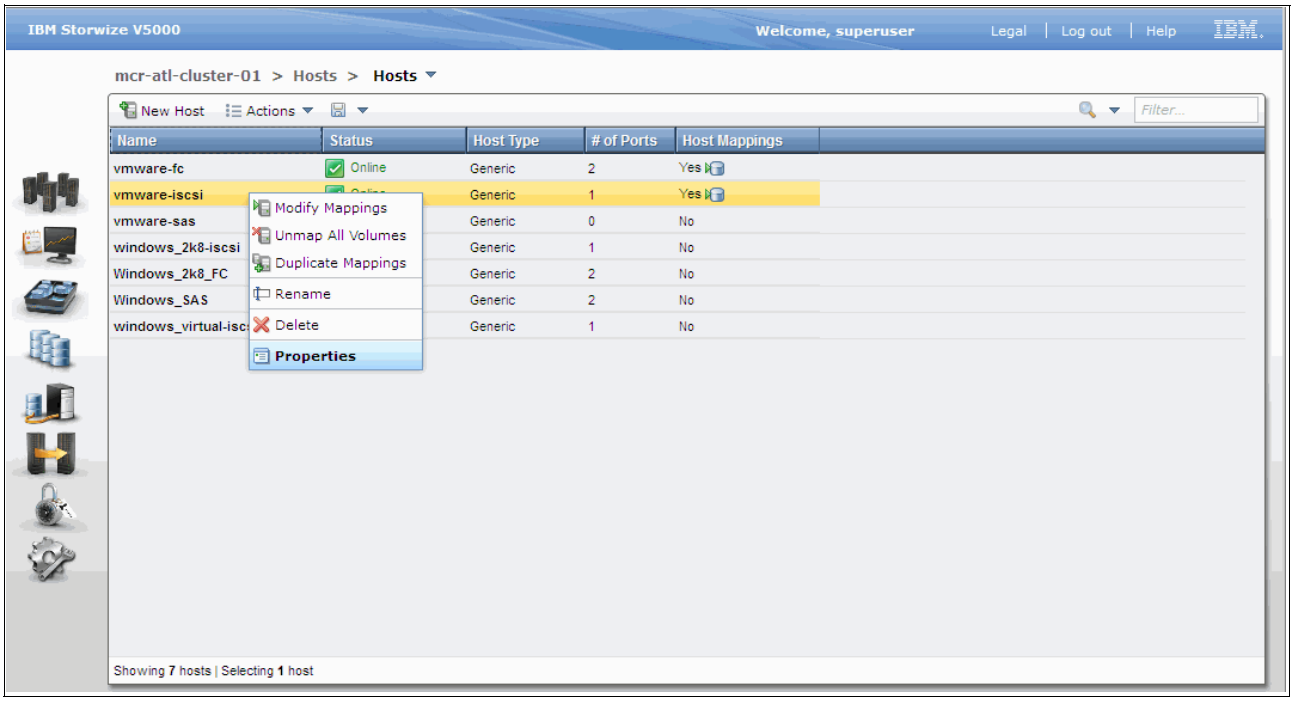

Figure 5-79 Select iSCSI ESX host properties

2. Browse to the **Mapped Volumes** tab, as shown in Figure 5-80.

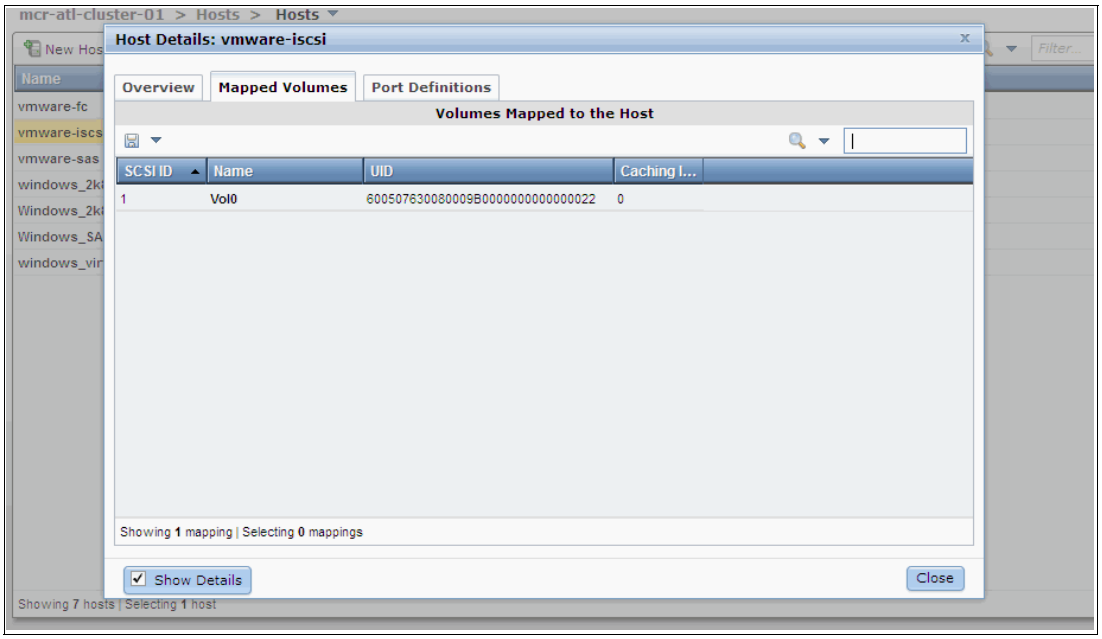

Figure 5-80 iSCSI ESX host properties

In the Host Details window, you see that there is one volume connected to the ESX iSCSI host that uses SCSI ID 1. The UID of the volume is also displayed.

3. Connect to your VMware ESX Server by using the vSphere Client. Browse to the **Configuration** tab and select **Storage Adapters**, as shown in Figure 5-81.

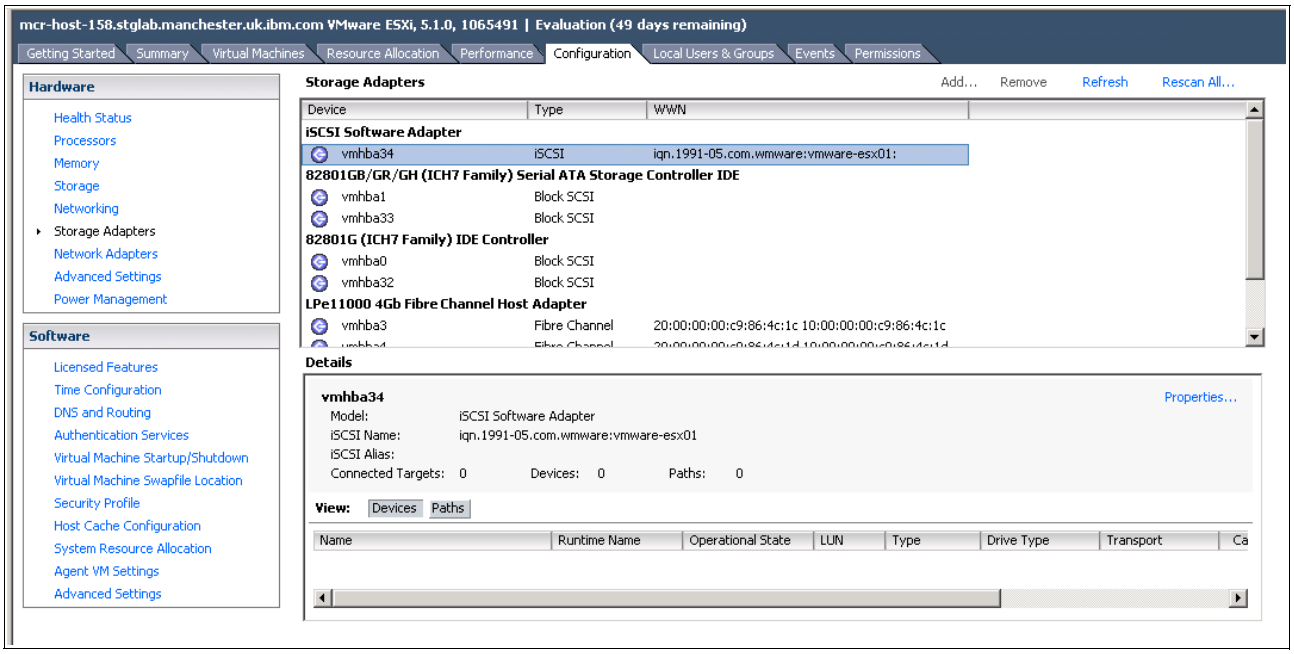

Figure 5-81 vSphere Client: Storage Adapters

4. Highlight the iSCSI Software Adapter and click **Properties**. The iSCSI initiator properties window opens. Select the **Dynamic Discovery** tab (as shown in Figure 5-82) and click **Add**.

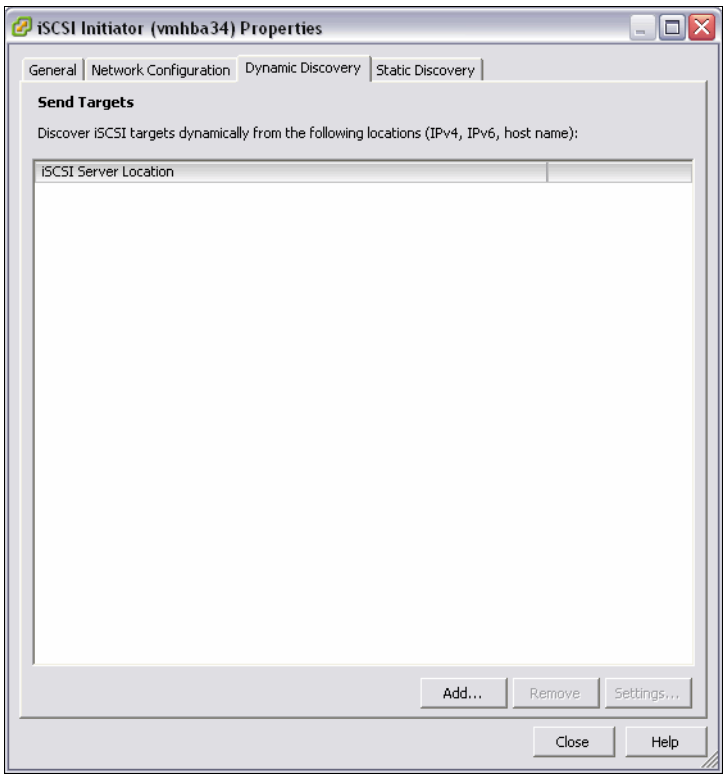

Figure 5-82 iSCSI Initiator properties

5. To add a target, enter the target IP address, as shown in Figure 5-83 on page 218. The target IP address is the iSCSI IP address of a node in the I/O Group from which you are mapping the iSCSI volume. Leave the IP port number at the default value of 3260 and click **OK**. The connection between the initiator and target is established.

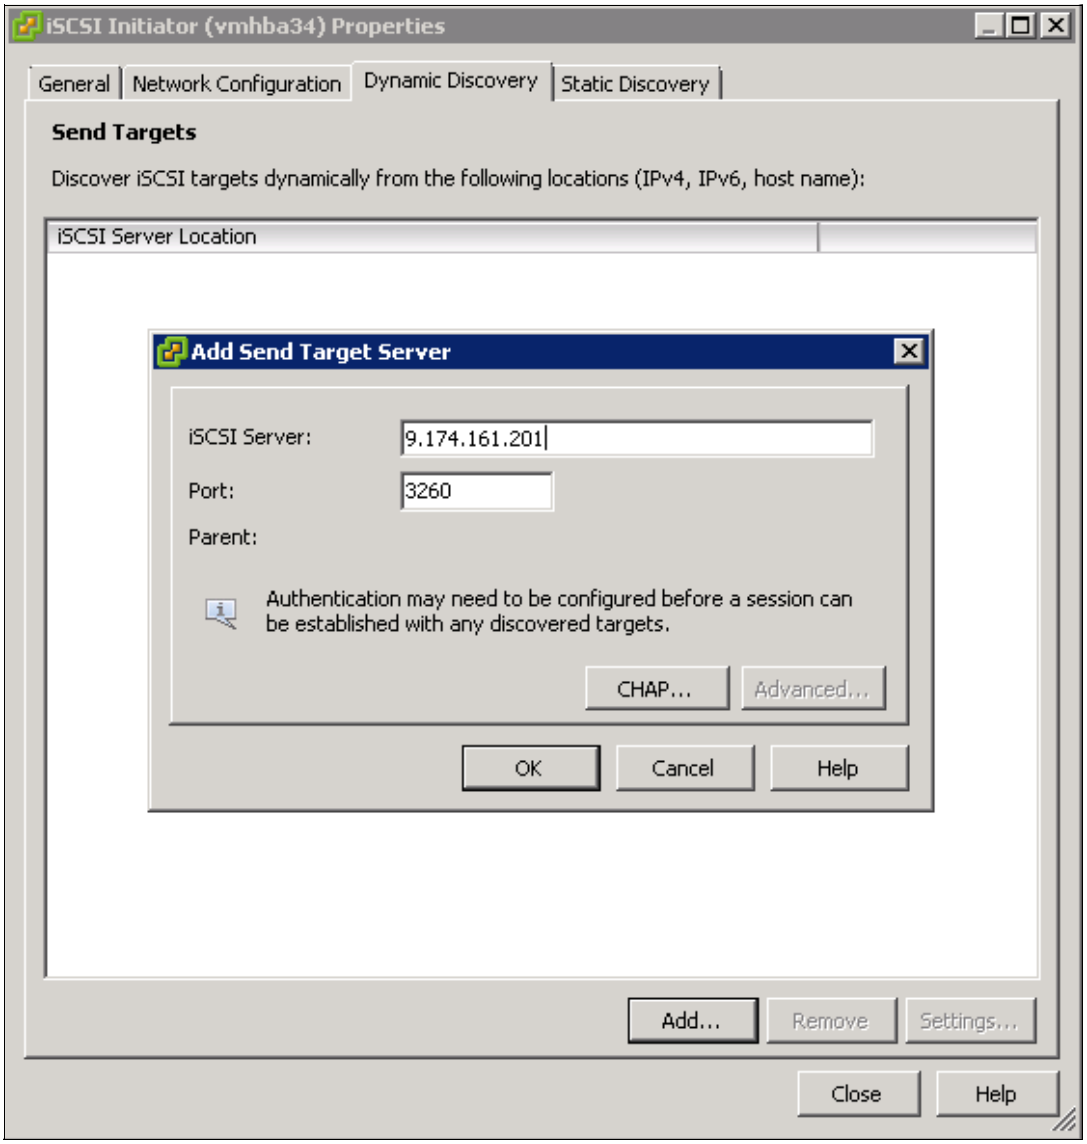

Figure 5-83 Enter a target IP address

Repeat this step for each IBM Storwize V5000 iSCSI port you want to use for iSCSI connections.

iSCSI IP addresses: The iSCSI IP addresses are different for the cluster and canister IP addresses. They are configured as described in Chapter 4, “Host configuration” on page 153.

6. After you add all the required ports, close the iSCSI Initiator properties by clicking **Close**, as shown in Figure 5-82 on page 217.

You are prompted to rescan for new storage devices. Confirm the scan by clicking **Yes**, as shown in Figure 5-84.

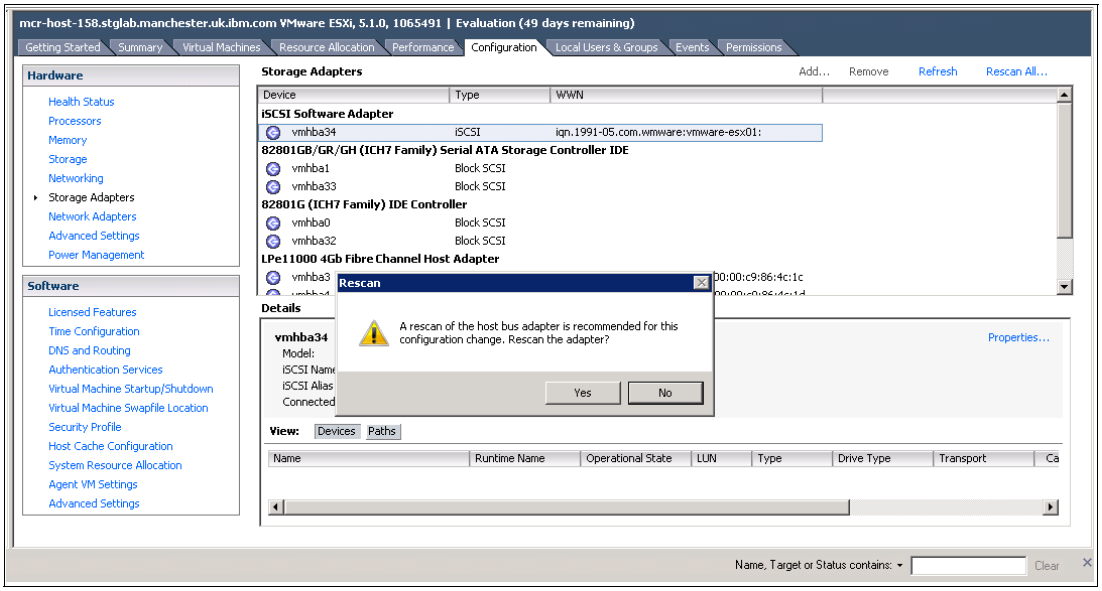

Figure 5-84 Confirm the rescan

7. Go to the storage view and click **Add Storage**. The Add Storage wizard opens, as shown in Figure 5-85. Select **Disk/LUN** and click **Next**.

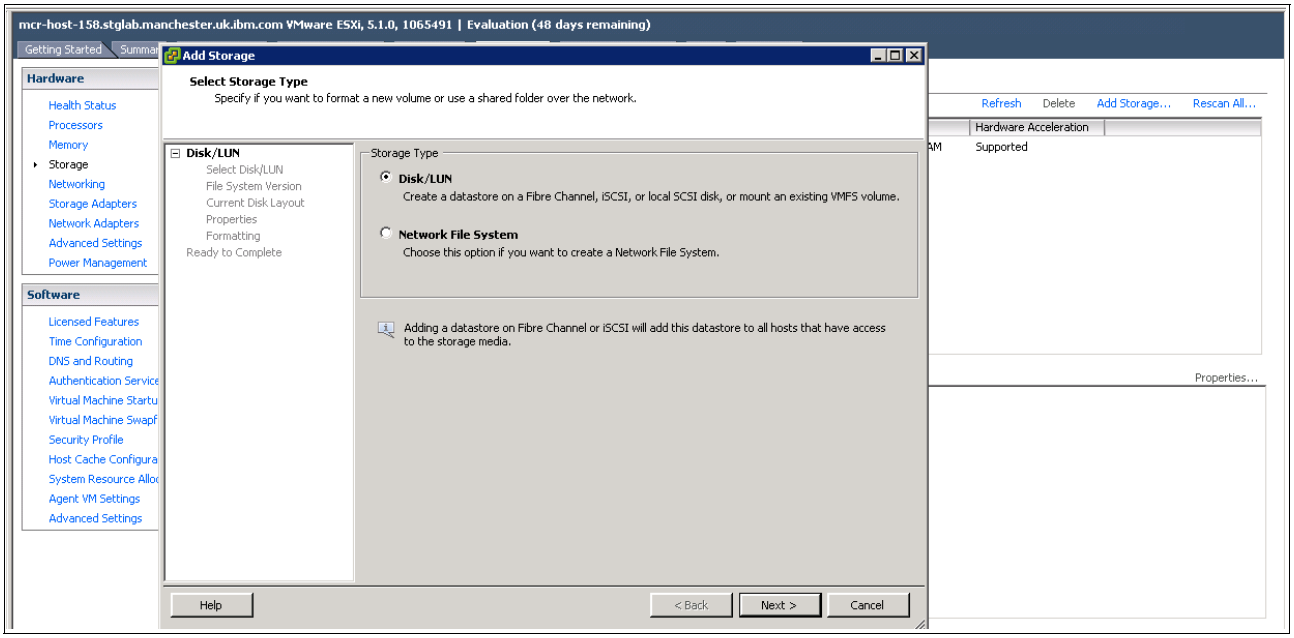

Figure 5-85 vSphere Client: Add Storage

8. The new iSCSI LUN is shown. Highlight it and click **Next**, as shown in Figure 5-86.

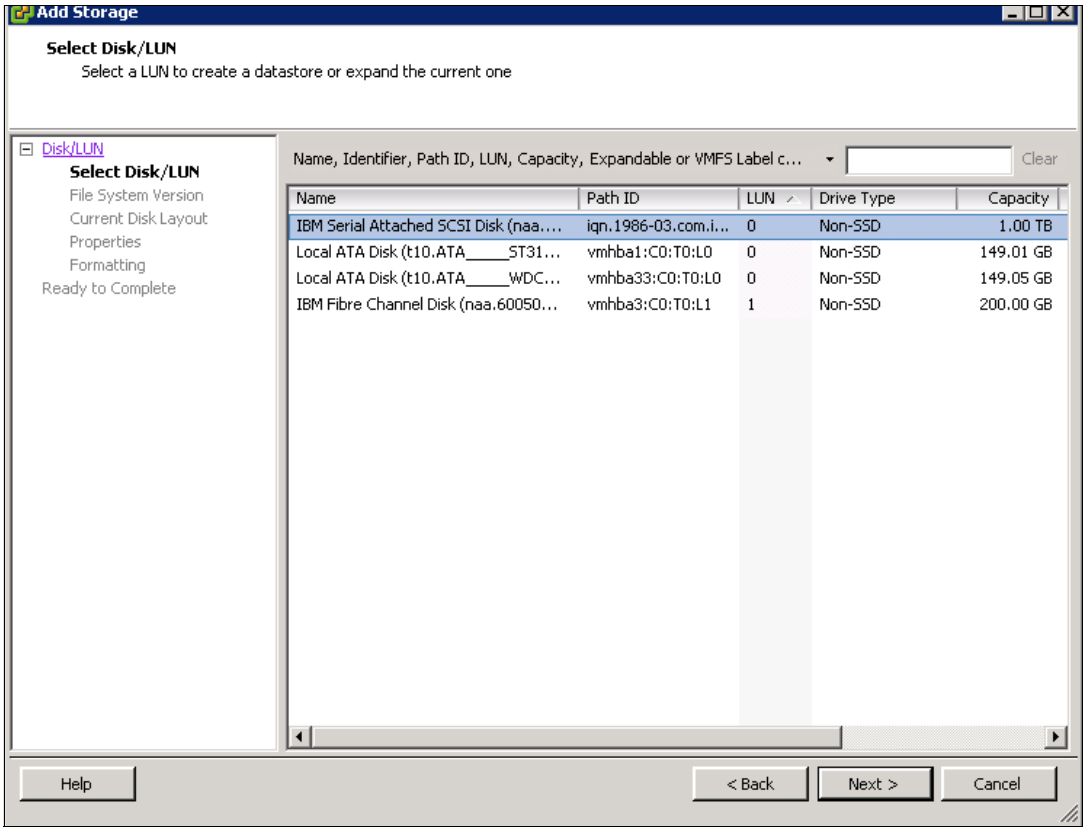

Figure 5-86 Select iSCSI LUN

9. Select a File System Version option. In our example, we selected **VMFS-5**, as shown in Figure 5-87.

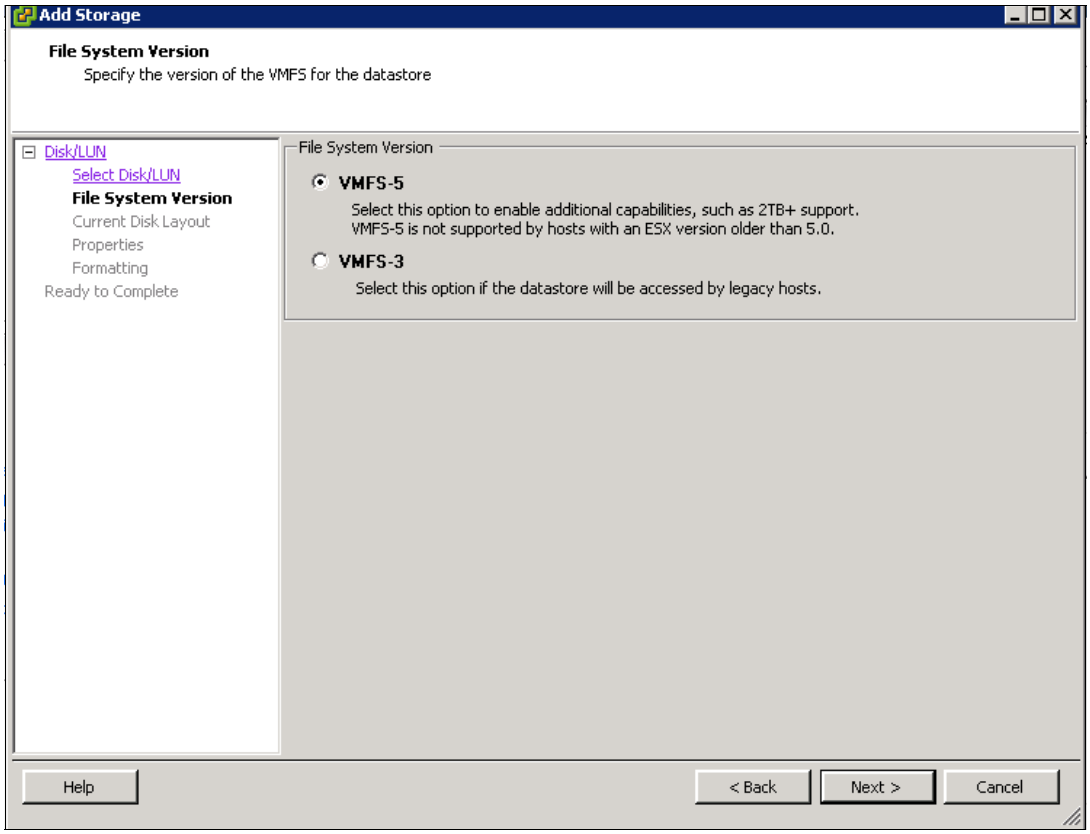

Figure 5-87 Select File System Version

10. Review the disk layout and click **Next**, as shown in Figure 5-88.

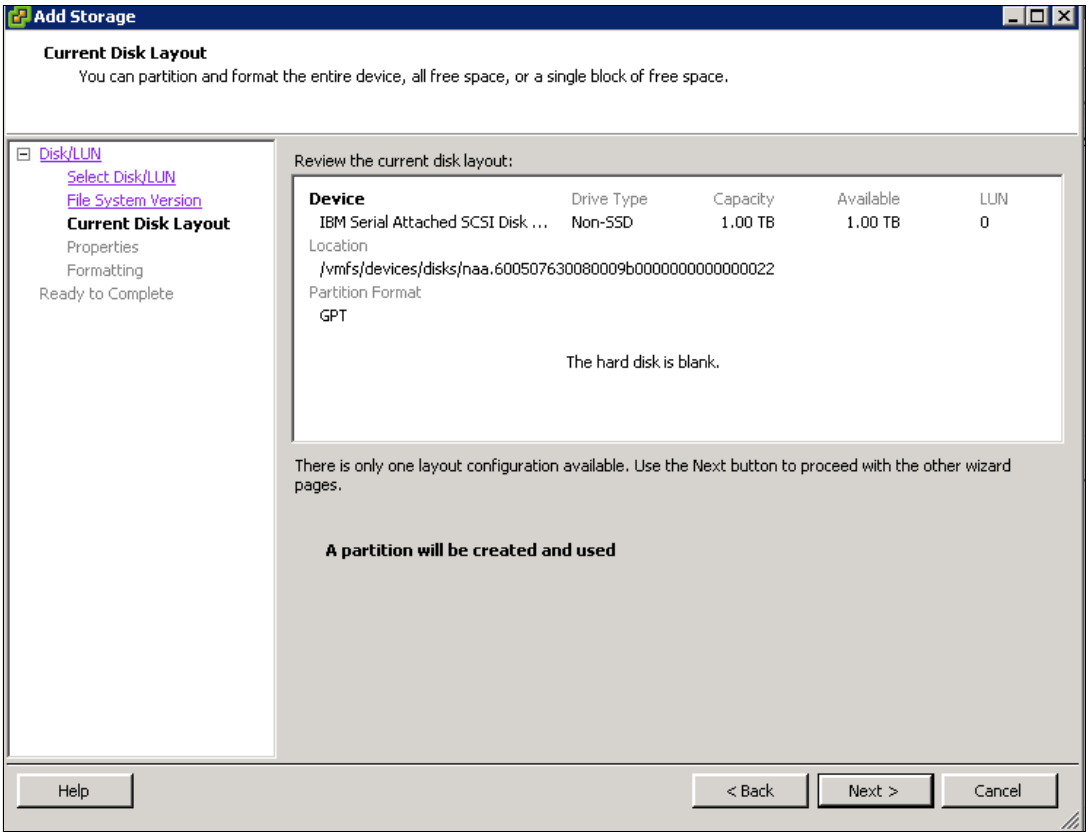

Figure 5-88 Current Disk Layout

11. Enter a name for the Datastore and click **Next**, as shown in Figure 5-89.

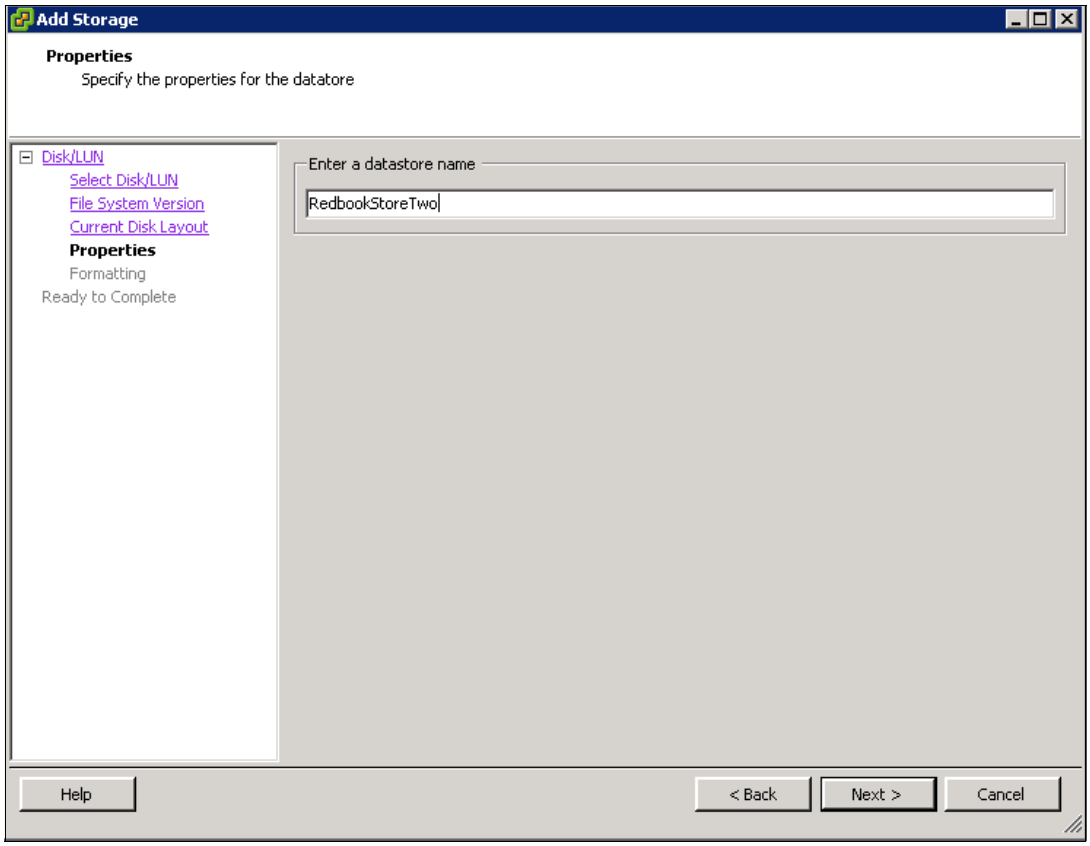

Figure 5-89 Enter a Datastore name

12. Select the Maximum available space and click **Next**, as shown in Figure 5-90.

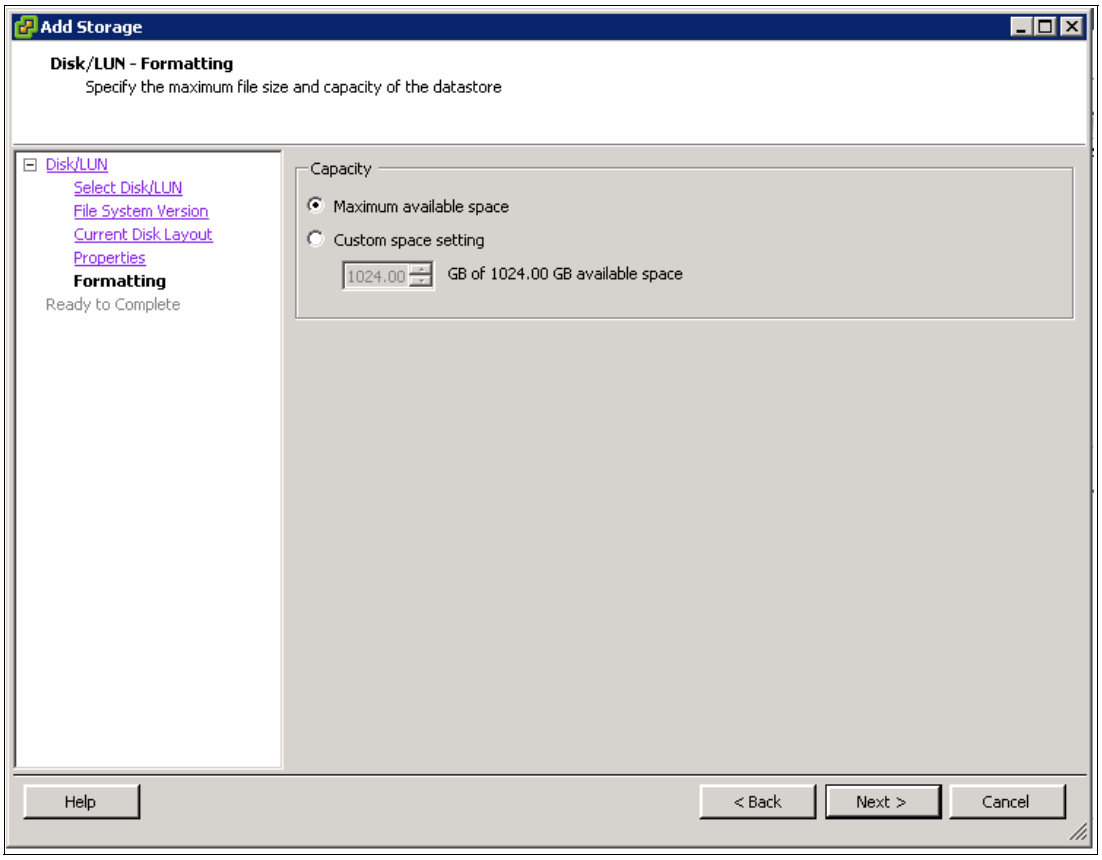

Figure 5-90 Capacity

13. Review your selections and click **Finish**, as shown in Figure 5-91.

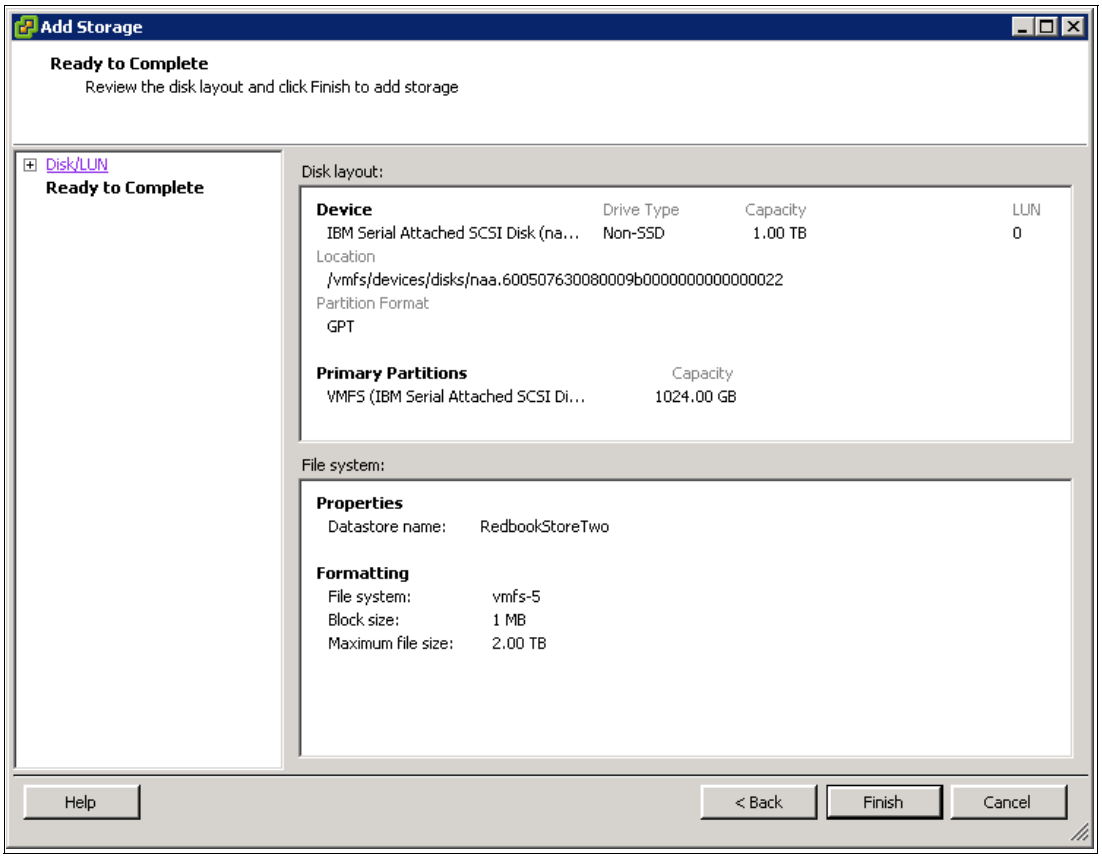

Figure 5-91 Finish the wizard

The process starts to add an iSCSI LUN, which can take a few minutes. After the task is complete, the new Datastore appears in the storage view, as shown in Figure 5-92.

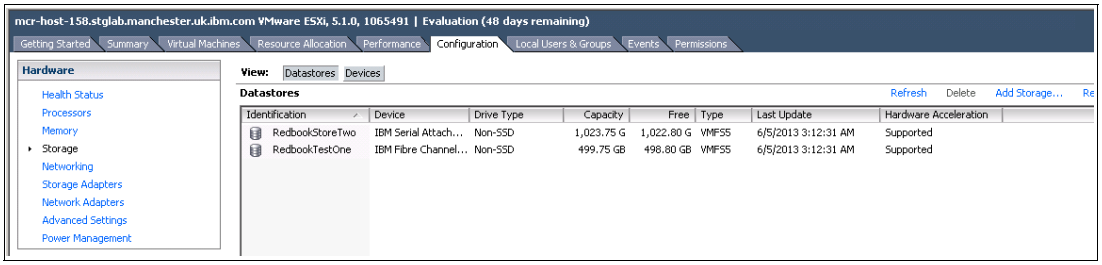

Figure 5-92 New Datastore available

14. Highlight the new Datastore and click **Properties** to open and review the Datastore settings, as shown in Figure 5-93.

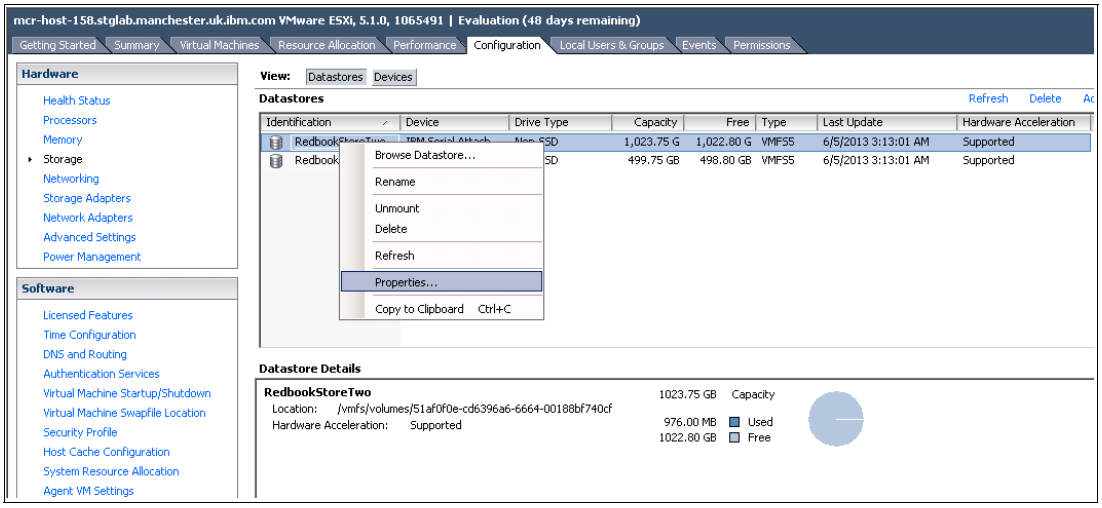

Figure 5-93 Datastore properties

15. Click **Manage Paths**, select **Round Robin** as the multipath policy (as shown in Figure 5-94), and click **Change**.

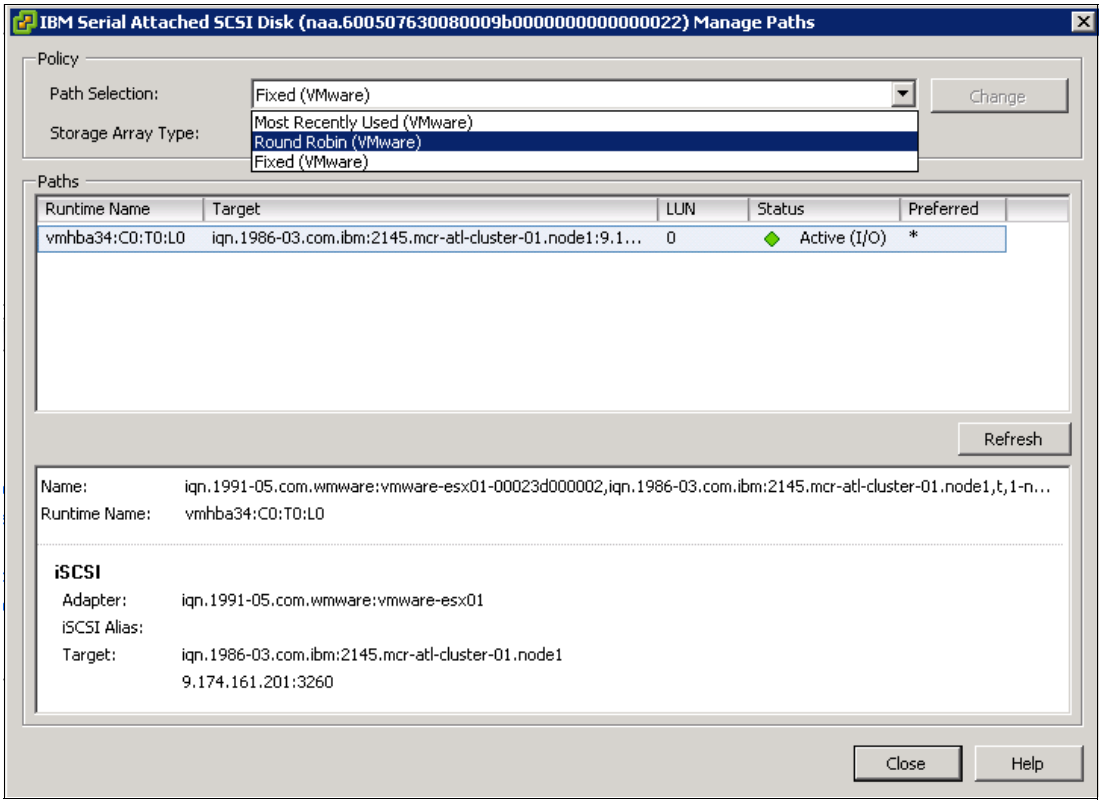

Figure 5-94 Change the multipath policy

16. Click **Close** twice to return to the storage view. The storage disk is available and ready to use for your VMware ESX server that uses an iSCSI attachment.

5.3.6 VMware ESX Direct SAS volume attachment

To perform VMware ESX Direct SAS attachment, complete the following steps:

1. Right-click your VMware ESX SAS host in the Hosts view and select **Properties,** as shown in Figure 5-95.

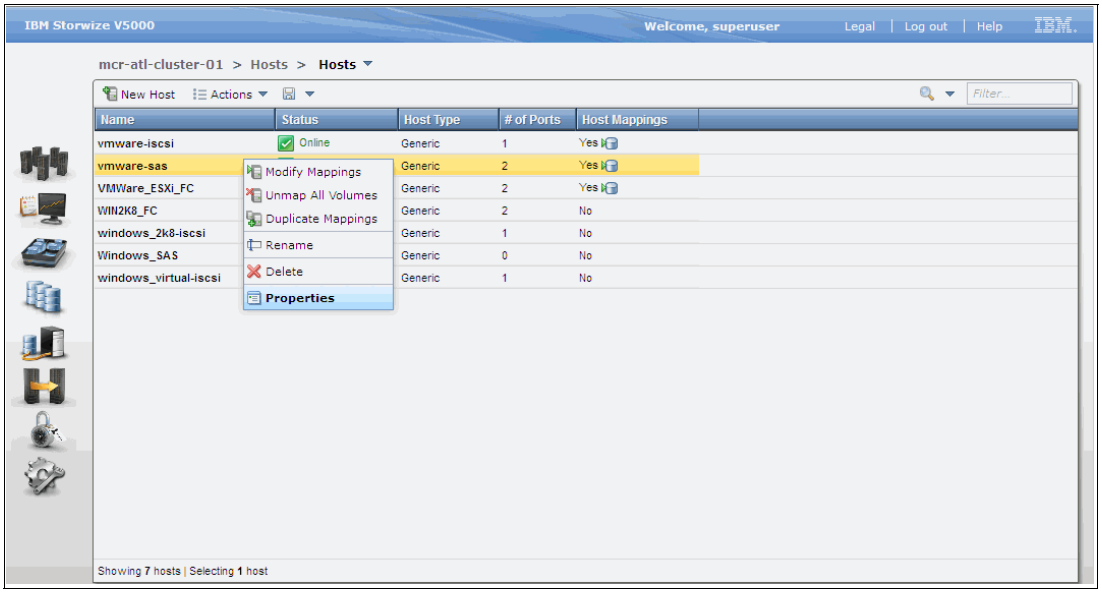

Figure 5-95 Example ESX SAS host

2. Browse to the **Mapped Volumes** tab, as shown in Figure 5-96.

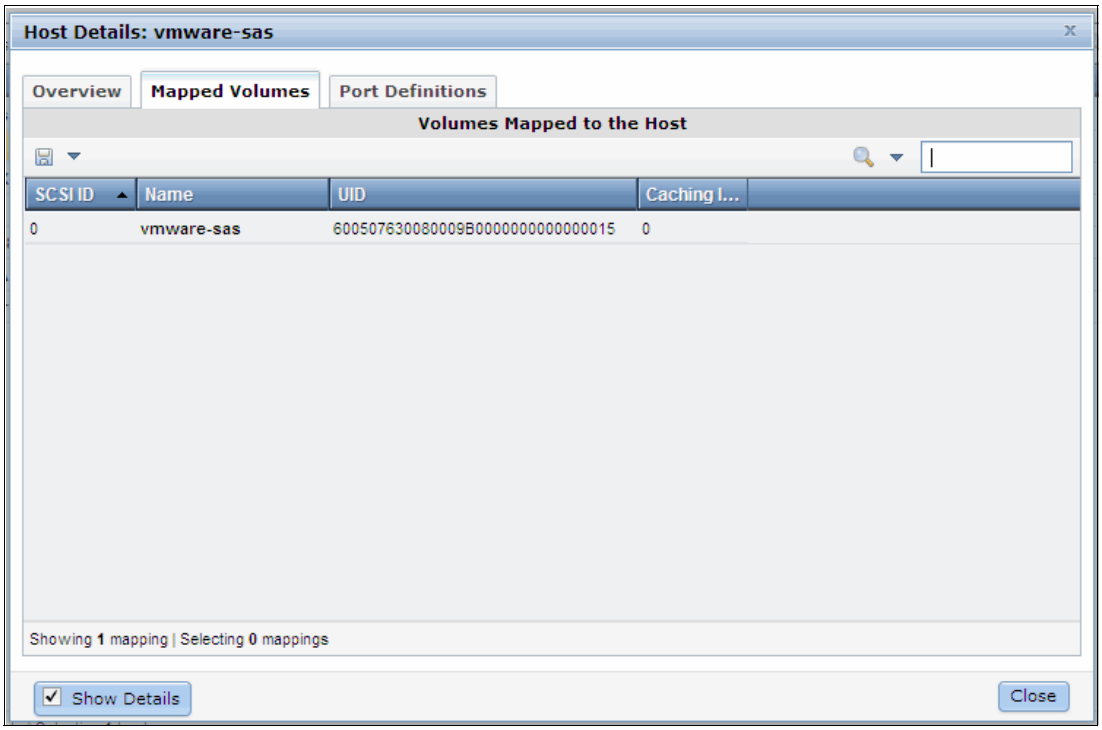

Figure 5-96 Mapped volumes to ESX SAS host

In the Host Details window, there is one volume connected to the ESX SAS host that uses SCSI ID 0. The UID of the volume is also displayed.

3. Connect to your VMware ESX Server by using the vSphere client. Browse to the Configuration tab and select **Storage Adapters**, as shown in Figure 5-97.

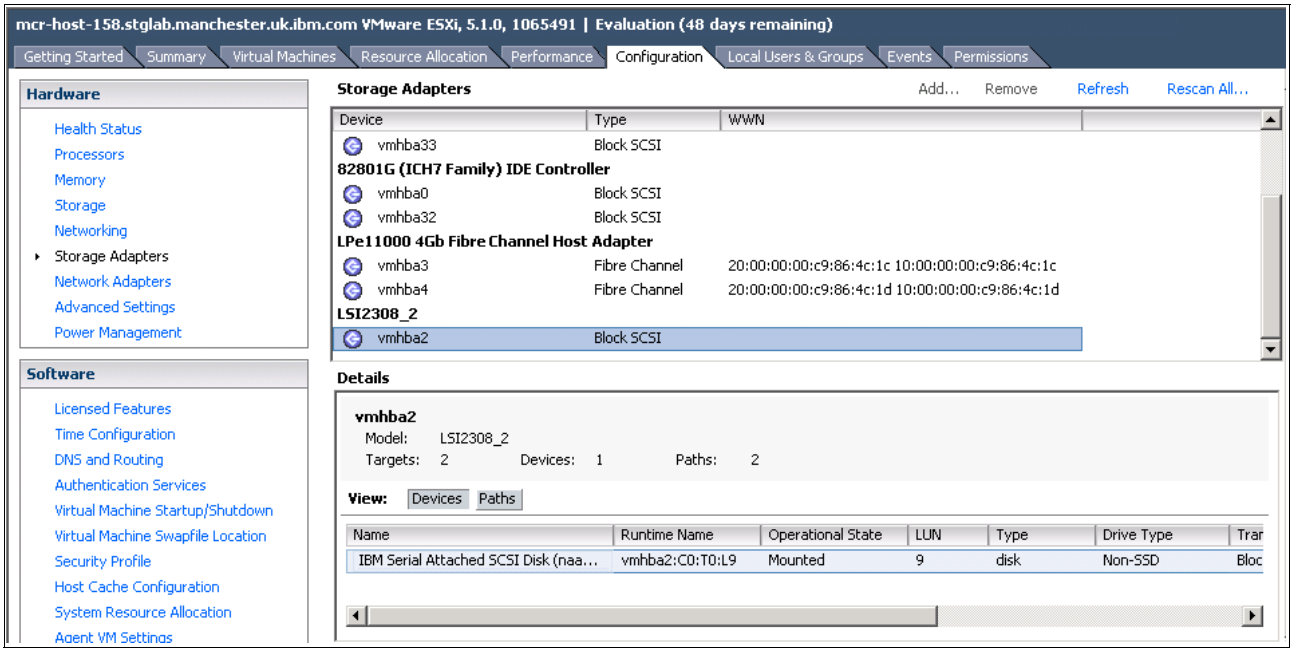

Figure 5-97 vSphere Client: Storage Adapters

4. Click **Rescan All...** in the upper right corner and click **OK** in the resulting pop-up window, as shown in Figure 5-98. This scans for new storage devices.

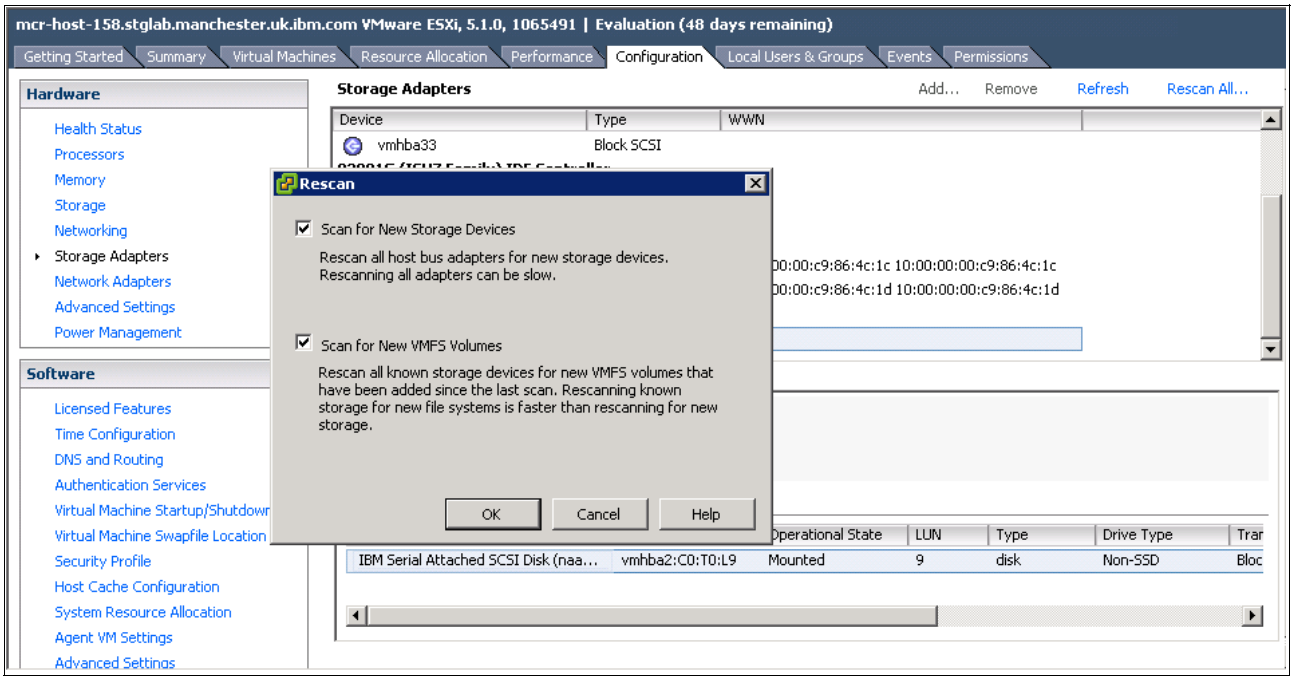

Figure 5-98 Rescan

The mapped volumes on the IBM Storwize V5000 should now appear against the SAS adapters.

5. Select **Storage** and click **Add Storage**, as shown in Figure 5-99.

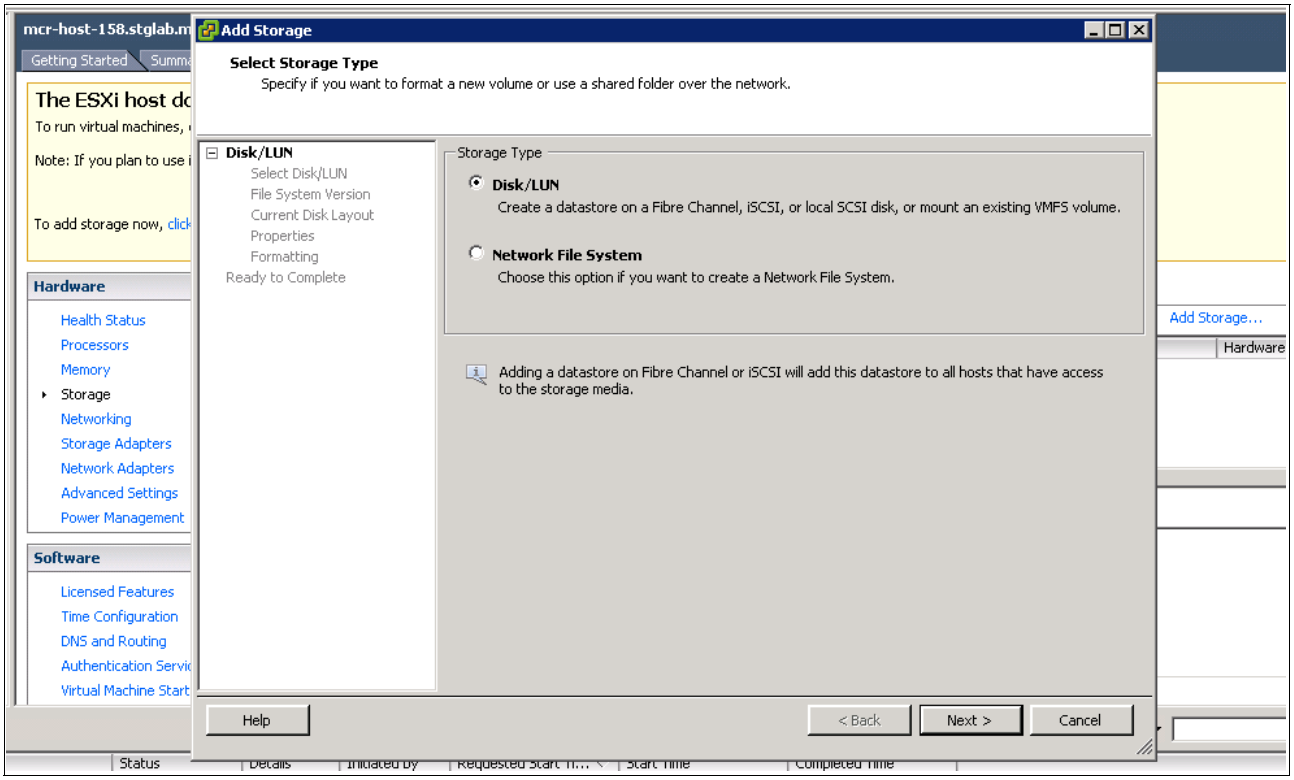

Figure 5-99 vSphere Client: Add Storage

6. The Add Storage wizard opens. Click **Select Disk/LUN** and click **Next**. The IBM Storwize V5000 disk appears, as shown in Figure 5-100. In our example, it is the SAS Disk. Highlight the disk and click **Next**.

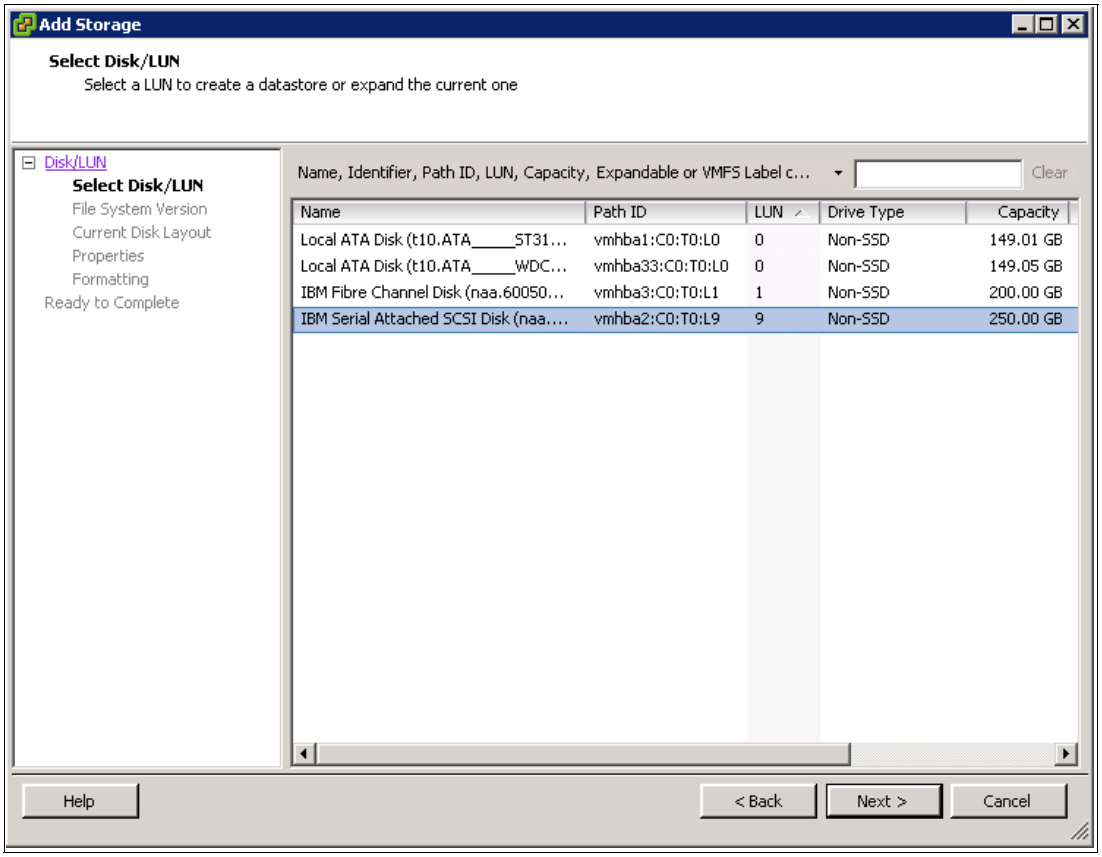

Figure 5-100 Select SAS Disk

7. Select a File System Version option. In our example, we selected **VMFS-5**, as shown in Figure 5-101.

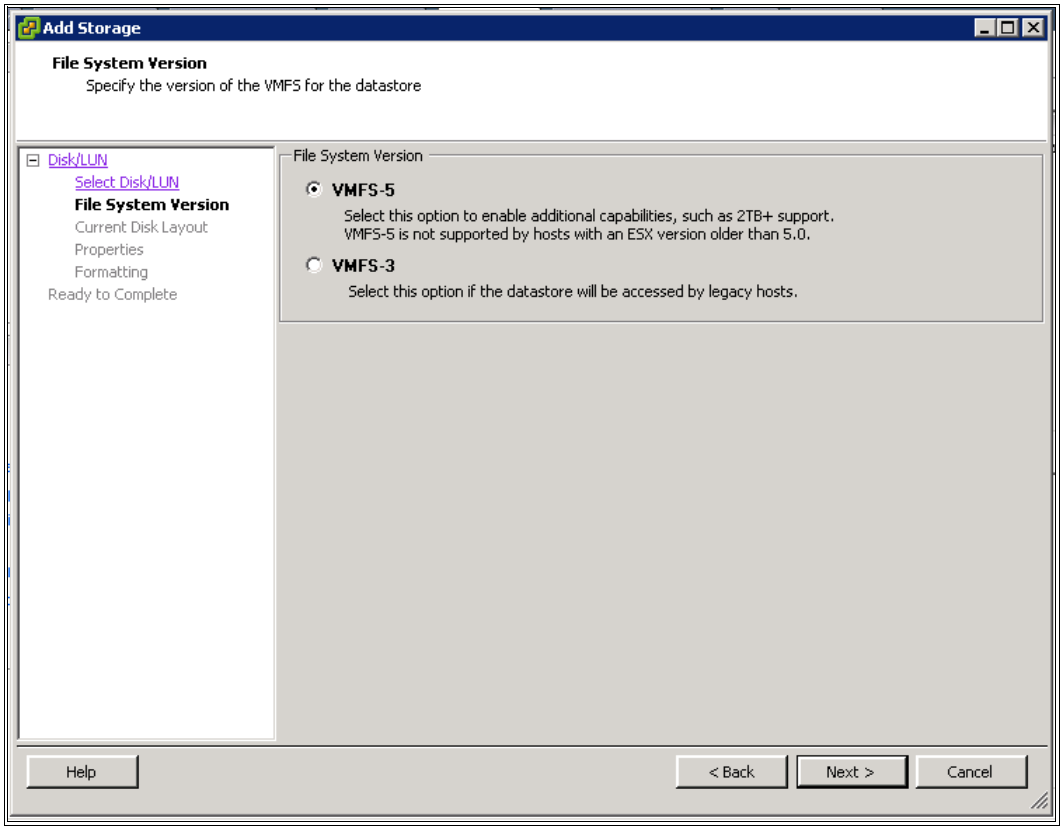

Figure 5-101 Select File System Version

8. Click **Next** to move through the wizard. A summary window of the current disk layout is shown, followed by the option to name the new Datastore. In our example, we chose `RedbookTestThree`, as shown in Figure 5-102.

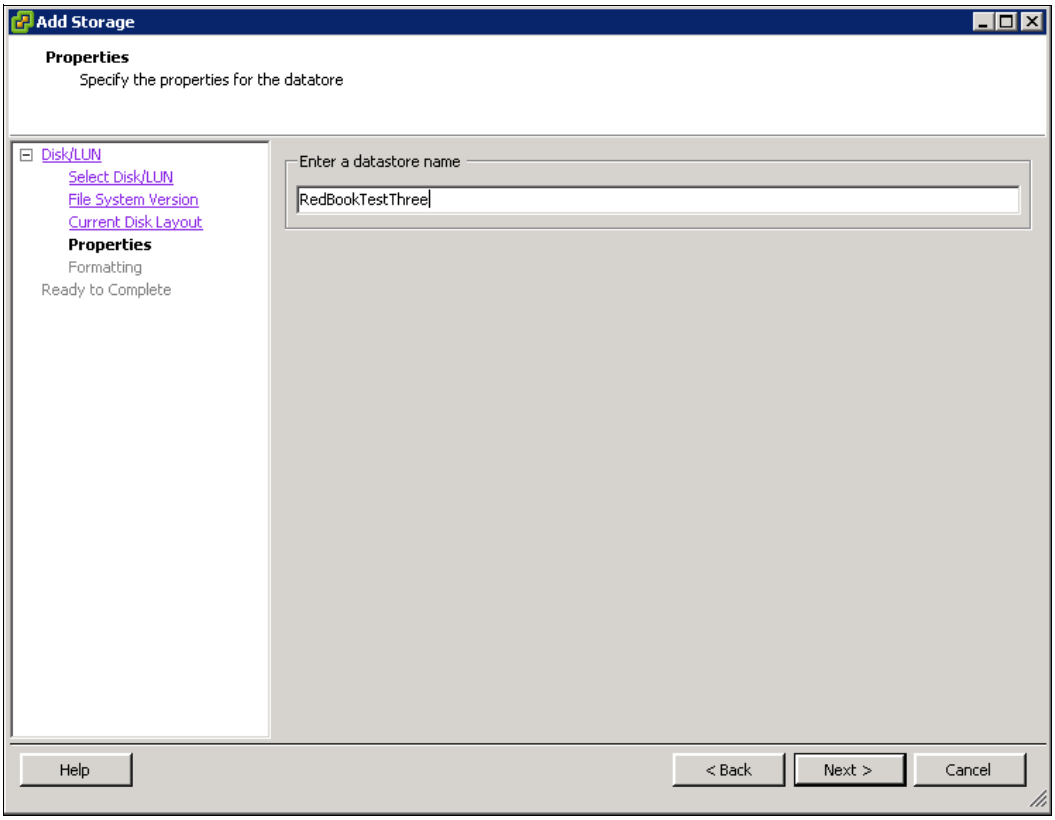

Figure 5-102 Adding Datastore name

9. Click **Next** and the final window presents the choice of creating the Datastore with the default maximum size of the volume or a proportion of it. After you click **Finish**, the wizard closes and you return to the storage view. In Figure 5-103, you see that the new volume was added to the configuration.

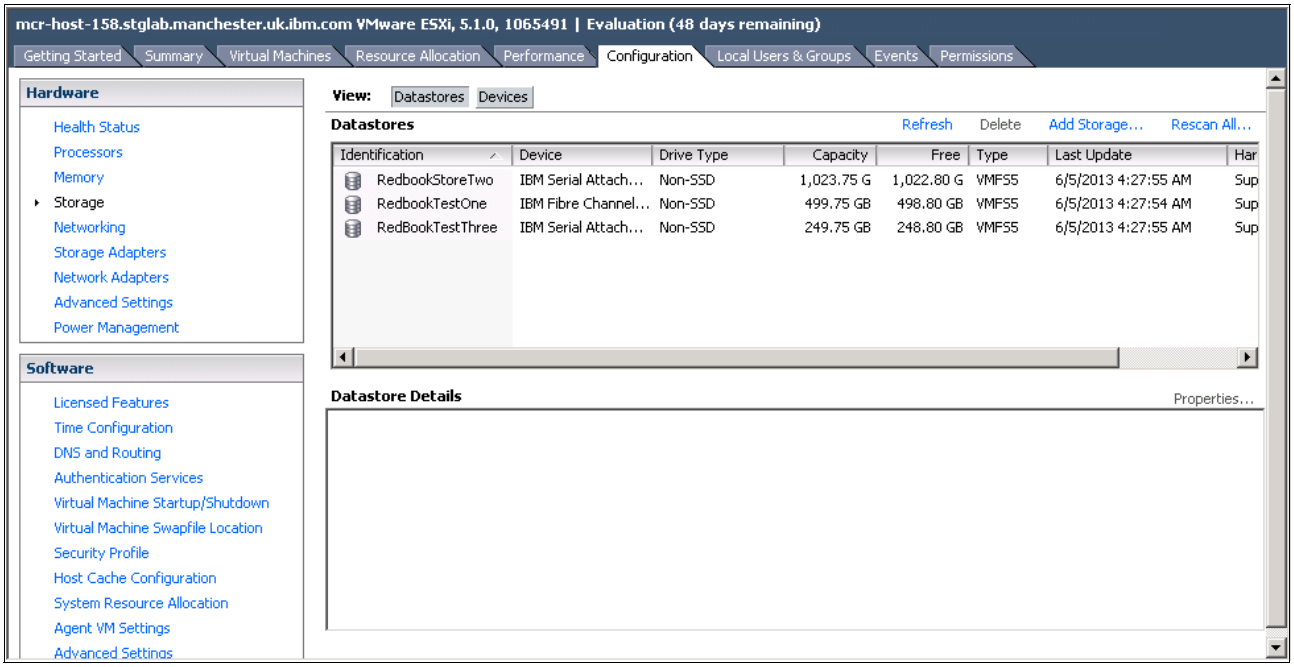

Figure 5-103 Add Storage task complete

10. Highlight the new Datastore and click **Properties** (as shown in Figure 5-104).

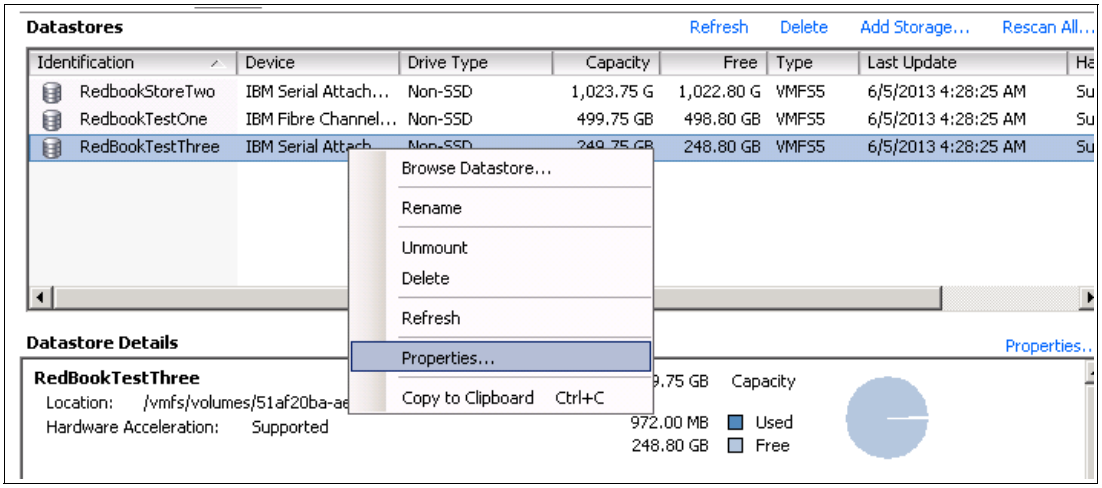

Figure 5-104 Datastore properties

The Datastore Properties window opens, as shown in Figure 5-105.

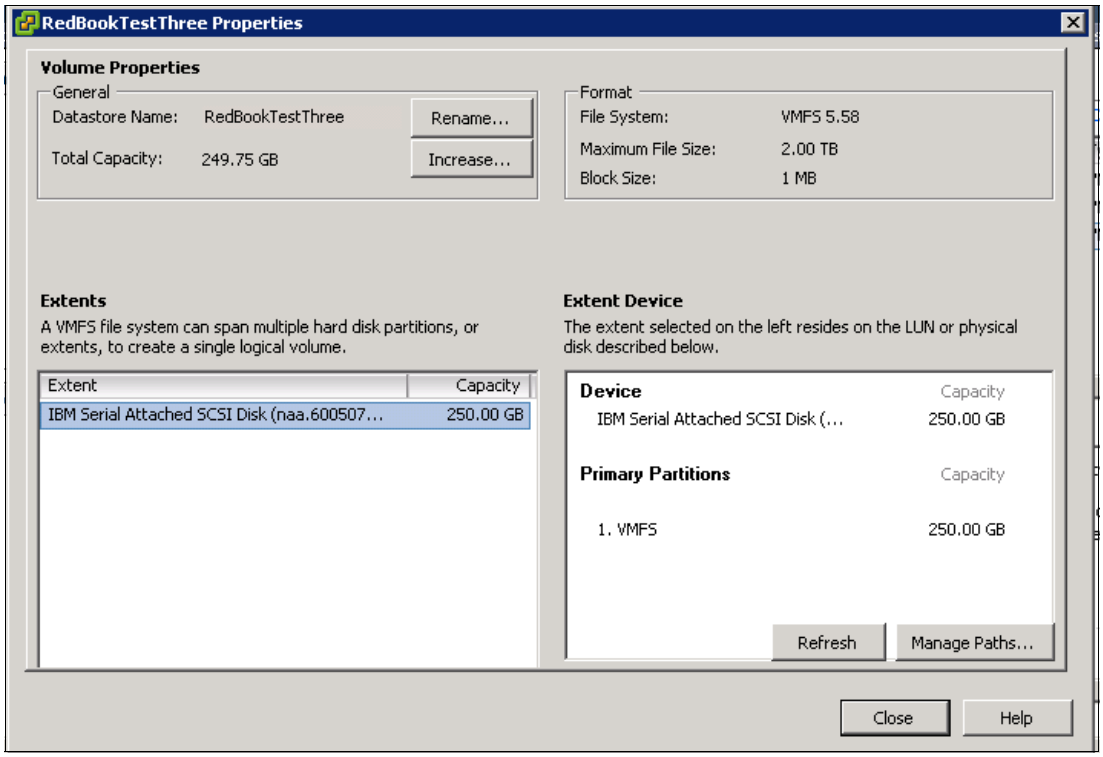

Figure 5-105 Datastore property details

11. Click **Manage Paths** to customize the multipath settings. Select **Round Robin** (as shown in Figure 5-106) and click **Change**.

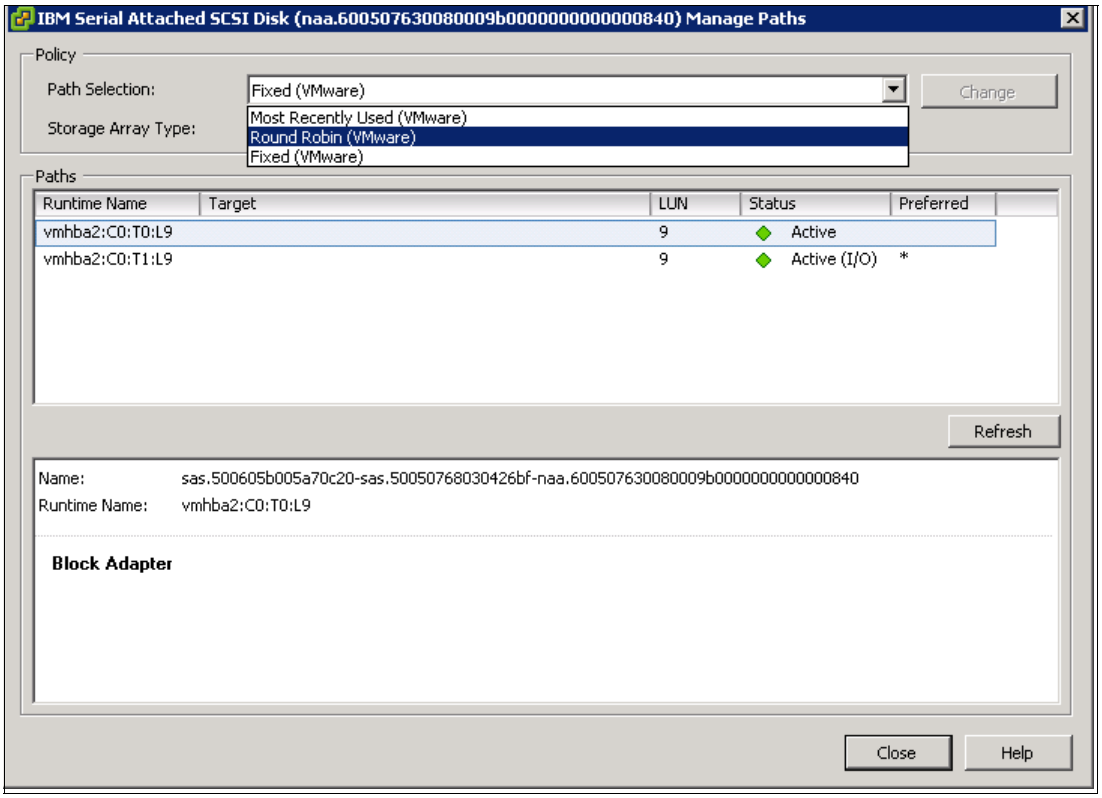

Figure 5-106 Select a Datastore multipath setting

When the change completes, click **Close** and the storage disk is available and ready to use with your VMware ESX server that uses Fibre Channel attachment.

6

Storage migration wizard

This chapter describes the steps of the storage migration wizard. The storage migration wizard is used to migrate data from older external storage systems to the internal capacity of the Storwize V5000. Migrating data from older storage systems to the Storwize V5000 storage system provides benefit from more functionality, such as, the easy-to-use GUI, internal virtualization, thin provisioning, and FlashCopy.

This chapter includes the following topics:

- Interoperability and compatibility
- Storage migration wizard
- Storage migration wizard example scenario

6.1 Interoperability and compatibility

Interoperability is an important consideration when a new storage system is set up in an environment that contains existing storage infrastructure. In this section, we describe how to check that the storage environment, the older storage system, and IBM Storwize V5000 are ready for the data migration process.

To ensure system interoperability and compatibility between all elements that are connected to the SAN fabric, check the proposed configuration with the IBM System Storage Interoperation Center (SSIC). SSIC can confirm whether the solution is supported and provide recommendations for hardware and software levels.

If the required configuration is not listed for support in the SSIC, contact your IBM marketing representative and a Request for Price Quotation (RPQ) for your specific configuration.

For more information about the IBM SSIC, see this website:

http://www.ibm.com/systems/support/storage/ssic/interoperability.wss

6.2 Storage migration wizard

The Storwize V5000 storage migration wizard simplifies the migration. The wizard features easy-to-follow panels that guide users through the entire process. This process involves external virtualization of the older storage system (in our example, an IBM DS3400) and performing an online migration. After data migration is complete, the older storage system is removed from Storwize V5000 control and can be retired.

6.2.1 External virtualization capability

To migrate data from an older storage system to the Storwize V5000, it is necessary to use the built-in external virtualization capability. This capability places external Fibre Channel connected Logical Units (LUs) under the control of the Storwize V5000. Control of the external LUs is established by using and following the storage migration wizard.

6.2.2 Overview of the storage migration wizard

An overview of the storage migration wizard process includes the following considerations:

- The older storage systems divide storage into many Small Computer System Interface (SCSI) LUs that are presented on a Fibre Channel SAN to hosts.
- I/O to the LUs is stopped and changes are made to the mapping of the storage system LUs and to the SAN fabric zoning so that the original LUs are presented directly to the Storwize V5000. The Storwize V5000 discovers the external LUs as unmanaged MDisks.
- The unmanaged MDisks are then imported to the Storwize V5000 as image mode MDisks and placed into a storage pool named MigrationPool_8192. This storage pool is now a logical container for the SAN-attached LUs.

- Image mode volumes are created from MigrationPool_8192. Each volume has a one-to-one mapping with an image mode MDisk. From a data perspective, the image mode volume represents the SAN-attached LUs exactly as it was before the import operation. The image mode volume is on the same physical drives of the older storage system and the data remains unchanged. The Storwize V5000 is presenting active images of the SAN-attached LUs.
- The hosts have the older storage system multipath device driver removed and are then configured for Storwize V5000 attachment. Further zoning changes are made for host-to-V5000 SAN connections. The Storwize V5000 hosts are defined with worldwide port names (WWPNs) and the volumes are mapped. After the volumes are mapped, the hosts discover the Storwize V5000 volumes through a host rescan device or reboot operation.
- Storwize V5000 volume mirror operations are then initiated. The image mode volumes are mirrored to generic volumes. The generic volumes are from user-nominated internal storage pools. The mirrors are online migration tasks, which means a defined host can access and use the volumes during the mirror synchronization process.
- After the mirror operations are complete, the migrations are finalized by the user. The finalization process is seamless and it removes the volume mirror relationships and the image mode volumes. The older storage system LUs are now migrated and the Storwize V5000 control of the old LUs can be removed.

6.2.3 Storage migration wizard tasks

The storage migration wizard is designed for the easy and nondisruptive migration of data from an older storage system to the internal capacity of the Storwize V5000.

This section describes the following storage migration wizard tasks:

- Avoiding data loss
- Accessing the storage migration wizard
- Step 1: Before you begin
- Step 2: Prepare environment for migration
- Step 3: Map storage
- Step 4: Migrating MDisks
- Step 5: Configure hosts
- Step 6: Map volumes to hosts
- Step 7: Select storage pool
- Step 8: Finish the storage migration wizard
- Finalize migrated volumes

Avoiding data loss

The risk of losing data when the storage migration wizard is used correctly is low. However, it is prudent to avoid potential data loss by creating a backup of all the data that is stored on the hosts, the older storage systems, and the Storwize V5000 before the wizard is used.

Accessing the storage migration wizard

Select **System Migration** in the Pools menu to open the System Migration panel. The System Migration panel provides access to the storage migration wizard and displays the migration progress information, as shown in Figure 6-1.

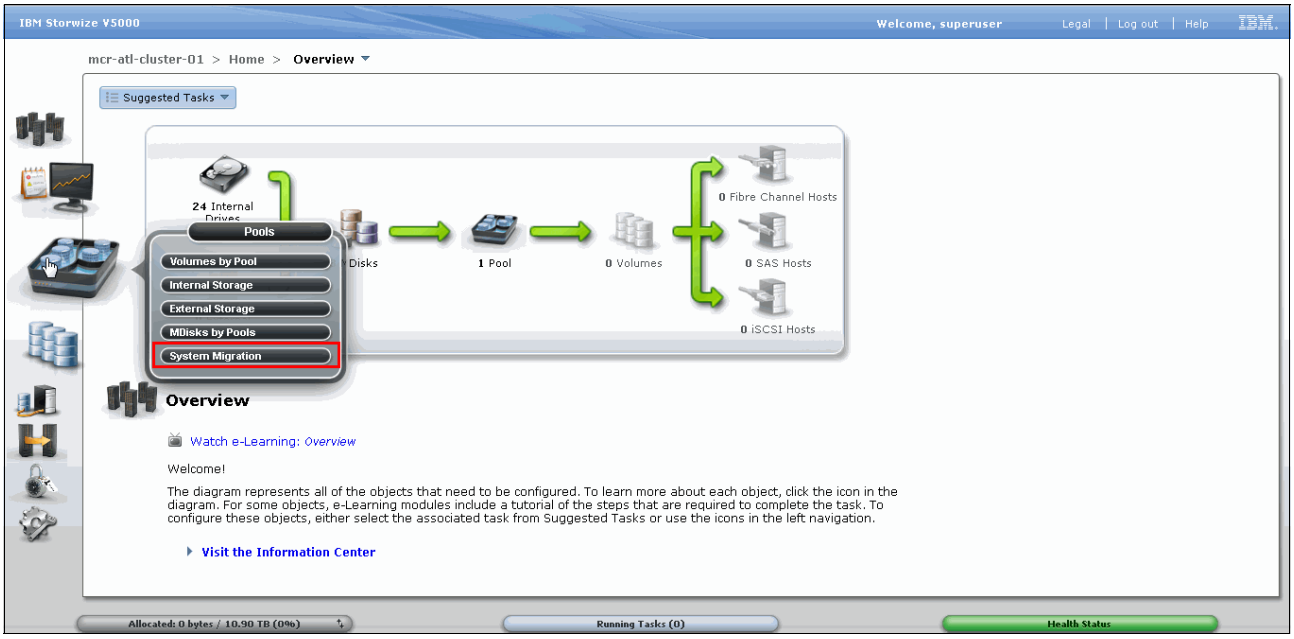

Figure 6-1 Pools menu

Click **Start New Migration** and the storage migration wizard is started. Figure 6-2 shows the System Migration panel.

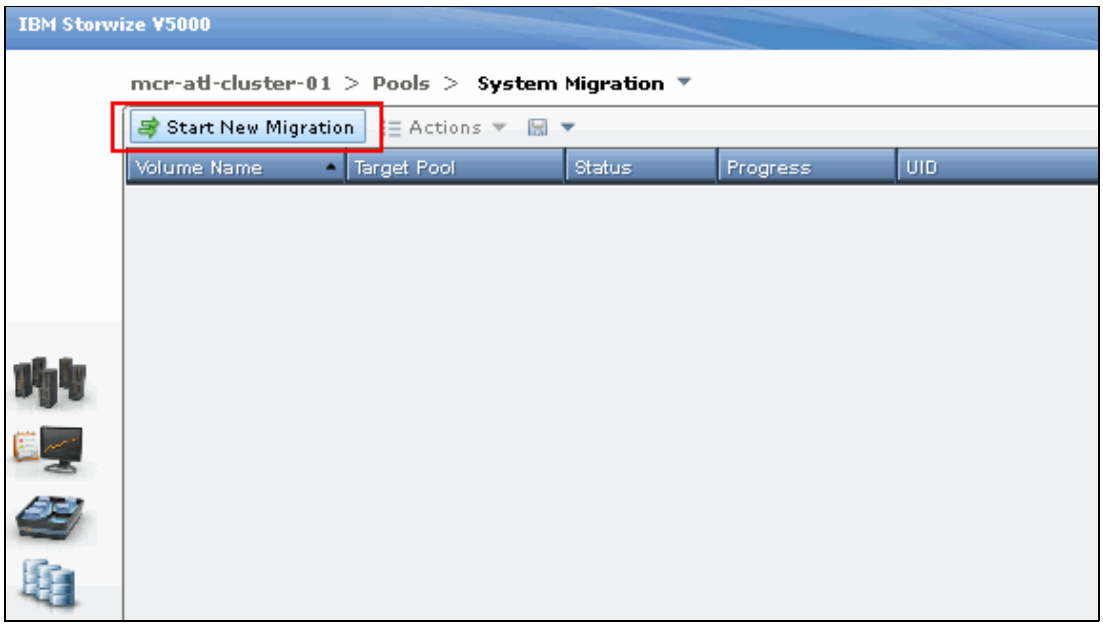

Figure 6-2 System Migration panel

Step 1: Before you begin

Follow step 1 of the storage migration wizard in which the restrictions and prerequisites are described. Read and select each restriction and prerequisite that applies to the planned migration, as shown in Figure 6-3.

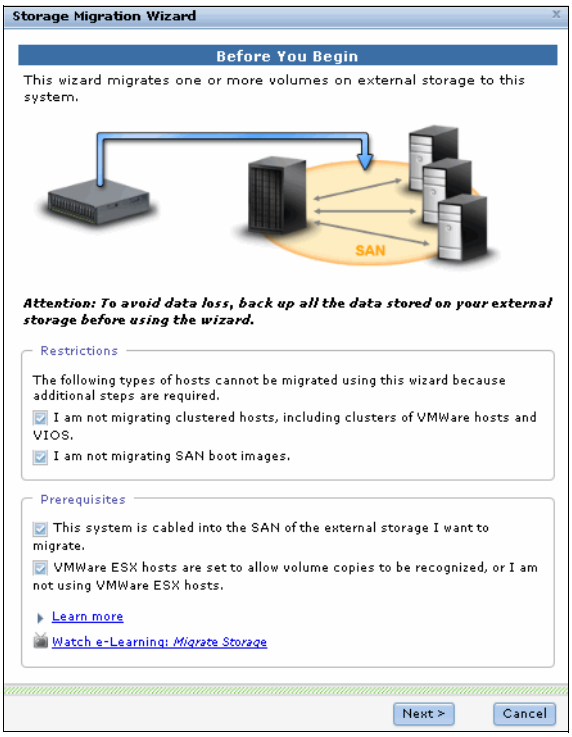

Figure 6-3 Step 1 of the storage migration wizard

Restrictions

Confirm that the following conditions are met:

- You are not using the storage migration wizard to migrate cluster hosts, including clusters of VMware hosts and Virtual I/O Servers (VIOS).
- You are not using the storage migration wizard to migrate SAN Boot images.

If the restrictions options cannot be selected, the migration must be performed outside of this wizard because more steps are required. For more information, see the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp

The VMware ESX Storage vMotion feature might be an alternative for migrating VMware clusters. For more information, see this website:

http://www.vmware.com/products/vmotion/overview.html

Prerequisites

Confirm that the following prerequisites apply:

- Make sure that the Storwize V5000, older storage system, hosts, and Fibre Channel ports are physically connected to the SAN fabrics.
- If there are VMware ESX hosts involved in the data migration, make sure that the VMware ESX hosts are set to allow volume copies to be recognized. For more information, see the VMware ESX product documentation at this website:

  `http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html?`

If all options can be selected, click **Next** to continue. In all other cases, Next cannot be selected and the data must be migrated without use of this wizard. Figure 6-4 shows step 1 of the storage migration wizard with all restrictions satisfied and prerequisites met.

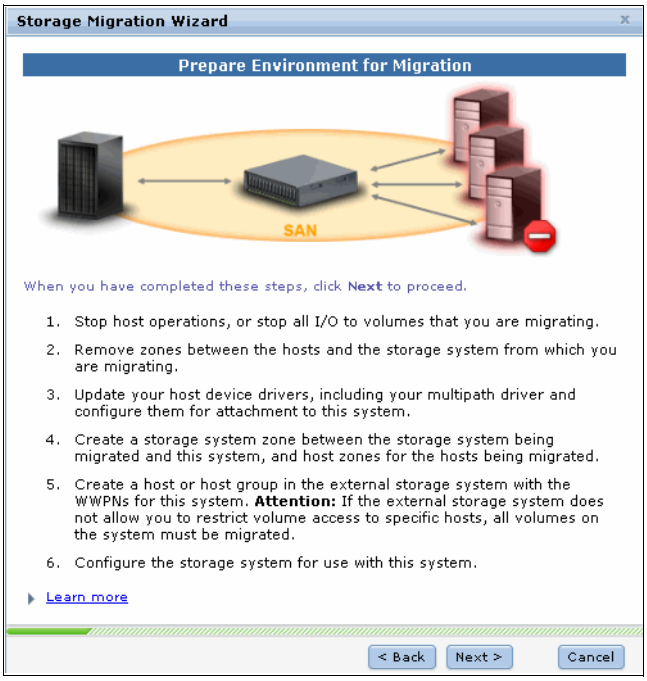

Figure 6-4 Prepare environment

Step 2: Prepare environment for migration

Follow step 2 of the wizard carefully. When all of the required tasks are complete, click **Next** to continue. Figure 6-4 shows the Prepare Environment for Migration panel.

Step 3: Map storage

Follow step 3 of the wizard and click **Next** to continue. Record all of the details carefully because the information can be used in later panels. Table 6-1 shows an example table for capturing the information that relates to older storage system LUs.

Table 6-1 Example table for capturing external LU information

LU name	Controller	Array	SCSI ID	Host name	Capacity
MCRPRDW2K801	DS3400_01	Array_01	0	MCRPRDW2K8	50 GB
MCRPRDW2K802	DS3400_01	Array_01	1	MCRPRDW2K8	200 GB
MCRPRDLNX01	DS3400_01	Array_02	0	MCRPRDLNX	100 GB
MCRPRDLNX02	DS3400_01	Array_02	1	MCRPRDLNX	300 GB

SCSI ID: Record the SCSI ID of the LUs to which the host is originally mapped. Some operating systems do not support changing the SCSI ID during the migration.

Table 6-2 shows an example table for capturing host information.

Table 6-2 Example table for capturing host information

Host Name/ LU Names	Adapter / Slot / Port	WWPN	HBA F/W	HBA Device Driver	Operating System	V5000 Multipath Software
MCRPRDW2K8	QLE2562 / 2 / 1	21000024FF2D0BE8	2.10	9.1.9.25	W2K8 R2 SP1	SDDDSM 2.4.3.1-2
MCRPRDW2K8	QLE2562 / 2 / 2	21000024FF2D0BE9	2.10	9.1.9.25	W2K8 R2 SP1	SSDDSM 2.4.3.1-2
MCRPRDLNX	LP10000 / 0 / 1	10000000C1234A56	2.72a2	8.2.0.63.3p	RHEL5	Device Mapper
MCRPRDLNX	LP10000 / 1 / 1	10000000C6789A01	2.72a2	8.2.0.63.3p	RHEL5	Device Mapper

Figure 6-5 shows the Map Storage panel.

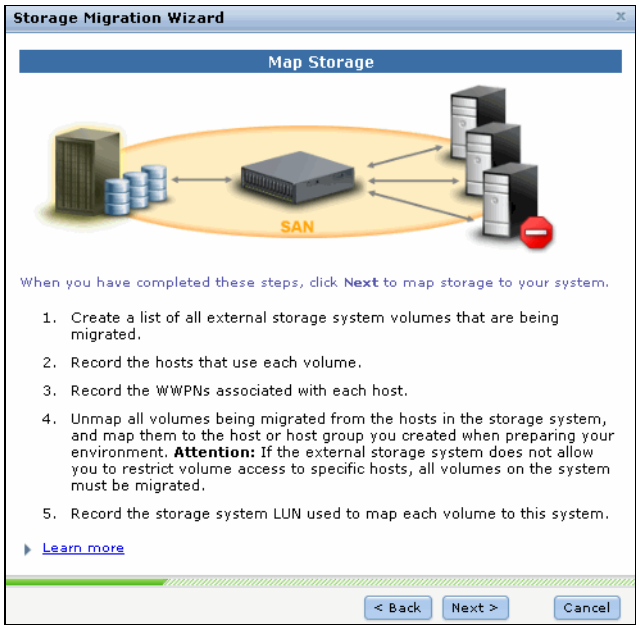

Figure 6-5 Map Storage panel

The Storwize V5000 runs the discover devices task. After the task is complete, click **Close** to continue. Figure 6-6 on page 245 shows the results of the Discover Devices task.

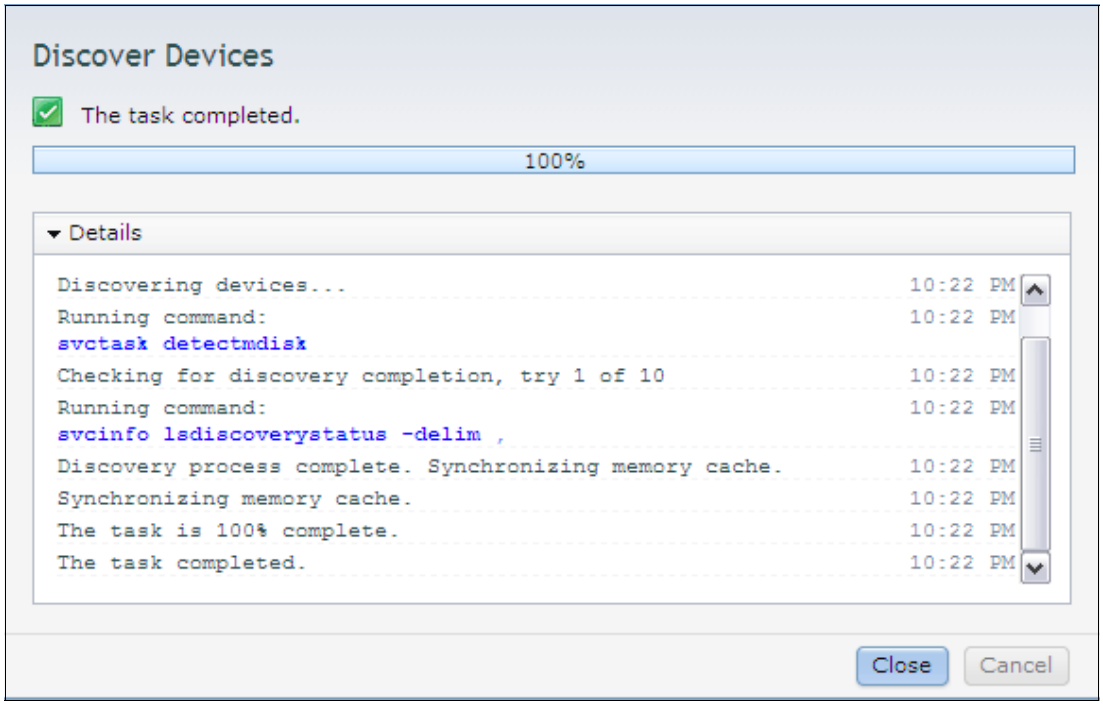

Figure 6-6 Discover Devices task

Step 4: Migrating MDisks

Follow step 4 of the wizard and select the MDisks that are to be migrated and then click **Next** to continue. Figure 6-7 shows the Migrating MDisks panel.

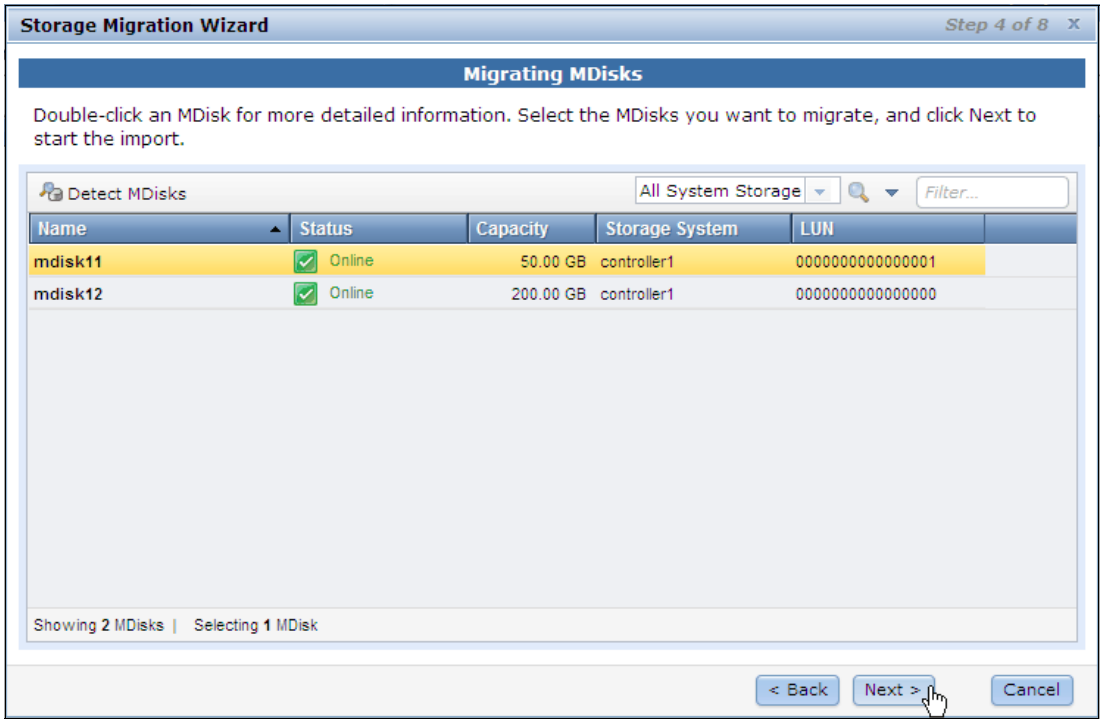

Figure 6-7 Migrating MDisks panel

MDisk selection: Select only the MDisks that are applicable to the current migration plan. After step 8 of the current migration completes, another migration plan can be started to migrate any remaining MDisks.

The Storwize V5000 runs the Import MDisks task. After the task is complete, click **Close** to continue. Figure 6-8 shows the result of the Import MDisks task.

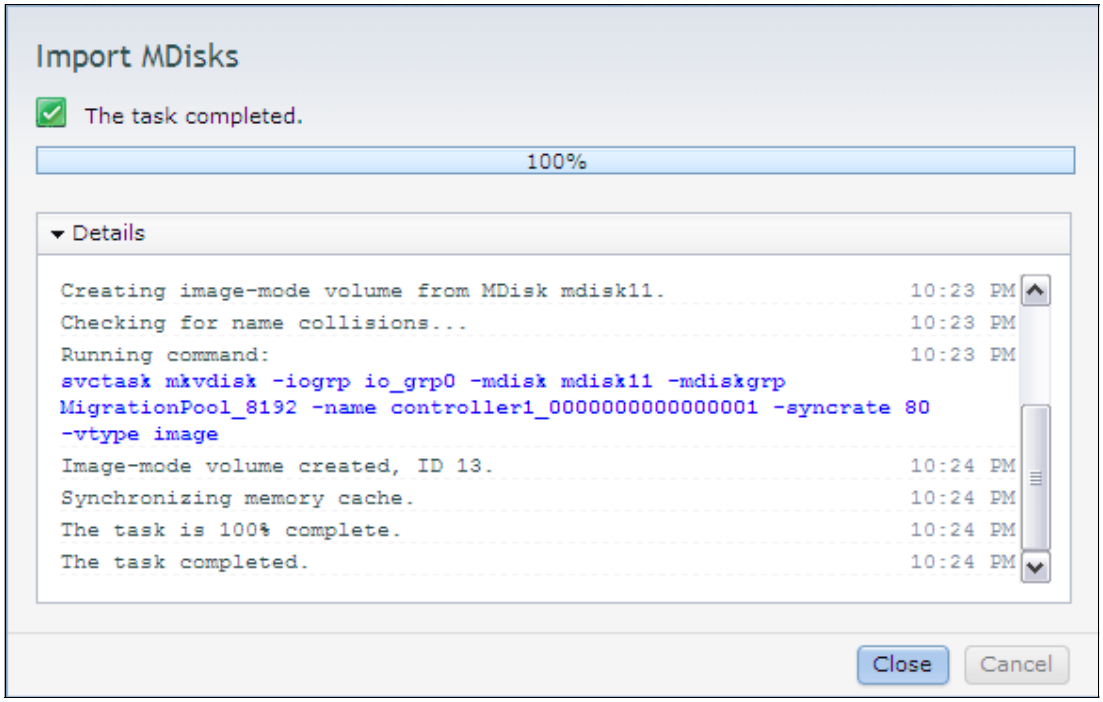

Figure 6-8 Import MDisks task

Step 5: Configure hosts

Follow step 5 of the wizard to select or configure new hosts, as required. Click **Next** to continue. Figure 6-9 shows the Configure Hosts panel.

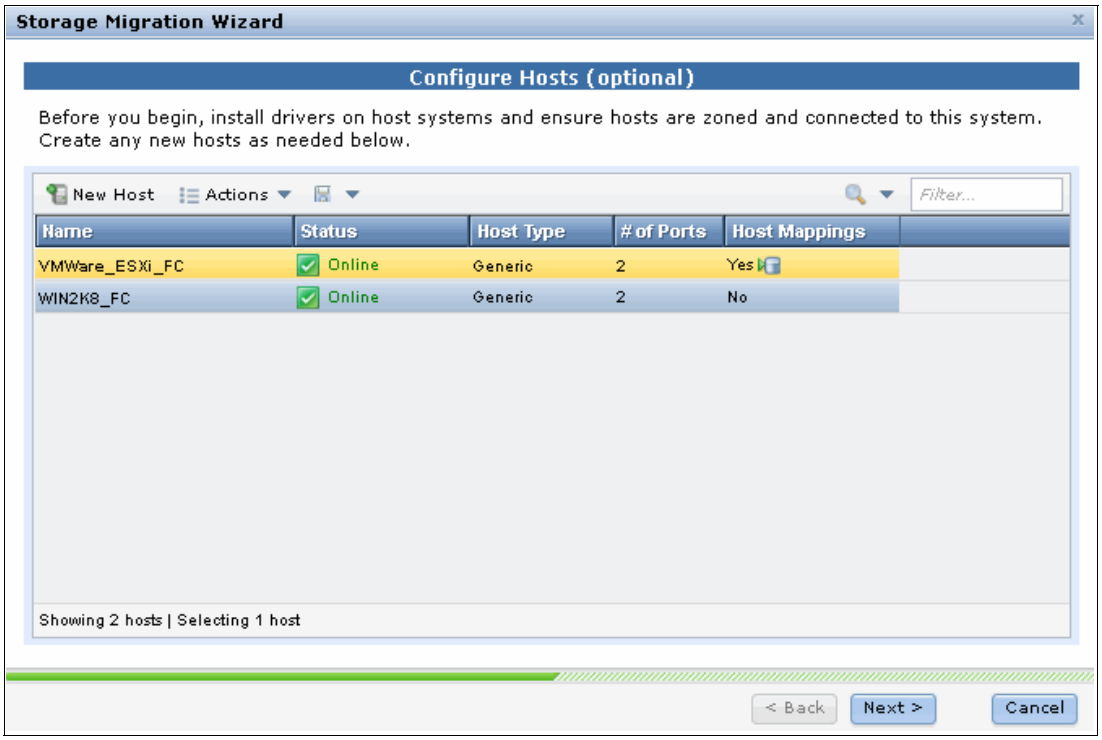

Figure 6-9 Configure Hosts panel

Important: It is not mandatory to select the hosts now. The actual selection of the hosts occurs in the next step, Map Volumes to Hosts. However, take this opportunity to cross-check the hosts that have data to be migrated by highlighting them in the list before you click Next.

Step 6: Map volumes to hosts

Follow step 6 of the wizard to select the newly migrated volume. Click **Map to Host.** Figure 6-10 on page 248 shows the Map Volumes to Hosts panel.

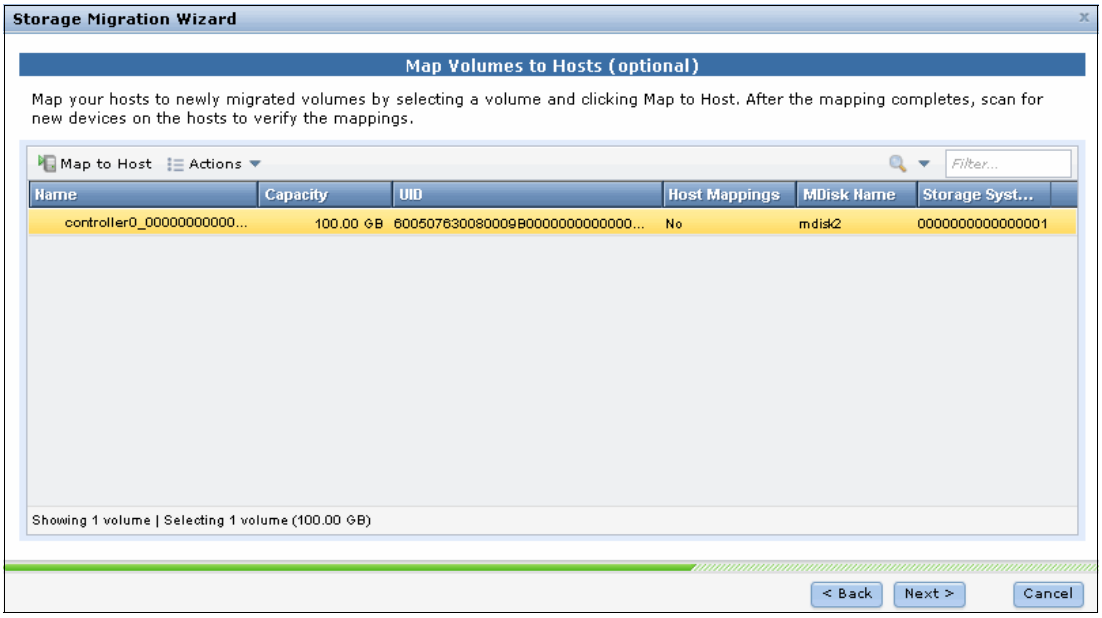

Figure 6-10 Map Volumes to Hosts panel

The image mode volumes are listed and the names of the image mode volumes are assigned automatically by the Storwize V5000 storage system. The names can be changed to reflect something more meaningful to the user by selecting the volume and clicking **Rename** in the Actions menu.

Names: The names of the image mode volumes must begin with a letter. The name can be a maximum of 63 characters. The following valid characters can be used:

- Uppercase letters (A - Z)
- Lowercase letters (a - z)
- Digits (0 - 9)
- Underscore (_)
- Period (.)
- Hyphen (-)
- Blank space

The names must not begin or end with a space.

A Host drop-down menu is displayed. Select the required host and the Modify Host Mappings panel is opened, in which the Choose a Host menu is available, as shown in Figure 6-11.

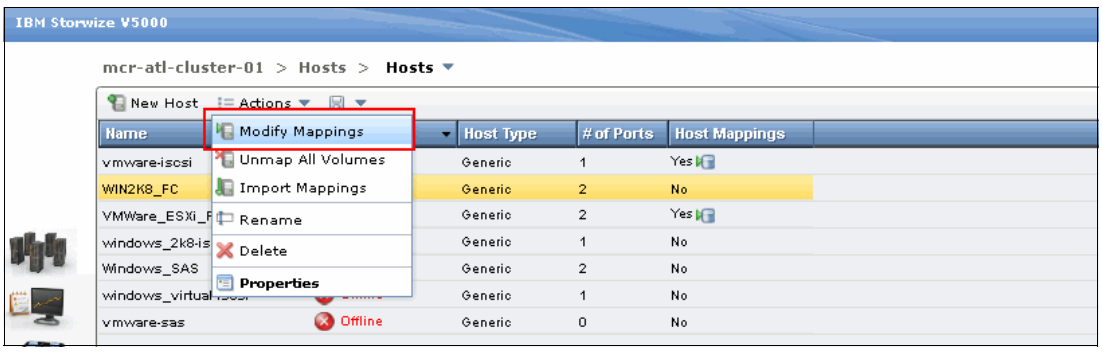

Figure 6-11 Modify Mappings

The MDisks highlighted in step 6 of the wizard are shown in yellow in the Modify Host Mappings panel. The yellow highlighting means that the volumes are not yet mapped to the host. Click **Edit SCSI ID** and modify as required. The SCSI ID should reflect the same SCSI ID as was recorded in step 3. Click **Map Volumes**. Figure 6-12 shows the Modify Host Mappings panel.

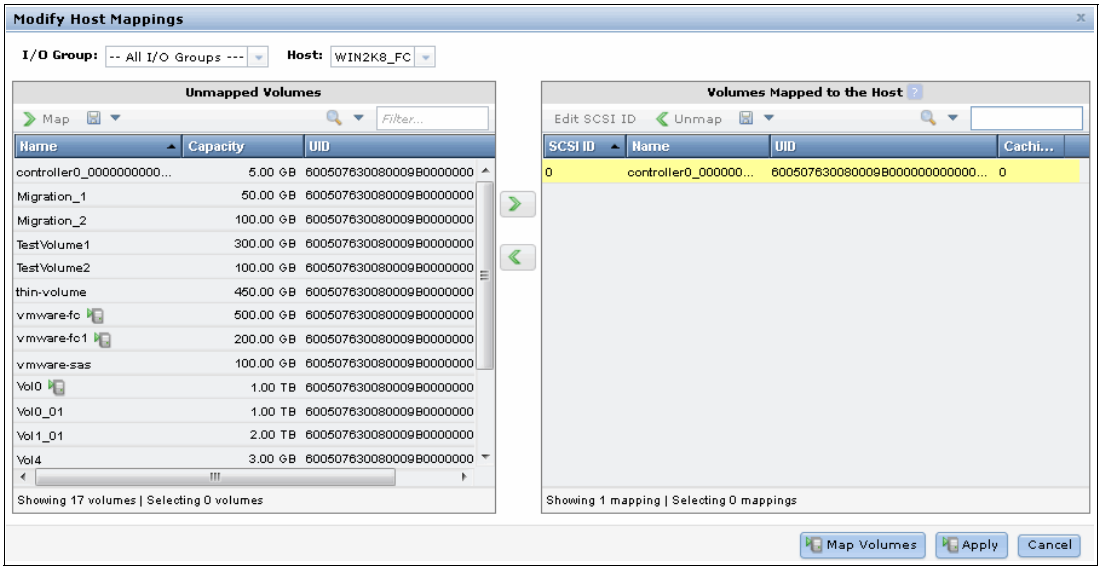

Figure 6-12 Modify Host Mappings panel

The Storwize V5000 runs the modify mappings task. After the task is complete, the volume is mapped to the host. Click **Close** to continue. Figure 6-13 shows the Modify Mappings task.

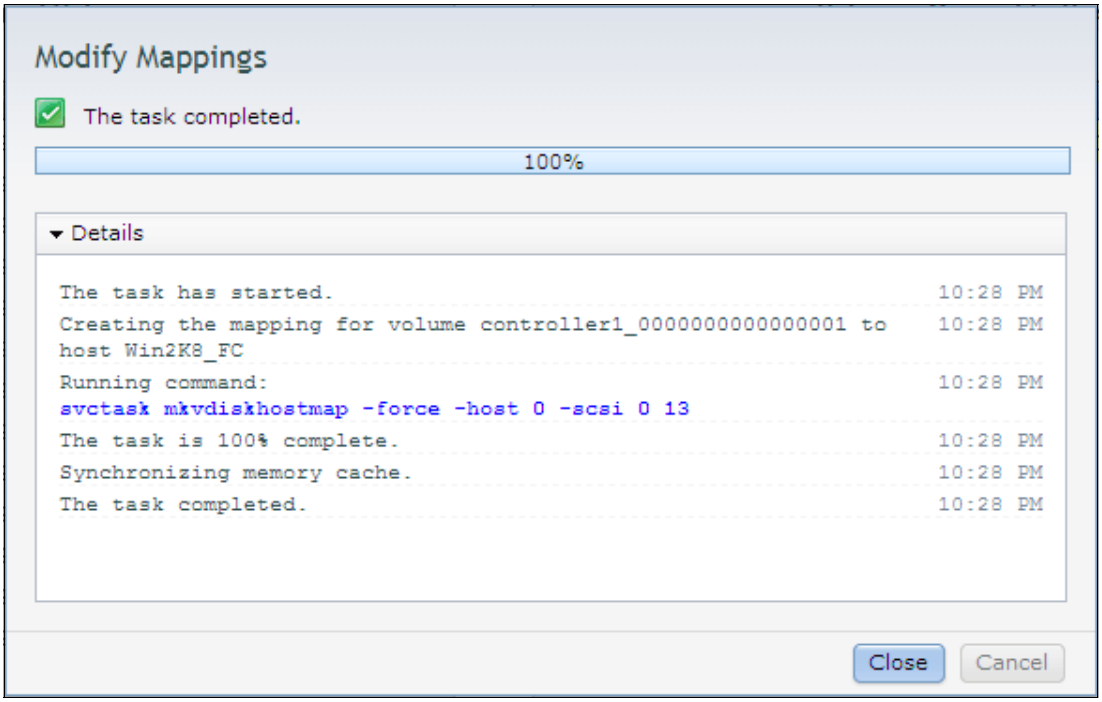

Figure 6-13 Modify Mappings task

The Map Volumes to Hosts panel is displayed again. Verify that the migrated volumes now have Yes in the Host Mappings column. Click **Next** to continue. Figure 6-14 shows the Map Volumes to Hosts panel with Yes in the Host Mappings column.

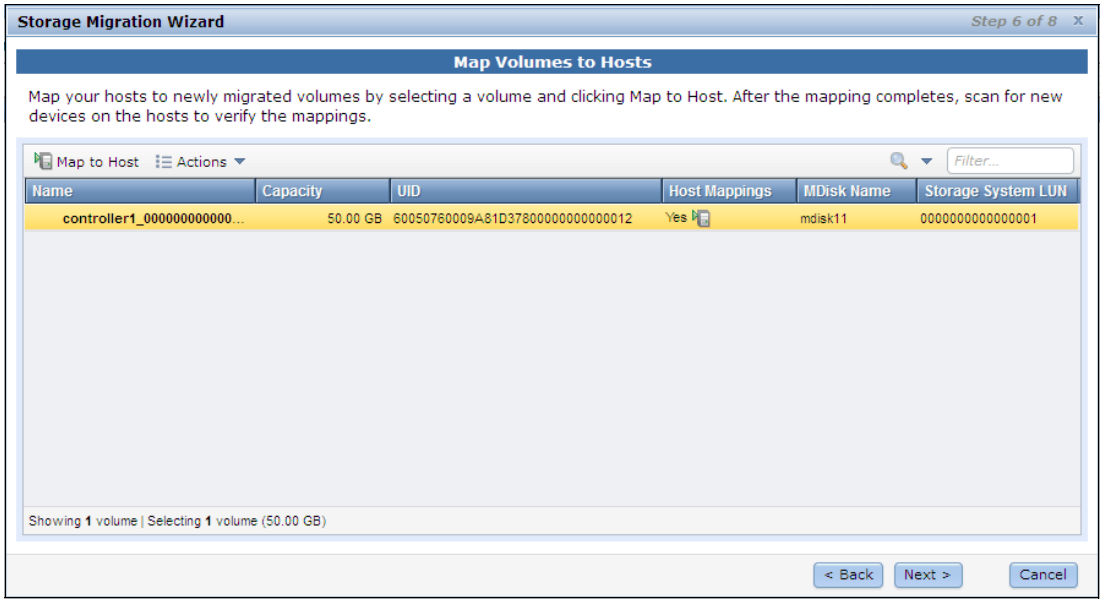

Figure 6-14 Map Volumes to Hosts panel that shows Yes in the Host Mappings column

Scan for new devices on the hosts to verify the mapping. The disks are now displayed as IBM 2145 Multi-Path disk devices. This disk device type is common for the IBM Storwize disk family and the IBM SAN Volume Controller.

Step 7: Select storage pool

Follow step 7 of the wizard to select an internal storage pool. Click **Next** to continue. The destination storage pool of the data migration is an internal storage pool of the Storwize V5000. Ensure that there is enough space in the selected storage pool to accommodate the migrated volume. The migration task runs in the background and results in a copy of the data is placed on the MDisks in the selected storage pool.

The process uses the volume mirroring function that is included with the Storwize V5000. When the process completes, the volumes feature points to the new copy on the internal storage pool that is selected and the older storage system. Figure 6-15 on page 251 shows the Select a Pool panel.

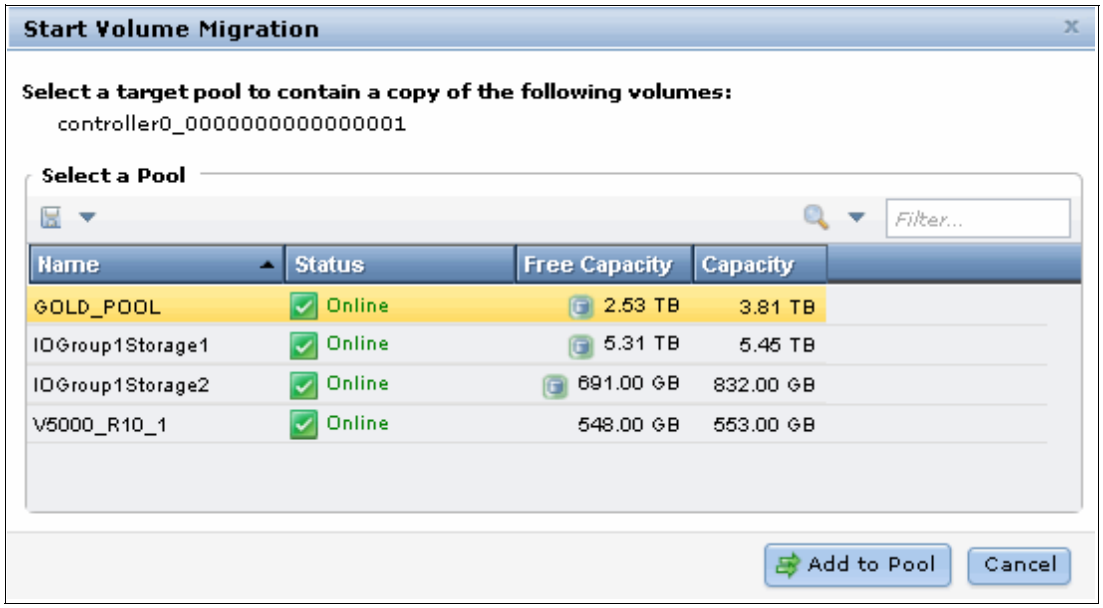

Figure 6-15 Select a Pool panel

The Storwize V5000 runs the start migration task. After the task is complete, click **Close** to continue. Figure 6-16 shows the result of the Start Migration task.

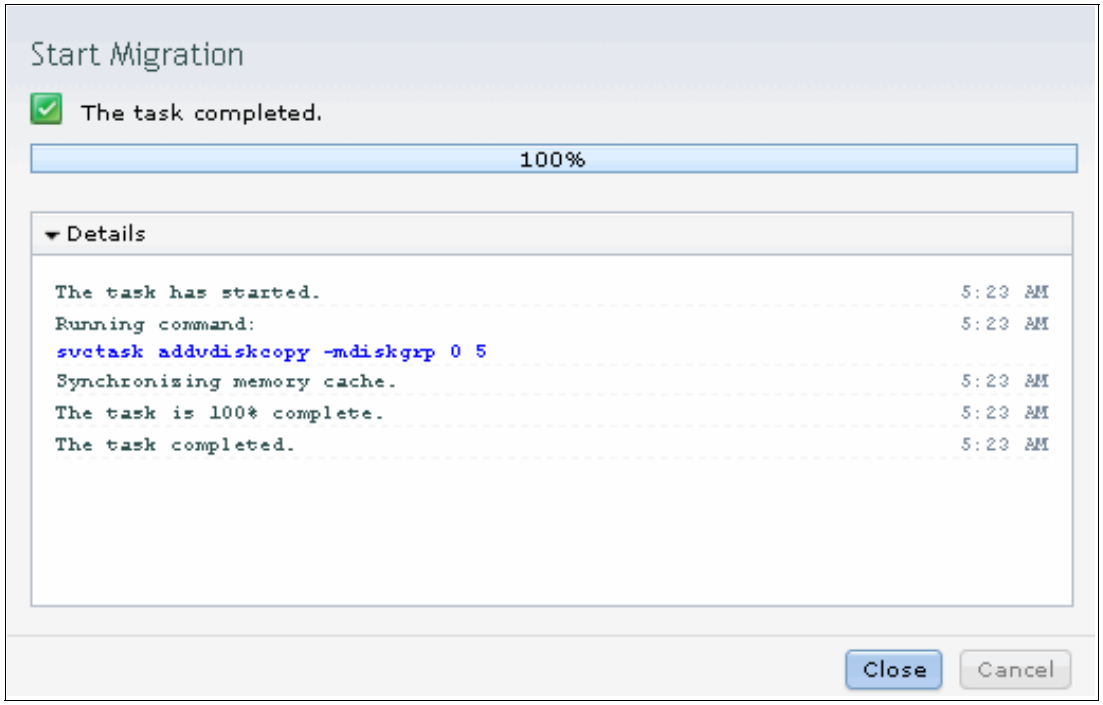

Figure 6-16 Start Migration task

Step 8: Finish the storage migration wizard

Follow step 8 of the wizard and click **Finish**, as shown in Figure 6-17.

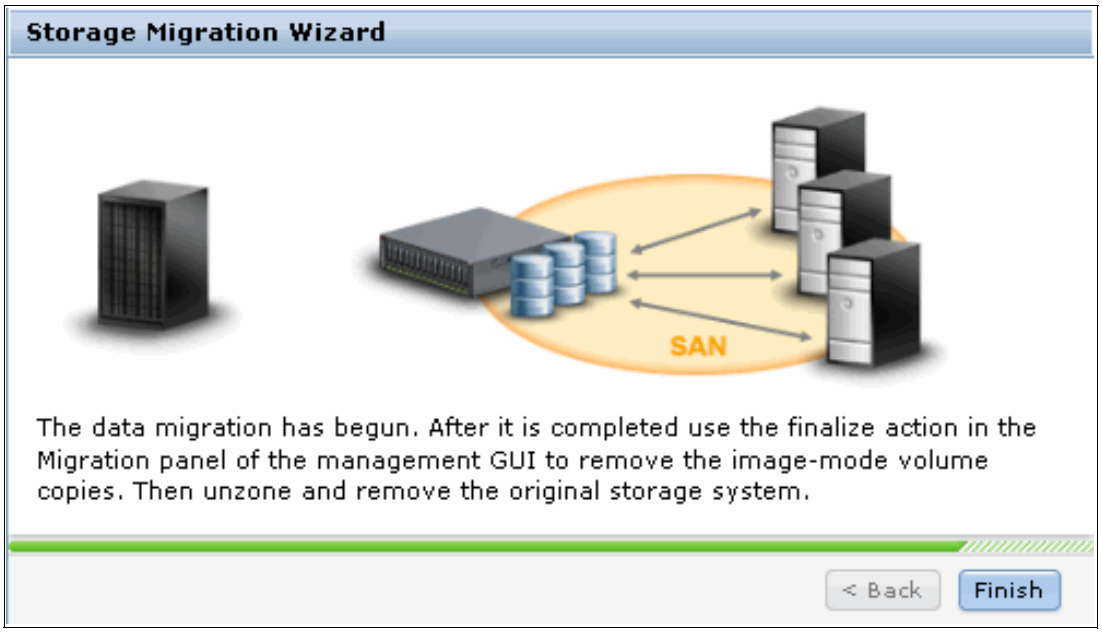

Figure 6-17 Step 8 of the storage migration wizard

The end of the storage migration wizard is not the end of the data migration process. The data migration is still in progress. A percentage indication of the migration progress is displayed in the System Migration panel, as shown in Figure 6-18.

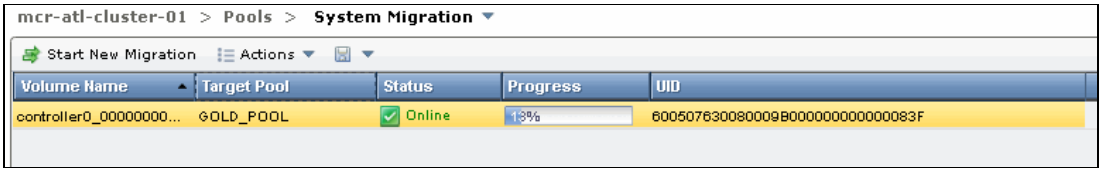

Figure 6-18 Storage Migration panel with a migration in progress

Finalize migrated volumes

When the migration completes, select the Migration and right-click **Finalize**. Verify that the volume names and the number of migrations and click **OK**. The image mode volumes are deleted and the associated image mode MDisks from the migration storage pool are removed. The status of those image mode MDisks is then unmanaged. When the finalization completes, the data migration to the IBM Storwize V5000 is done. Remove zoning and retire the older storage system.

6.3 Storage migration wizard example scenario

This section describes an example scenario that provides some details that relate to the attachment and verification tasks that are associated with running the storage migration wizard.

6.3.1 Storage migration wizard example scenario description

The example scenario shows the introduction of a Storwize V5000 to an environment that contains existing storage infrastructure, which includes a SAN fabric, a Windows 2008 host, and an IBM DS3400 storage system.

The Windows 2008 host has existing data on the disks of an IBM DS3400 storage system. That data must be migrated to the internal storage of the Storwize V5000. The Windows 2008 host has a dual port QLogic Host Bus Adapter (HBA) type QLE2562. Each of the Fibre Channel switches is the IBM 2498-24B type. There are two host disks to be migrated: devices Disk 1 and Disk 2. Figure 6-19 shows the Windows 2008 Disk management panel. The two disks feature defined volumes. The volume labels are Migration 1 (G: drive) and Migration 2 (H: drive).

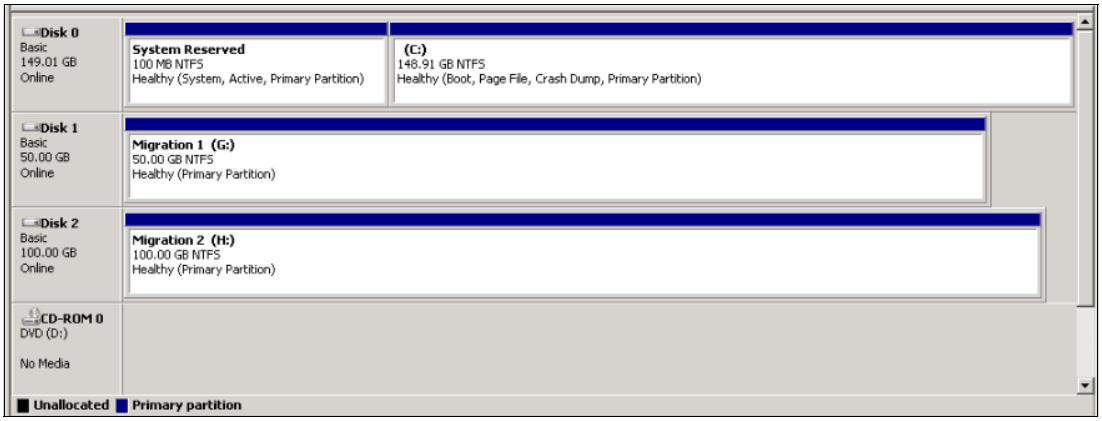

Figure 6-19 Windows 2008 disk management panel

The two disks to be migrated are on the IBM DS3400 storage system. Therefore, the disk properties display the disk device type as an IBM1726-4xx FAStT disk device. To show this disk attribute, right-click the disk to show the menu and then select **Properties**, as shown in Figure 6-20.

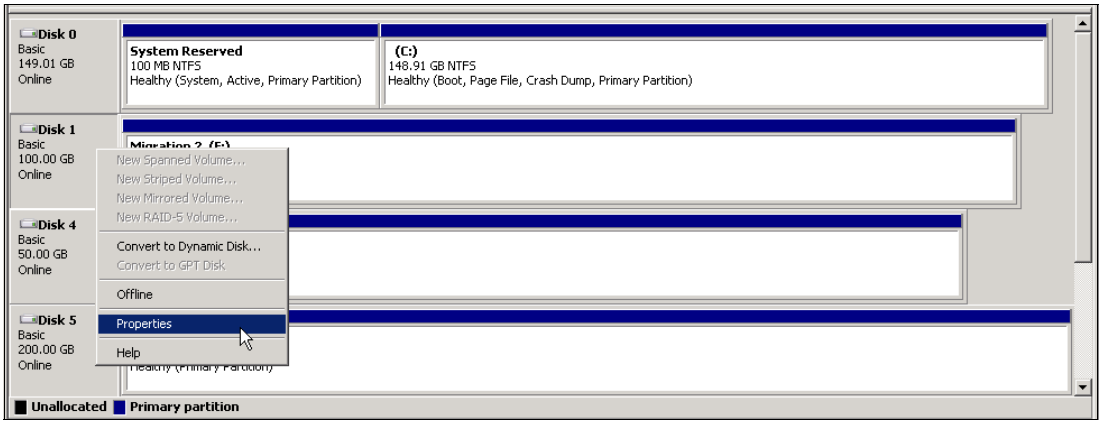

Figure 6-20 Display properties of disk before migration

After the disk properties panel is opened, the General tab shows the disk device type. Figure 6-21 shows the General tab in the Windows 2008 Disk Properties window.

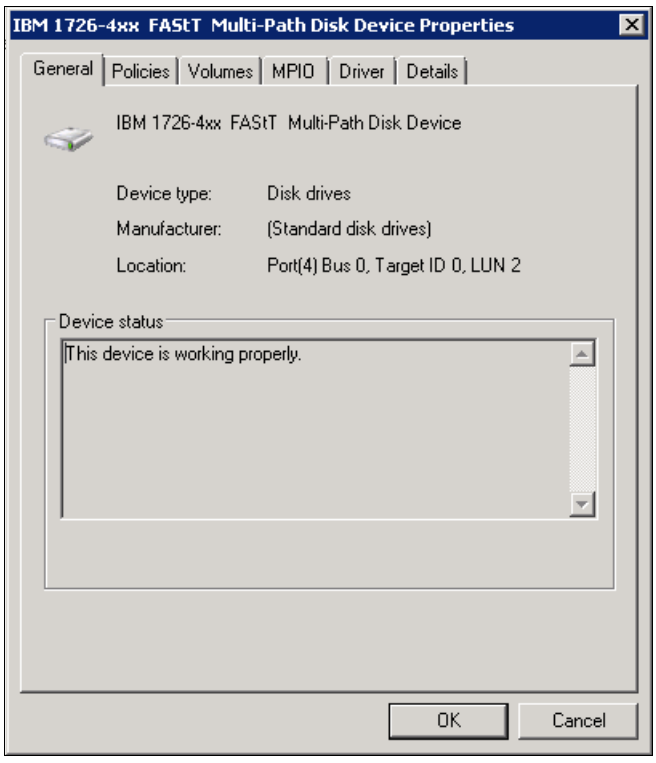

Figure 6-21 Windows 2008 Disk Properties: General tab

Perform this task on all disks before the migration and then the same check can be done after the disks are presented from the Storwize V5000. After the Storwize V5000 mapping and host rescan, the disk device definitions are changed to IBM 2145 Multi-Path disk device and it is confirmed that the disks are under Storwize V5000 control.

Example scenario Fibre Channel cabling layout

To provide more information about the example migration, Figure 6-22 shows the example scenario Fibre Channel cabling layout. The Host, IBM DS3400, and Storwize V5000 are cabled into a dual SAN fabric configuration. The connection method that is shown can provide improved availability through fabric and path redundancy and improved performance through workload balancing.

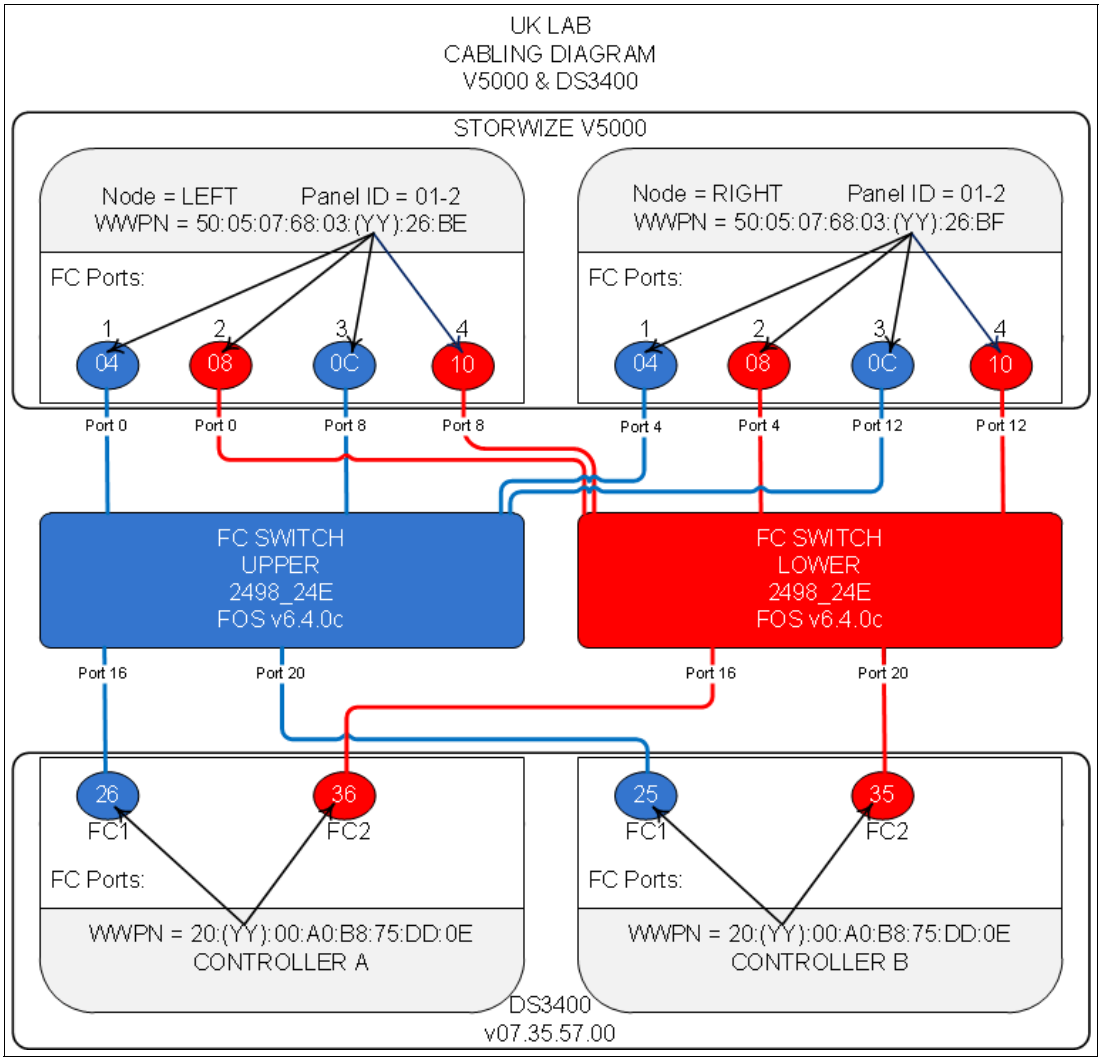

Figure 6-22 Example scenario Fibre Channel cabling layout

6.3.2 Using the storage migration wizard for example scenario

This section provides an overview of the storage migration tasks that are performed when the storage migration wizard is used for the example scenario. A more detailed perspective also is provided to assist users that require more information.

Overview of storage migration wizard tasks for example scenario

The following steps provide an overview of the wizard tasks for our example scenario:

1. Search the IBM SSIC for scenario compatibility.
2. Back up all of the data that is associated with the host, DS3400, and Storwize V5000.

3. Start New Migration to open the wizard on the Storwize V5000.
4. Follow Step 1 in the wizard before you begin.
5. Follow Step 2 in the wizard to prepare the environment for migration, including the following steps:
   a. Stop host operations or stop all I/O to volumes that you are migrating.
   b. Remove zones between the hosts and the storage system from which you are migrating. Remove Host-to-DS3400 zones on SAN.
   c. Update your host device drivers, including your multipath driver and configure them for attachment to this system. Complete the steps that are described in 4.2.1, “Windows 2008 R2: Preparing for FC attachment” on page 155 to connect to Storwize V5000 that uses Fibre Channel.

      Pay attention to the following tasks:

      i. Make sure that the latest OS service pack and test fixes are applied to your Microsoft server.
      ii. Use the latest firmware and driver levels on your host system.
      iii. Install and HBA or HBAs on the Windows server that uses the latest BIOS and drivers.
      iv. Connect the FC Host Adapter ports to the switches.
      v. Configure the switches (zoning).
      vi. Configure the HBA for hosts that are running Windows.
      vii. Set the Windows timeout value.
      viii. Install the Subsystem Device Driver Device Specific Module (SDDDSM) multipath module.
   d. Create a storage system zone between the storage system that is migrated and this system and host zones for the hosts that are migrated.

      Pay attention to the following tasks:

      i. Locate the WWPNs for Host.
      ii. Locate WWPNs for IBM DS3400.
      iii. Locate WWPNs for Storwize V5000.
      iv. Define port aliases definitions on SAN.
      v. Add V5000-to-DS3400 zones on SAN.
      vi. Add Host-to-V5000 zones on SAN.
   e. Create a host or host group in the external storage system with the WWPNs for this system.

      Important: If you cannot restrict volume access to specific hosts by using the external storage system, all volumes on the system must be migrated.

      Add Storwize V5000 host group on DS3400
   f. Configure the storage system for use with this system.

      Follow the IBM Storwize V5000 Version 6.4.1 Information Center for DS3400 configuration recommendations.
6. Follow Step 3 of the wizard to map storage, including the following steps:
   a. Create a list of all external storage system volumes that are migrated.

      Create a DS3400 LU table.

b. Record the hosts that use each volume.

   Create Host table.

c. Record the WWPNs associated with each host.

   Add WWPNs to Host table.

d. Unmap all volumes that are migrated from the hosts in the storage system and map them to the host or host group that you created when your environment was prepared.

   Important: If you cannot restrict volume access to specific hosts by using the external storage system, all volumes on the system must be migrated.

   Move LUs from Host to Storwize V5000 Host Group on DS3400.

e. Record the storage system LUN that is used to map each volume to this system.

   Update the DS3400 LU table.

7. Follow Step 4 of the wizard to migrate MDisks. Select discovered MDisk on Storwize V5000.
8. In Step 5 of the wizard, configure hosts by completing the following steps:
   a. Create Host on Storwize V5000.
   b. Select Host on Storwize V5000.
9. In Step 6 of the wizard, map volumes to hosts by completing the following steps:
   a. Map volumes to Host on Storwize V5000.
   b. Verify that disk device type is now 2145 on Host.
   c. SDDDSM datapath query commands on Host.
10. In Step 7 of the wizard, select the storage pool. Select internal storage pool on Storwize V5000.
11. Finish the storage migration wizard.
12. Finalize the migrated volumes.

Detailed view of the storage migration wizard for the example scenario

The following steps provide an overview of the wizard tasks for our example scenario:

1. Search the IBM SSIC for scenario compatibility.
2. Back up all of the data that is associated with the host, DS3400, and Storwize V5000.
3. Start New Migration to open the wizard on the Storwize V5000, as shown in Figure 6-23.

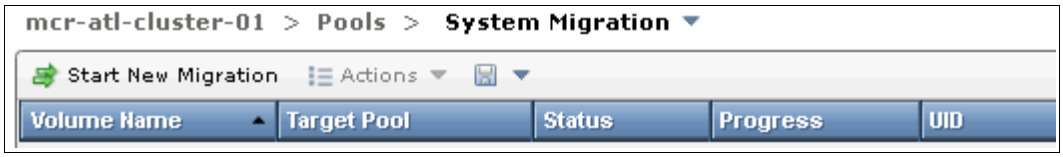

Figure 6-23 Start new migration

4. Follow step 1 of the wizard and select all of the restrictions and prerequisites, as shown in Figure 6-24. Click **Next** to continue.

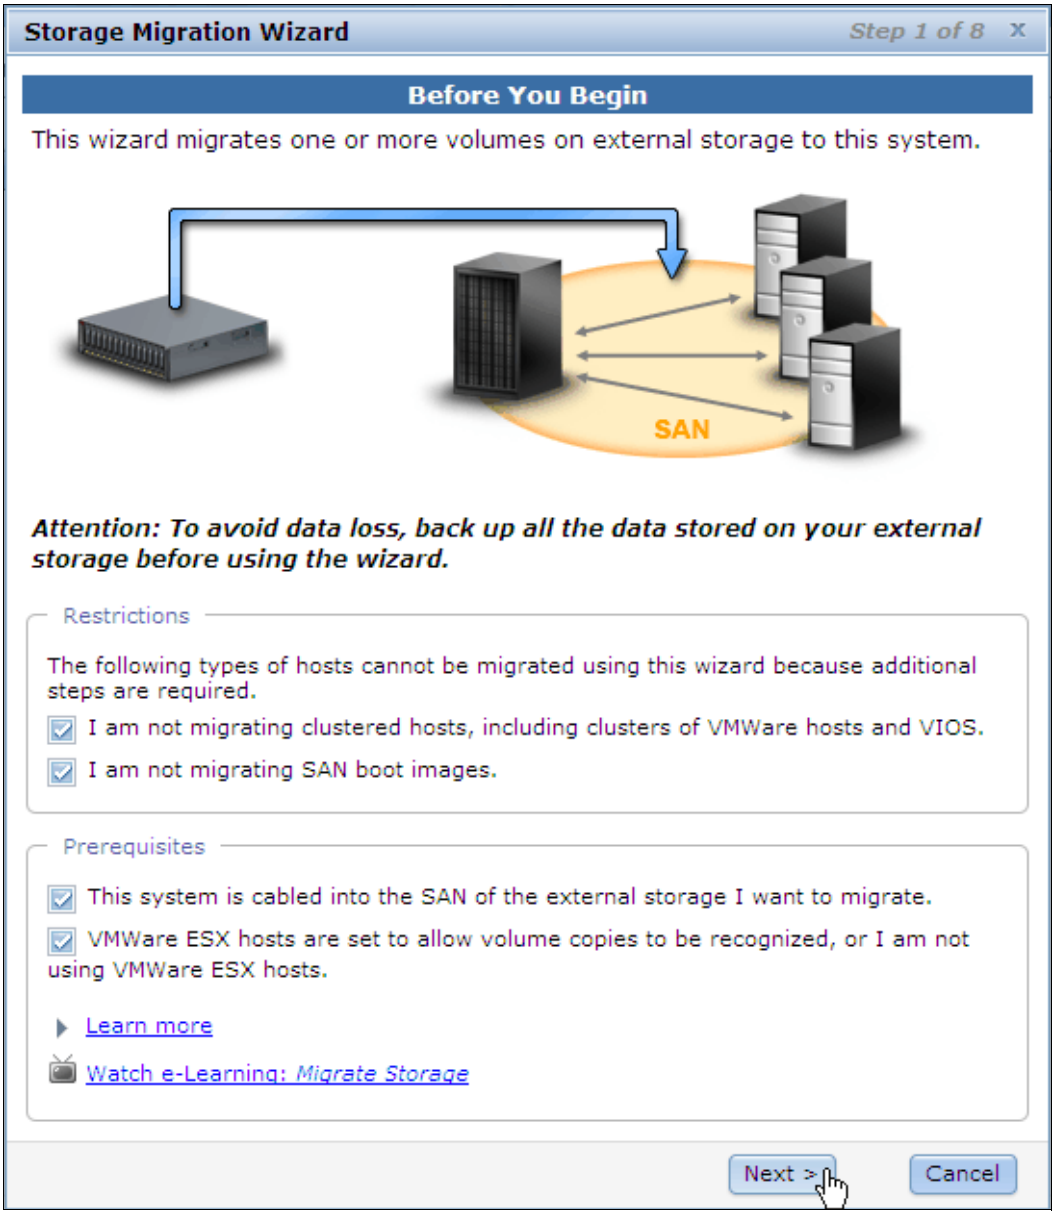

Figure 6-24 Storage Migration wizard: Step 1

5. Follow step 2 the wizard, as shown in Figure 6-25. Complete all of the steps before you continue.

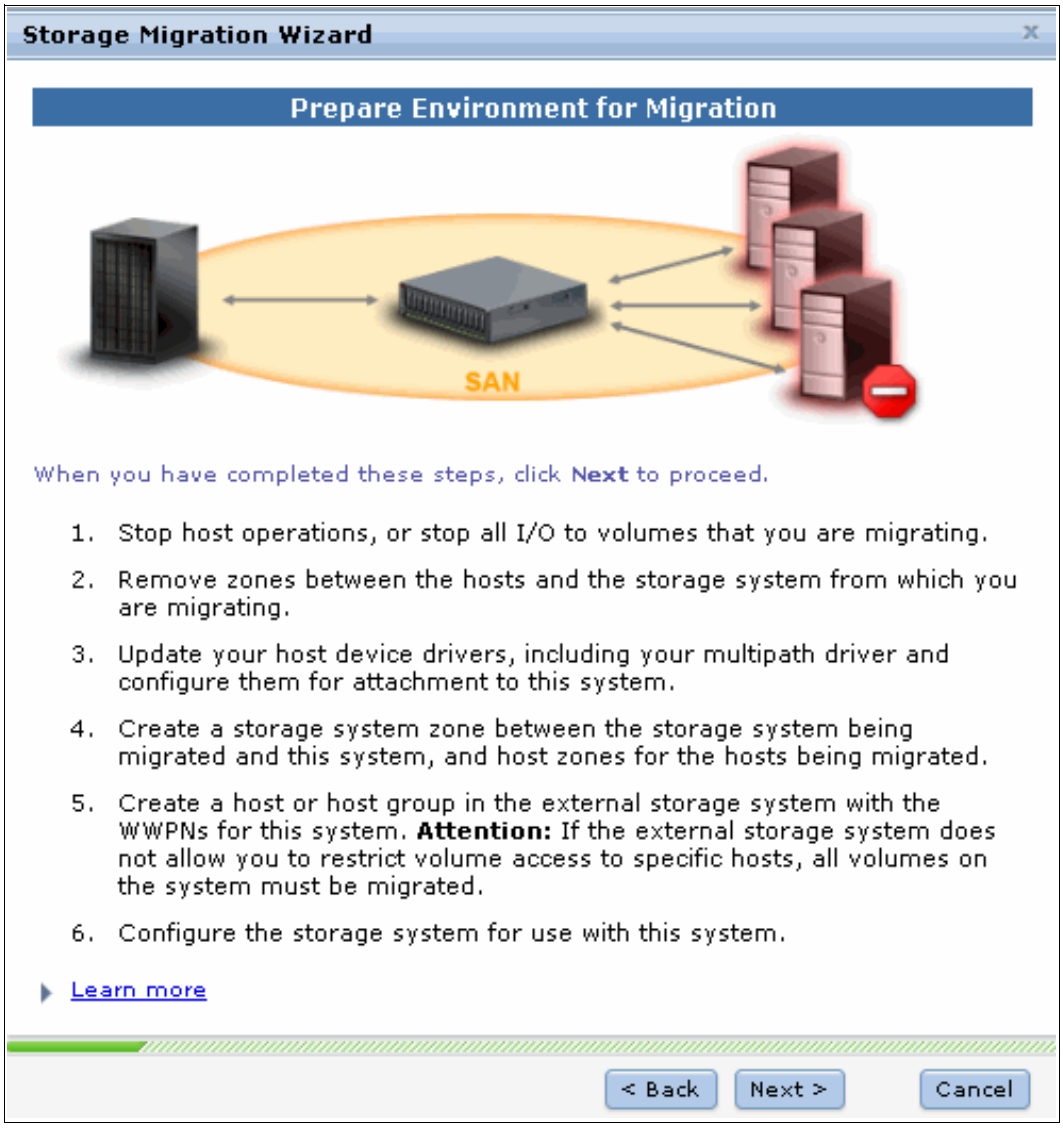

Figure 6-25 Storage Migration wizard: Step 2

Pay attention to the following tasks:

a. Stop host operations or stop all I/O to volumes that you are migrating.

b. Remove zones between the hosts and the storage system from which you are migrating.

c. Update your host device drivers (including your multipath driver) and configure them for attachment to this system. Complete the steps that are described in 4.2.1, “Windows 2008 R2: Preparing for FC attachment” on page 155 to prepare a Windows host to connect to Storwize V5000 by using Fibre Channel.

   Pay attention to the following tasks during this process:

   i. Make sure that the latest OS service pack and test fixes are applied to your Microsoft server.

   ii. Use the latest firmware and driver levels on your host system.

iii. Install HBAs on the Windows server that uses the latest BIOS and drivers.

iv. Connect the FC Host Adapter ports to the switches.

v. Configure the switches (zoning).

vi. Configure the HBA for hosts that are running Windows.

vii. Set the Windows timeout value.

viii. Install the multipath module.

d. Create a storage system zone between the storage system that is migrated and this system and host zones for the hosts that are migrated.

To perform this step, locate the WWPNs of the host, IBM DS3400, and Storwize V5000, then create an alias for each port to simplify the zone creation steps.

Important: A WWPN is a unique identifier for each Fibre Channel port that is presented to the SAN fabric.

Locating the HBA WWPNs on the Windows 2008 host

See the original IBM DS3400 Host definition to locate the WWPNs of the host's dual port QLE2562 HBA. To complete this task, open the IBM DS3400 Storage Manager and click the **Modify** tab, as shown in Figure 6-26. Select **Edit Host Topology** to show the host definitions.

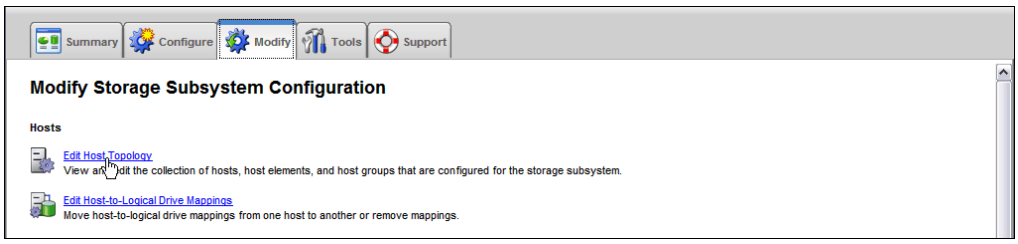

Figure 6-26 IBM DS3400 modify tab: Edit Host Topology

Figure 6-27 shows the IBM DS3400 storage manager host definition and the associated WWPNs.

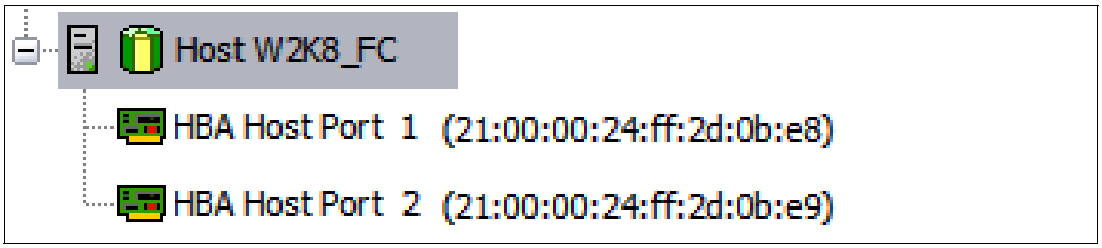

Figure 6-27 IBM DS3400 host definition

Record the WWPNs for alias, zoning, and the Storwize V5000 New Host task.

Important: Alternatively, the QLogic SAN Surfer application for the QLogic HBAs or the SAN fabric switch reports can be used to locate the WWPNs of the host.

Locating the controller WWPNs on the IBM DS3400

The IBM DS3400 Storage Manager can provide the controller WWPNs through the Storage Subsystem Profile. Open the IBM DS3400 Storage Manager, click **Support**, and select **View Storage Subsystem Profile.** Figure 6-28 shows the IBM DS3400 Storage Manager Support tab.

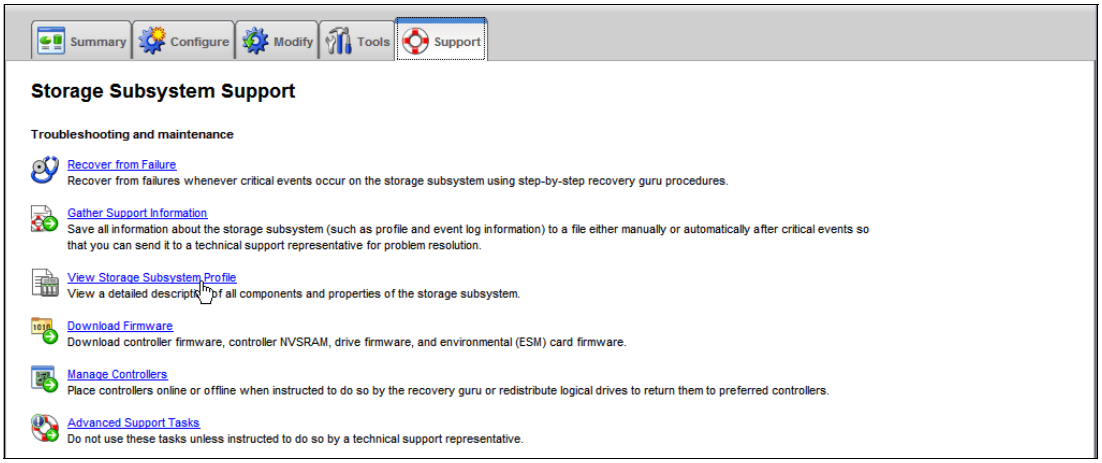

Figure 6-28 Storage Subsystem Support profile

Click the **Controllers** tab to show the WWPNs for each controller. Figure 6-29 shows the IBM Ds3400 storage manager Storage Subsystem Profile.

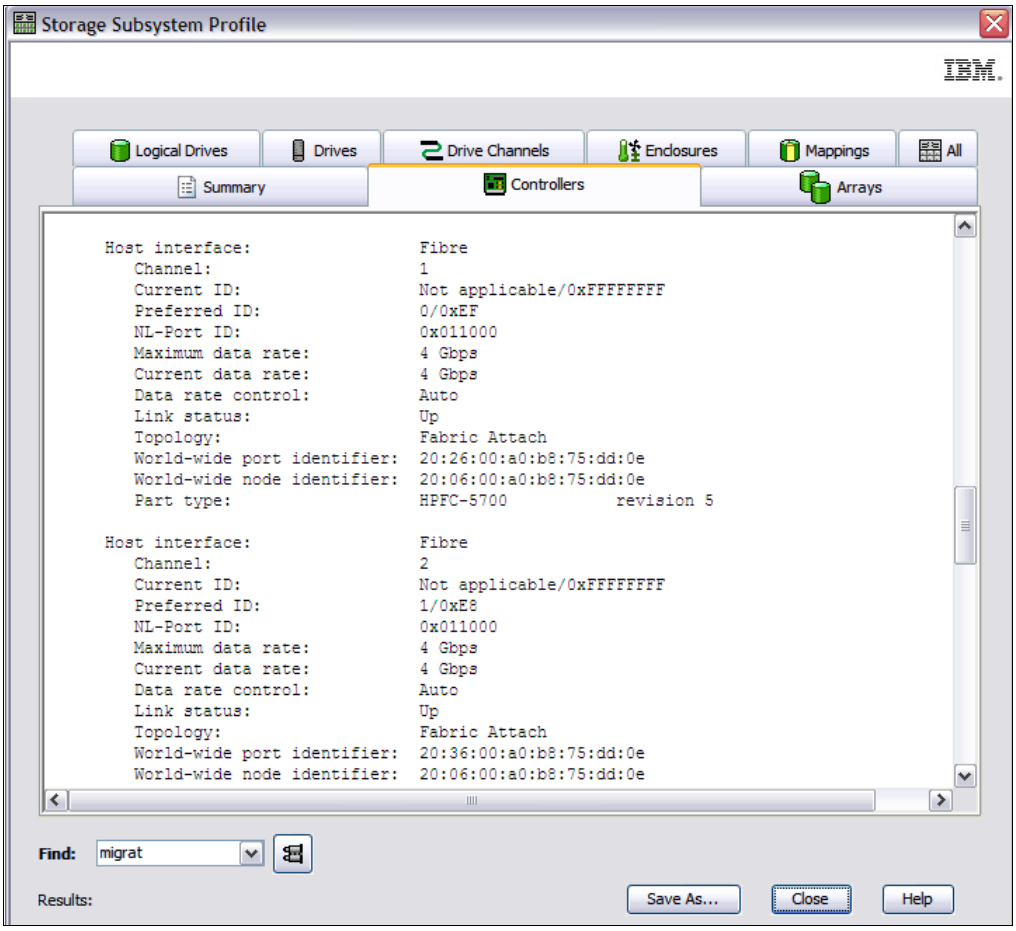

Figure 6-29 Storage Subsystem Profile: Controller WWPNs

Locating node canister WWPNs on the Storwize V5000

To locate the WWPNs for the Storwize V5000 node canisters, expand the control enclosure section and select the canister from the System Details panel. Scroll down to Ports to see the associated WWPNs. Figure 6-30 shows the Storwize V5000 System Details panel with the WWPNs shown when you click **IBM-Storwize-V5000** → **Enclosure 1** → **Canister 1**.

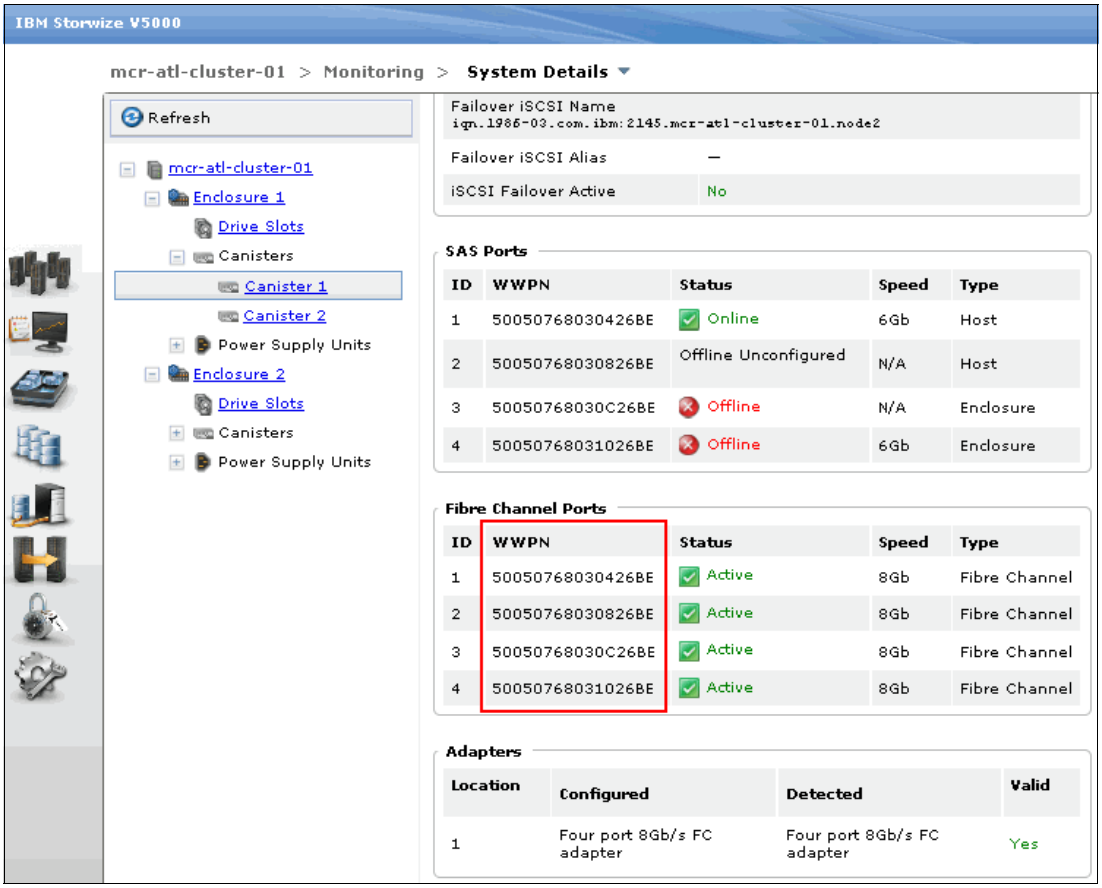

Figure 6-30 Storwize V5000 node canister WWPNs information

WWPN: The WWPN consists of eight bytes (two digits per byte). In Figure 6-30, the third byte pair in the listed WWPNs are 04, 08, 0C, and 10. They are the differing bytes for each WWPN only. Also, the last two bytes in the listed example of 04BF are unique for each node canister. Taking note of these types of patterns can help when you are zoning or troubleshooting SAN issues.

Example scenario Storwize V5000 and IBM DS3400 WWPN diagram

Each port on the Storwize V5000 and IBM DS3400 has a unique and persistent WWPN. This configuration means if an HA in the storage system is replaced, the new HA presents the same WWPNs as the old HA. This configuration means that if you understand the WWPN of a port, you can match it to the storage system and the Fibre Channel port. Figure 6-31 on page 264 shows the relationship between the device WWPNs and the Fibre Channel ports for the Storwize V5000 and the IBM DS3400 that are used in the example scenario, as shown in Figure 6-31 on page 264.

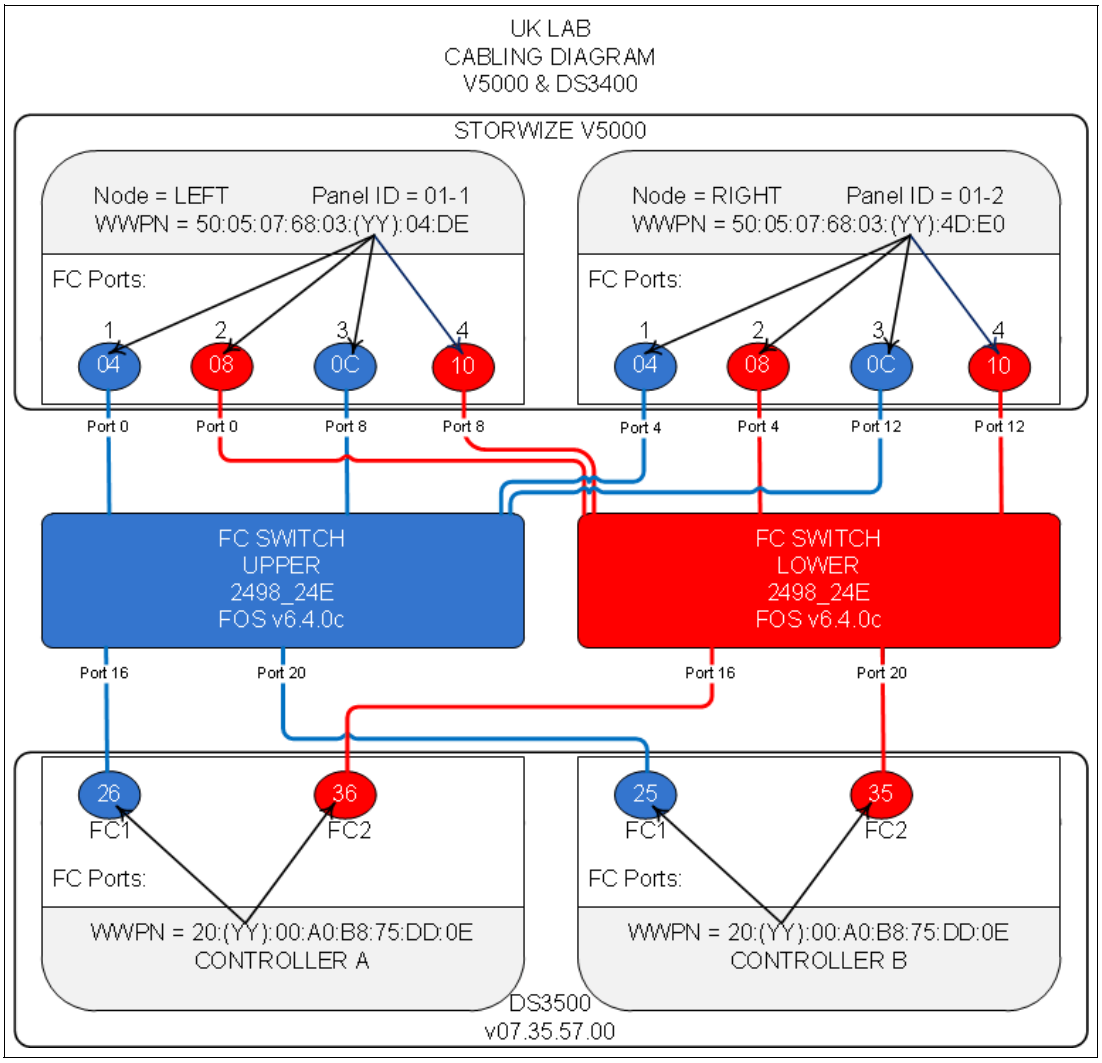

Figure 6-31 Example scenario Storwize V5000 and IBM DS3400 WWPN location diagram

Zoning: Defining aliases on the SAN fabrics

Now that the WWPNs for Storwize V5000, IBM DS3400, and Windows 2008 host are located, you can define the WWPN aliases on the SAN fabrics for the Storwize V5000. Aliases for the DS3400 and Windows 2008 host also can be created, if necessary. Aliases can simplify the zone creation process. Create an alias name for each interface, then add the WWPN.

Aliases can contain the FC Switch Port to which the device is attached, or the attached device's WWPN. In this example scenario, WWPN-based zoning is used instead of port-based zoning. Either method can be used; however, it is best not to intermix the methods and keep the zoning configuration consistent throughout the fabric.

When WWPN-based zoning is used, be mindful when host HBA cards are replaced because occasions can occur when a new HBA card contains new WWPNs and, as a consequence, the previously defined aliases must be modified to match the new card. This situation is not the case for IBM Storage Systems because they use persistent WWPNs, which means that the WWPNs remain unchanged after an HA card is replaced.

Figure 6-31 on page 264 shows the alias definitions:

```
Storwize V5000 ports connected to SAN Fabric A:
      alias= V5000_Canister_Left_Port1 wwpn= 50:05:07:68:03:04:26:BE
      alias= V5000_Canister_Left_Port3 wwpn= 50:05:07:68:03:0C:26:BE
      alias= V5000_Canister_Right_Port1 wwpn= 50:05:07:68:03:04:26:BF
      alias= V5000_Canister_Right_Port3 wwpn= 50:05:07:68:03:0C:26:BF
Storwize V5000 ports connected to SAN Fabric B:
      alias= V5000_Canister_Left_Port2 wwpn= 50:05:07:68:03:08:26:BE
      alias= V5000_Canister_Left_Port4 wwpn= 50:05:07:68:03:10:26:BE
      alias= V5000_Canister_Right_Port2 wwpn= 50:05:07:68:03:08:26:BF
      alias= V5000_Canister_Right_Port4 wwpn= 50:05:07:68:03:10:26:BF
IBM DS3400 ports connected to SAN Fabric A:
      alias= DS3400_CTRLA_FC1 wwpn= 20:26:00:A0:B8:75:DD:0E
      alias= DS3400_CTRLB_FC1 wwpn= 20:27:00:A0:B8:75:DD:0E
IBM DS3400 ports connected to SAN Fabric B:
      alias= DS3400_CTRLA_FC2 wwpn= 20:36:00:A0:B8:75:DD:0E
      alias= DS3400_CTRLB_FC2 wwpn= 20:37:00:A0:B8:75:DD:0E
Window 2008 HBA port connected to SAN Fabric A:
      alias= W2K8_HOST_P2 wwpn= 21:00:00:24:FF:2D:0B:E9
Window 2008 HBA port connected to SAN Fabric B:
      alias= W2K8_HOST_P1 wwpn= 21:00:00:24:FF:2D:0B:E8
```

Zoning: Defining the V5000-to-DS3400 zones on the SAN fabrics

Define the V5000-to-DS3400 zones on the SAN fabrics. The best way to zone DS3400-to-V5000 connections is to ensure that the IBM DS3400 controllers are not in the same zone. The zoning configuration that is provided shows the two zones per fabric that are necessary to ensure that the IBM DS3400 controllers are not in the same zone. Also, all Storwize V5000 node canisters must detect the same ports on IBM DS3400 storage system.

See Figure 6-31 on page 264 and the previously defined SAN aliases for the following zones definitions:

```
FABRIC A
Zone name= ALL_V5000_to_DS3400_CTRLA_FC1:
               DS3400_CTRLA_FC1
               V5000_Canister_Left_Port1
               V5000_Canister_Left_Port3
               V5000_Canister_Right_Port1
               V5000_Canister_Right_Port3
Zone name= ALL_V5000_to_DS3400_CTRLB_FC1:
               DS3400_CTRLB_FC1
               V5000_Canister_Left_Port1
               V5000_Canister_Left_Port3
               V5000_Canister_Right_Port1
               V5000_Canister_Right_Port3
FABRIC B
Zone name= ALL_V5000_to_DS3400_CTRLA_FC2:
               DS3400_CTRLA_FC2
               V5000_Canister_Left_Port2
               V5000_Canister_Left_Port4
               V5000_Canister_Right_Port2
               V5000_Canister_Right_Port4
Zone name= ALL_V5000_to_DS3400_CTRLB_FC2:
               DS3400_CTRLB_FC2
               V5000_Canister_Left_Port2
```

```
                V5000_Canister_Left_Port4
                V5000_Canister_Right_Port2
                V5000_Canister_Right_Port4
```

Zoning: Defining the Host-to-V5000 zones on the SAN fabrics

Define the Host-to-V5000 zones on each of the SAN fabrics. Zone each Host HBA port with one port from each node canister. This configuration provides four paths to the Windows 2008 host. SDDDSM is optimized to use four paths. See Figure 6-22 on page 255 and the previously defined SAN aliases for the following host zone definitions:

```
FABRIC A
Zone name= W2K8_HOST_P2_to_V5000_Port1s:
                W2K8_HOST_P2
                V5000_Canister_Left_Port1
                V5000_Canister_Right_Port1
FABRIC B
Zone name= W2K8_HOST_P1_to_V5000_Port2s:
                W2K8_HOST_P1
                V5000_Canister_Left_Port2
                V5000_Canister_Right_Port2
```

Important: The configuration of an intra-cluster zone (V5000-to-V5000) on each fabric is recommended. Place all Storwize V5000 port aliases from each node canister into the one zone on each of the fabrics. This configuration provides further resilience by providing another communication path between each of the node canisters.

Create a host or host group in the external storage system with the WWPNs for this system.

Important: If you cannot restrict volume access to specific hosts by using the external storage system, all volumes on the system must be migrated.

To complete this step, an IBM DS3400 Host Group is defined for the Storwize V5000, which contains two hosts. Each host is a node canister of the Storwize V5000.

Creating an IBM DS3400 Host Group

To define a new Host Group for the Storwize V5000 by using the DS3400 Storage Manager, click the **Configure** tab and then select **Create Host Group** to open the Create Host Group panel, as shown in Figure 6-32.

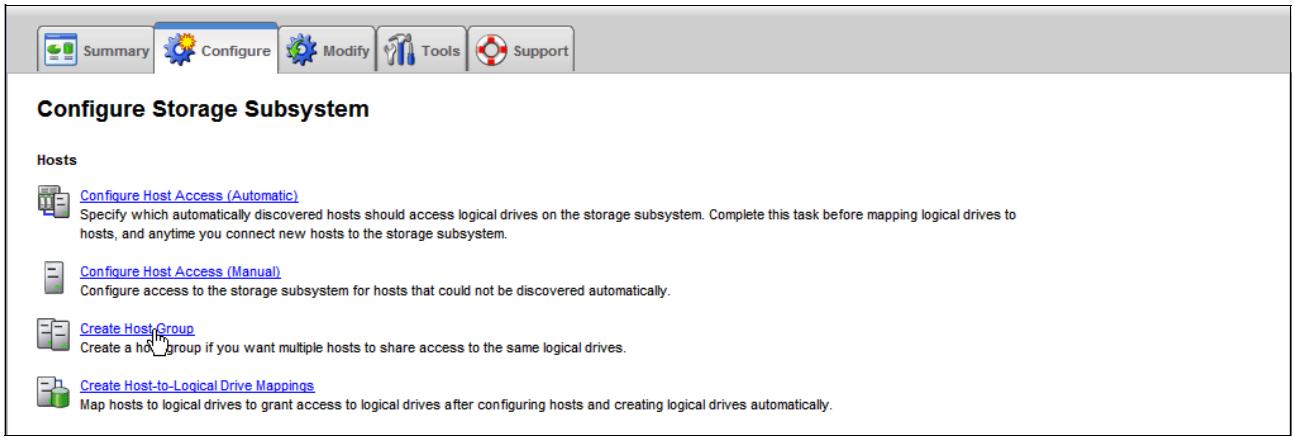

Figure 6-32 Configure Storage Subsystem

By using the IBM DS3400 Storage Manager, create a Host Group that is named `Storwize_V5000`. Figure 6-33 shows the IBM DS3400 Create Host Group panel.

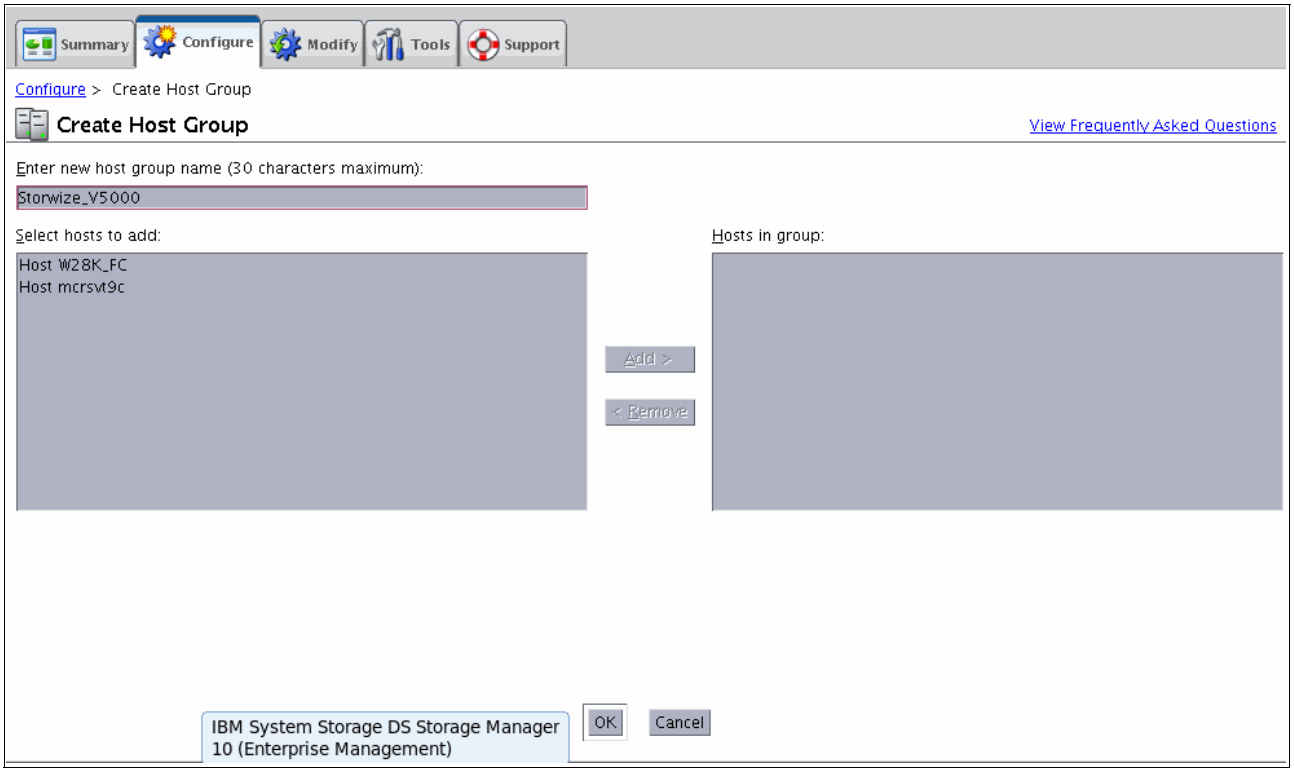

Figure 6-33 IBM DS3400 Create Host Group panel

Creating IBM DS3400 hosts

By using the IBM DS3400 Storage Manager, create a Host for each node canister of the Storwize V5000. To define a new Host by using the DS3400 Storage Manager, click the **Configure** tab and then select **Configure Host-Access (Manual)** to open the configure host access panel, as shown in Figure 6-34.

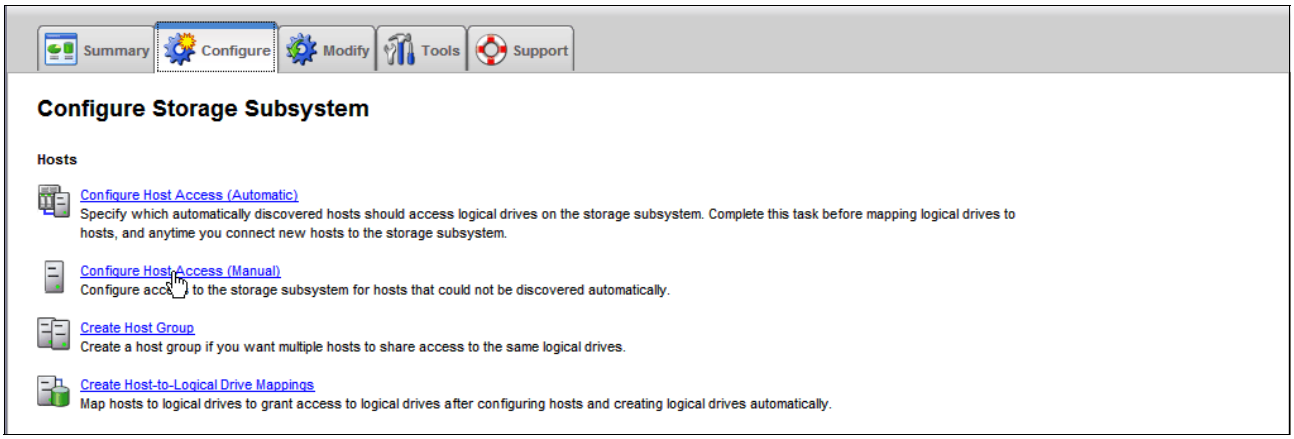

Figure 6-34 Selecting Configure Host-Access (Manual) option

Provide a name for the host and ensure that the selected host type is `IBM TS SAN VCE`. The name of the host should be easily recognizable, such as, `Storwize_V5000_Canister_Left` and `Storwize_V5000_Canister_Right`. Click **Next** to continue. Figure 6-35 shows the IBM DS3400 storage manager Configure Host Access (Manual) panel.

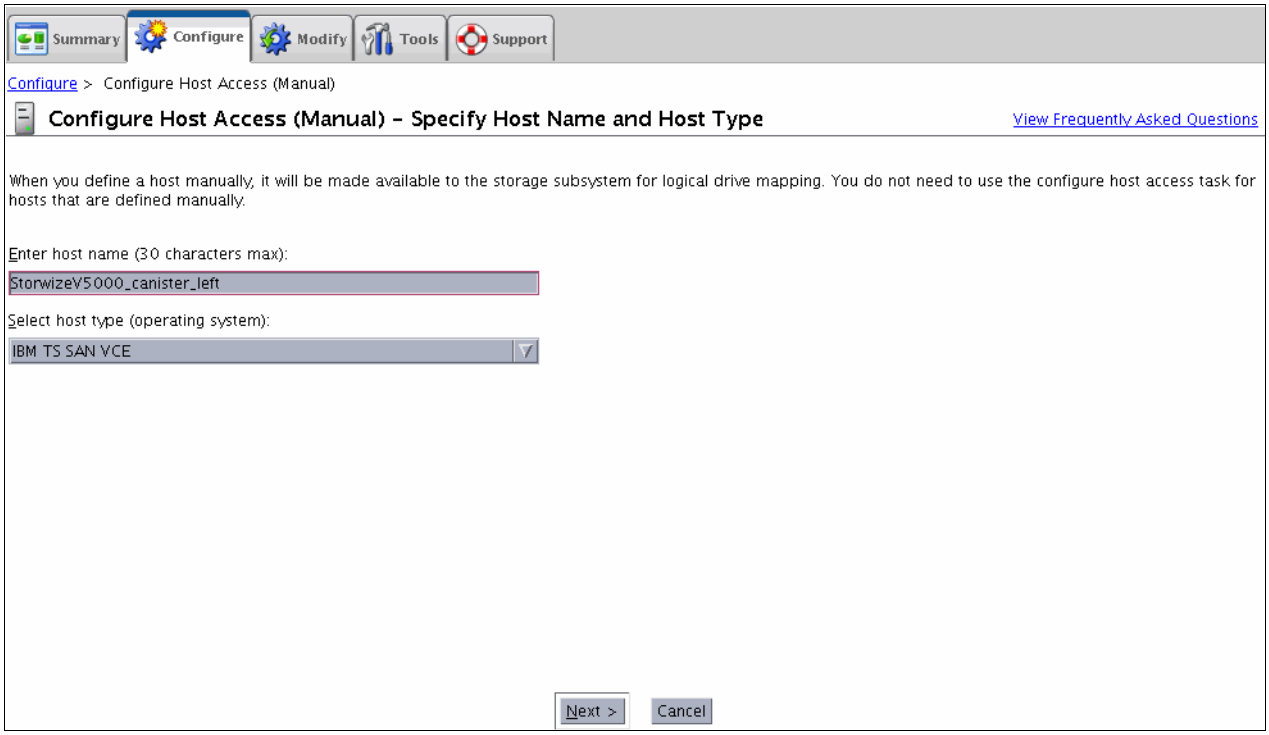

Figure 6-35 IBM DS3400 storage manager Configure tab: Configure host

The node canister's WWPNs are automatically discovered and must be matched to the canister's host definition. Select each of the four WWPNs for the node canister and then click **Add >**. The selected WWPN moves to the right side of the panel, as shown in Figure 6-36.

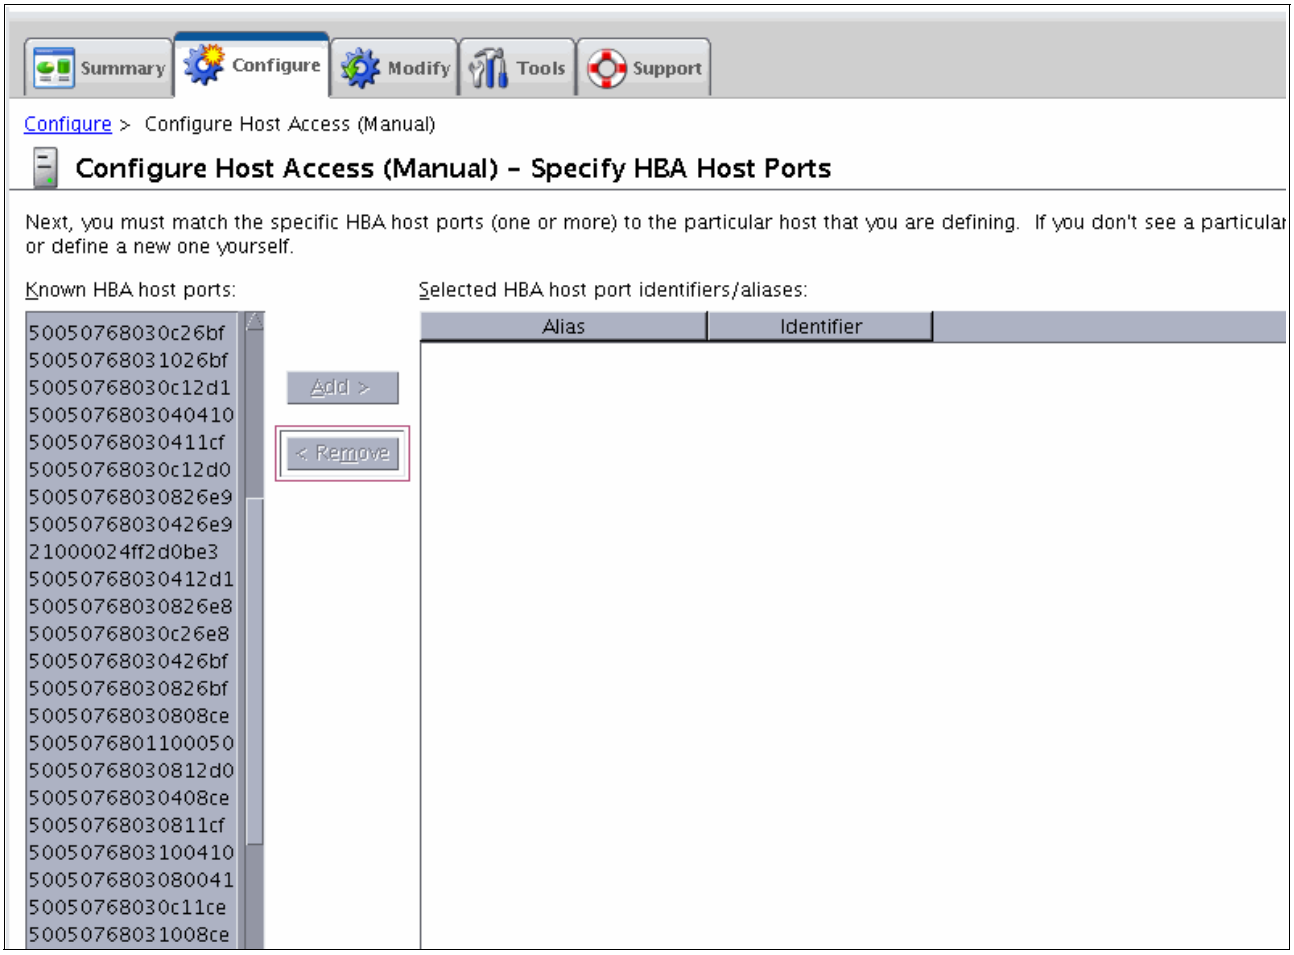

Figure 6-36 IBM DS3400 Specify HBA Host Ports panel

Click **Edit** to open the Edit HBA Host Port panel, as shown in Figure 6-37.

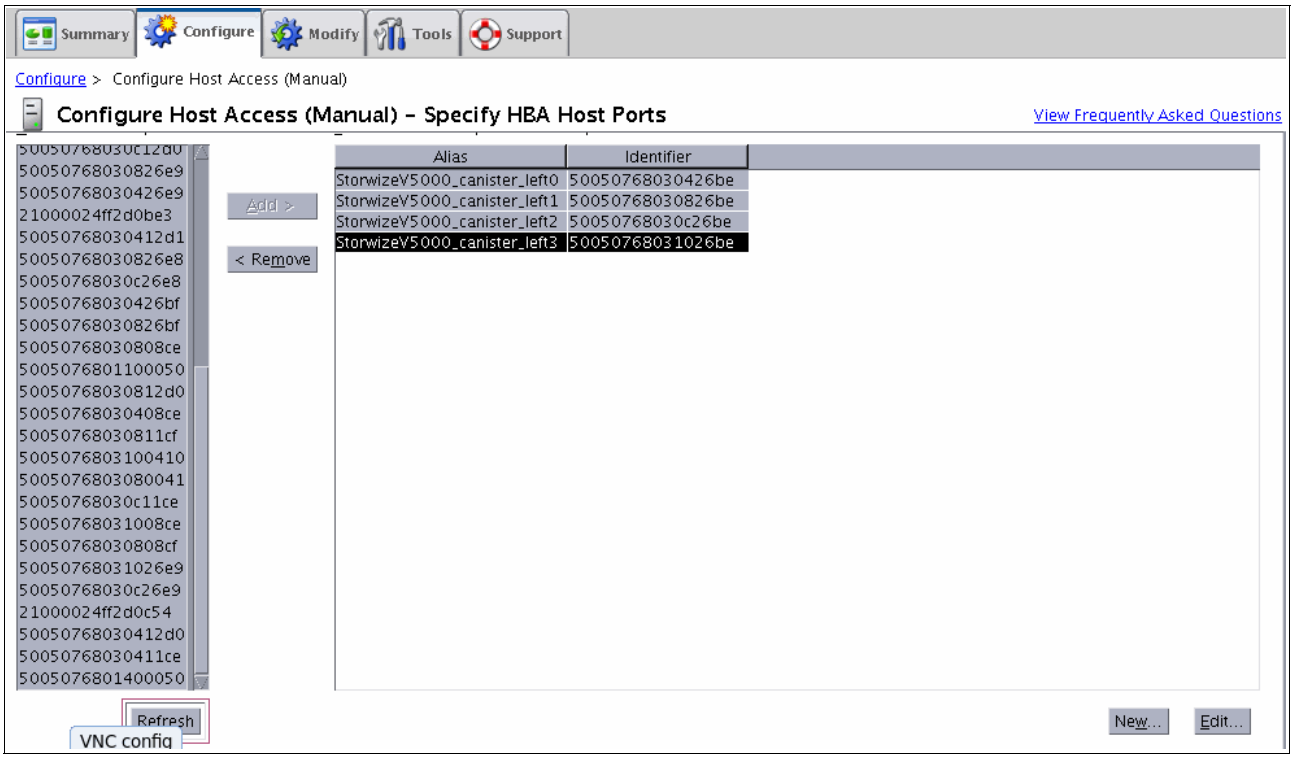

Figure 6-37 IBM DS3400 storage manager specifying HBA host ports: Edit alias

Enter a meaningful alias for each of the WWPNs, such as, `V5000_Canister_Left_P1`, as shown in Figure 6-38. To ensure that the information was added correctly, see Figure 6-31 on page 264 and the previously defined SAN fabric aliases.

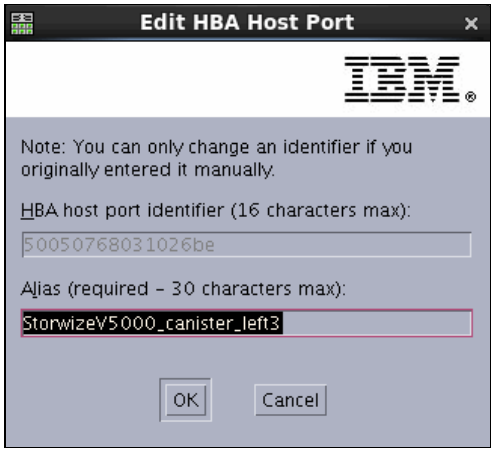

Figure 6-38 IBM DS3400 Edit HBA Host Port panel

After the four ports for the node canister with the meaningful aliases are added to the node canister host definition, click **Next** to continue. Figure 6-39 shows the node canister WWPNs that are added to the host definition on the IBM DS3400 Specify HBA Host Ports panel.

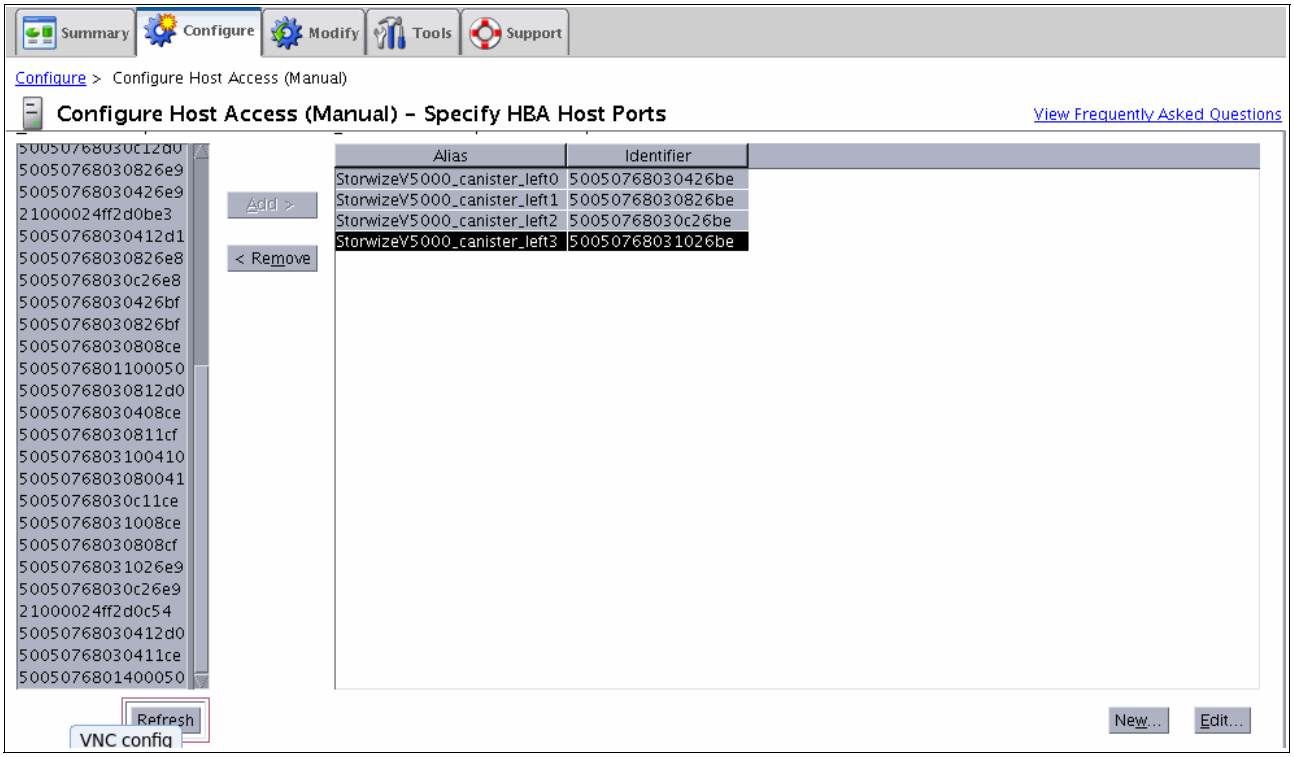

Figure 6-39 IBM DS3400 Specify HBA Host Ports panel

Select **Yes** to allow the host to share access with other hosts for the same logical drives. Ensure that the existing Host Group is selected and shows the previously defined Storwize_V5000 host group. Click **Next** to continue. Figure 6-40 shows the IBM DS3400 Specify Host Group panel.

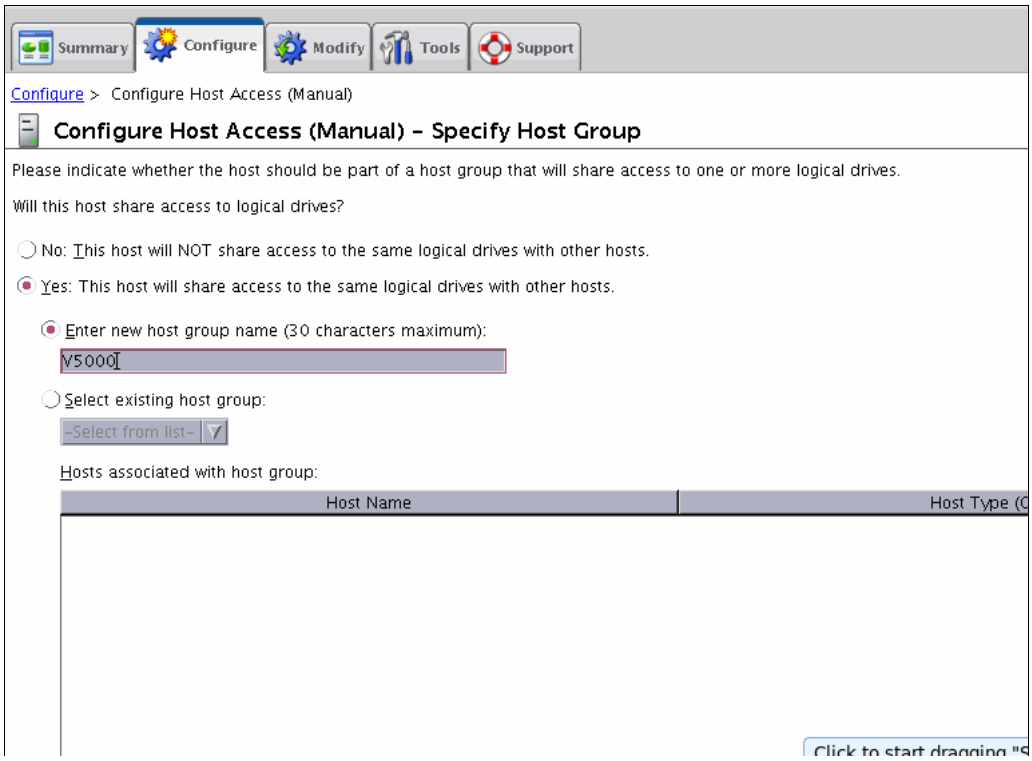

Figure 6-40 IBM DS3400 Specify Host Group panel

A summary panel of the defined host and its associated host group is displayed. Cross-check and confirm the host definition summary, and then click **Finish**, as shown in Figure 6-41.

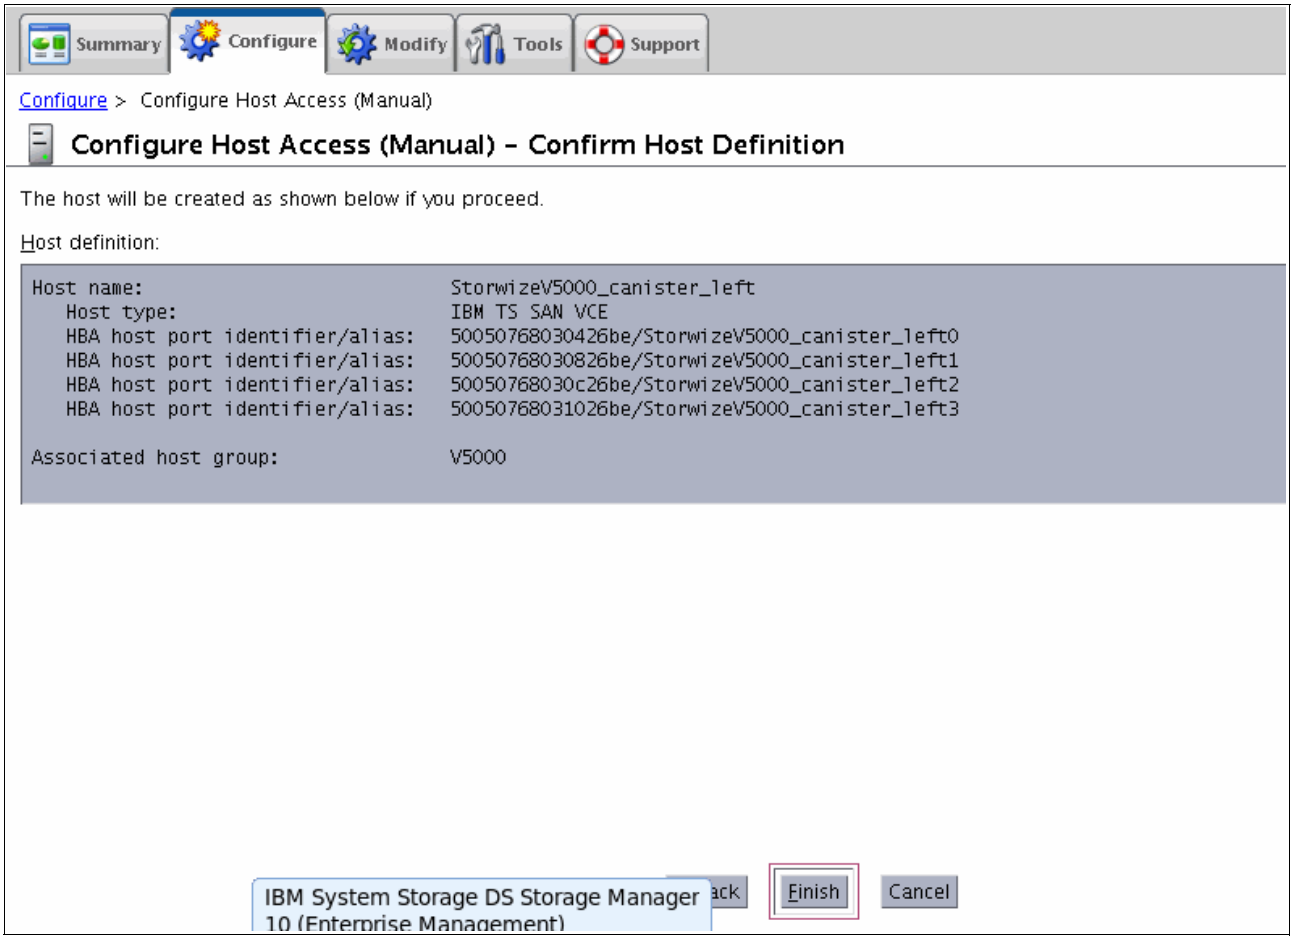

Figure 6-41 IBM DS3400 Confirm Host Definition panel

A host definition must be created for the other node canister. The host definition also is associated to the Host Group Storwize_V5000. To configure the other node canister, complete the steps that are described in “Creating IBM DS3400 hosts” on page 267.

The node canister Host definitions are logically contained in the Storwize_V5000 Host Group. After both node canister hosts are created, confirm the host group configuration by reviewing the IBM DS3400 host topology tree. To access the host topology tree, use the IBM DS3400 storage manager, click the **Modify** tab and select **Edit Host Topology**, as shown in Figure 6-42.

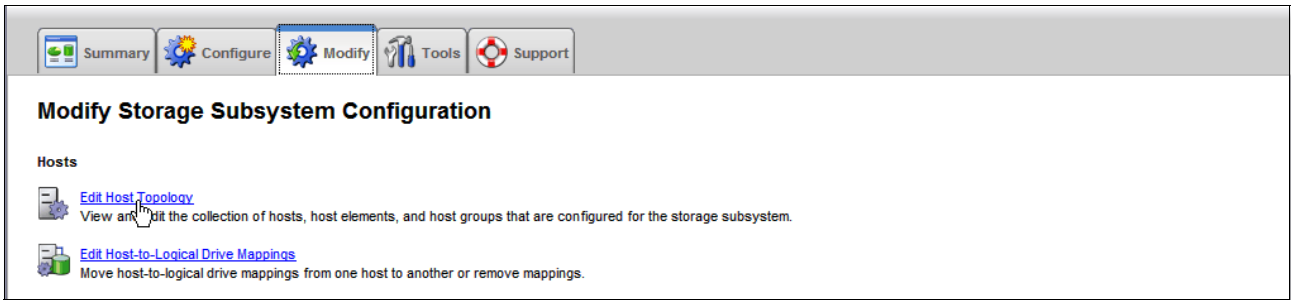

Figure 6-42 Selecting the Edit Host Topology option

Figure 6-43 shows the host topology of the defined Storwize_V5000 Host Group with both of the created node canister hosts, as seen through the DS3400 Storage Manager software.

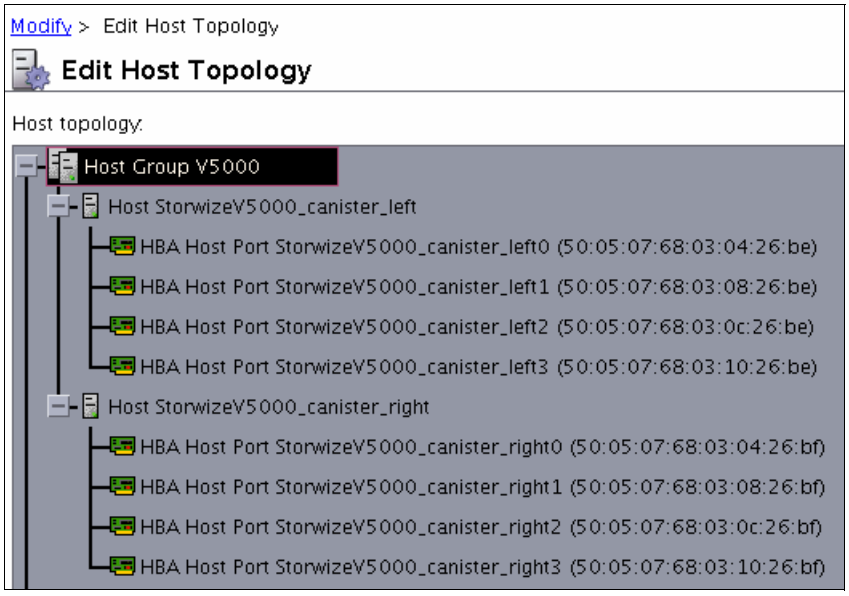

Figure 6-43 IBM DS3400 host group definition for the Storwize V5000

Configure the storage system for use with this system.

See the IBM Storwize V5000 Version 6.4.1 Information Center for DS3400 configuration recommendations at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Fsvc_configdiskcontrollersovr_22n9uf.html

Now that the environment is prepared, return to step 2 of the Storage Migration wizard in the Storwize V5000 GUI and click **Next** to continue, as shown in Figure 6-44.

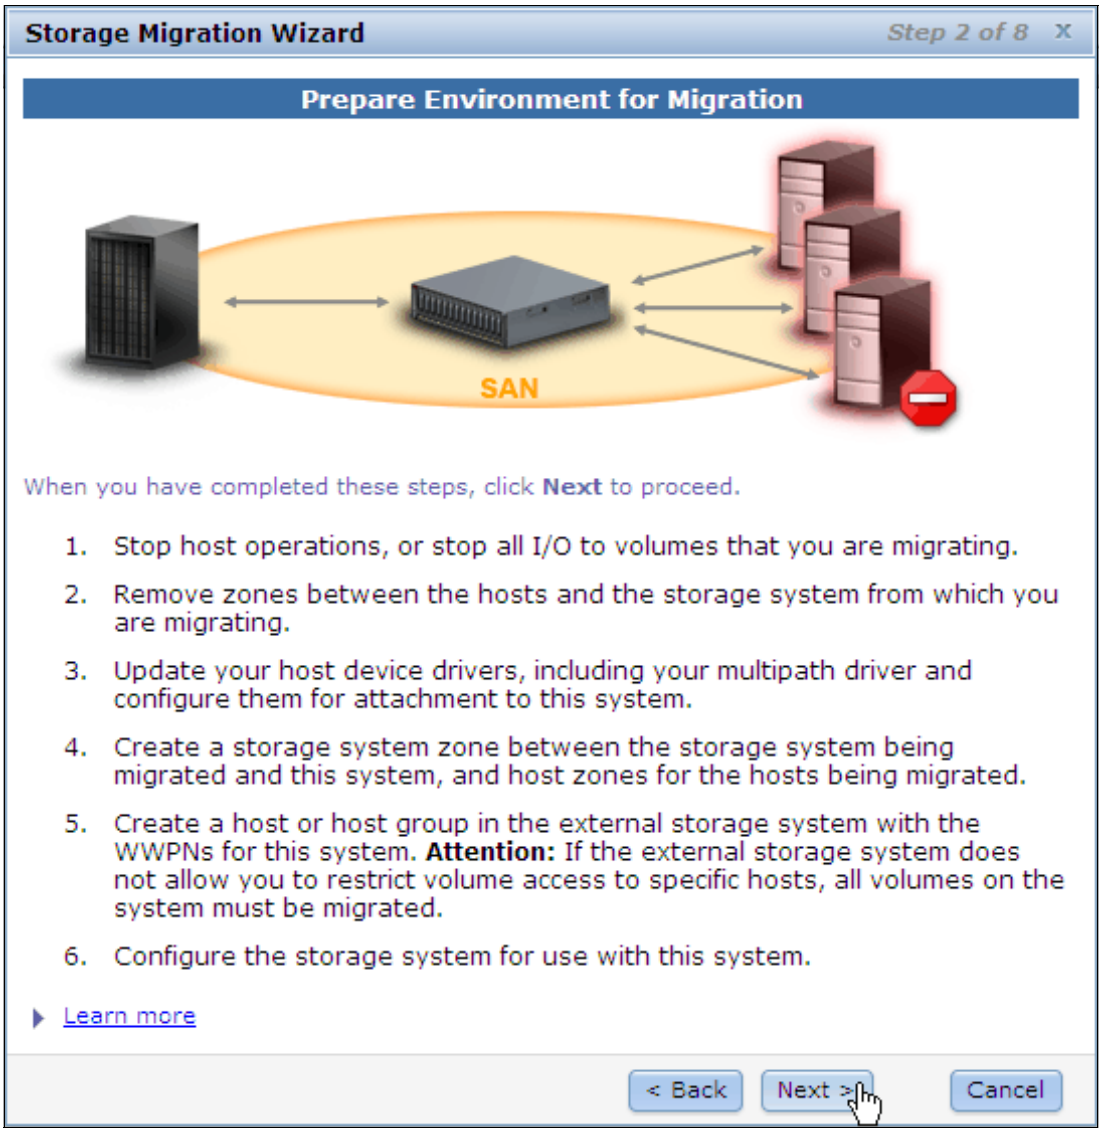

Figure 6-44 Step 2 of the Storage Migration wizard

Follow step 3 of the Storage Migration wizard and map the storage, as shown in Figure 6-45.

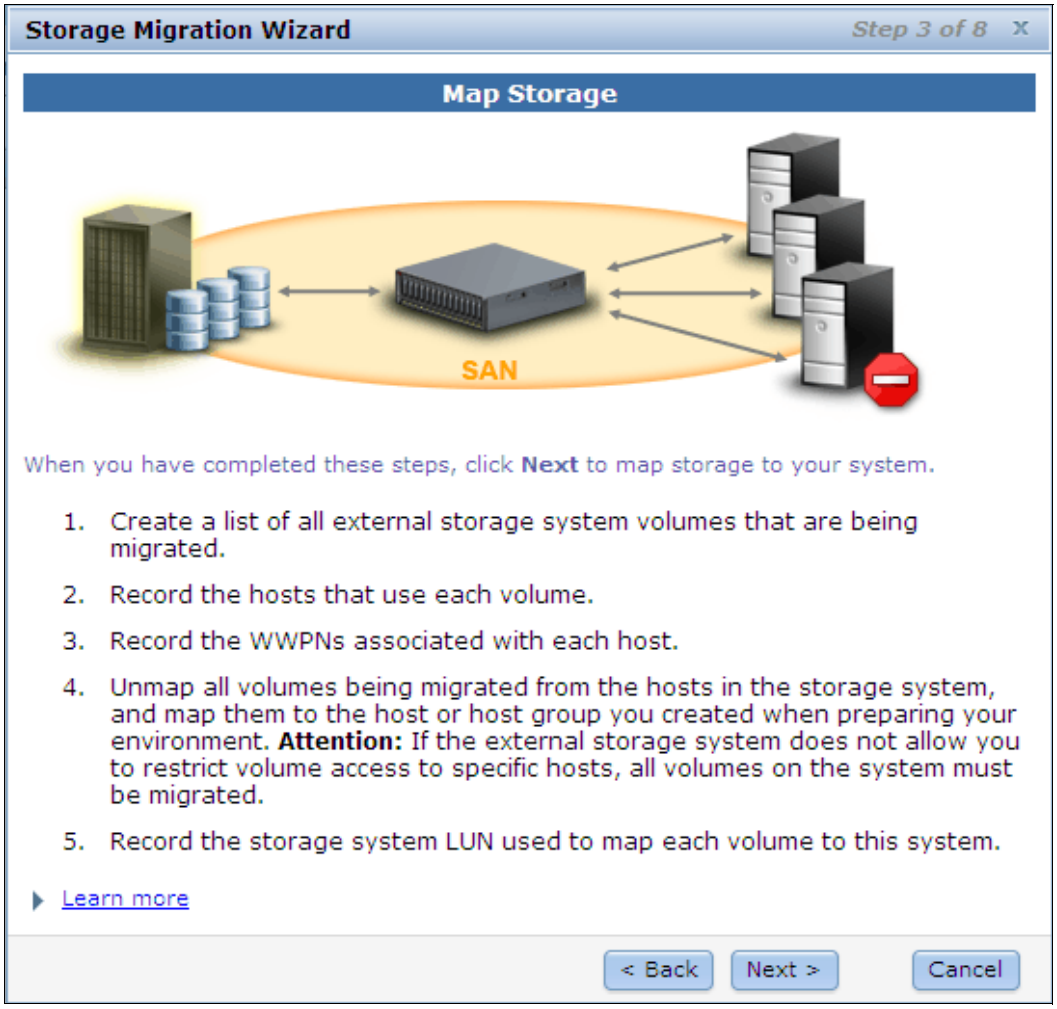

Figure 6-45 Step 3 of the Storage Migration wizard

Create a list of all external storage system volumes that are being migrated. Record the hosts that use each volume.

Table 6-3 shows a list of the IBM DS3400 LUs that were migrated and the host that uses them.

Table 6-3 List of the IBM DS3400 logical units that are migrated and hosted

LU name	Controller	Array	SCSI ID	Host name	Capacity
Migration_1	DS3400	Array 1	0	W2K8_FC	50 GB
Migration_2	DS3400	Array 3	1	W2K8_FC	100 GB

Record the WWPNs that are associated with each host.

The WWPNs that are associated to the host can be seen in Table 6-4. It is recommended that you record the HBA firmware, HBA device driver version, adapter information, operating system, and V5000 multi-path software version, if possible.

Table 6-4 WWPNs that are associated to the host

Host name	Adapter / Slot / Port	WWPNs	HBA F/W	HBA Device Driver	Operating System	V5000 Multipath Software
W2K8_FC	QLE2562 / 2 / 1	21000024FF2D0BE8	2.10	9.1.9.25	W2K8 R2 SP1	SDDDSM 2.4.3.1-2
	QLE2562 / 2 / 2	21000024FF2D0BE9				

Unmap all volumes that are migrated from the hosts in the storage system and map them to the host or host group that you created when your environment was prepared.

Important: If you cannot restrict volume access to specific hosts by using the external storage system, all volumes on the system must be migrated.

Change IBM DS3400 LU mappings

The LUs that are migrated are presented from the IBM DS3400 to the Windows 2008 host because of a mapping definition that was configured on the IBM DS3400. To modify the mapping definition so that the LUs are accessible only by the Storwize V5000 Host Group, a modify mapping operation must be completed. To modify the mapping on the IBM DS3400, click the **Modify** tab and select **Edit Host-to-Logical Drive Mappings**, as shown in Figure 6-46.

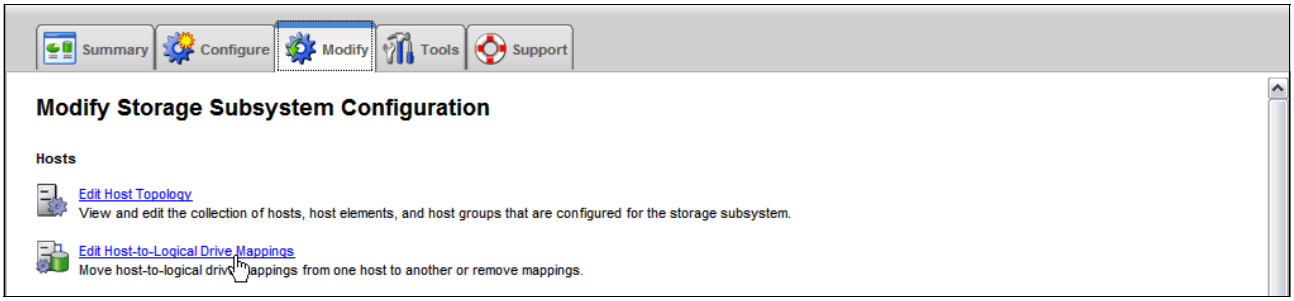

Figure 6-46 IBM DS3400 storage manager Modify tab

The IBM DS3400 logical drives are accessible by the Windows 2008 host. Figure 6-47 shows the IBM DS3400 logical drives mapping information before the change.

Figure 6-47 IBM DS3400 Logical drives mapping information before changes

To modify the mapping definition so that the LUs are accessible only by the Storwize V5000 Host Group, select **Change...** to open the Change Mapping panel and modify the mapping. This step ensures that the LU cannot be accessed from the Windows 2008 Host, as shown in Figure 6-48 on page 278.

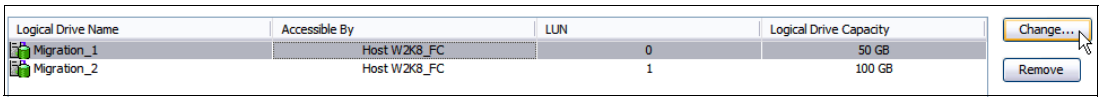

Figure 6-48 IBM DS3400 modify mapping panel: Change mapping

Select **Host Group Storewize_V7000** in the menu and ensure that the Logical Unit Number (LUN) remains the same. Record the LUN for later reference. Figure 6-49 shows the IBM DS3400 Change Mapping panel.

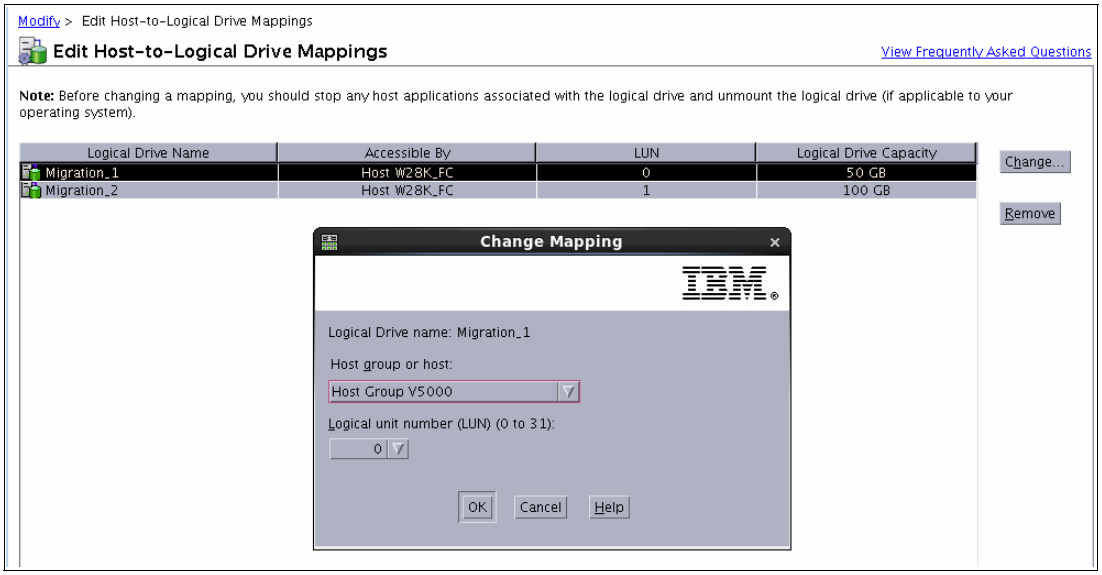

Figure 6-49 IBM DS3400 Change Mapping panel

Confirm the mapping change by selecting **Yes.** Figure 6-50 shows the Change Mapping confirmation panel.

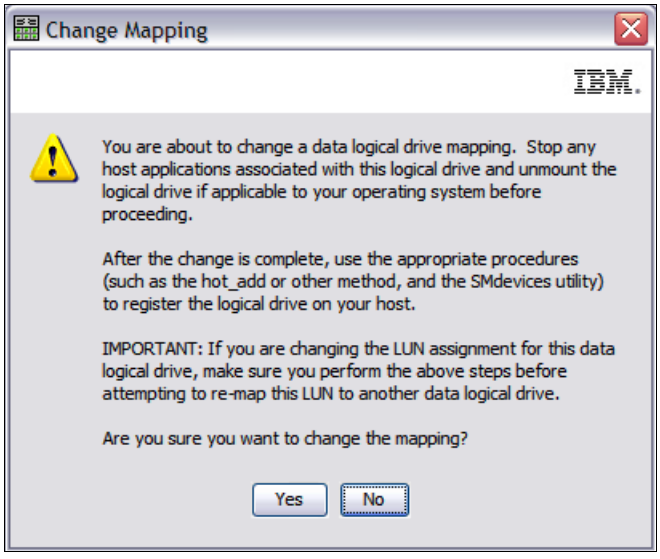

Figure 6-50 Change Mapping confirmation panel

Repeat the steps that are described in “Change IBM DS3400 LU mappings” on page 277 for each of the LUs that are migrated. Confirm that the Accessible By column now reflects the mapping changes. Figure 6-51 shows that both logical drives are now accessible by Host Group Storwize_V5000.

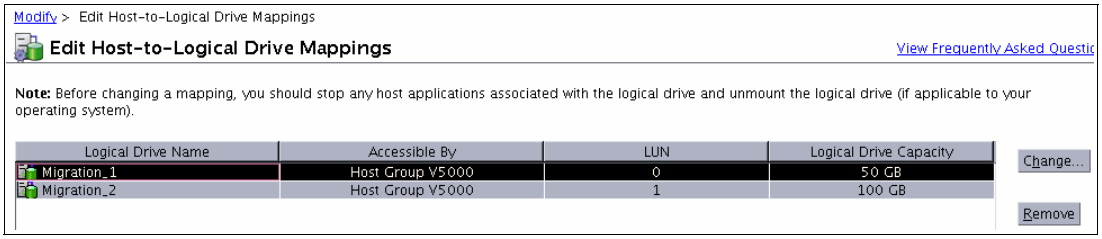

Figure 6-51 Edit Host-to-Logical Drive Mappings panel

Record the storage system LUN that is used to map each volume to this system.

The LUNs that are used to map the logical drives remained unchanged and can be found in Table 6-3 on page 276. Now that step 3 of the storage migration wizard is complete, click **Next** to show the Detect MDisks running task, as shown in Figure 6-52.

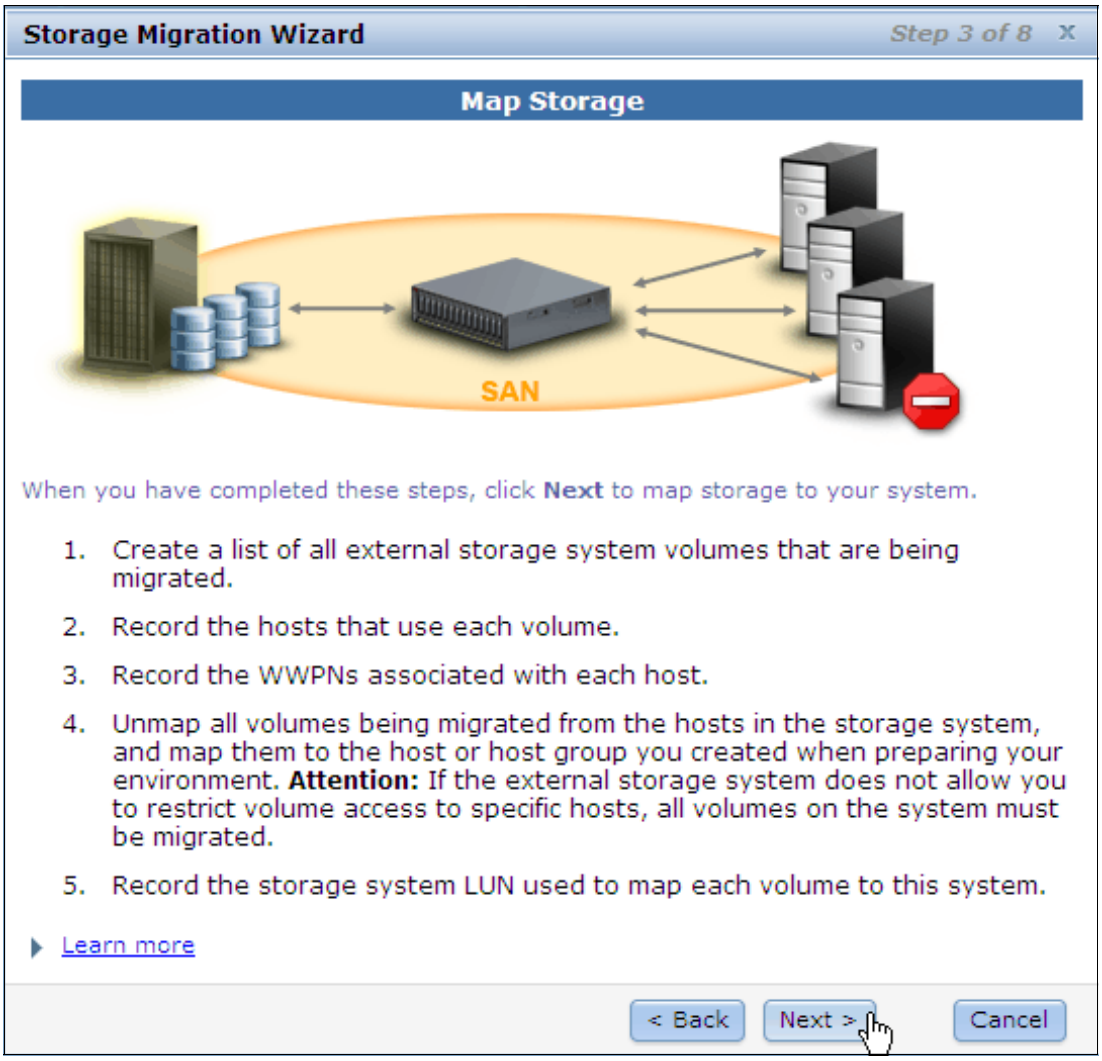

Figure 6-52 Step 3 of the Storage Migration wizard

After the Discover Devices running task is complete, select **Close** to show the step 4 of the wizard, as shown in Figure 6-53.

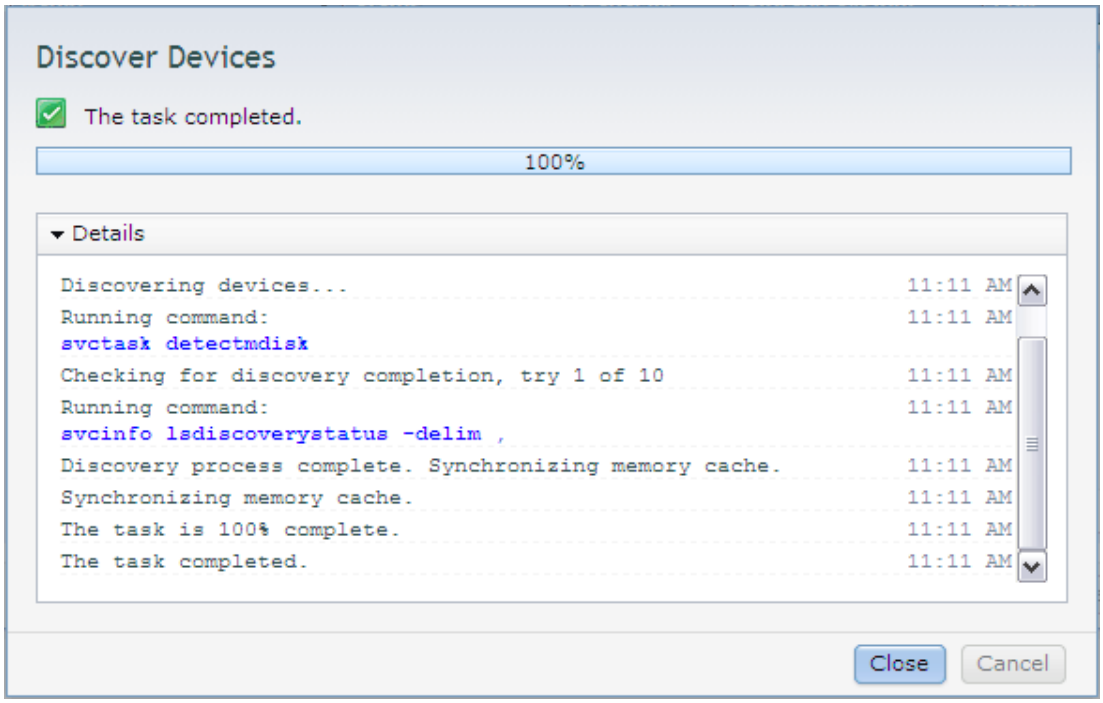

Figure 6-53 Discover Devices panel

Follow step 4 of the Storage Migration wizard, as shown in Figure 6-54. The MDisk name is allocated depending on the order of device discovery; mdisk0 in this case is LUN 1 and mdisk1 is LUN 0. There is an opportunity to change the MDisk names to something more meaningful to the user in later steps.

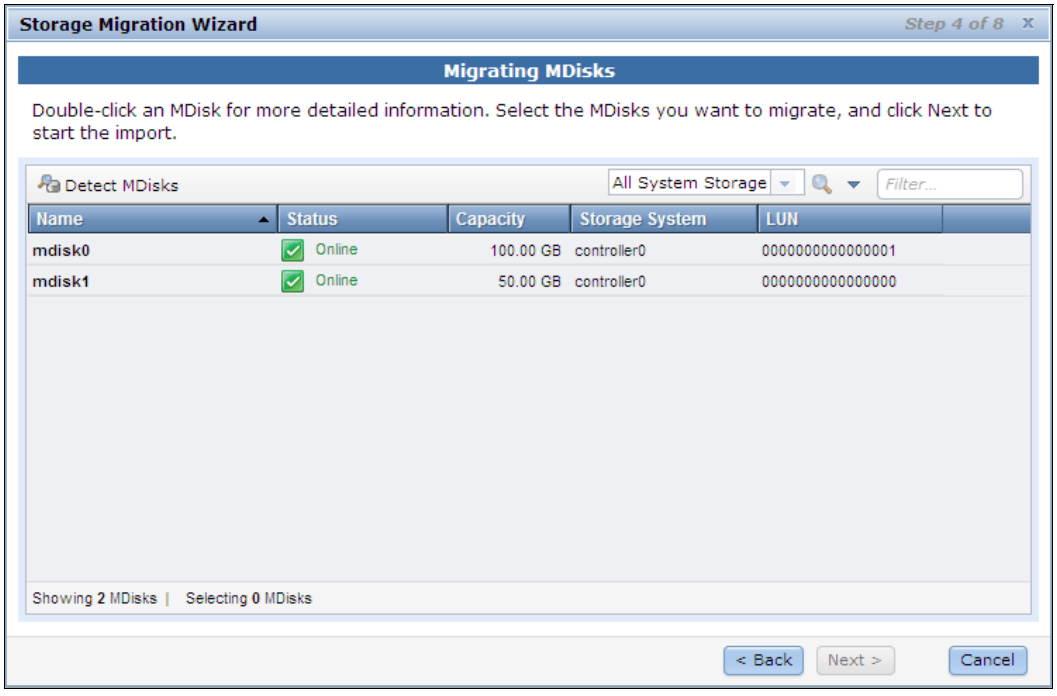

Figure 6-54 Step 4 of the Storage Migration wizard

Select the discovered MDisks and click **Next** to open the Import MDisks running task panel, as shown in Figure 6-55.

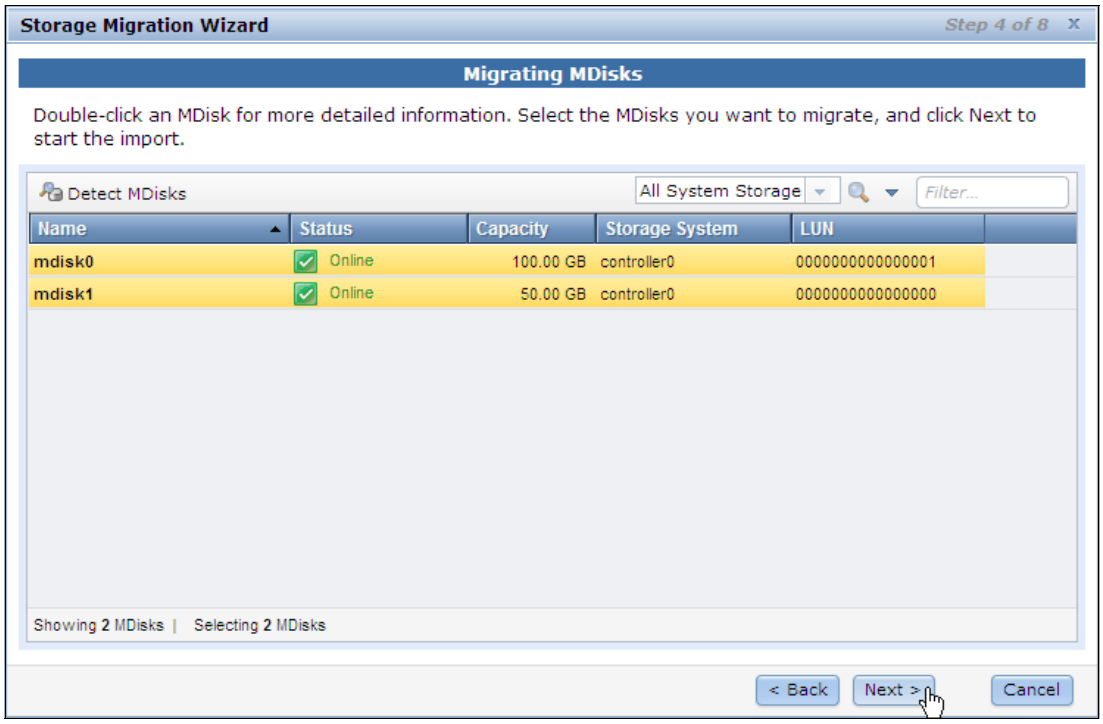

Figure 6-55 Selecting MDisk to migrate

After the Import MDisks running task is complete, select **Close** to open step 5 of the storage migration wizard, as shown in Figure 6-56.

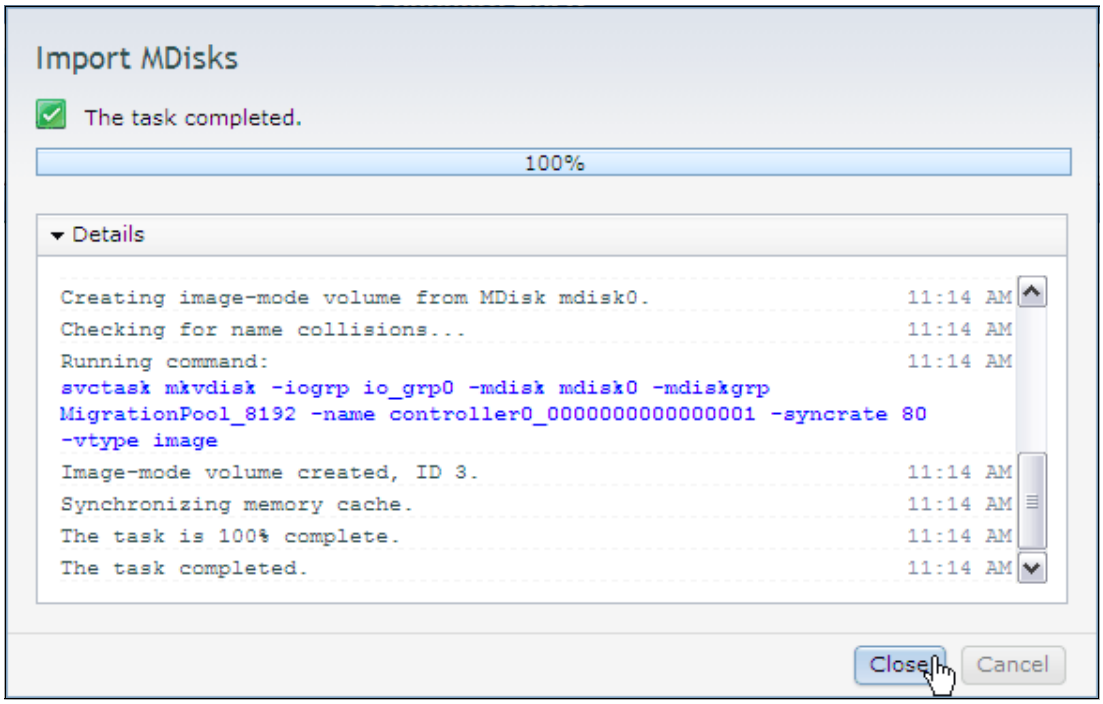

Figure 6-56 Import MDisks panel

Follow step 5 of the storage migration wizard, as shown in Figure 6-57. The Windows 2008 host is not yet defined in the Storwize V5000. Select **New Host** to open the Create Host panel.

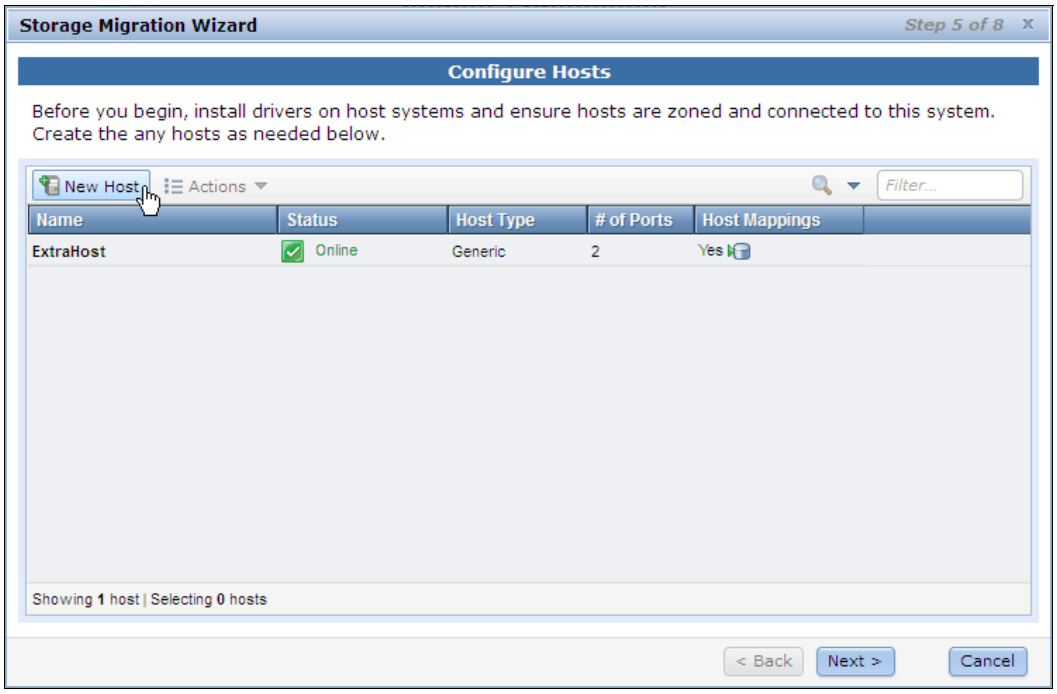

Figure 6-57 Selecting Create Host option

Enter a host name and select the WWPNs that were recorded earlier from the Fibre Channel ports menu. Select **Add Port to List** for each WWPN. Figure 6-58 on page 283 shows the Create Host panel.

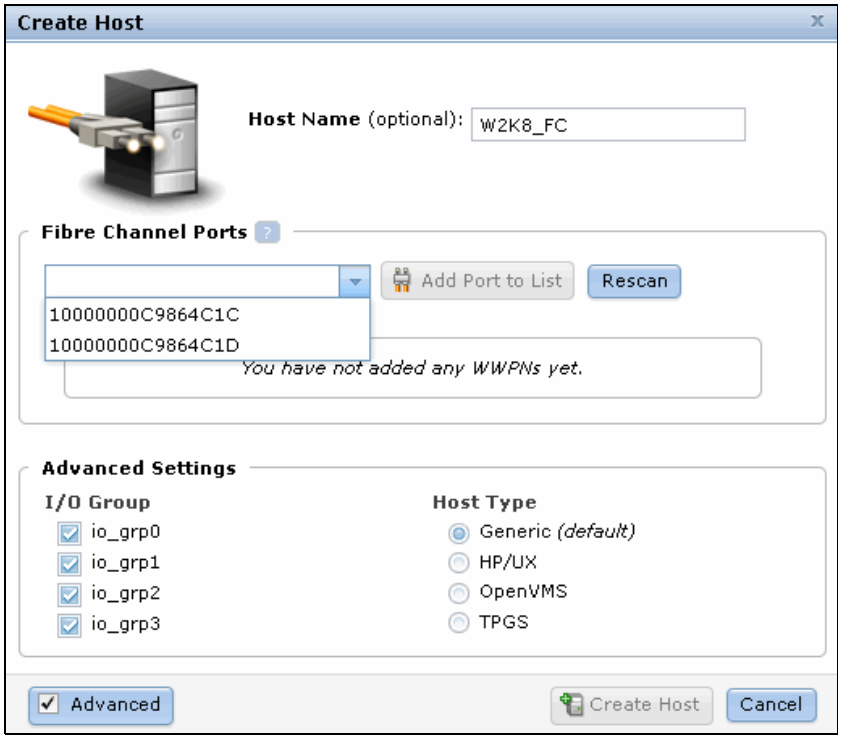

Figure 6-58 Select WWPNs

After all of the port definitions are added, click **Create Host** to open the Create Host running task. Figure 6-59 shows the Create Host panel with the required port definitions listed.

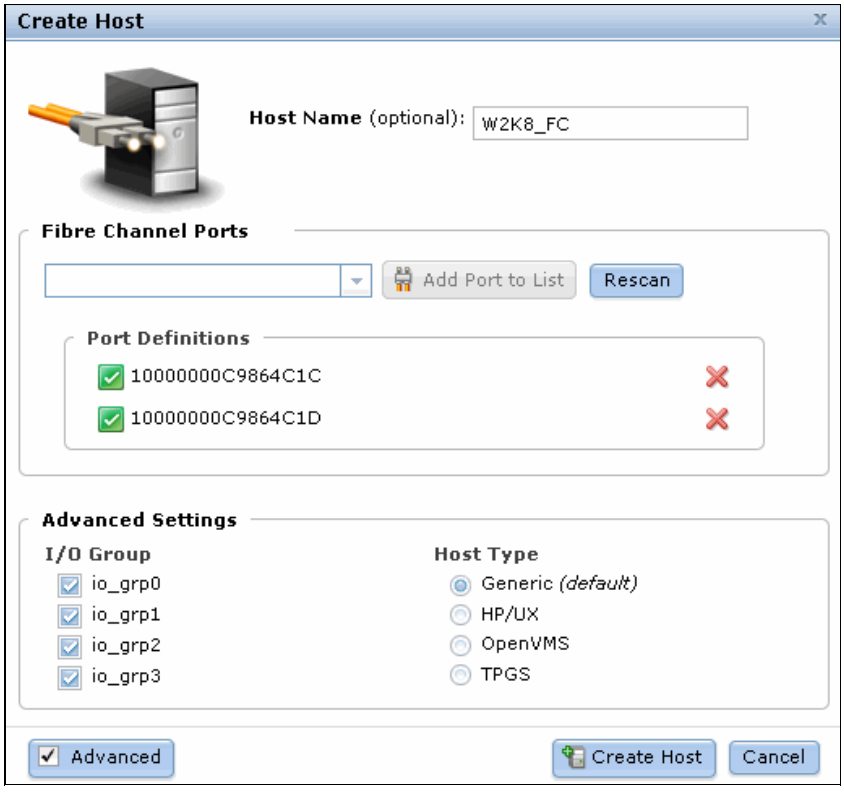

Figure 6-59 Required port definitions listed

After the Create Host running task is complete, select **Close** to reopen step 5 of the wizard, as shown in Figure 6-60.

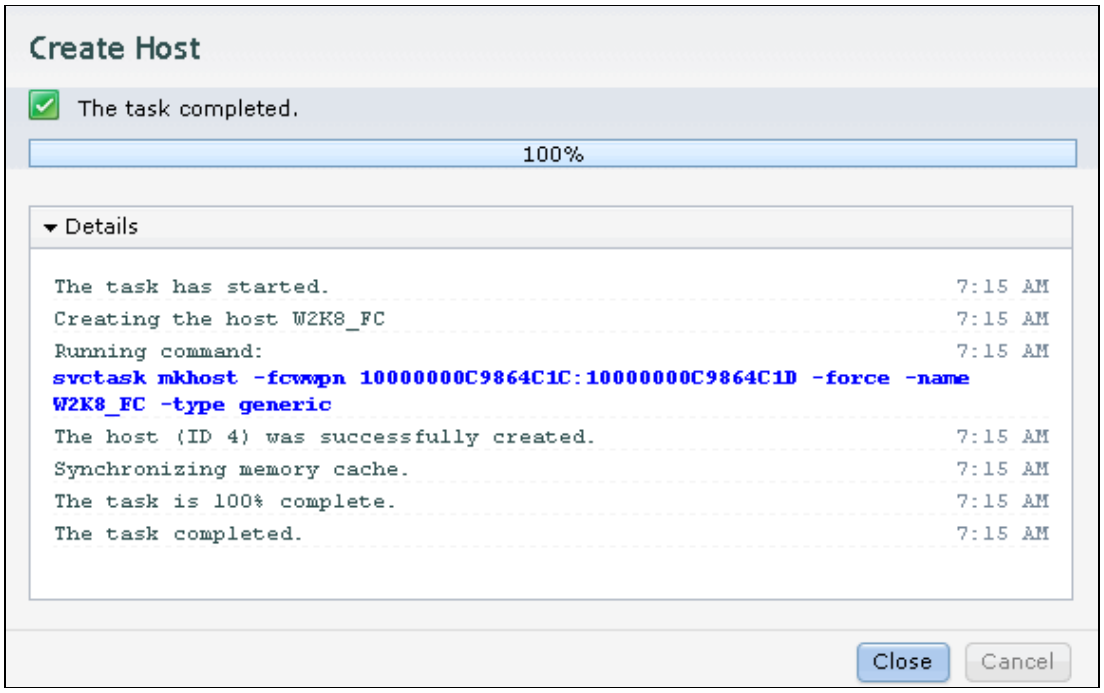

Figure 6-60 Create Host running task panel

From step 5 of the wizard, select the host that was configured and click **Next** to open step 6 of the wizard, as shown in Figure 6-61.

Figure 6-61 Select host

Important: It is not mandatory to select the hosts now. The actual selection of the hosts occurs in the next step. However, cross-check the hosts that have data that must be migrated by highlighting them in the list before you click **Next**.

Follow step 6 of the wizard. Rename the MDisks to reflect something more meaningful. Right-click the MDisk and select **Rename** to open the Rename Volume panel, as shown in Figure 6-62.

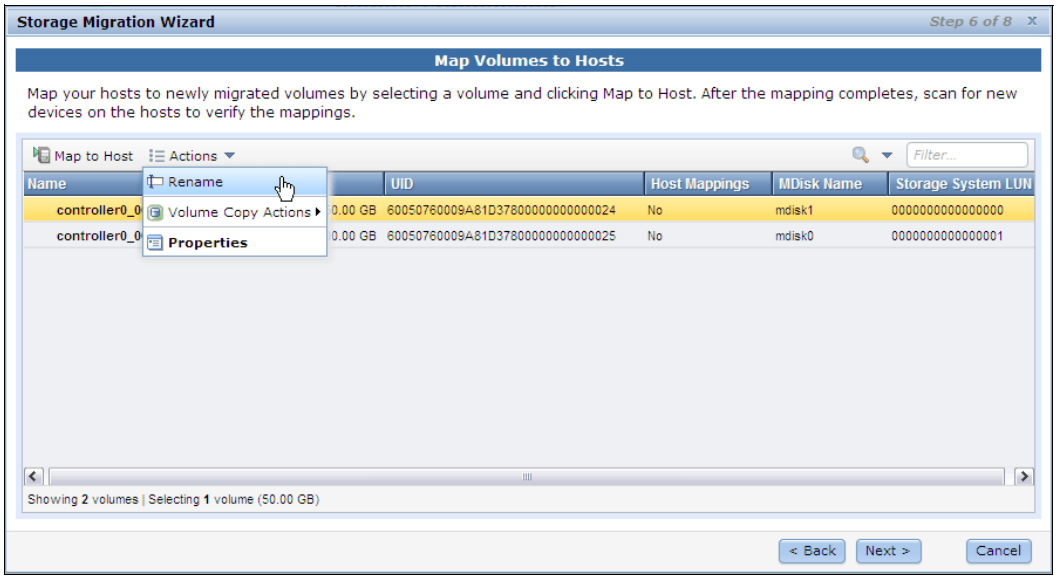

Figure 6-62 Step 6 of the Storage Migration wizard

The name that is automatically given to the image mode volume includes the controller and the LUN information. Use this information to determine an appropriate name for the volume. After the new name is entered, click **Rename** from the Rename Volume panel to start the rename running task. Rename both volumes. Figure 6-63 shows the Rename Volume panel.

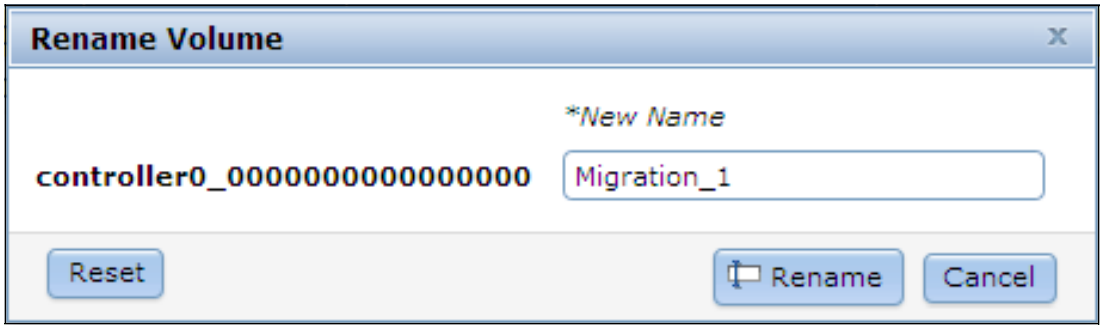

Figure 6-63 Rename volume panel

After the final rename running task is complete, click **Close** to reopen step 6 of the wizard, as shown in Figure 6-64.

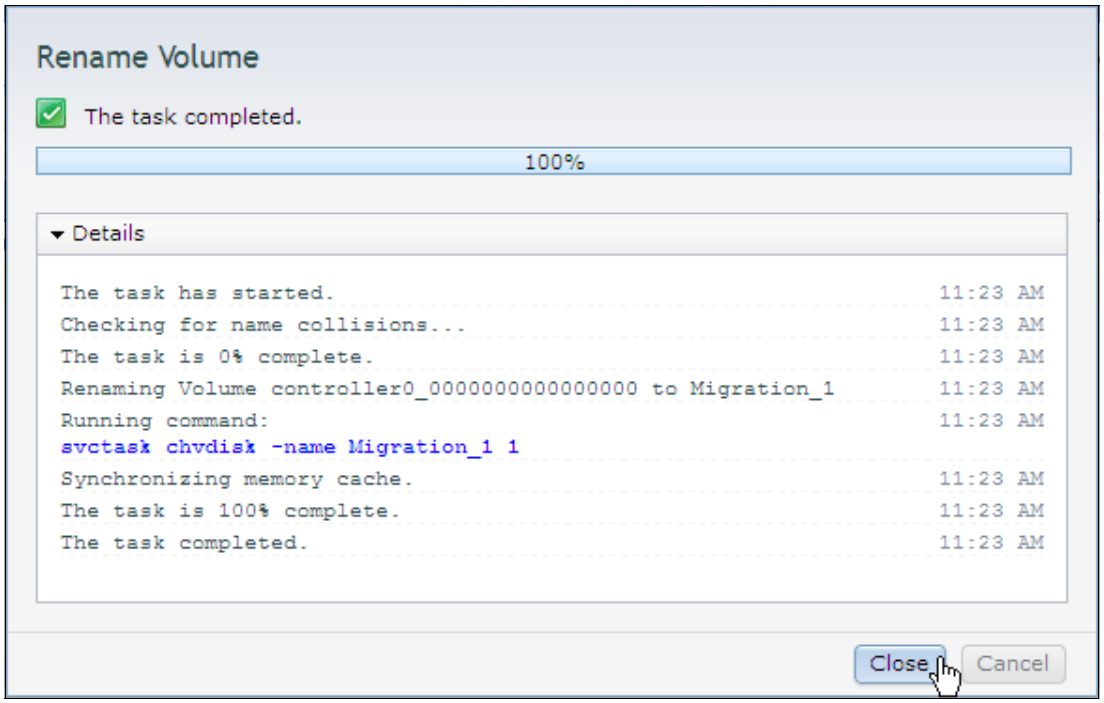

Figure 6-64 Rename Volume running task

From step 6 of the wizard, highlight the two MDisks and select **Map to Host** to open the Modify Host Mappings panel, as shown in Figure 6-65.

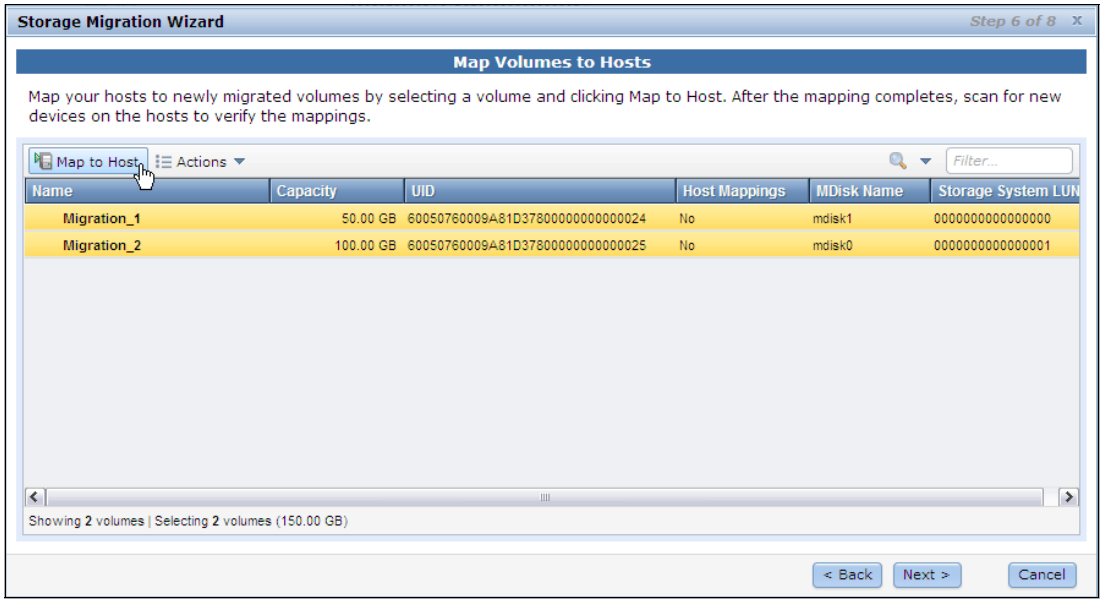

Figure 6-65 Renamed MDisks highlighted for mapping

Select the host from the menu on the Modify Host Mappings panel, as shown in Figure 6-66. The rest of the Modify Host Mappings panel opens.

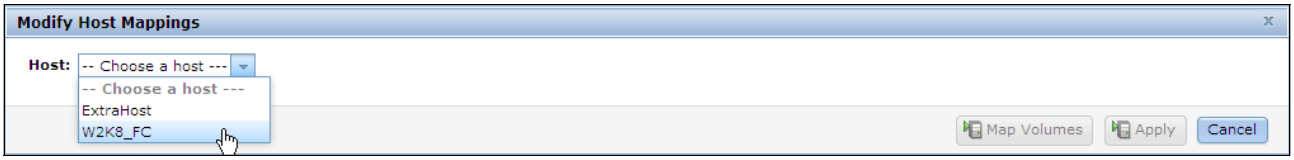

Figure 6-66 Modify Host Mappings panel

The MDisks that were highlighted in step 6 of the wizard are highlighted in yellow in the Modify Host Mappings panel. The yellow highlighting means that the volumes are not yet mapped to the host. Now is the time to edit the SCSI ID, if required. (In this case, it is not necessary.) Click **Map Volumes** to open the Modify Mappings running task, as shown in Figure 6-67.

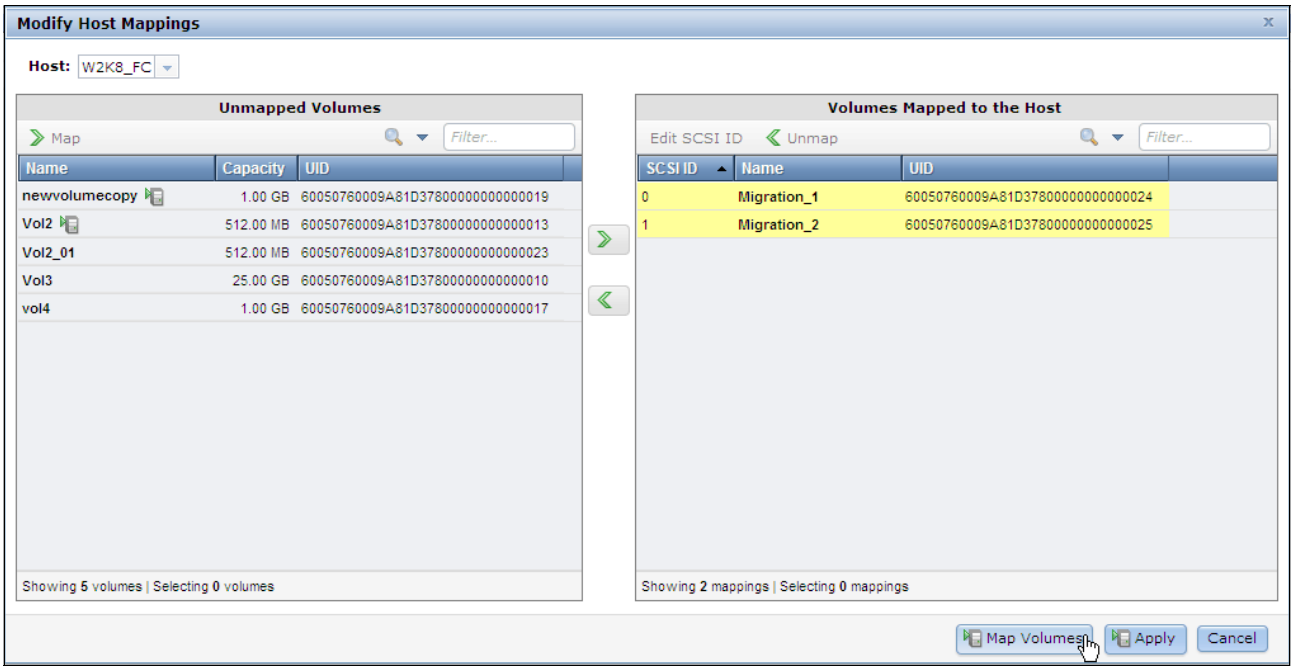

Figure 6-67 Modify Host Mappings panel

After the Modify Mappings running task is complete, select **Close** to reopen step 6 of the wizard, as shown in Figure 6-68.

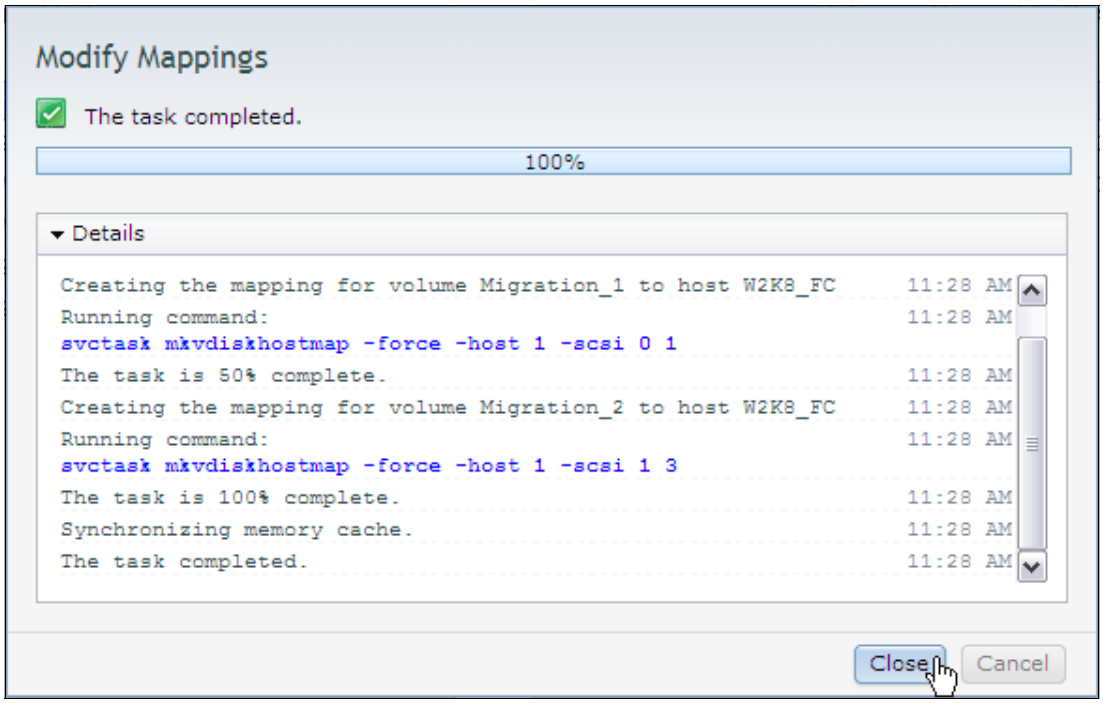

Figure 6-68 Modify Mappings running task

Confirm that the MDisks are now mapped by ensuring the Host Mappings column has a `Yes` listed for each MDisk, as shown in Figure 6-69.

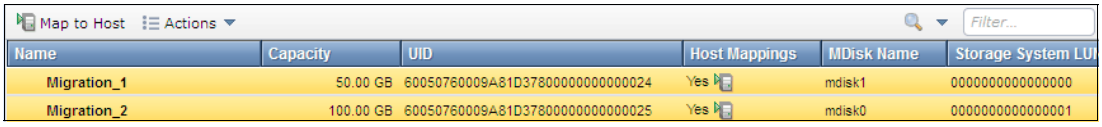

Figure 6-69 MDisks mapped

Verifying that migrated disk device type is now 2145 on the host

The migrated volumes are now mapped to the Storwize V5000 host definition. The migrated disks properties show the disk device type as an IBM 2145 Multi-Path disk device. To confirm that this information is accurate, right-click the disk to open the menu and select **Properties**, as shown in Figure 6-70.

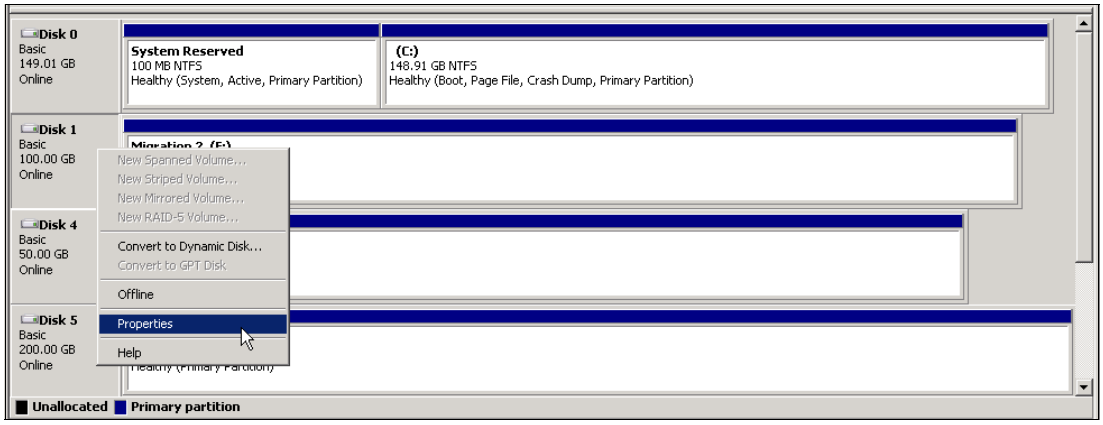

Figure 6-70 Display the disk properties from the Windows 2008 disk migration panel

After the disk properties panel is opened, the General tab shows the disk device type, as shown in Figure 6-71.

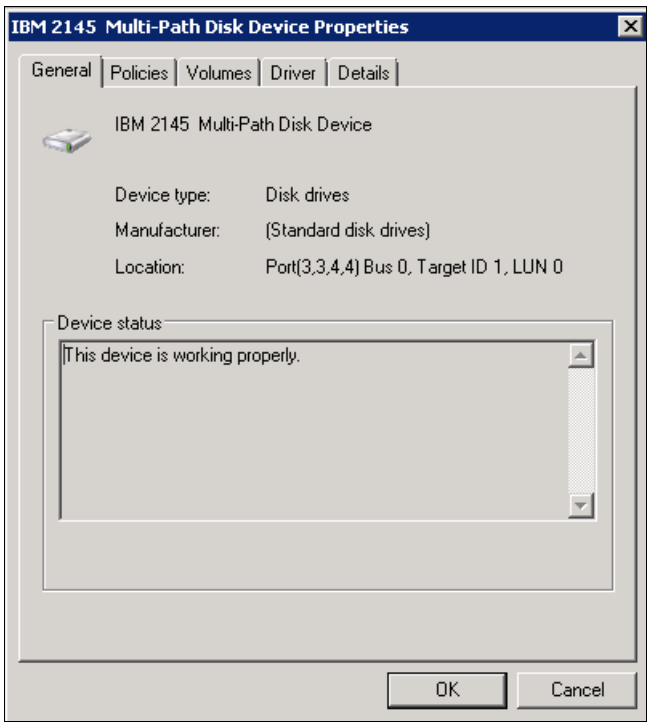

Figure 6-71 Windows 2008 properties General tab

The Storwize V5000 SDDDSM also can be used to verify that the migrated disk device is connected correctly. Open the SDDDSM command-line interface (CLI) to run the disk and adapter queries. As an example, on a Windows 2008 R2 SP1 host, click **Subsystem Device Driver DSM** to open the SSDDSM CLI window, as shown Figure 6-72 on page 290.

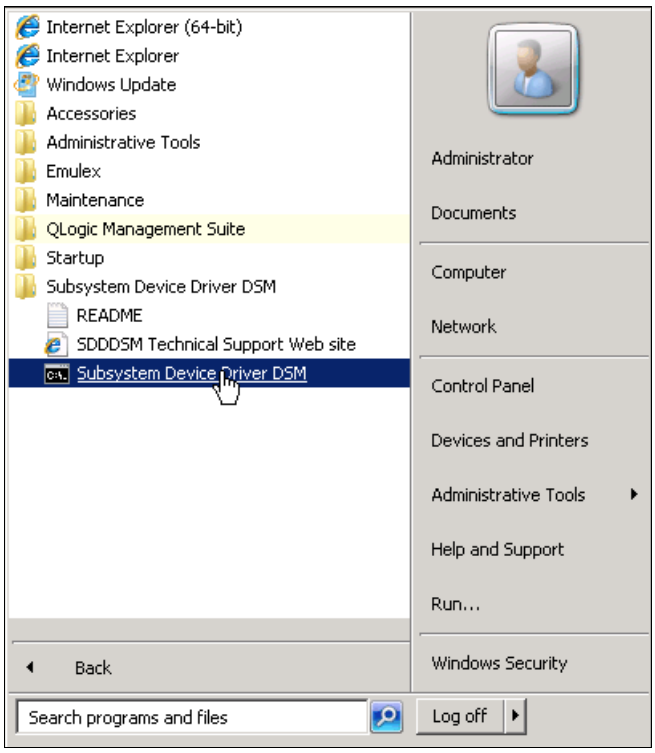

Figure 6-72 Windows 2008 R2 example: Open SSDDSM command line

The SDDDSM disk and adapter queries can be found in the SDDDSM user's guide. As an example on a Windows 2008 R2 SP1 host, useful commands to run include **`datapath query adapter`** and **`datapath query device`**. Example 6-1 shows the output of SDDDSM commands that were run on the Window 2008 host.

Example 6-1 Output from datapath query adapter and datapath query device SDDDSM commands

```
C:\Program Files\IBM\SDDDSM>datapath query adapter

Active Adapters :2

Adpt#            Name        State     Mode       Select     Errors  Paths   Active
    0  Scsi Port3 Bus0   NORMAL   ACTIVE        171          0      4        4
    1  Scsi Port4 Bus0   NORMAL   ACTIVE        174          0      4        4

C:\Program Files\IBM\SDDDSM>datapath query device

Total Devices : 2

DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 60050760009A81D37800000000000024
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port3 Bus0/Disk1 Part0     OPEN   NORMAL         90          0
    1     Scsi Port3 Bus0/Disk1 Part0     OPEN   NORMAL          0          0
    2     Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL         81          0
    3     Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL          0          0

DEV#:   1  DEVICE NAME: Disk2 Part0  TYPE: 2145       POLICY: OPTIMIZED
```

```
SERIAL: 60050760009A81D37800000000000025
============================================================================
Path#            Adapter/Hard Disk          State  Mode       Select      Errors
    0     Scsi Port3 Bus0/Disk2 Part0       OPEN   NORMAL         81           0
    1     Scsi Port3 Bus0/Disk2 Part0       OPEN   NORMAL          0           0
    2     Scsi Port4 Bus0/Disk2 Part0       OPEN   NORMAL         93           0
    3     Scsi Port4 Bus0/Disk2 Part0       OPEN   NORMAL          0           0
```

Use the SSDDSM output to verify that the expected number of devices, paths, and adapters are shown. Example 6-1 on page 290 shows the workload is balanced across each adapter and that there are four paths to the device. The **`datapath query device`** output shows two devices with `SERIALs: 6005070009A81D37800000000000024` and 6005070009A81D37800000000000025. The serial numbers can be cross-checked with the UID values that are now shown in step 6 of Storage Migration wizard, as shown in Figure 6-73.

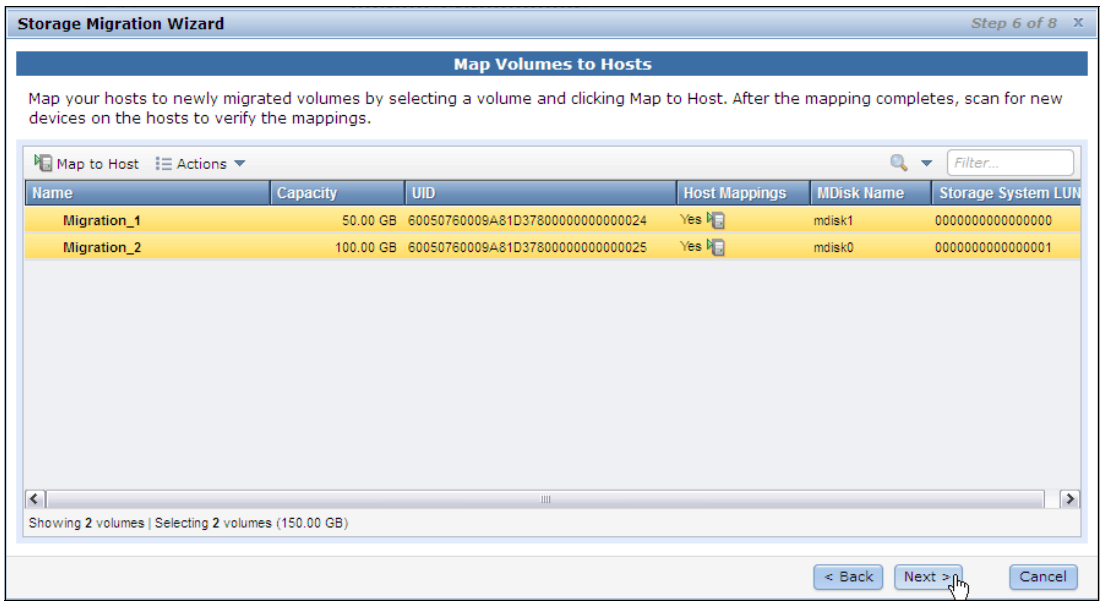

Figure 6-73 Mapped volumes and UIDs

From step 6 of the storage migration wizard, click **Next** to open step 7 of the wizard, as shown in Figure 6-73 on page 291.

Follow step 7 of the wizard. Highlight an internal storage pool and click **Next** to open the Start Migration running task panel, as shown in Figure 6-74.

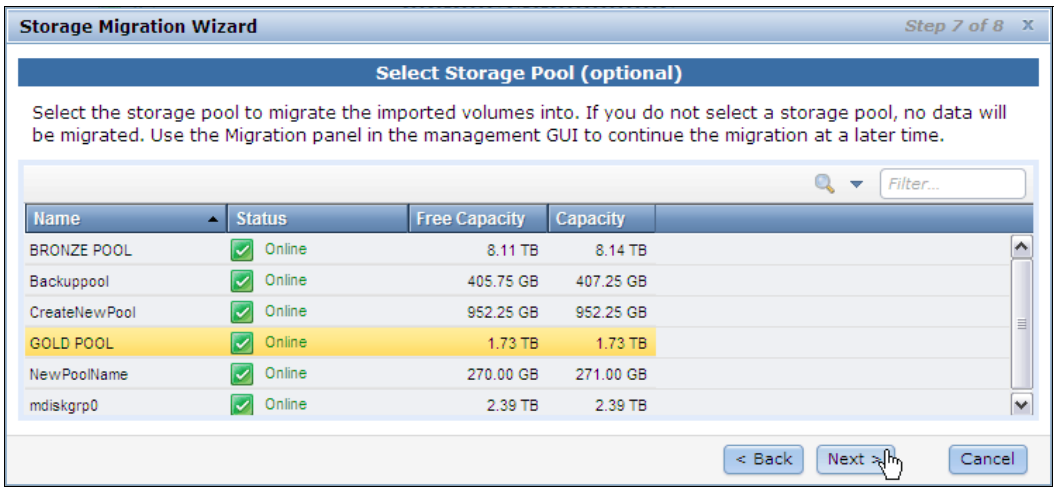

Figure 6-74 Select storage pool

After the Start Migration running task is complete, select **Close** to open step 8 of the storage migration wizard, as shown in Figure 6-75.

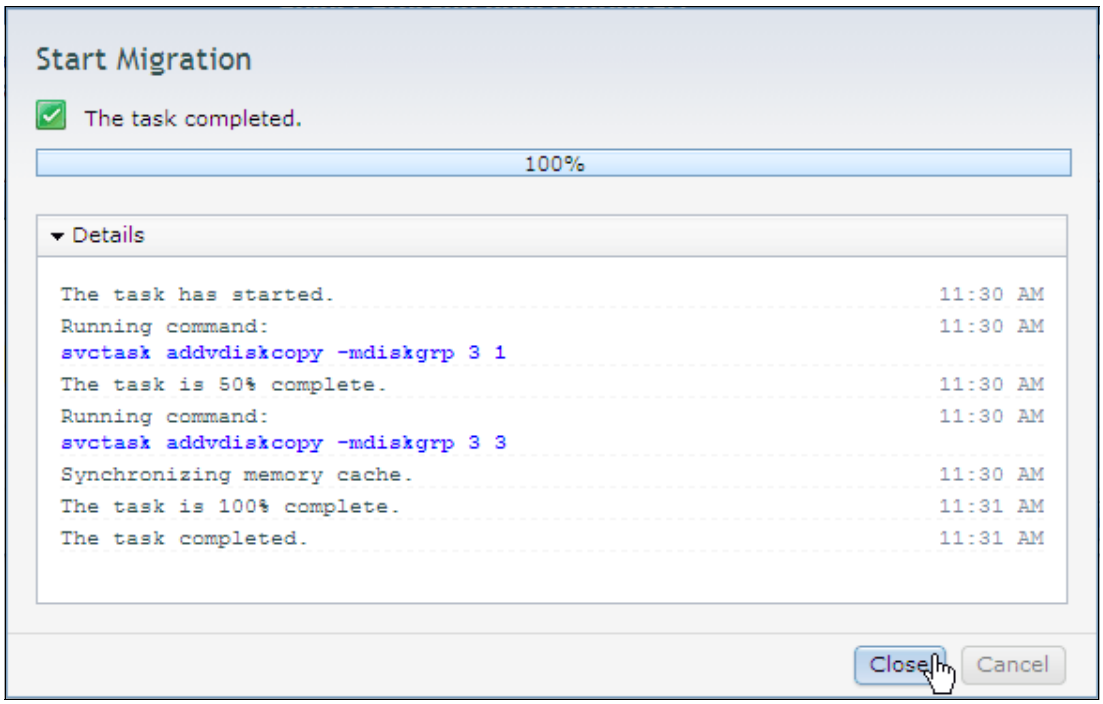

Figure 6-75 Start Migration completed task panel

Follow step 8 of the wizard and click **Finish** to open the System Migration panel, as shown in Figure 6-76.

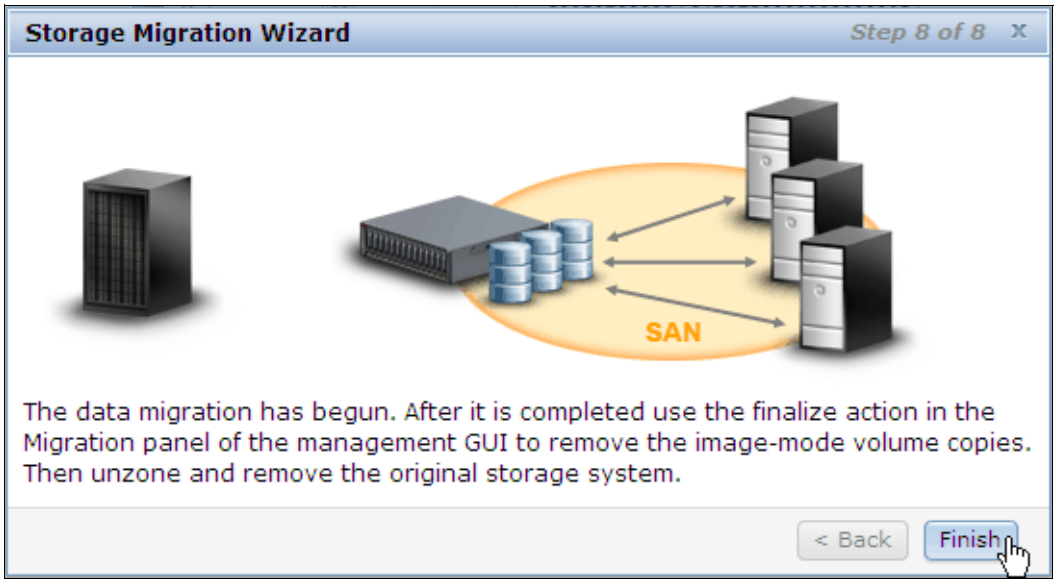

Figure 6-76 Step 8 of the Storage Migration wizard

The end of the Storage Migration wizard is not the end of the data migration process. The data migration is still in progress. A percentage indication of the migration progress is shown in the System Migration panel, as shown in Figure 6-77.

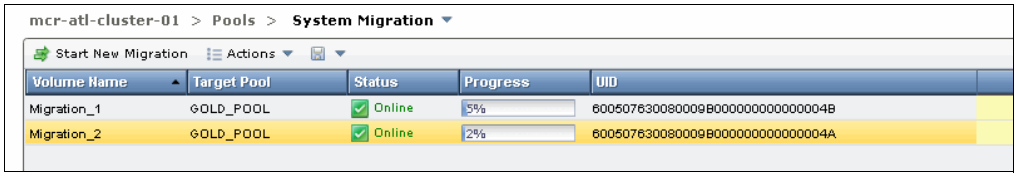

Figure 6-77 Migration progress indicators

Finalize the volume migrations. When the volume migrations are complete, select the volume migration instance and right-click **Finalize** to open the Finalize Volume Migrations panel, as shown in Figure 6-78.

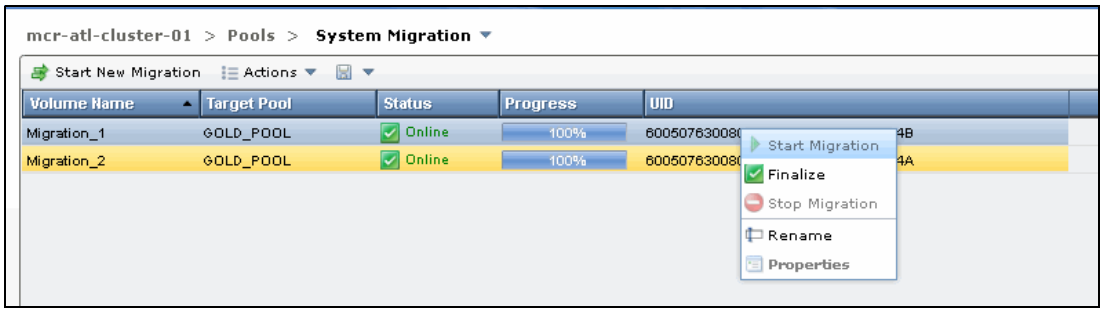

Figure 6-78 Finalize Volume Migrations

From the Finalize Volume Migrations panel, verify the volume names and the number of migrations and click **OK**, as shown in Figure 6-79.

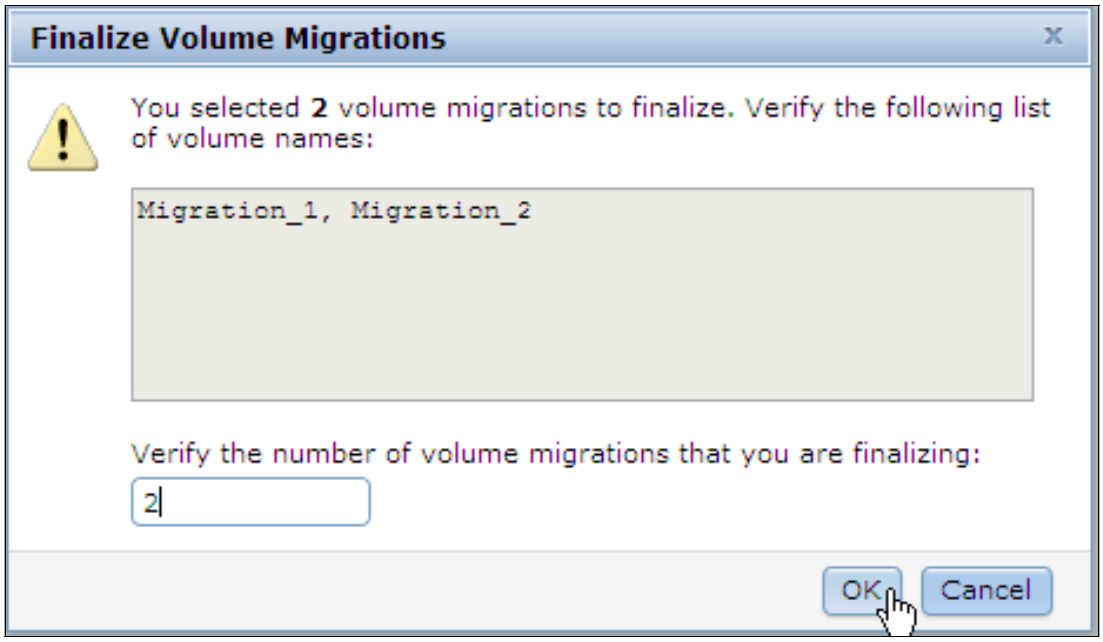

Figure 6-79 Finalize Volume Migrations panel

The image mode volumes are deleted and the associated image mode MDisks are removed from the migration storage pool. The status of those image mode MDisks is unmanaged. When the finalization completes, the data migration to the IBM Storwize V5000 is done. Remove the DS3400-to-V5000 zoning and retire the older storage system.

7

Storage pools

This chapter describes how IBM Storwize V5000 manages physical storage resources. All storage resources that are under IBM Storwize V5000 control are managed by using *storage pools*. Storage pools make it easy to dynamically allocate resources, maximize productivity, and reduce costs. Advanced internal storage, Managed Disks (MDisks), and storage pool management are covered in this chapter; external storage is covered in Chapter 11, “External storage virtualization” on page 547.

Storage pools can be configured through the Easy Setup wizard when the system is first installed, as described in Chapter 2, “Initial configuration” on page 27.

All available drives are configured based on recommended configuration preset values for the RAID level and drive class. The recommended configuration uses all the available drives to build arrays that are protected with the appropriate number of spare drives.

The management GUI also provides a set of presets to help you configure for different RAID types. You can tune storage configurations slightly that are based on best practices. The presets vary according to how the drives are configured. Selections include the drive class, the preset from the list that is shown, whether to configure spares, whether to optimize for performance or capacity, and the number of drives to provision.

This chapter includes the following topics:

- Working with internal drives
- Configuring internal storage
- Working with MDisks on internal and external storage
- Working with storage pools

Default extent size: The IBM Storwize V5000 GUI has a default extent size value of 1 GB when you define a new storage pool. This is a change in the IBM Storwize code v7.1 (earlier versions of code used a default extent size of 256 MB).

The GUI cannot change the extent size; therefore, creating storage pools with a different extent size must be done via the command-line interface (CLI) by using the **mkmdiskgrp** and **mkarray** commands.

7.1 Working with internal drives

This section describes how to configure the internal storage disk drives by using different RAID levels and optimization strategies.

The IBM Storwize V5000 storage system provides an Internal Storage window for managing all internal drives. The Internal Storage window can be accessed by opening the Overview window, clicking the **Internal Drives** function icon, and then clicking **Pools**, as shown in Figure 7-1.

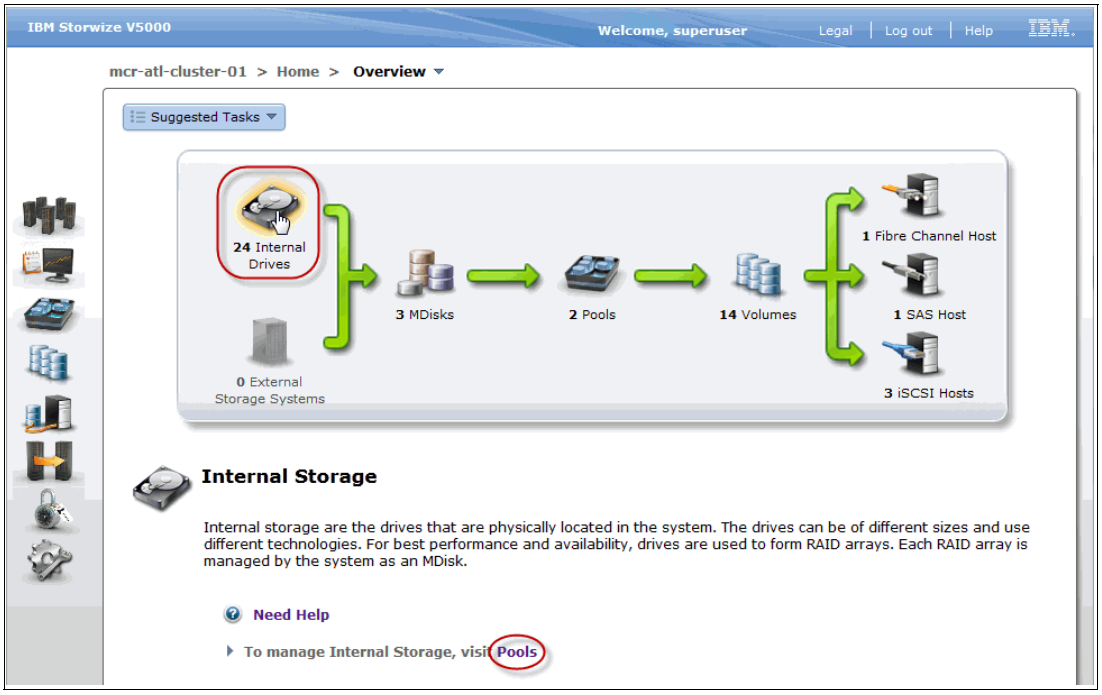

Figure 7-1 Internal Storage via Home Overview

An alternative way to access the Internal Storage window is by clicking the **Pools** icon on the left side of the window and selecting **Internal Storage**, as shown in Figure 7-2 on page 297.

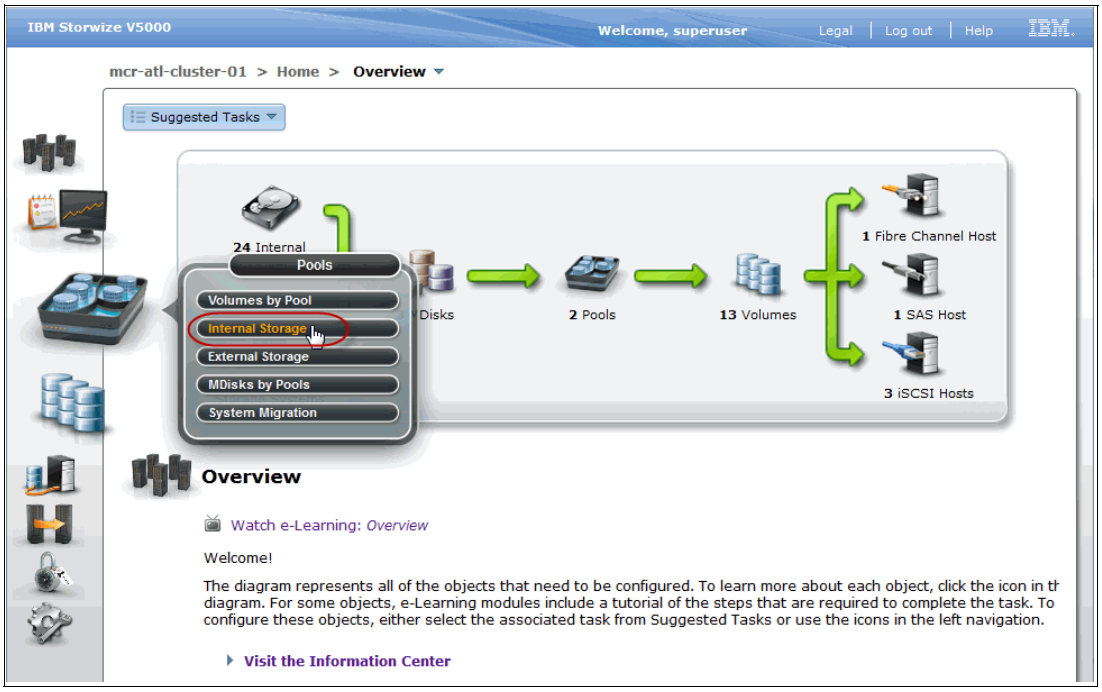

Figure 7-2 Internal Storage Details via Pools icon

7.1.1 Internal Storage window

The Internal Storage window (as shown in Figure 7-3) provides an overview of the internal drives that are installed in the IBM Storwize V5000 storage system. Selecting **All Internal** in the Drive Class Filter shows all of the drives that are installed in the managed system, including attached expansion enclosures. Alternatively, you can filter the drives by their type or class; for example, you can choose to show only serial-attached SCSI (SAS), Serial Advanced Technology Attachment (SATA), or solid-state drives (SSDs).

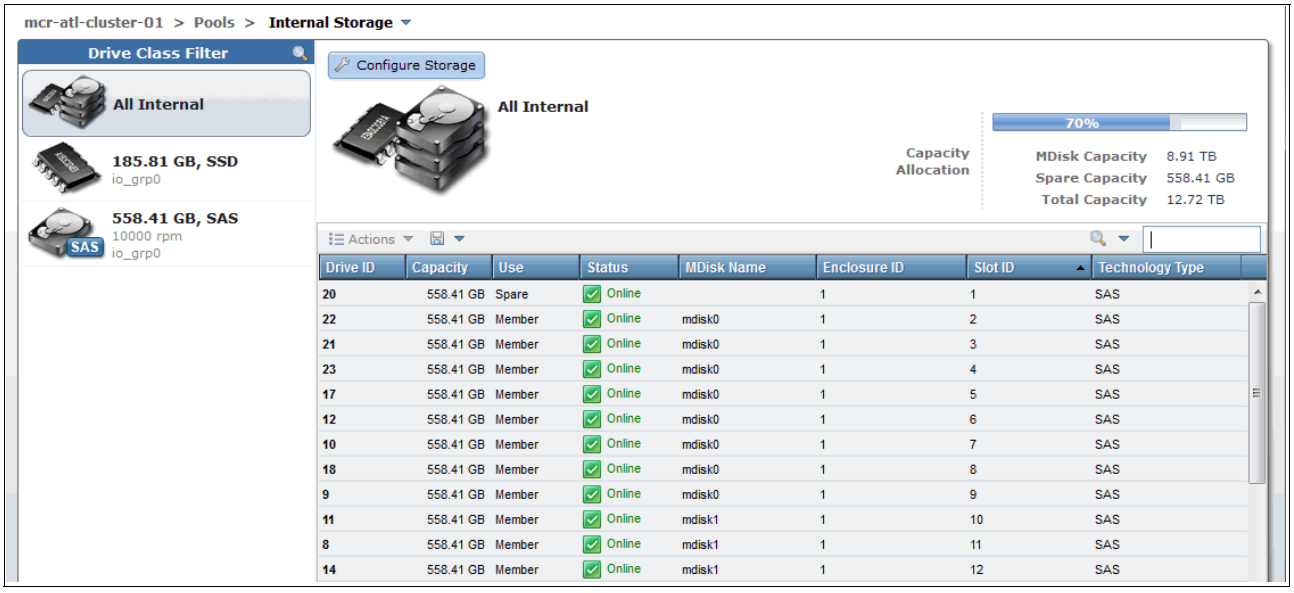

Figure 7-3 Internal storage window

On the right side of the Internal Storage window, the selected type of internal disk drives is listed. By default, the following information also is listed:

- Logical drive ID
- Drive's capacity
- Current type of use (unused, candidate, member, spare, or failed)
- Status (online, offline, and degraded)
- MDisk's name that the drive is a member of
- Enclosure ID that it is installed in
- Physical Drive Slot ID of the enclosure in which it is installed

The default sort order is by enclosure ID (this default can be changed to any other column by left-clicking the column header). To toggle between ascending and descending sort order, left-click the column header again.

More details can be shown (for example, the drive's Technology Type) by right-clicking the blue header bar of the table, which opens the selection panel, as shown in Figure 7-4.

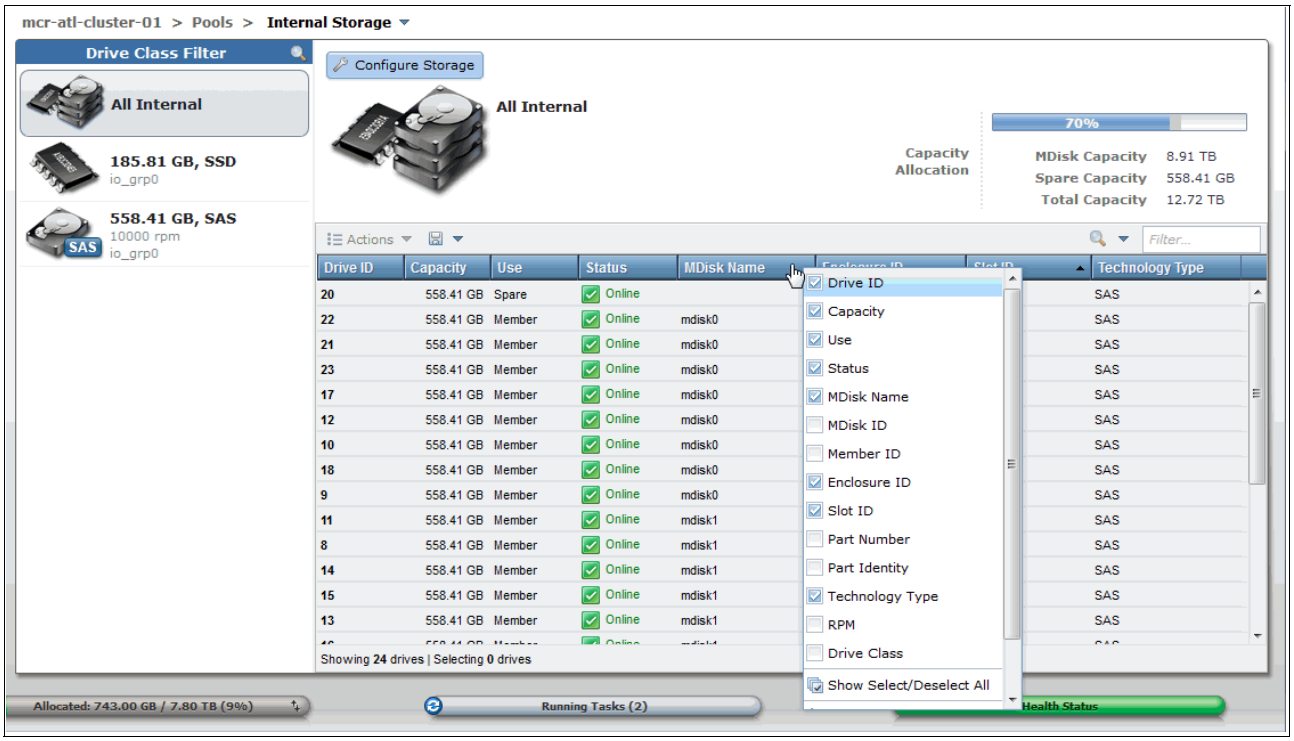

Figure 7-4 Internal storage window details selection

You also can find the internal storage capacity allocation indicator in the upper right corner. The Total Capacity shows the overall capacity of the internal storage that is installed in the IBM Storwize V5000 storage system. The MDisk Capacity shows the internal storage capacity that is assigned to the MDisks. The Spare Capacity shows the internal storage capacity that is used for hot spare disks.

The percentage bar that is shown in Figure 7-5 indicates how much capacity is allocated.

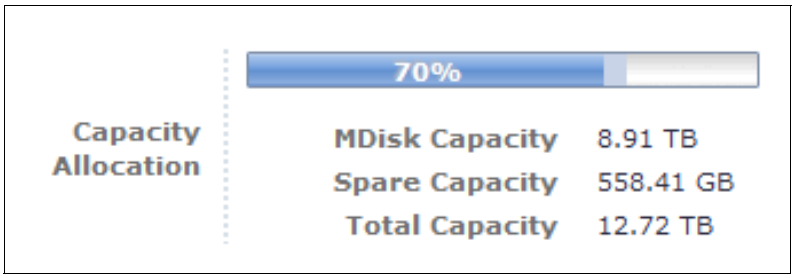

Figure 7-5 Internal storage allocation indicator

7.1.2 Actions on internal drives

There are a number of actions that can be performed on the internal drives when you select them and right-click or click the **Actions** drop-down menu, as shown in Figure 7-6.

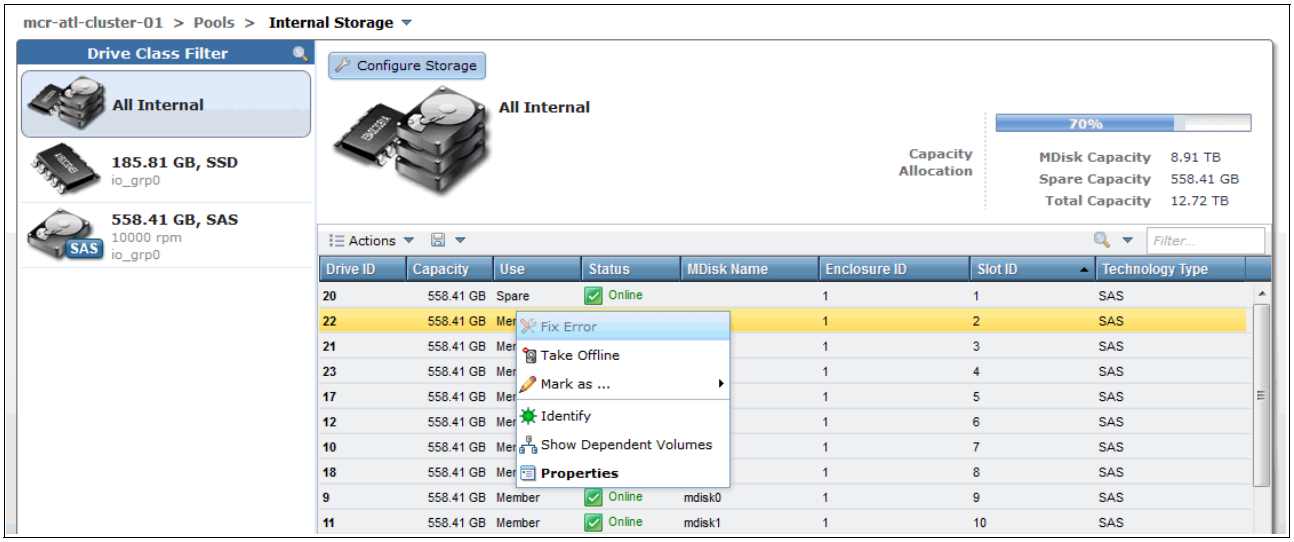

Figure 7-6 Internal drive Actions menu

Fix Error

The Fix Error action starts the Directed Maintenance Procedure (DMP) for a defective drive. For more information, see Chapter 12, “RAS, monitoring, and troubleshooting” on page 559.

Take Drive Offline window

The internal drives can be taken offline when there are problems with the drives. A confirmation window opens, as shown in Figure 7-7 on page 300.

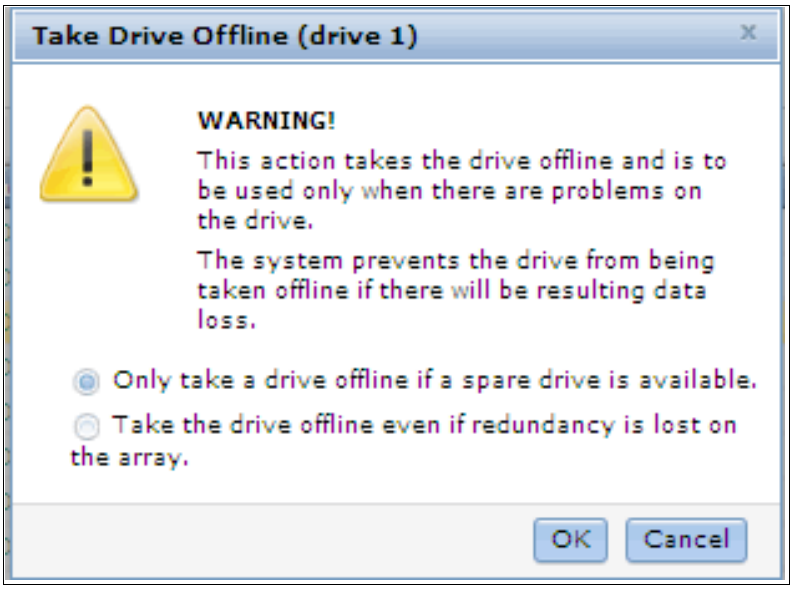

Figure 7-7 Take internal drive offline warning

A drive should be taken offline only if a spare drive is available. If the drive fails (as shown in Figure 7-8), the MDisk (of which the failed drive is a member) remains online and a hot spare is automatically reassigned.

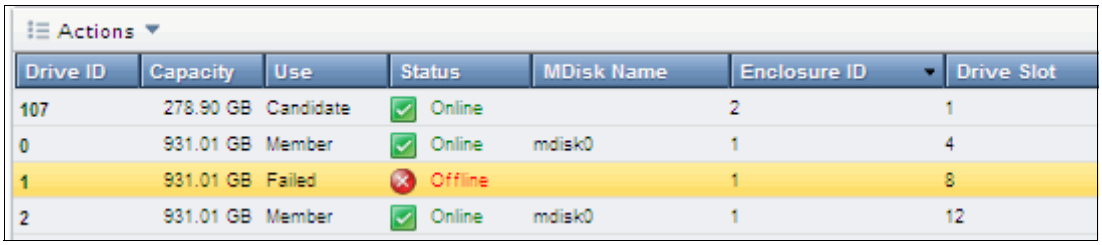

Figure 7-8 Internal drive taken offline

If no sufficient spare drives are available and one drive must be taken offline, the second option for no redundancy must be selected. This option results in a degraded MDisk, as shown in Figure 7-9.

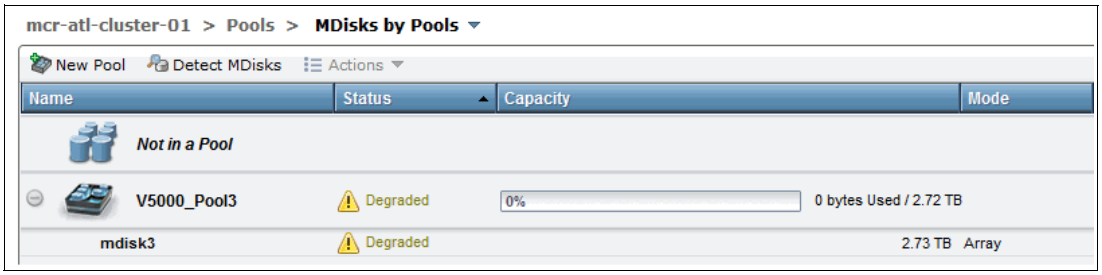

Figure 7-9 Internal drive that is failed with MDisk degraded

The IBM Storwize V5000 storage system prevents the drive from being taken offline if there might be data loss as a result. A drive cannot be taken offline (as shown in Figure 7-10) if no suitable spare drives are available and, based on the RAID level of the MDisk, drives are already offline.

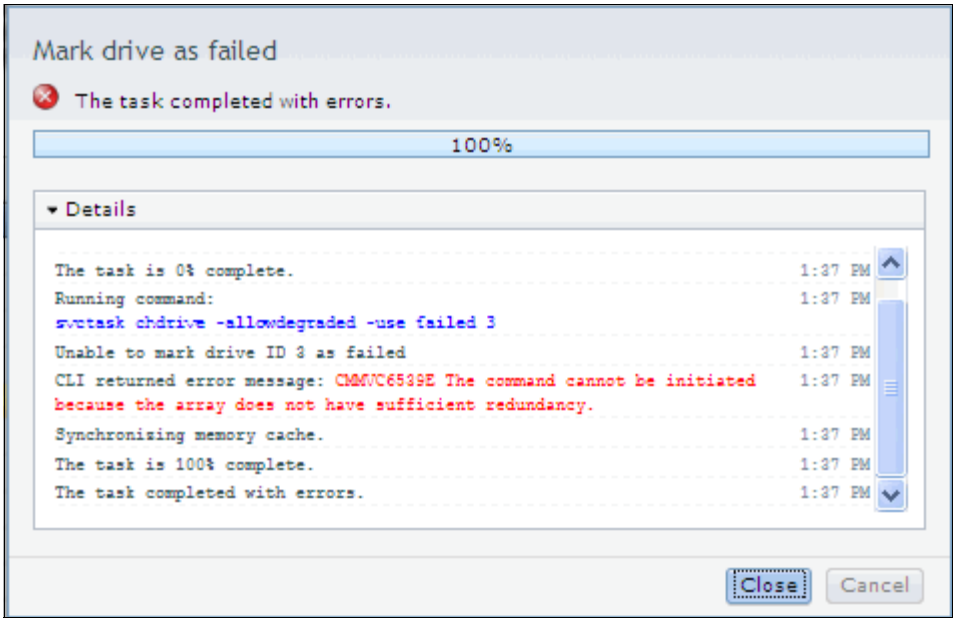

Figure 7-10 Internal drive offline not allowed because of insufficient redundancy

Example 7-1 shows how to use the **chdrive** CLI command to set the drive to failed.

Example 7-1 The use of the chdrive command to set drive to failed

```
chdrive -use failed driveID
chdrive -use failed -allowdegraded driveID
```

Mark as

The internal drives in the IBM Storwize V5000 storage system can be assigned to the following usage roles, as shown in Figure 7-11 on page 302:

- Unused: The drive is not in use and cannot be used as a spare.
- Candidate: The drive is available for use in an array.
- Spare: The drive can be used as a hot spare, if required.

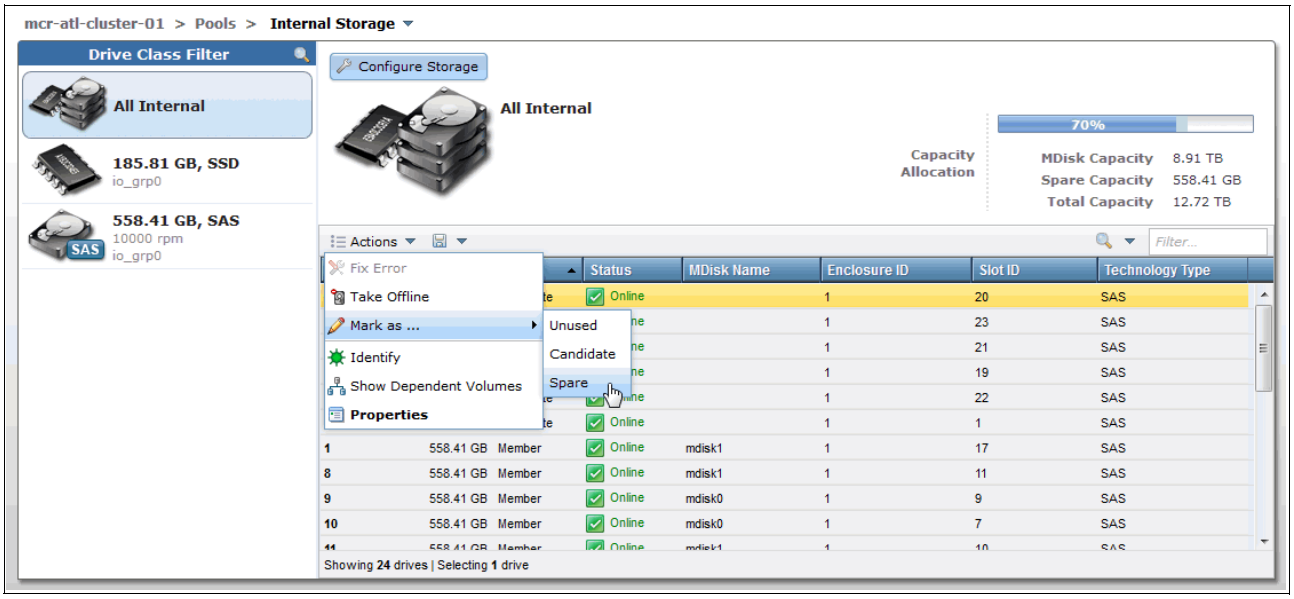

Figure 7-11 Internal drive Mark as... option

The new role that can be assigned depends on the current drive usage role. These dependencies are shown in Figure 7-12.

		To				
		Unused	Candidate	Failed	Member	Spare
From	Unused	allowed	allowed	no option		not allowed
	Candidate	allowed	allowed			allowed
	Failed	allowed	allowed			not allowed
	Member	No change on member drives				
	Spare	not allowed	allowed	no option		allowed

Figure 7-12 Internal drive usage role table

Identify

Use the Identify action to turn on the LED light so that you can easily identify a drive that must be replaced or that you want to troubleshoot. The panel that is shown in Figure 7-13 on page 303 appears when the LED is on.

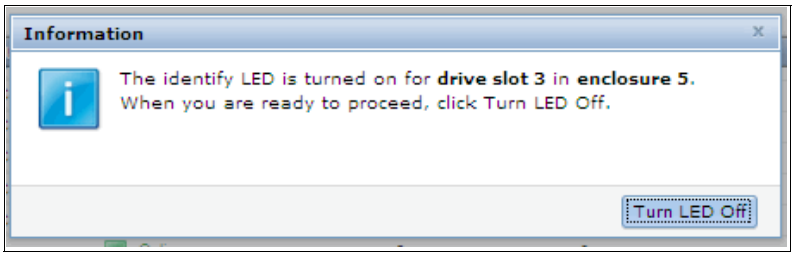

Figure 7-13 Internal drive identification

Click **Turn LED Off** when you are finished.

Example 7-2 shows how to use the **chenclosureslot** command to turn on and off the drive LED.

Example 7-2 The use of the chenclosureslot command to turn on and off drive LED

```
chenclosureslot -identify yes/no -slot slot enclosureID
```

Show Dependent Volumes

Clicking **Show Dependent Volumes** shows you volumes that are dependent on the selected drive. Volumes are dependent on a drive only when the underlying disks or MDisks are in a degraded or inaccessible state and removing further hardware causes the volume to go offline. This condition is true for any RAID 0 MDisk or if the associated MDisk is degraded already.

Use the Show Dependent Volumes option before you perform maintenance to determine which volumes are affected.

Important: A lack of listed dependent volumes does not imply that there are no volumes created that use this drive.

Figure 7-14 shows an example if no dependent volumes are detected for this specific drive.

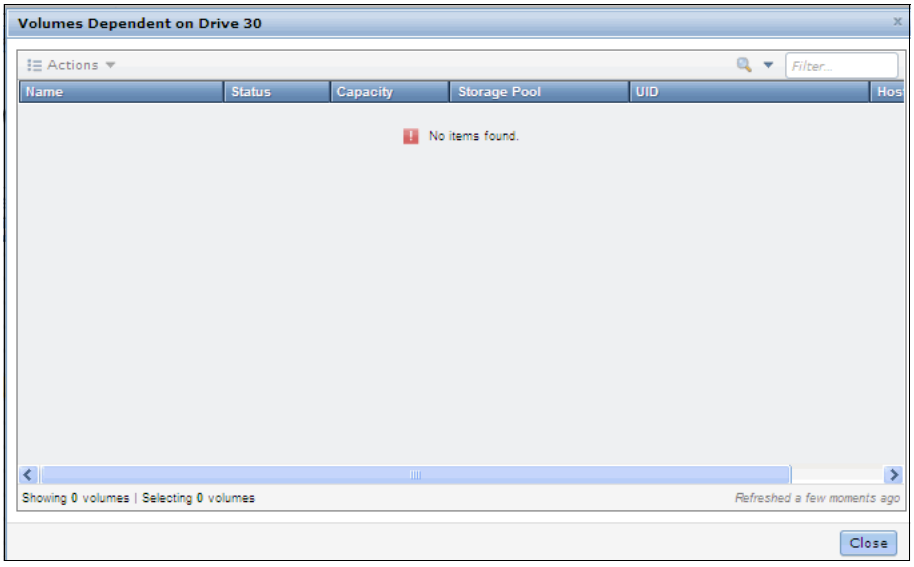

Figure 7-14 Internal drive no dependent volume

Figure 7-15 shows the list of dependent volumes for a drive when its underlying MDisk is in a degraded state.

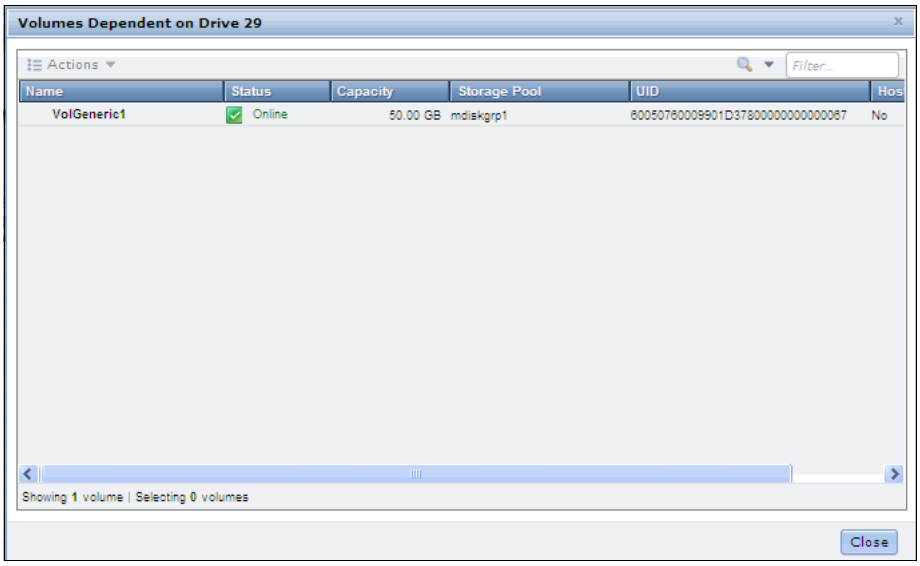

Figure 7-15 Internal drive with dependent volume

Example 7-3 shows how to view dependent volumes for a specific drive by using the CLI.

Example 7-3 Command to view dependent Vdisks for a specific drive

```
lsdependentvdisks -drive driveID
```

Properties

Clicking **Properties** (as shown in Figure 7-16) in the Actions menu or double-clicking the drive provides the vital product data (VPD) and the configuration information. The Show Details option was selected to show more information.

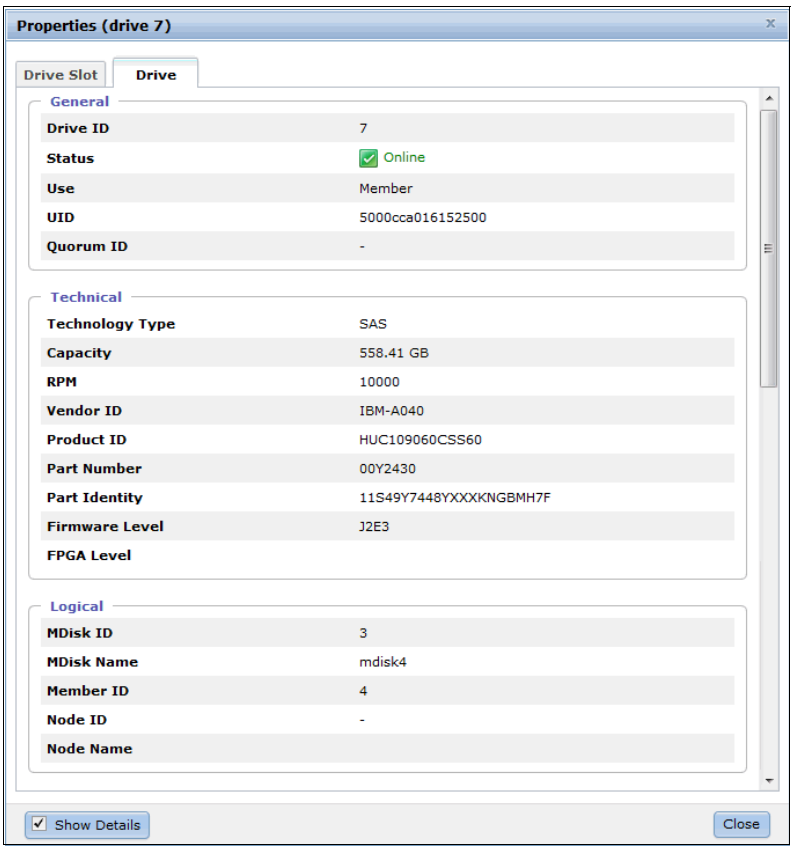

Figure 7-16 Internal drives properties: Part1

If the Show Details option is not selected, the technical information section is reduced, as shown in Figure 7-17.

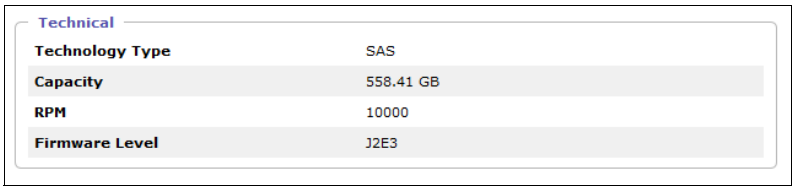

Figure 7-17 Internal drives properties no details

A tab for the Drive Slot is available in the Properties panel (as shown in Figure 7-18) to get specific information about the slot of the selected drive.

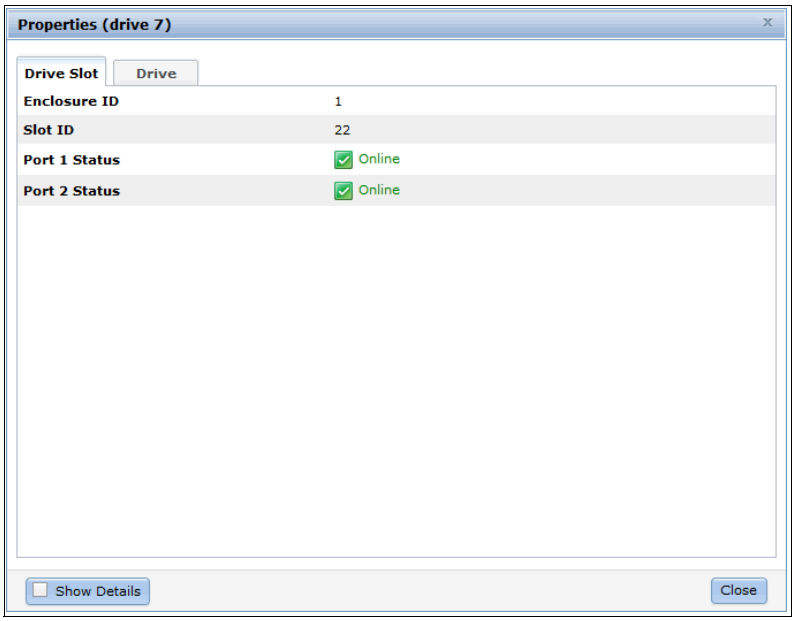

Figure 7-18 Internal drive properties slot

Example 7-4 shows how to use the **lsdrive** command to display configuration information and drive VPD.

Example 7-4 The use of the lsdrive command to display configuration information and drive VPD

```
lsdrive driveID
```

7.2 Configuring internal storage

The internal storage of an IBM Storwize V5000 can be configured into MDisks and pools by using the system setup wizard during the initial configuration. For more information, see Chapter 2, “Initial configuration” on page 27.

The decision that is shown in Figure 7-19 must be made when a IBM Storwize V5000 is configured.

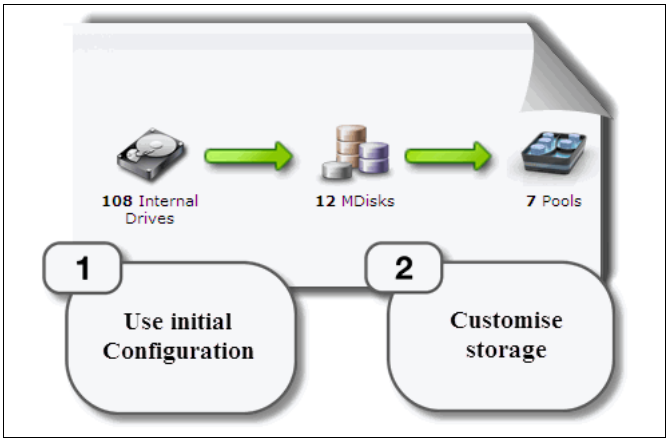

Figure 7-19 Decision to customize storage configuration

The decision choices include the following meanings:

- Use initial configuration

  During system setup, all available drives can be configured based on the RAID configuration presets. The setup creates MDisks and pools but does not create volumes.

  If this automated configuration fits your business requirement, it is recommended that this configuration is kept.

- Customize storage configuration

  A storage configuration might be customized for the following reasons:

  – The automated initial configuration does not meet customer requirements.

  – More storage was attached to the IBM Storwize V5000 and must be integrated into the existing configuration.

7.2.1 RAID configuration presets

RAID configuration presets are used to configure internal drives that are based on recommended values for the RAID level and drive class. Each preset has a specific goal for the number of drives per array and the number of spare drives to maintain redundancy.

Table 7-1 on page 308 describes the presets that are used for SSDs for the IBM Storwize V5000 storage system.

Table 7-1 SSD RAID presets

Preset	Purpose	RAID level	Drives per array goal	Drive count (Min - Max)	Spare drive goal
SSD RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1
SSD RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1
SSD RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	2 - 16 (even)	1
SSD RAID 1	Protects against at least one drive failure. All data is mirrored on two array members.	1	2	2	1
SSD RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0
SSD Easy Tier	Mirrors data to protect against drive failure. The mirrored pairs are spread between storage pools to be used for the Easy Tier function.	10	2	2 - 16 (even)	1

SSD RAID instances: In all SSD RAID instances, drives in the array are balanced across enclosure chains, if possible.

Table 7-2 describes the RAID presets that are used for hard disk drives (HDDs) for the IBM Storwize V5000 storage system.

Table 7-2 HDD RAID presets

Preset	Purpose	RAID level	Drives per array goal	Drive count (Min - Max)	Spare goal	Chain balance
Basic RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1	All drives in the array are from the same chain wherever possible.
Basic RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1	All drives in the array are from the same chain wherever possible.
Basic RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	2 - 16 (evens)	1	All drives in the array are from the same chain wherever possible.

Preset	Purpose	RAID level	Drives per array goal	Drive count (Min - Max)	Spare goal	Chain balance
Balanced RAID 10	Protects against at least one drive or enclosure failure. All data is mirrored on two array members. The mirrors are balanced across the two enclosure chains.	10	8	2 - 16 (evens)	1	Exactly half of the drives are from each chain.
RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0	All drives in the array are from the same chain wherever possible.

7.2.2 Customizing initial storage configuration

If the initial storage configuration does not meet the requirements, pools must be deleted. Select the Pool navigator in the GUI and click **Pools** → **MDisks by Pools**. Select and right-click the pool and then select **Delete Pool**, as shown in Figure 7-20.

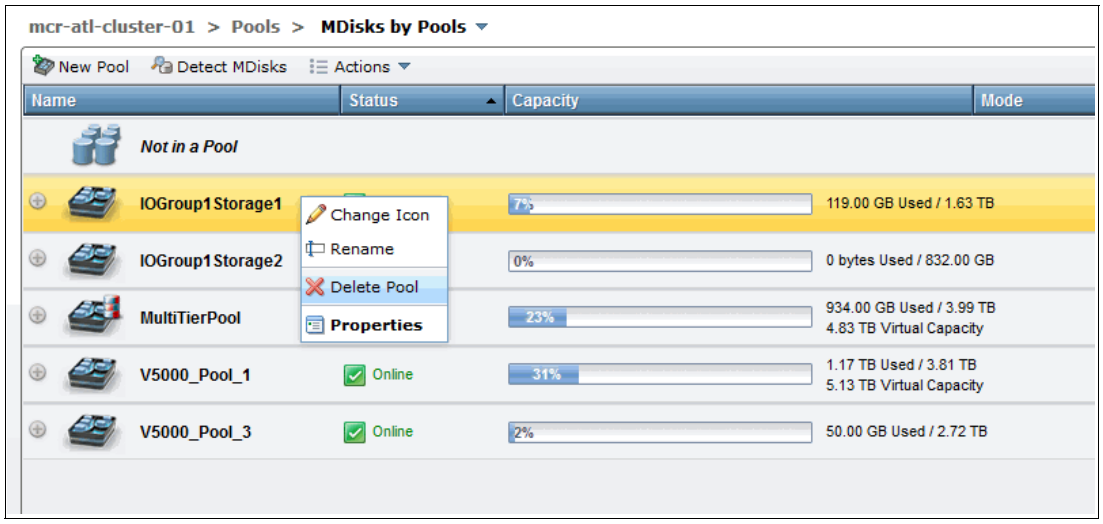

Figure 7-20 Delete selected pool

The option for deleting the volume, host mappings, and MDisks must be selected so that all associated drives are marked as a candidate for deletion, as shown in Figure 7-21.

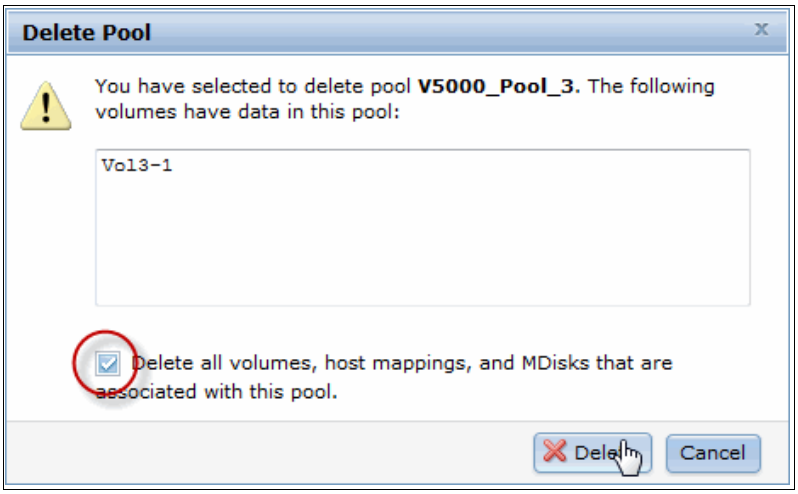

Figure 7-21 Delete pool confirmation

These drives now can be used for a different configuration.

Important: When a pool is deleted, data that is contained within any volume that is provisioned from this pool is deleted.

7.2.3 Creating an MDisk and pool

To configure internal storage for use with hosts, click **Pools** → **Internal Storage** and then click **Configure Storage**, as shown in Figure 7-22.

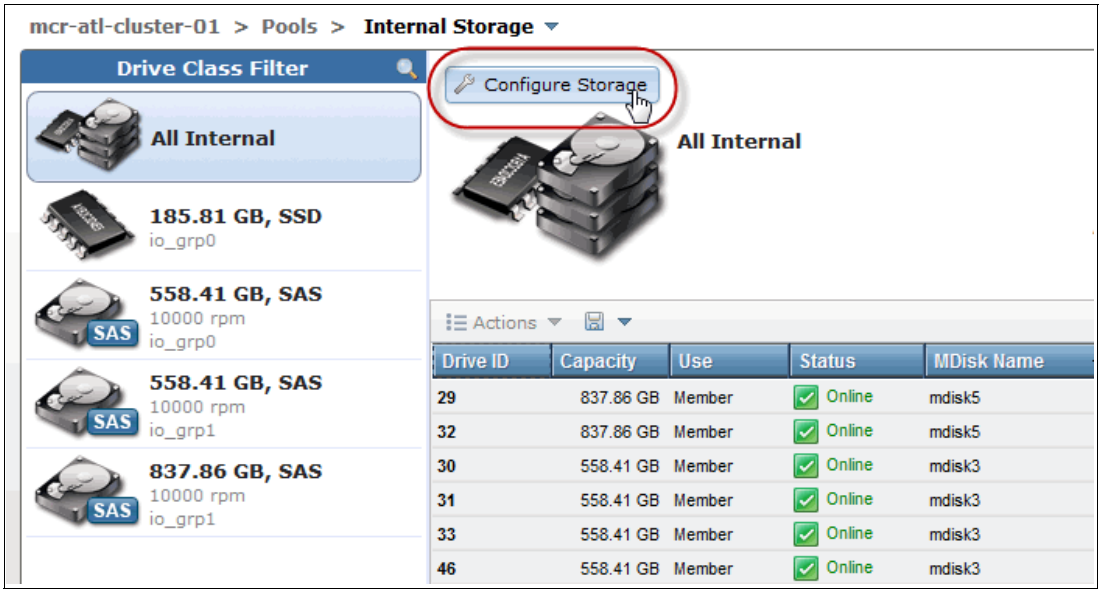

Figure 7-22 Click Configure Storage

A configuration wizard opens and guides you through the process of configuring internal storage. The wizard shows all internal drives, their status, and their use. The status shows whether it is Online, Offline, or Degraded. The Use status shows if a drive is Unused, a Candidate for configuration, a Spare, a Member of a current configuration, or Failed. Figure 7-23 shows an example in which 15 drives are available for configuration.

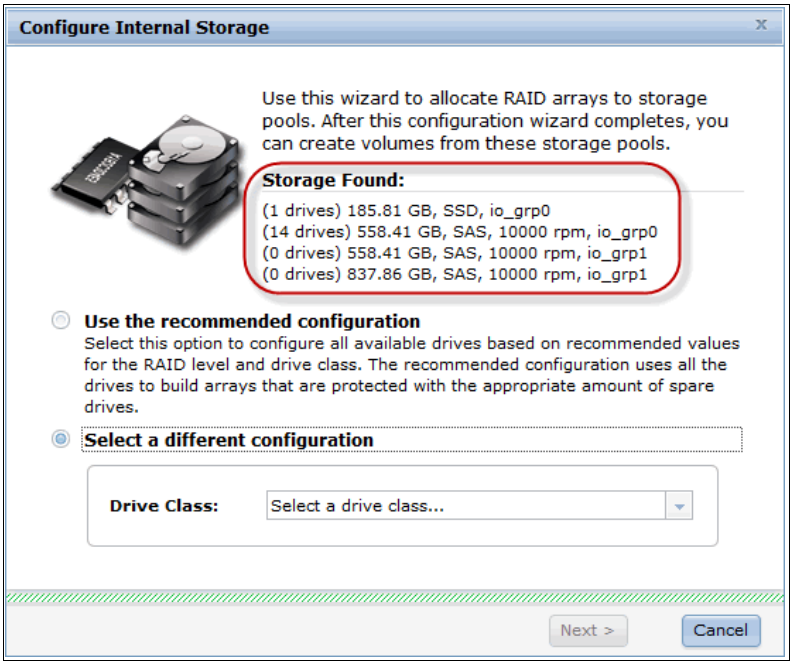

Figure 7-23 Available drives for new MDisk

If there are internal drives with a status of Unused, a window opens, which gives the option to include them in the RAID configuration, as shown in Figure 7-24.

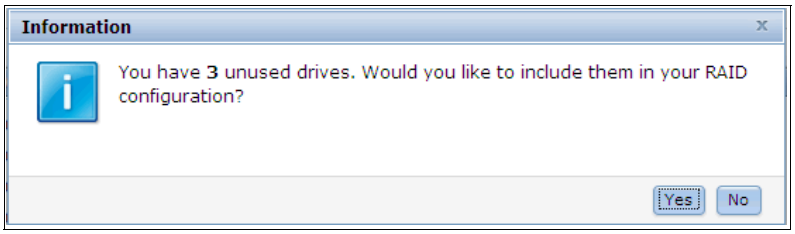

Figure 7-24 Unused drives warning

When the decision is made to include the drives into the RAID configuration, their status is set to Candidate, which also makes them available for a new MDisk.

The use of the storage configuration wizard simplifies the initial disk drive setup and offers the following options:

- Use the recommended configuration
- Select a different configuration

Selecting **Use the recommended configuration** guides you through the wizard that is described in 7.2.4, “Using the recommended configuration” on page 312. Selecting **Select a different configuration** uses the wizard that is described in 7.2.5, “Selecting a different configuration” on page 314.

7.2.4 Using the recommended configuration

As shown in Figure 7-25, when you click **Use the recommended configuration**, the wizard offers a recommended storage configuration at the bottom of the window.

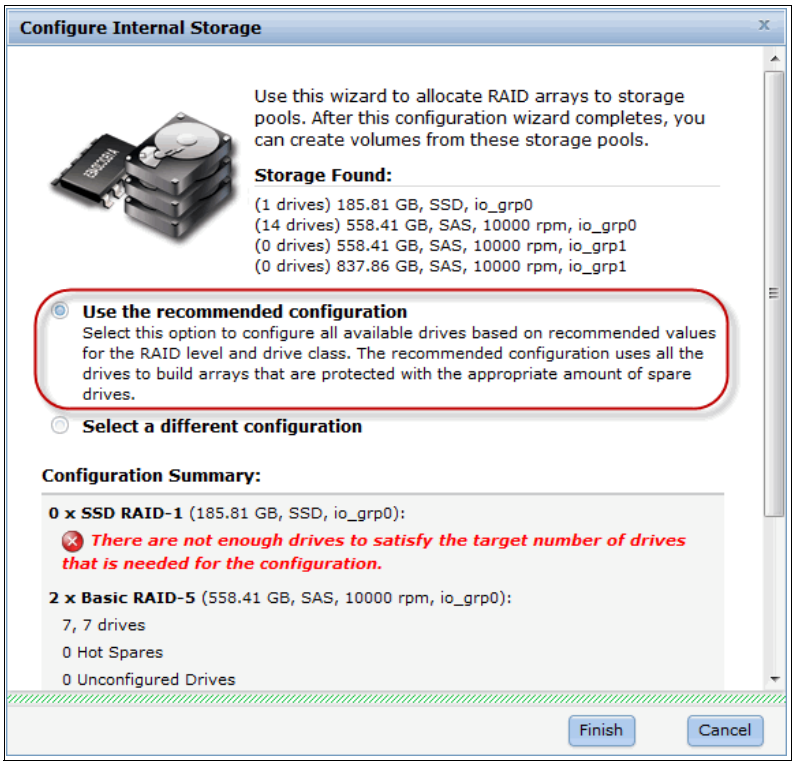

Figure 7-25 The recommended configuration

The following recommended RAID presets for different drive classes are available:

- SSD EasyTier or RAID 1 for SSDs
- Basic RAID 5 for SAS drives
- Basic RAID 6 for Nearline SAS drives

Figure 7-25 shows a sample configuration with 1x SSD and 14x SAS drives. The Configuration Summary shows a warning that there are insufficient SSDs installed to satisfy the RAID 1 SSD preset (two drives are required to do this), plus a third drive for a hot spare.

By using the recommended configuration, spare drives are also automatically created to meet the spare goals according to the preset chosen; one spare drive is created out of every 24 disk drives of the same drive class on a single chain. Spares are not created if sufficient spares are already configured.

Spare drives in the IBM Storwize V5000 are *global spares*, which means that any spare drive that has at least the same capacity as the drive to be replaced can be used in any array. Thus, an SSD array with no SSD spare available uses an HDD spare instead.

If the proposed configuration meets your requirements, click **Finish**, and the system automatically creates the array MDisks with a size according to the chosen RAID level.

Storage pools also are automatically created to contain the MDisks with similar performance characteristics, including the consideration of RAID level, number of member drives, and drive class.

Important: This option adds new MDisks to an existing storage pool when the characteristics match. If this is not what is required, the Select a different configuration option should be used.

After an array is created, the Array MDisk members are synchronized with each other through a background initialization process. The progress of the initialization process can be monitored by clicking the icon at the left of the Running Tasks status bar and selecting the initialization task to view the status, as shown in Figure 7-26.

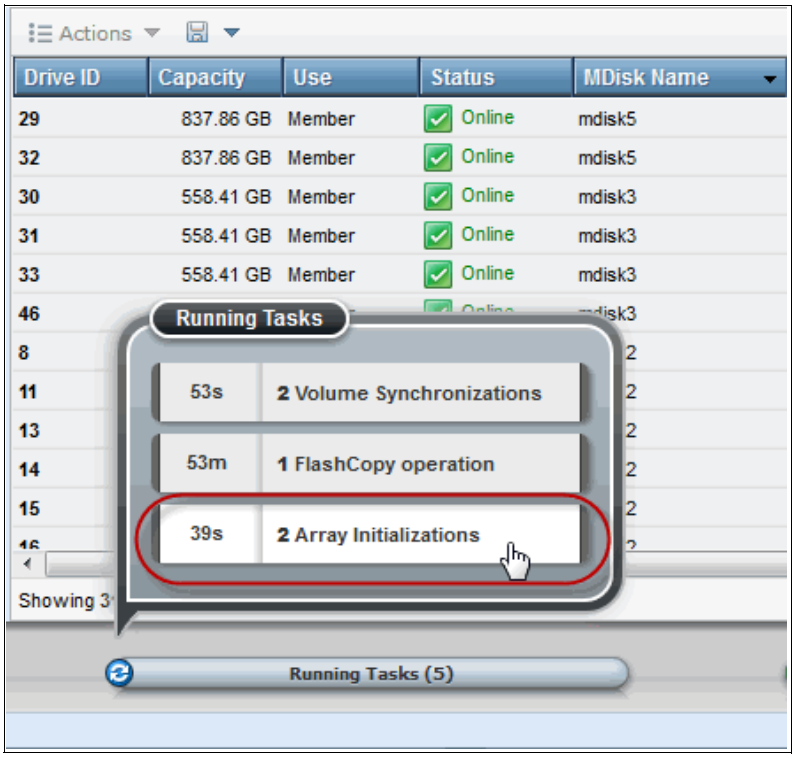

Figure 7-26 Running task panel

Click the taskbar to open the progress window, as shown in Figure 7-27. The array is available for I/O during this process. The initialization does not affect the availability because of possible member drive failures.

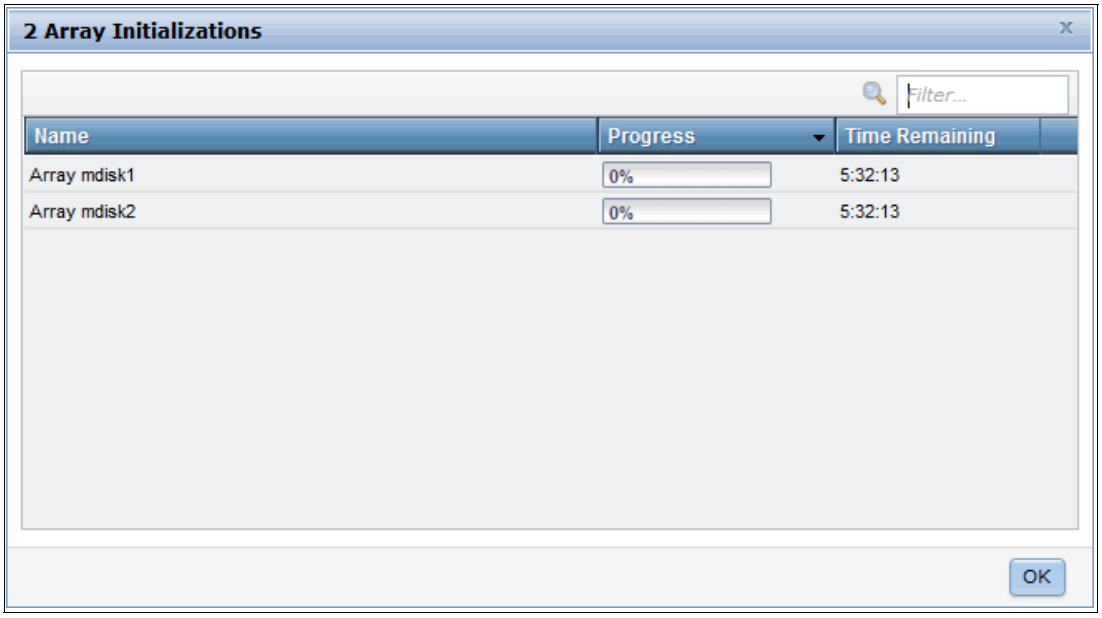

Figure 7-27 Initialization progress view

7.2.5 Selecting a different configuration

The Select a different configuration option offers a more flexible way to configure the internal storage as compared to the Use the recommended configuration preset in terms of drive selection, RAID level, and storage pool to be used.

Only one drive class (RAID configuration) can be allocated at a time.

Complete the following steps to select a different configuration:

1. Choose drive class and RAID preset.

   The drive class selection list contains each drive class that is available for configuration, as shown in Figure 7-28 on page 315.

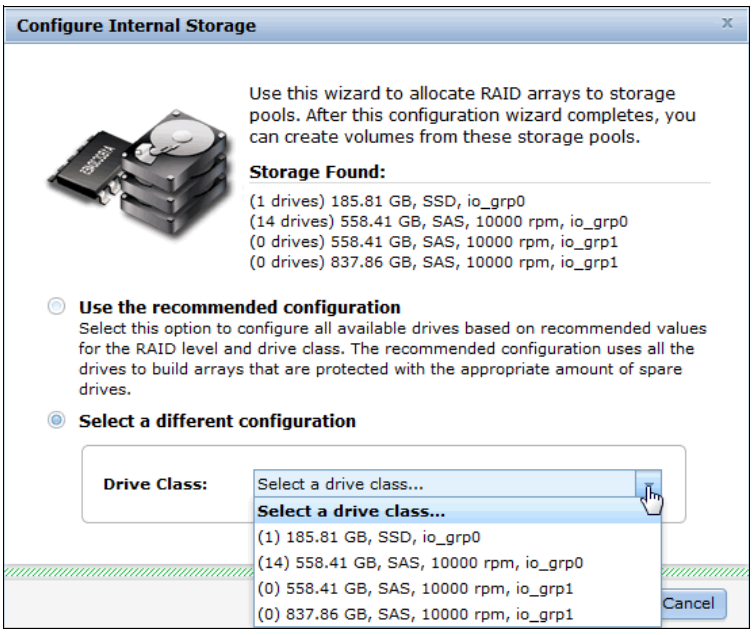

Figure 7-28 Select drive class for new configuration

2. Click **Next** and select the appropriated RAID preset, as shown in Figure 7-29.

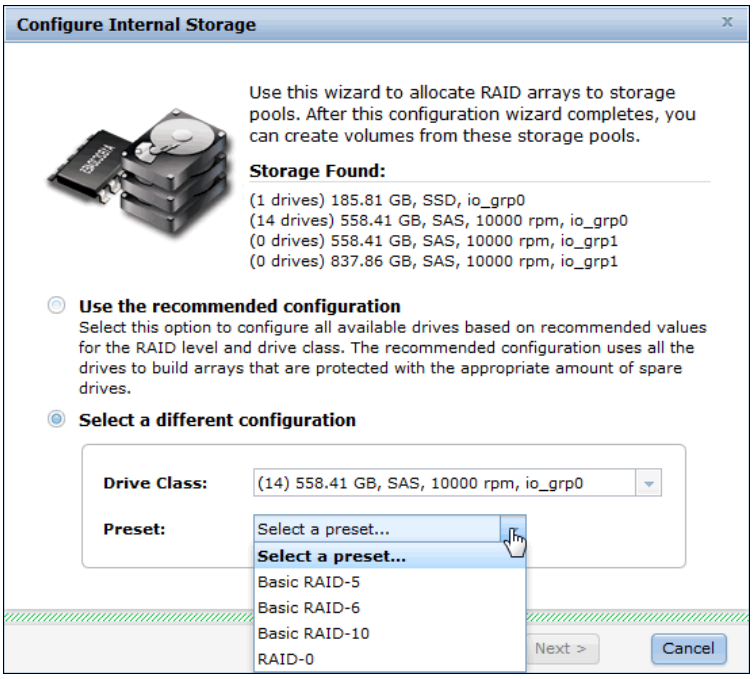

Figure 7-29 Select the RAID preset

3. Define the RAID attributes.

   You can tune RAID configurations slightly that are based on best practices. Selections include the configuration of spares, optimization for performance, optimization for capacity, and the number of drives to provision.

Each IBM Storwize V5000 preset has a specific goal for the number of drives per array. For more information, see the Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp

Table 7-3 shows the RAID goal widths.

Table 7-3 RAID goal width

RAID level	HDD goal width	SSD goal width
0	8	8
5	8	9
6	12	10
10	8	8

The following RAID configurations are available:

– Optimize for Performance

Optimizing for performance creates arrays with the same capacity and performance characteristics. The RAID goal width (as shown in Table 7-3) must be met for this target. In a performance optimized setup, the IBM Storwize V5000 provisions eight physical disk drives in a single array MDisk, except for the following situations:

- RAID 6 uses 12 disk drives.
- SSD Easy Tier uses two disk drives.

Hence, creating an Optimized for Performance configuration is only possible if there are enough drives available to match your needs.

As a consequence, all arrays with similar physical disks feature the same performance characteristics. Because of the defined presets, this setup might leave drives unused. The remaining unconfigured drives can be used in another array.

Figure 7-30 shows an example in which not all of the provisioned drives can be used in a performance optimized configuration (six drives remain).

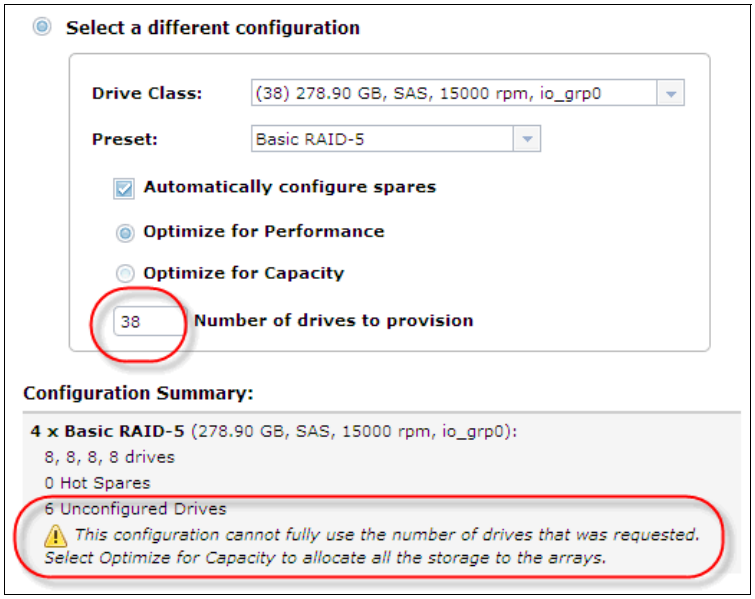

Figure 7-30 Optimization for performance failed

Figure 7-31 shows that the number of drives is not enough to satisfy the needs of the configuration.

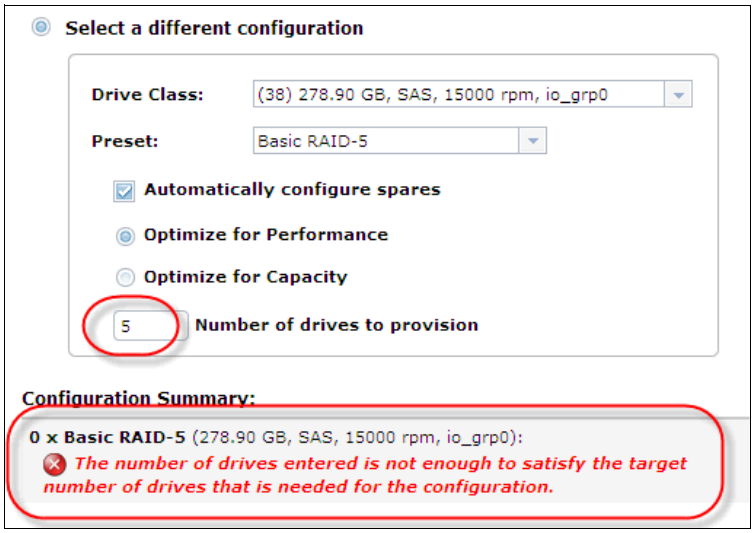

Figure 7-31 Not enough drives for performance optimization

Figure 7-32 shows that there are a suitable number of drives to configure performance optimized arrays.

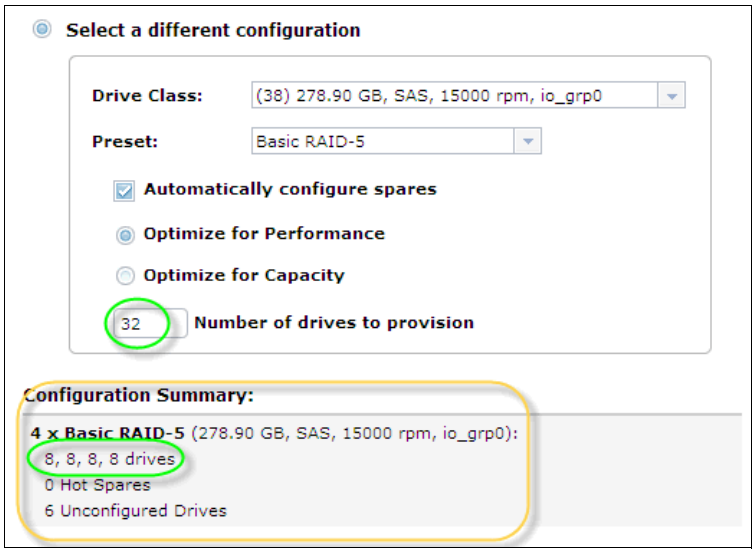

Figure 7-32 Arrays match performance goals

Four RAID 5 arrays where built and all provisioned drives are used.

– Optimize for Capacity

Optimizing for capacity creates arrays that allocate all the drives that are specified in the Number of drives to provision field. This option results in arrays of different capacities and performance. The number of drives in each MDisk does not vary by more than one drive, as shown in Figure 7-33 on page 318.

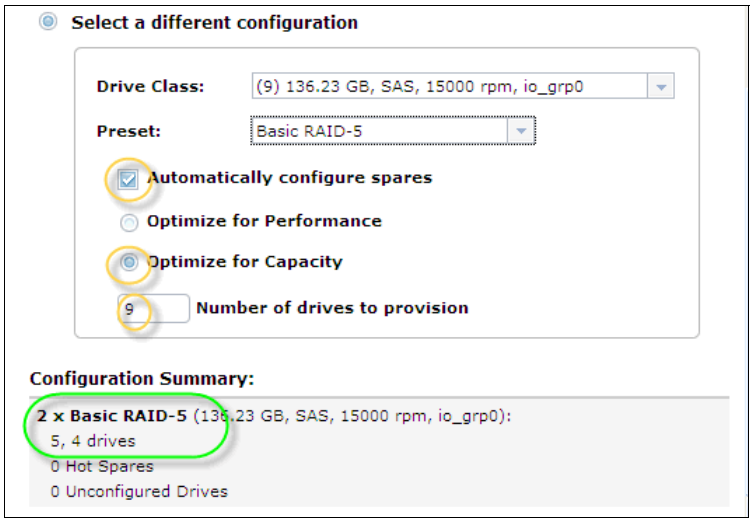

Figure 7-33 Capacity optimized configuration

4. Storage pool assignment.

   Choose whether an existing pool must be expanded or whether a pool is created for the configuration, as shown in Figure 7-34.

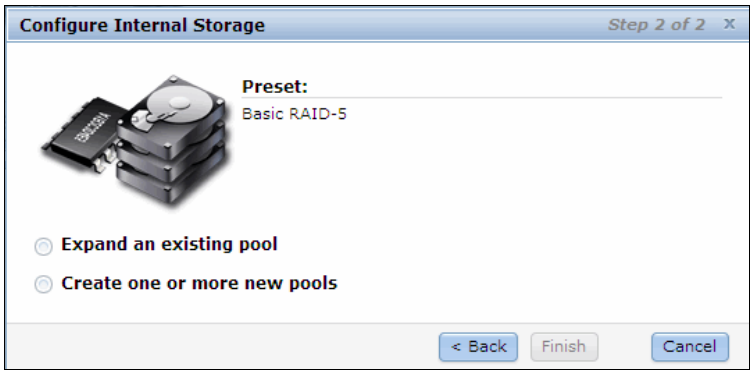

Figure 7-34 Storage pool selection

Complete the following steps to expand or create a pool:

a. Expand an existing pool.

   When an existing pool is to be expanded, you can select an existing storage pool that does not contain MDisks, or a pool that contains MDisks with the same performance characteristics (which is listed automatically), as shown in Figure 7-35 on page 319.

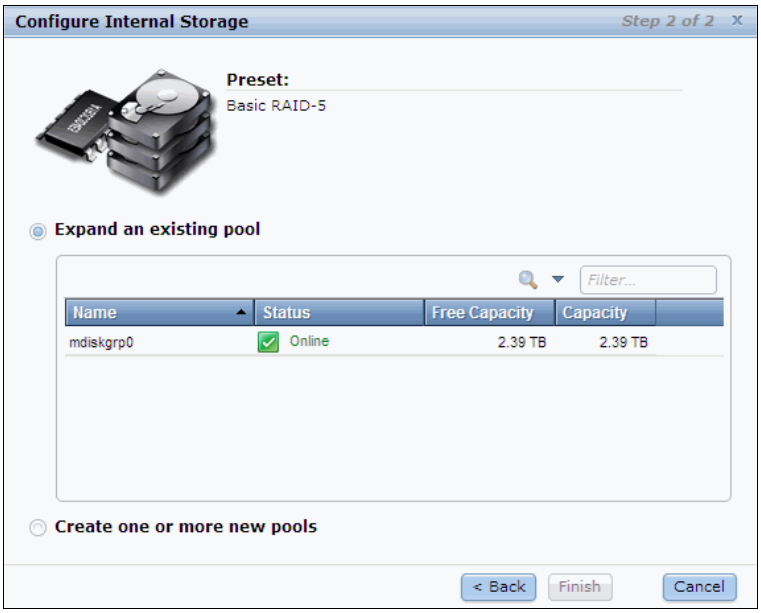

Figure 7-35 List of matching storage pool

b. Create one or more pools.

Alternatively, a storage pool is created by entering the required name, as shown in Figure 7-36.

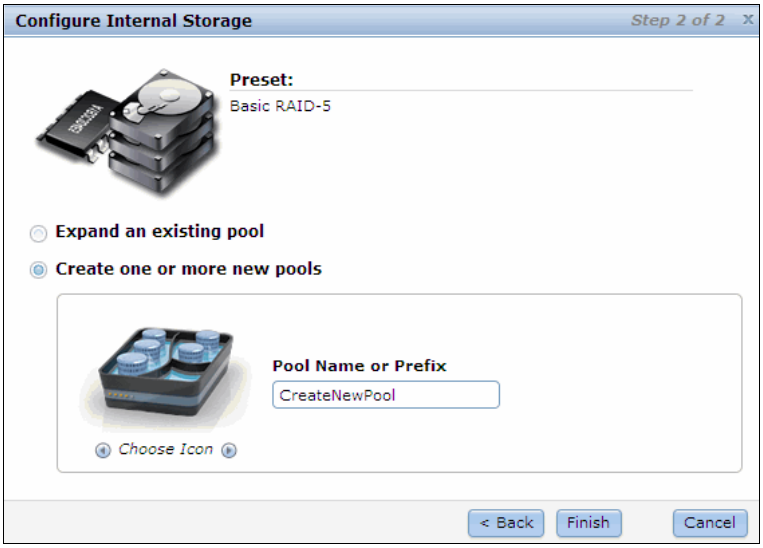

Figure 7-36 Create new pool

All drives are initialized when the Configuration wizard is finished.

7.3 Working with MDisks on internal and external storage

After the configuration is complete for the internal storage, you can find the MDisks that were created on the internal drives in the MDisks by Pools window.

You can access the MDisks window by clicking **Home** → **Overview** and then clicking the **MDisks** function icon. In the extended help information window, click **Pools**, as shown in Figure 7-37.

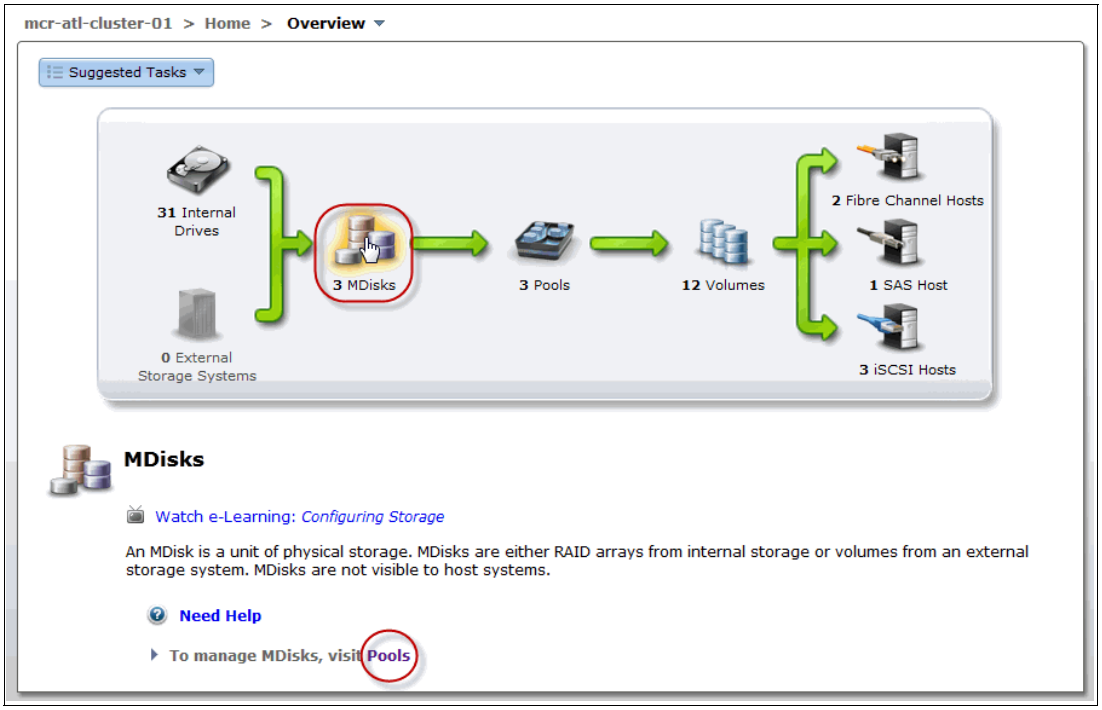

Figure 7-37 MDisk from Overview window

An alternative way to access the MDisks window is by using the **Pools** function icon and selecting **MDisk by Pools**, as shown in Figure 7-38.

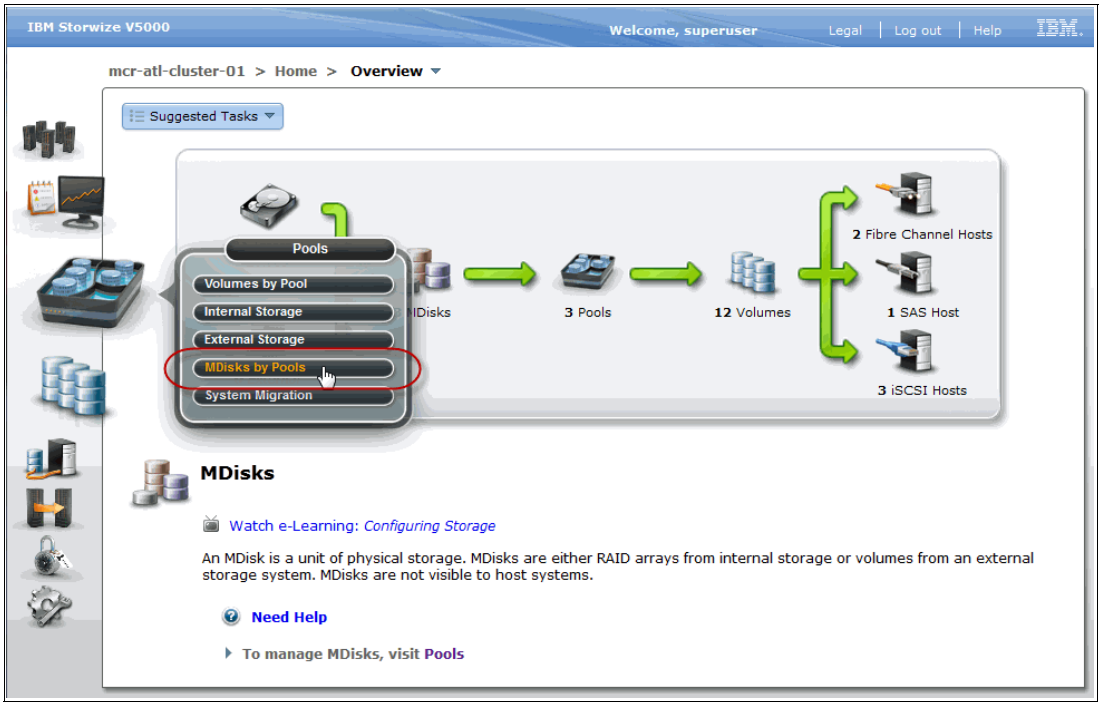

Figure 7-38 MDisk from Pools icon

By using the MDisks by Pools window, you can manage all MDisks that are made up of internal and external storage. Figure 7-39 shows internal and externally virtualized MDisks. In this example, the MDisks that are associated with the storage system (DS3400) are externally virtualized on an IBM DS3400 system.

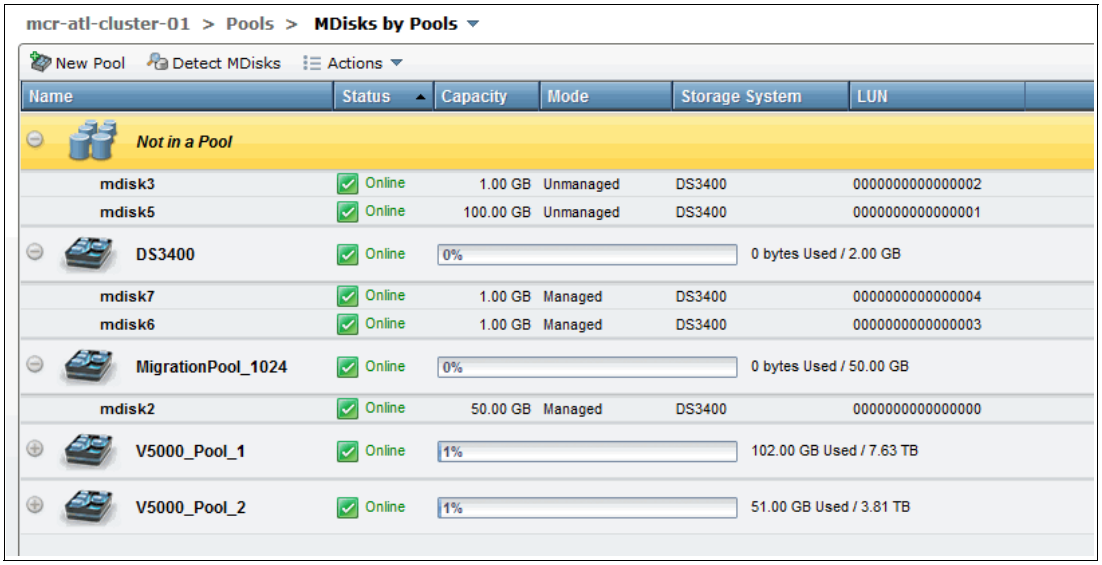

Figure 7-39 MDisks by Pools window

The window provides the following information:

- MDisk name
- Status
- Capacity
- Mode
- Name of the storage pool it belongs to
- Name of the backing storage system for MDisk on external storage
- MDisk's LUN ID from external storage systems
- Assigned storage tier

In IBM Storwize V5000, an MDisk features the following MDisk modes:

- Array

  Array mode MDisks are constructed from internal drives by using the RAID functionality. Array MDisks are always associated with storage pools.

- Unmanaged

  Logical Unit Numbers (LUNs) that are presented by external storage systems to IBM Storwize V5000 are discovered as unmanaged MDisks. The MDisk is not a member of any storage pools, which means it is not used by the IBM Storwize V5000 storage system.

- Managed

  Managed MDisks are LUNs that are presented by external storage systems to an IBM Storwize V5000 that are assigned to a storage pool and provide extents so that volumes can use it. Any data that was on these LUNs when they are imported is lost.

- Image

  Image MDisks are LUNs presented by external storage systems to an IBM Storwize V5000 and assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. For more information, see Chapter 6, "Storage migration wizard" on page 237.

For more information about attaching, zoning, and presenting external storage to the IBM Storwize V5000, see Chapter 11, "External storage virtualization" on page 547.

Externally virtualized storage can be used on an IBM Storwize V5000 in one of the following ways:

- Create empty LUNs on the external storage that is seen as unmanaged MDisks when they are presented to the IBM Storwize V5000. These MDisks can then be added to existing or new storage pools. If existing LUNs are used, any data on these LUNs is lost.
- Use existing LUNs on the external storage that is seen as unmanaged MDisks when they are presented to the IBM Storwize V5000. These MDisks can then be imported into an existing storage pool or a storage pool that is created. Any data on these LUNs is preserved.

7.3.1 Adding Externally Virtualized MDisks to storage pools

By adding unmanaged MDisks to a pool, their status changes to Managed MDisks. Managed MDisks can belong to only one pool. Unmanaged MDisks can be added to a newly created pool or to an existing pool to expand its capacity. Pools are commonly used to group MDisks from the same storage subsystem.

A pool can be created in the MDisks by Pools window by clicking the **New Pool** icon. Assign a name to the pool and choose an icon, if wanted, as shown in Figure 7-40 on page 323.

Existing data: If there is existing data on the unmanaged MDisks that you must preserve, do not use the Add to Pool feature because this action deletes data. Use the Import feature instead, which is described in 7.3.2, “Importing externally virtualized MDisks to storage pools” on page 326.

Figure 7-40 Create Pool: Part 1

By using the **Create Pool** window (as shown in Figure 7-41), you can include unmanaged MDisks in the new pool. Several filter options at the top of the window with which you can limit the selection by storage subsystem, capacity, and so on. Several MDisks can be selected by pressing the Ctrl or Shift keys while you click the MDisks that are listed. Also, the Detect MDisks icon starts a SAN discovery for finding recently attached external storage systems.

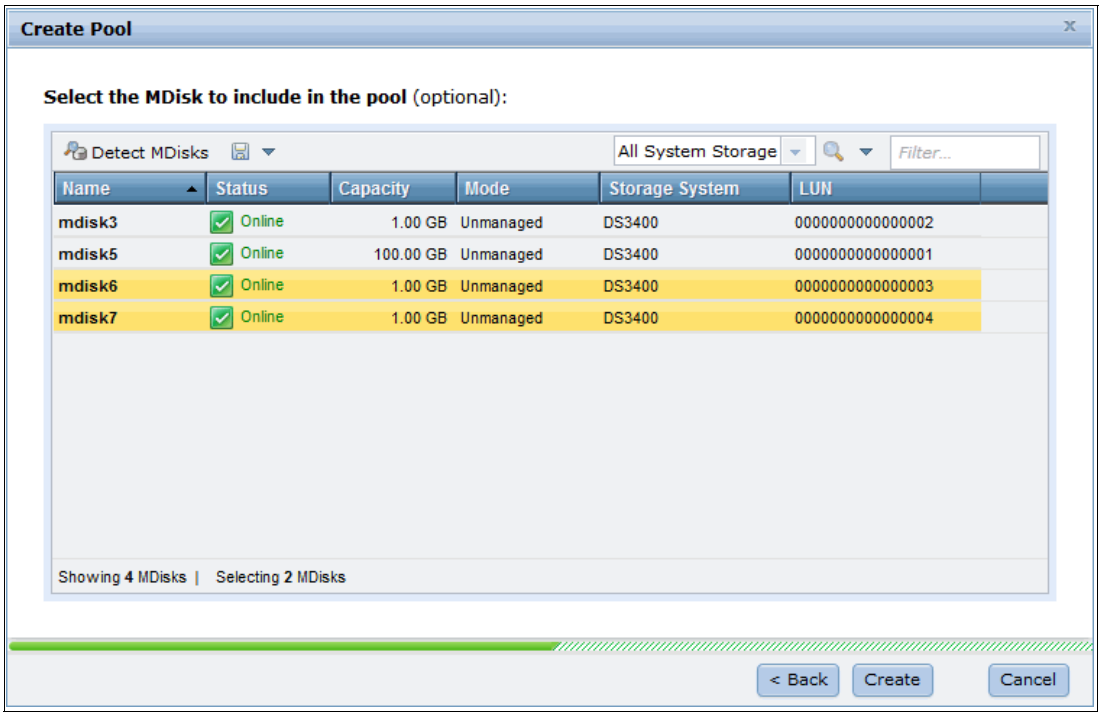

Figure 7-41 Create Pool: Part 2

To add unmanaged MDisks to an existing pool, select the MDisk from the Not in a Pool section, click **Actions** → **Add to Pool**, as shown in Figure 7-42 on page 324.

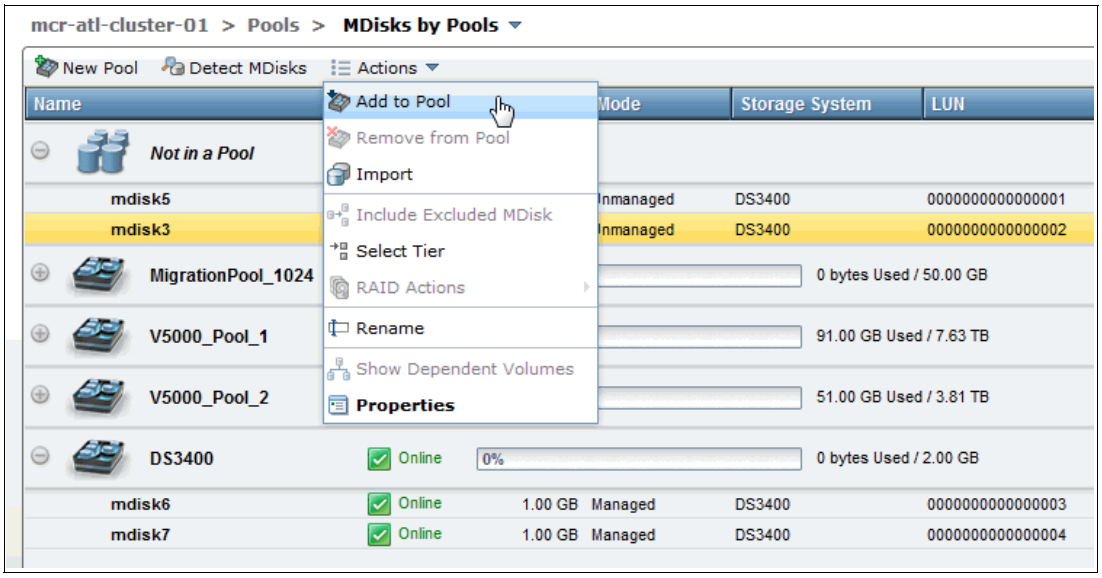

Figure 7-42 Add an unmanaged MDisk to a storage pool

Existing data: If there is existing data on the unmanaged MDisks that you must preserve, do not select Add to Pool on this LUN because this action deletes the data. Use the Import feature instead, which is described in 7.3.2, “Importing externally virtualized MDisks to storage pools” on page 326.

Choose the storage pool to which you want to add the MDisk and click **Add to Pool**, as shown in Figure 7-43.

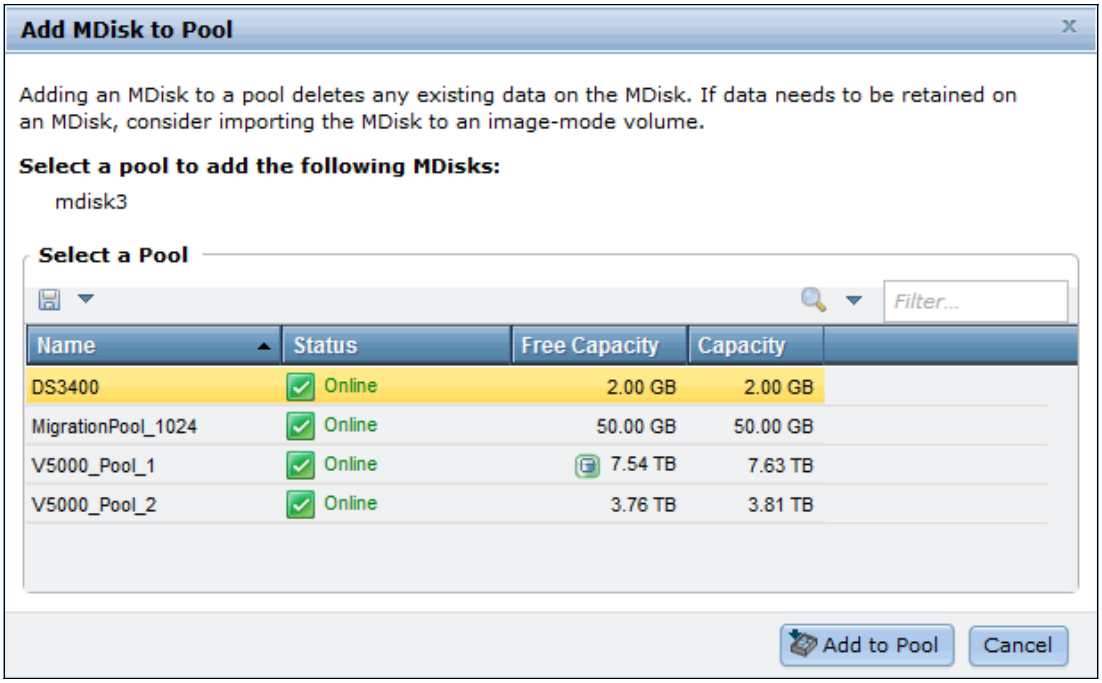

Figure 7-43 Add MDisk to pool

After the IBM Storwize V5000 system completes this action, the MDisk is shown in the pool to which it was added, as shown in Figure 7-44.

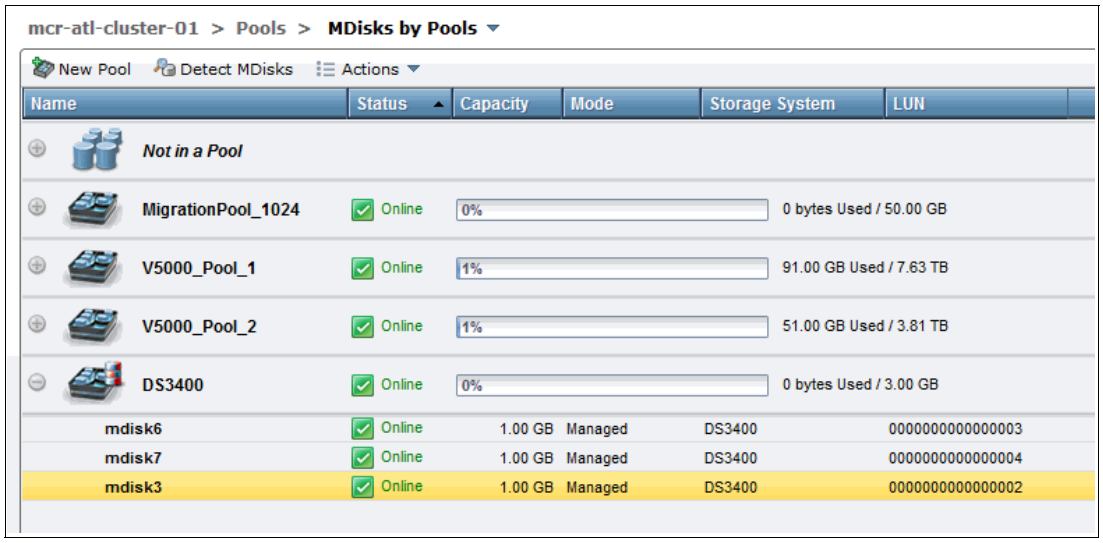

Figure 7-44 MDisk added to pool

In some cases, you might want to remove MDisks from storage pools to reorganize your storage allocation. You can remove MDisks from storage pools by selecting the MDisks and clicking **Remove from Pool** from the Actions drop-down menu, as shown in Figure 7-45.

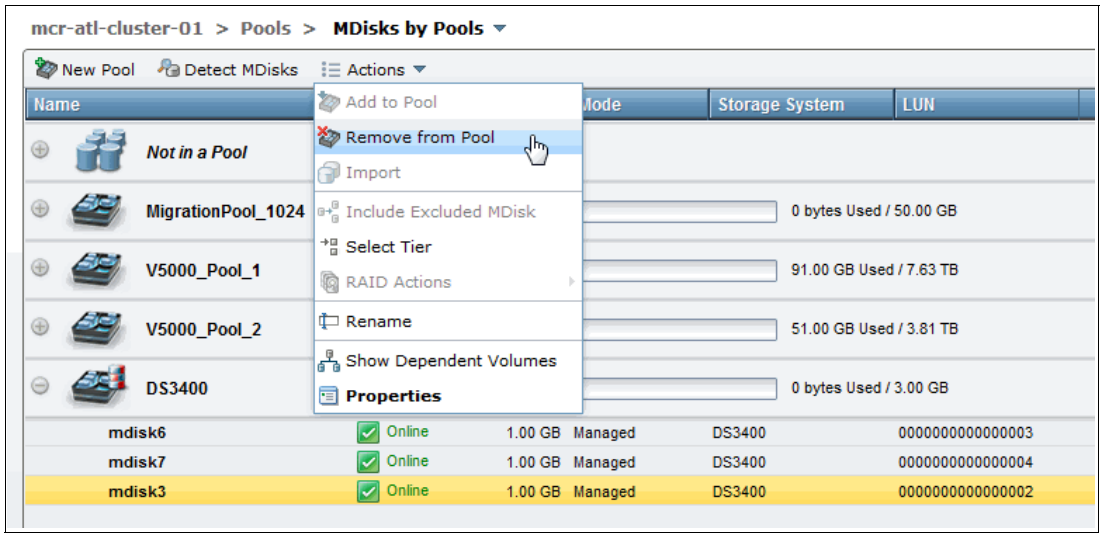

Figure 7-45 Remove an MDisk from the storage pool

You must confirm the number of MDisks that you want to remove, as shown in Figure 7-46 on page 326. If you have data on the MDisks and you still must remove the MDisks from the pool, select the **Remove the MDisk from the storage pool even if it has data on it. The system migrates the data to other MDisks in the pool** option.

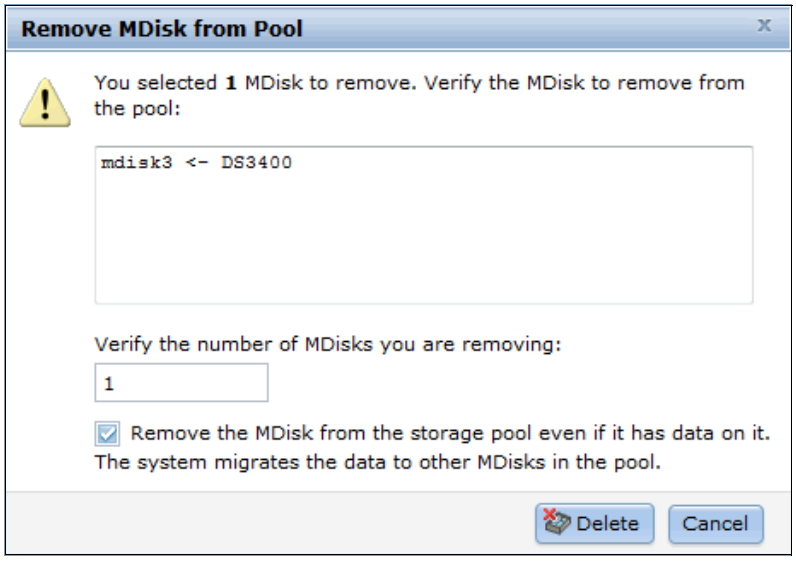

Figure 7-46 Confirm the removal of MDisk from the pool

Available capacity: Make sure that you have enough available capacity left in the storage pool for the data on the MDisks to be removed.

After you click **Delete**, data migration from the selected MDisk starts. You can find the migration progress in the Running Tasks status indicator, as shown in Figure 7-47.

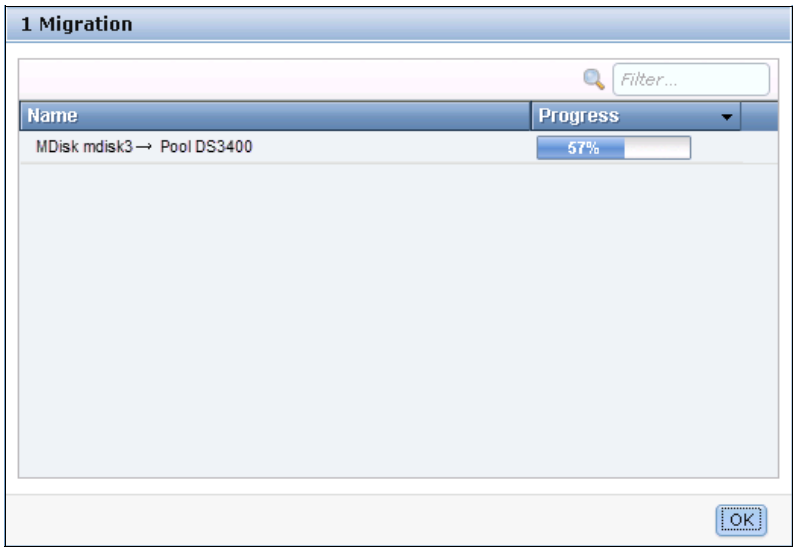

Figure 7-47 Data migration progress when MDisks are removed from the pool

7.3.2 Importing externally virtualized MDisks to storage pools

LUNs that are hosted on external storage systems can be imported into IBM Storwize V5000 storage. Hosts are used to be directly attached to these external storage systems. The hosts can continue to use their storage that is now presented through the IBM Storwize V5000.

To achieve this configuration, the existing external LUNs must be imported as an image-mode volume by using the Import option. This action is possible for unmanaged MDisks only. Those disks must not be added to a pool, as described in 7.3.1, “Adding Externally Virtualized MDisks to storage pools” on page 322.

If the Import option is used and no existing storage pool is chosen, a temporary *migration pool* is created to hold the new image-mode volume. This image-mode volume has a direct block-for-block translation from the imported MDisk to the volume and existing data is preserved.

Figure 7-48 shows an example of how to import an unmanaged MDisk. Select the unmanaged MDisk and click **Import** from the Actions drop-down menu.

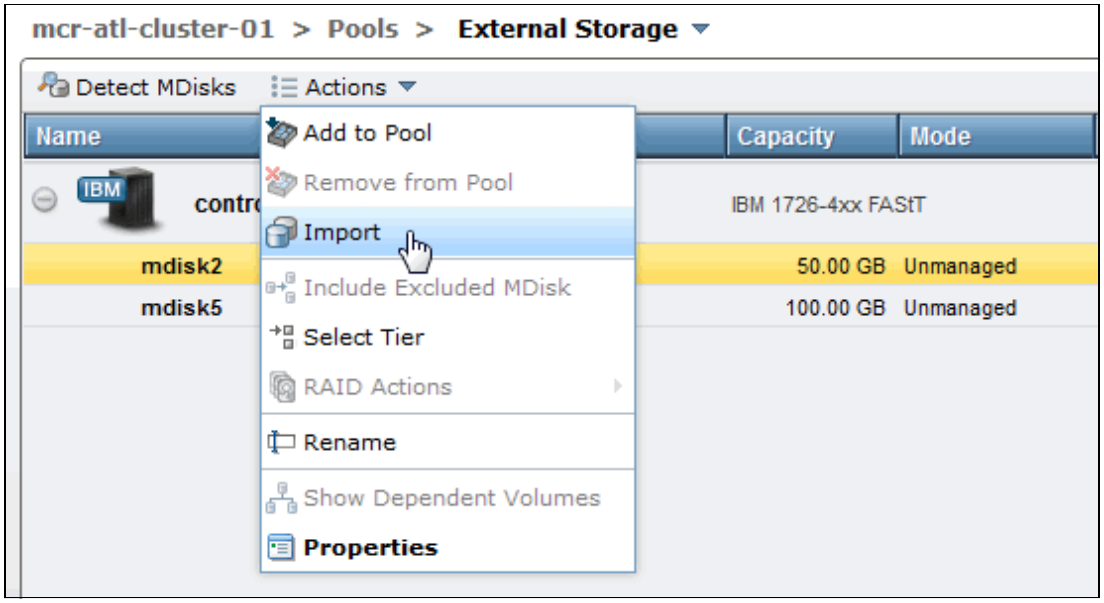

Figure 7-48 Import MDisk

As shown in Figure 7-49, the Import wizard starts and then guides you through the import process.

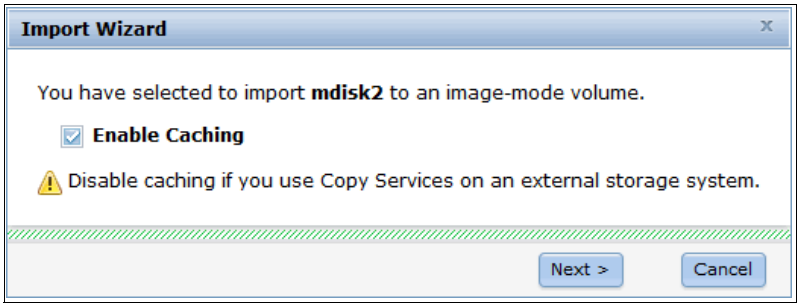

Figure 7-49 Import wizard: Step 1

In step 1 of the Import wizard, caching for the volume can be disabled; it is enabled by default.

Clear the Enable Caching option if you use copy services on the external storage system that is hosting the LUN. It is a best practice to use the copy services of IBM Storwize V5000 for virtualized volumes. For more information about virtualizing external storage, see in Chapter 11, “External storage virtualization” on page 547. For more information about exporting volumes, see in Chapter 8, “Advanced host and volume administration” on page 349.

Figure 7-50 shows step 2 of the Import wizard, which includes the option to import the MDisk into an existing pool or a temporary pool.

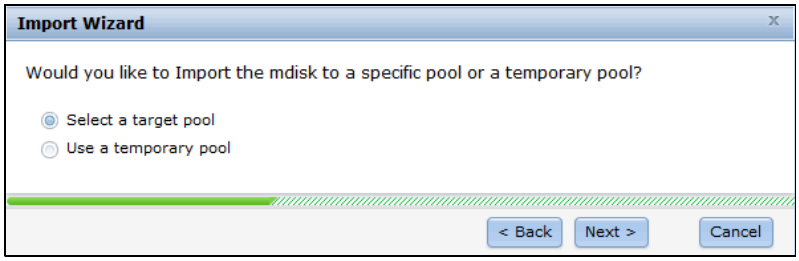

Figure 7-50 Import wizard: Step 2

If you select the option to import the MDisk to an existing pool, click **Next** and you see step 3 of the Import wizard (as shown in Figure 7-51), which includes the option to choose an existing destination storage pool (only pools with sufficient available capacity are listed). The actual data migration begins after the MDisk is imported successfully.

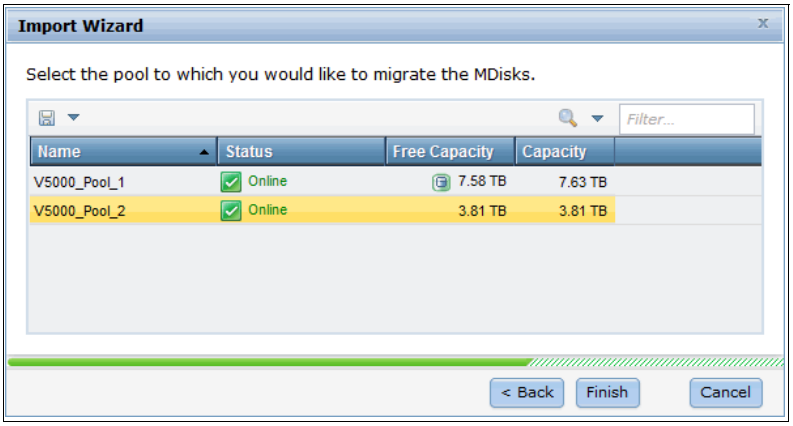

Figure 7-51 Import wizard: Step 3

You can check the migration progress in the Running Tasks status indicator (as shown in Figure 7-52) or by clicking **Pools** → **System Migration**.

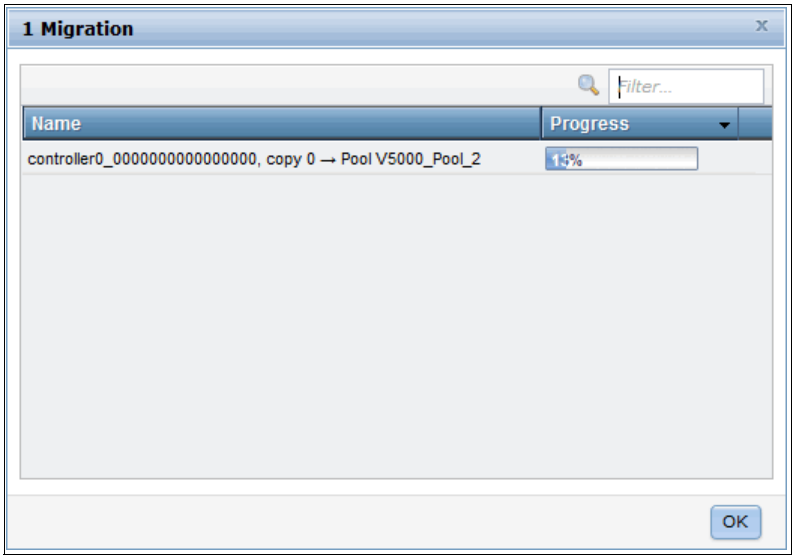

Figure 7-52 Migration progress in the status indicator of Running Tasks

After the migration completes, you can find the volume in the chosen destination pool, as shown in Figure 7-53.

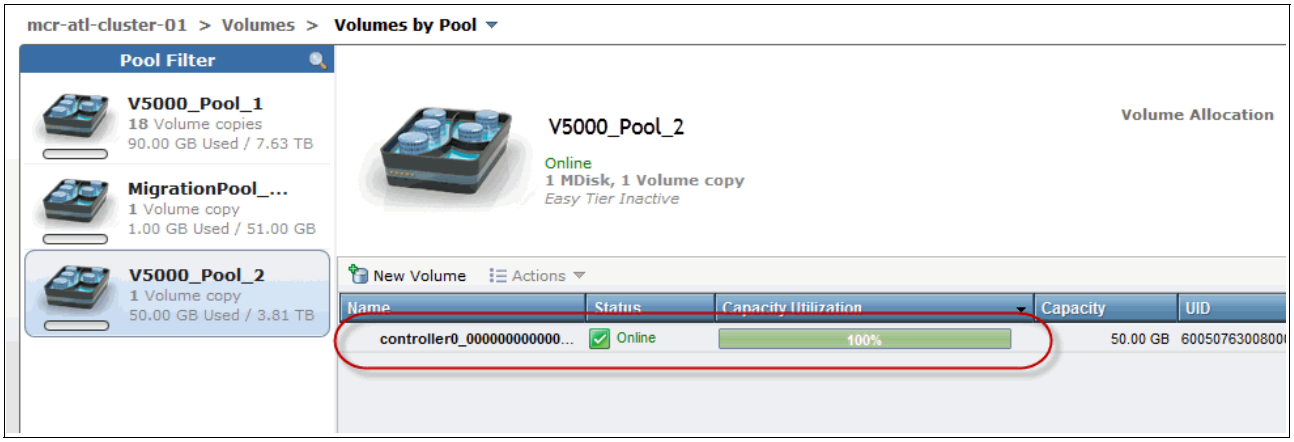

Figure 7-53 Volume migrated to destination pool

All data is migrated off the source MDisk to MDisks in the destination storage pool. The source MDisk changed its status to managed and is associated with an automatically created migration pool. It can be used as a regular MDisk to host volumes, as shown in Figure 7-54.

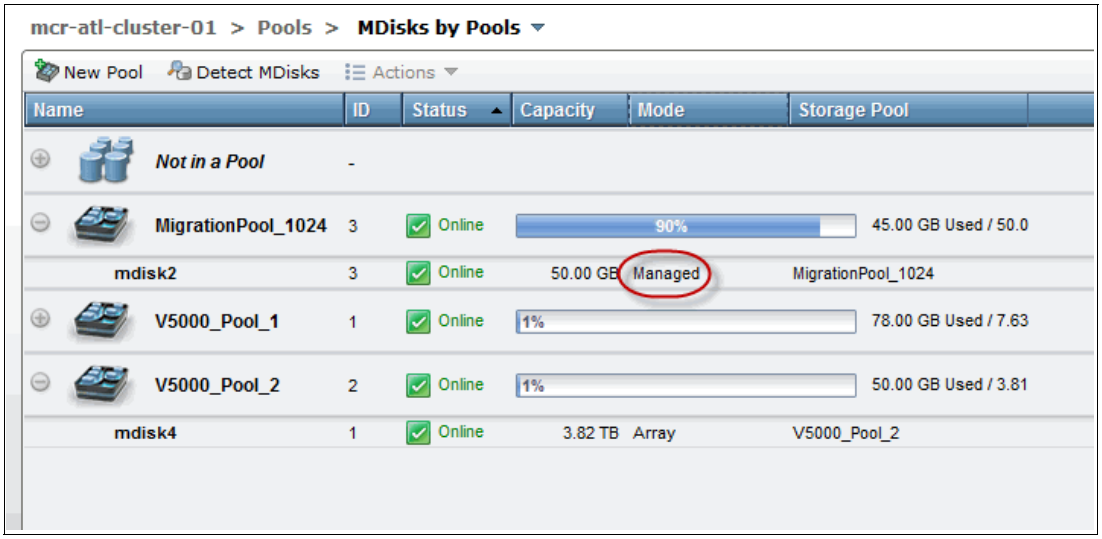

Figure 7-54 MDisk mode that is changed to managed

If you selected the Use a temporary Pool option, the MDisk is imported in step 2 of the Import wizard. The window that is shown in Figure 7-55 opens in which you can specify the extent size of the temporary pool. If you are planning to manually migrate this MDisk to a different pool later, choose the extent size to match that pool.

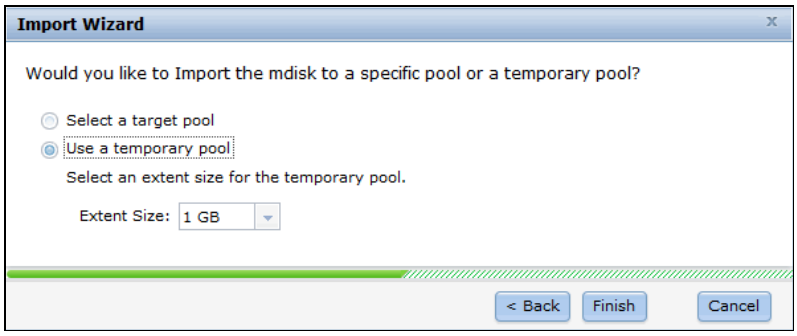

Figure 7-55 Import MDisk to a temporary pool

The imported MDisk remains in its temporary storage pool as an image mode volume, as shown in Figure 7-56 on page 331.

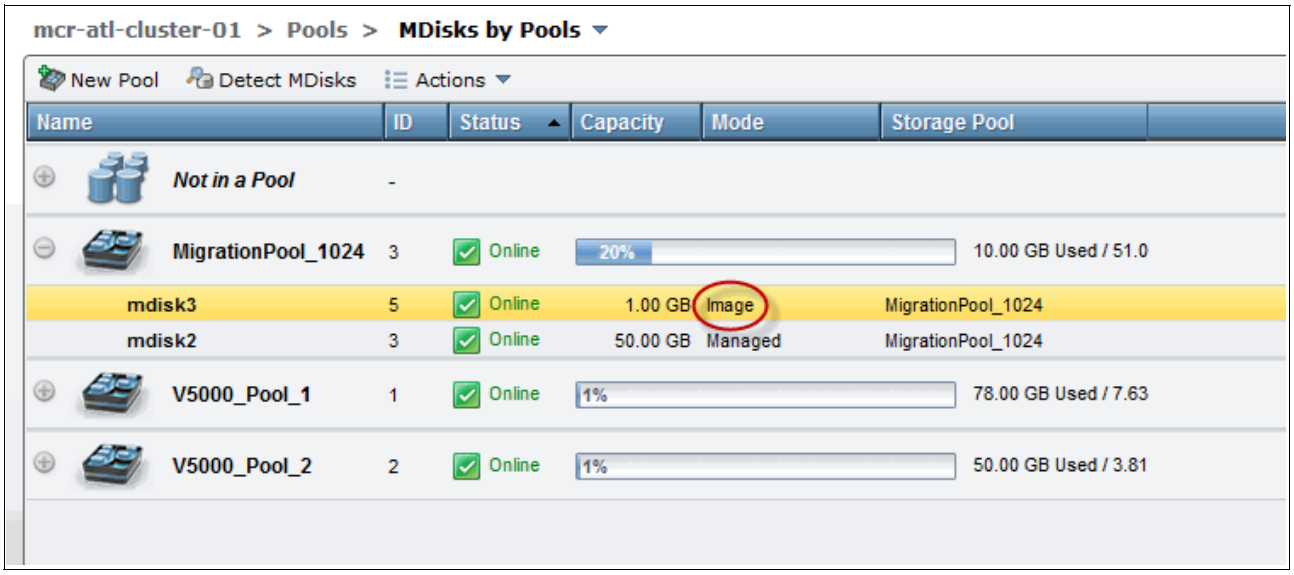

Figure 7-56 MDisk after import

If needed, the image mode volume can be migrated manually into a different pool by selecting **Migration to Another Pool** or **Volume Copy Actions**. For more information about volume actions, see Chapter 5, “I/O Group basic volume configuration” on page 161.

Alternatively, the migration into another pool can be done by clicking **Pools** → **System Migration**. For more information about migration, see Chapter 6, “Storage migration wizard” on page 237.

Any imported MDisk that was not migrated into a pool is listed under **Pools** → **System Migration**, as shown in Figure 7-57.

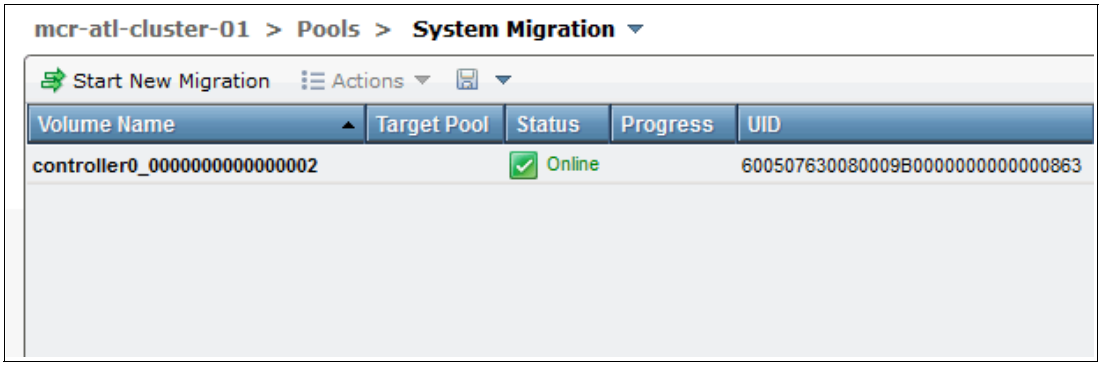

Figure 7-57 Imported MDisk in the System Migration window

This feature is normally used as a vehicle to migrate data from existing external LUNs into storage pools that are located internally or externally on the IBM Storwize V5000. You should not use image mode volumes as a long-term solution for reasons of performance and reliability.

To migrate an image mode volume into a regular storage pool, select the volume to be migrated and click **Actions** → **Migrate to Another Pool**. Choose the required target storage pool to migrate the data into and click **Migrate**, as shown in Figure 7-58 on page 332.

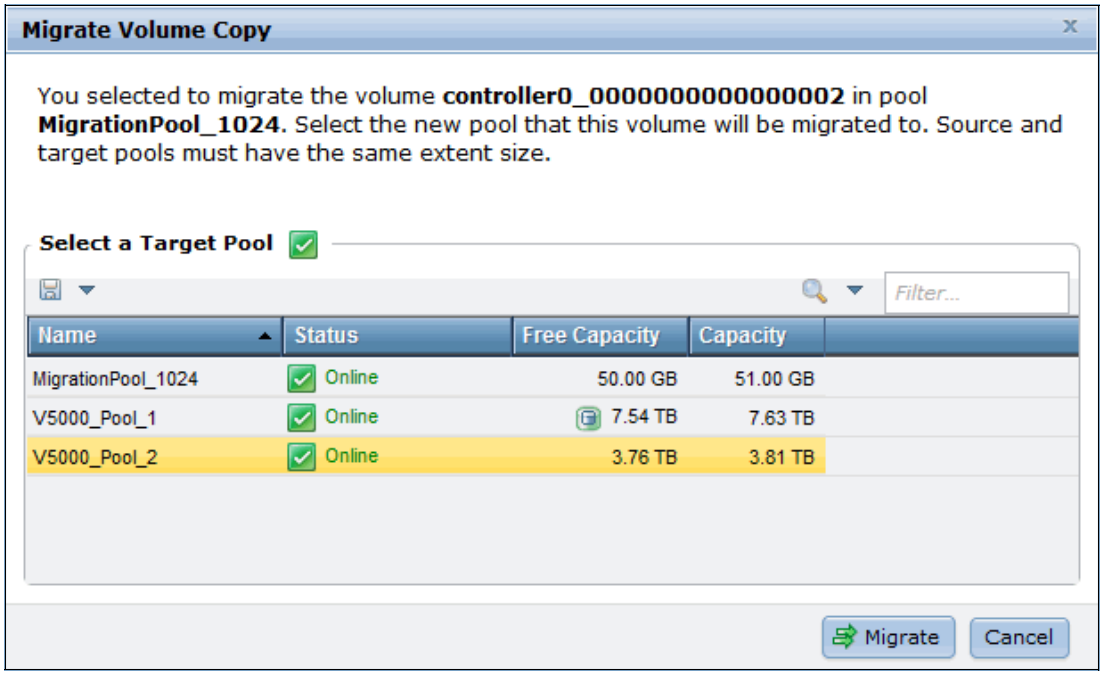

Figure 7-58 Migrate Image Mode Volume into a regular storage pool

The migration internally uses the volume copy function, which creates a second copy of the existing volume in the chosen target pool. For more information about the volume copy function, see Chapter 8, “Advanced host and volume administration” on page 349.

The original volume copy on the image mode MDisk is deleted and the newly created copy is kept.

7.3.3 MDisk by Pools panel

The MDisks by Pools panel (as shown in Figure 7-59) displays information about all MDisks made of internal and external storage. The MDisks are categorized by the pools to which they are attached.

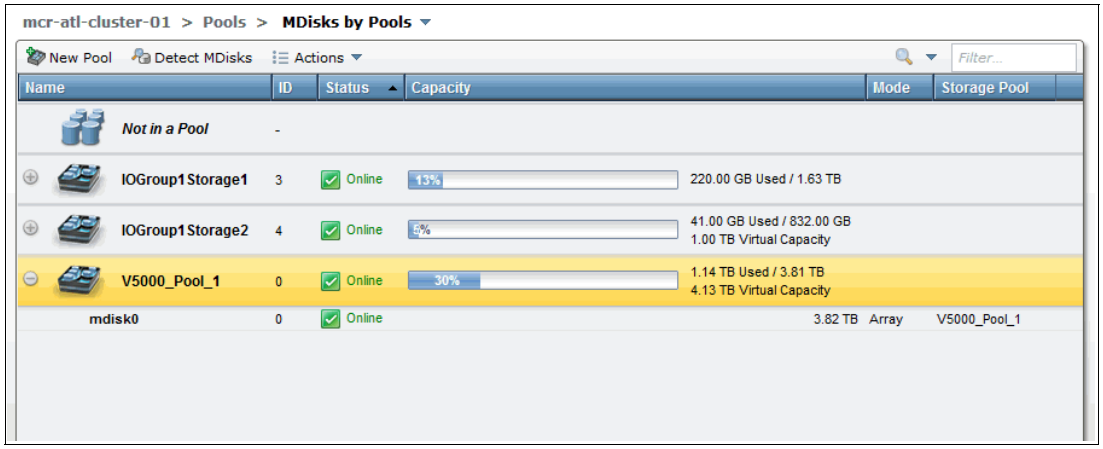

Figure 7-59 MDisk by Pool window

The following default information is provided:

- Name

  The MDisk or the storage pool name that is provided during the configuration process.

- ID

  The MDisk or storage pool ID that is automatically assigned during the configuration process.

- Status

  The status of the MDisk and storage pool. The following statuses are possible:

  – Online

    All MDisks are online and performing optimally.

  – Degraded

    One MDisk is in degraded state (for example, missing SAS connection to enclosure of member drives or a failed drive with no spare available). As shown in Figure 7-60, the pool also is degraded.

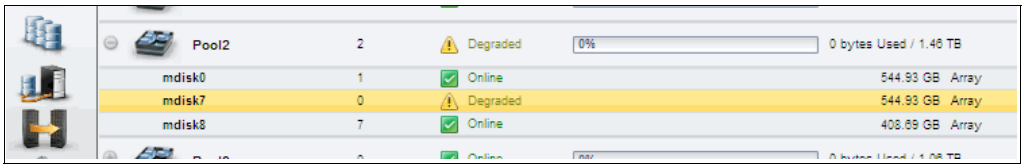

Figure 7-60 One degraded MDisk in pool

  – Offline

    One or more MDisks in a pool are offline. The pool (Pool3) also changes to offline, as shown in Figure 7-61.

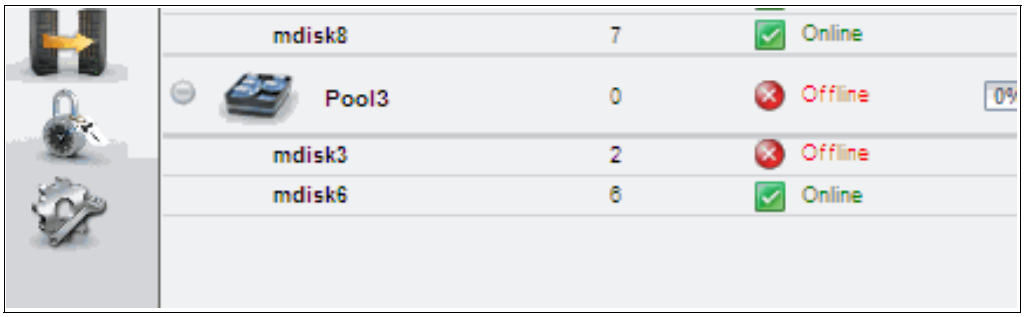

Figure 7-61 Offline MDisk in a pool

- Capacity

  The capacity of the MDisk. The capacity is shown for the storage pool, which is the total of all the MDisks in this storage pool. The usage of the storage pool is represented by a bar and the number.

- Mode

  An MDisk features the following modes:

  – Array

    Array mode MDisks are constructed from internal drives by using the RAID functionality. Array MDisks are always associated with storage pools.

– Unmanaged

  LUNs that are presented by external storage systems to IBM Storwize V5000 are discovered as unmanaged MDisks. The MDisk is not a member of any storage pools, which means it is not used by the IBM Storwize V5000 storage system.

– Managed

  Managed MDisks are LUNs that are presented by external storage systems to an IBM Storwize V5000 that are assigned to a storage pool and provide extents so that volumes can use it. Any data that was on these LUNs when they are imported is lost.

– Image

  Image MDisks are LUNs that are presented by external storage systems to an IBM Storwize V5000 and assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This status is an intermediate status of the migration process and is described in Chapter 6, “Storage migration wizard” on page 237.

- Storage Pool

  The name of the storage pool to which the MDisk belongs.

For more information about how to attach external storage to an IBM Storwize V5000 storage system, see in Chapter 11, “External storage virtualization” on page 547.

The CLI command **`lsmdiskgrp`** returns a concise list or a detailed view of the storage pools that are visible to the system, as shown in Example 7-5.

Example 7-5 CLI command lsmdiskgrp

```
lsmdiskgrp
lsmdiskgrp mdiskgrpID
```

7.3.4 RAID action for MDisks

Internal drives in the IBM Storwize V5000 are managed as Array mode MDisks, on which several RAID actions can be performed. Select the appropriate Array MDisk by clicking **Pools** → **MDisks by Pools**, and then click **Actions** → **RAID Actions**, as shown in Figure 7-62.

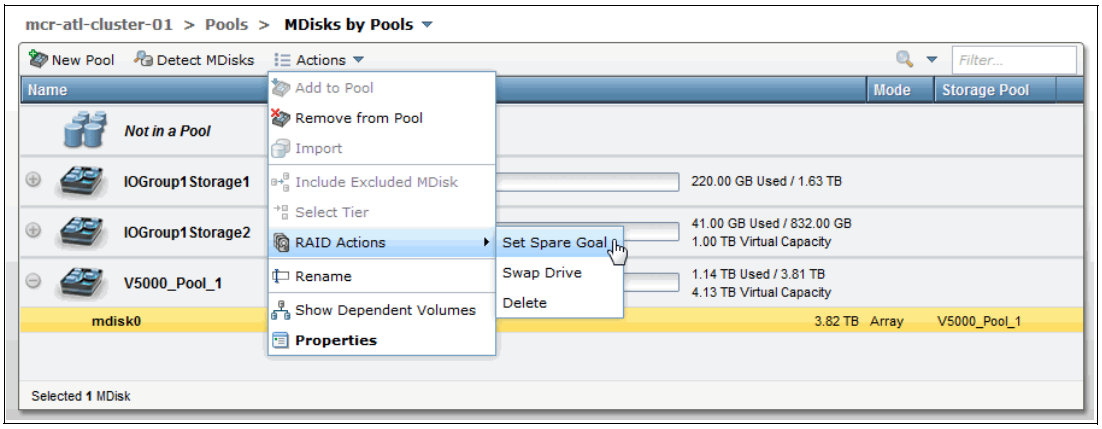

Figure 7-62 MDisk RAID actions

You can choose the following RAID actions:

- Set Spare Goal

  Figure 7-63 shows how to set the number of spare drives that are required to protect the array from drive failures.

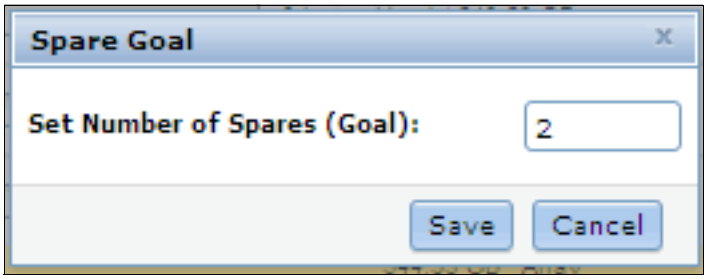

Figure 7-63 MDisk set spare goal

The alternative CLI command is shown in Example 7-6.

Example 7-6 CLI command to set spares

```
charray -sparegoal mdiskID goal
```

If the number of drives that are assigned as Spare does not meet the configured spare goal, an error is logged in the event log that reads: “Array MDisk is not protected by sufficient spares.” This error can be fixed by adding more drives as spares. During the internal drive configuration, spare drives are automatically assigned according to the chosen RAID preset’s spare goals, as described in 7.2, “Configuring internal storage” on page 307.

- Swap Drive

  The Swap Drive action can be used to replace a drive in the array with another drive with the status of Candidate or Spare. This action is used to replace a drive that failed, or is expected to fail soon; for example, as indicated by an error message in the event log. Select an MDisk that contains the drive to be replaced and click **RAID Actions** → **Swap Drive**. In the Swap Drive window, select the member drive that is replaced (as shown in Figure 7-64 on page 336) and click **Next**.

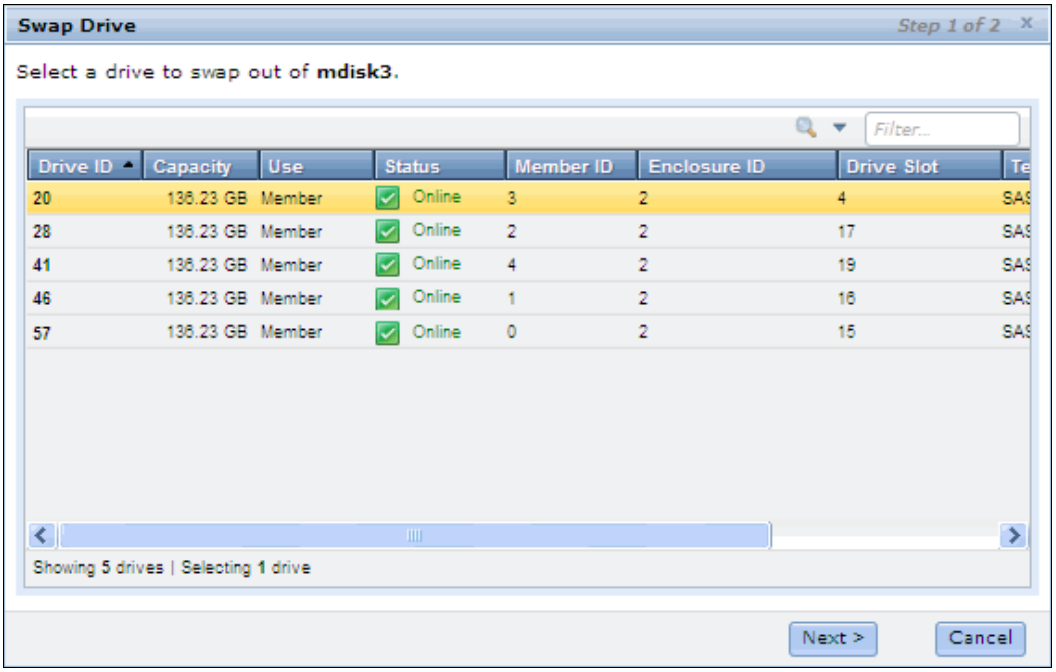

Figure 7-64 MDisk swap drive: Step 1

In step 2 (as shown as Figure 7-65), a list of suitable drives is presented. One drive must be selected to swap into the MDisk. Click **Finish**.

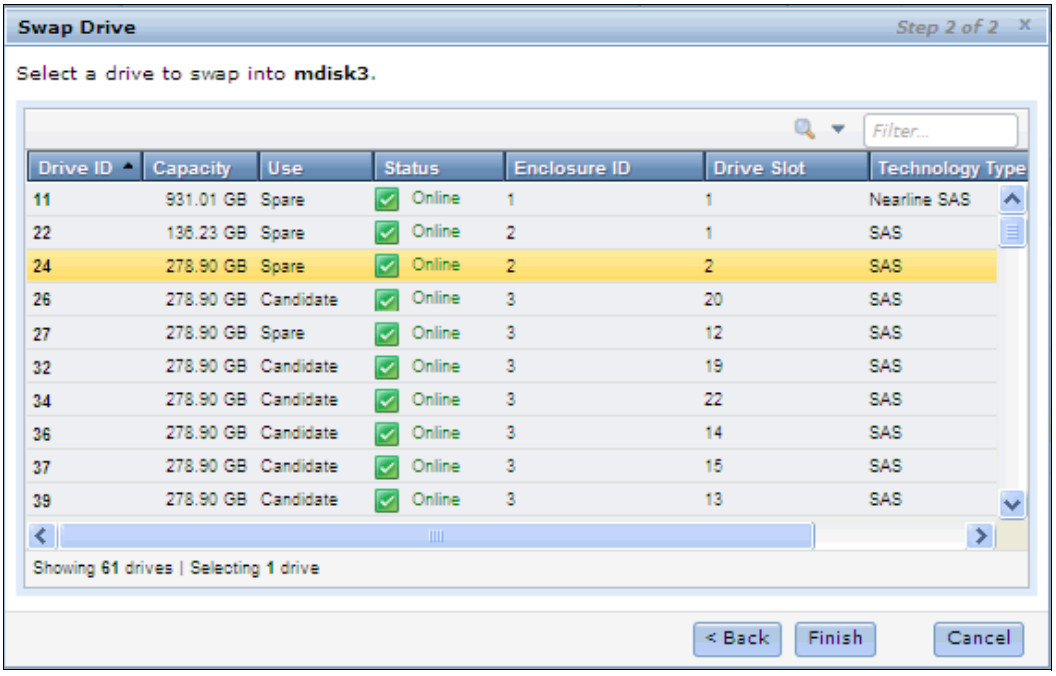

Figure 7-65 MDisk swap drive: Step 2

The exchange process starts and then runs in the background. The volumes on the affected MDisk remain accessible.

If the GUI process is not used for any reason, the CLI command in Example 7-7 on page 337 can be run.

Example 7-7 CLI command to swap drives

```
charraymember -balanced -member oldDriveID -newdrive newDriveID mdiskID
```

- Delete

An Array MDisk can be deleted by clicking **RAID Actions** → **Delete**. To select more than one MDisk, press Ctrl+left-mouse click. A confirmation is required by entering the correct number of MDisks to be deleted, as shown in Figure 7-66. You must confirm the number of MDisks that you want to delete. If there is data on the MDisk, it can be deleted only by tagging the option Delete the RAID array MDisk even if it has data on it. The system migrates the data to other MDisks in the pool.

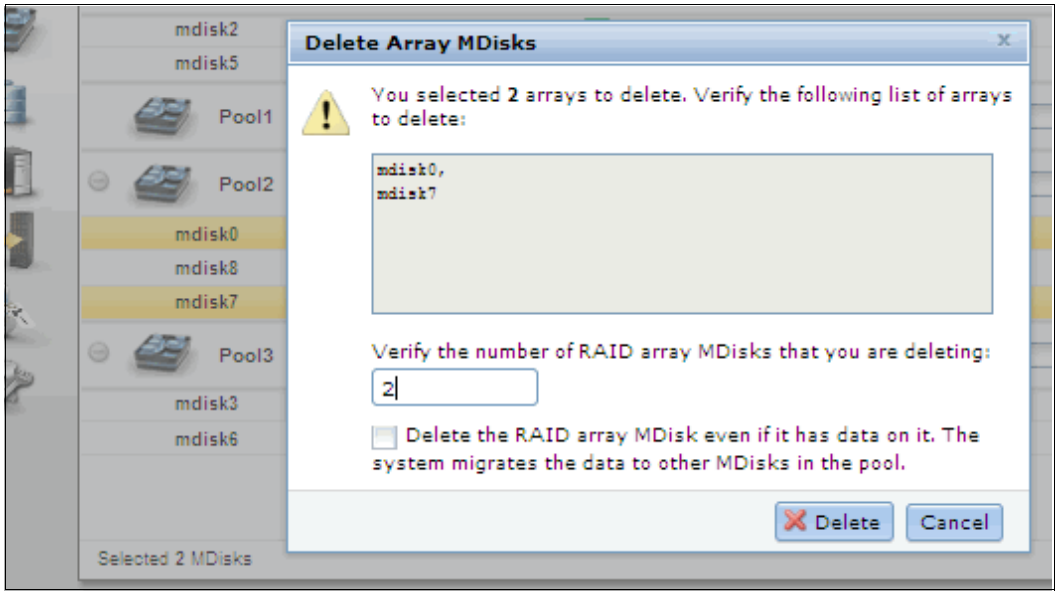

Figure 7-66 MDisk delete confirmation

Data that is on MDisks is migrated to other MDisks in the pool if enough space is available on the remaining MDisks in the pool.

Available capacity: Make sure that you have enough available capacity left in the storage pool for the data on the MDisks to be removed.

After an MDisk is deleted from a pool, its former member drives return to candidate mode.

The alternative CLI command to delete MDisks is shown in Example 7-8.

Example 7-8 CLI command to delete MDisk

```
rmmdisk -mdisk list -force mdiskgrpID
```

If all the MDisks of a storage pool were deleted, the pool remains as an empty pool with 0 bytes of capacity, as shown in Figure 7-67.

Figure 7-67 Empty storage pool after MDisk deletion

7.3.5 Selecting the drive tier for externally virtualized MDisks

The IBM Storwize V5000 Easy Tier feature is described in Chapter 9, “Easy Tier” on page 411. In this section, we show how to adjust the tier settings.

The following tiers are available:

- Generic SSD tier for storage that is made of SSDs, which is the faster-performing storage.
- Generic HDD tier for everything else.

Internal drives have their tier assigned automatically by the IBM Storwize V5000. MDisks on external storage systems are assigned the generic HDD tier by default. This setting can be changed manually by the user. To assign a specific tier to an MDisk, click **Pools** → **MDisks by Pool** and click **Select Tier** from the Actions drop-down menu, as shown in Figure 7-68.

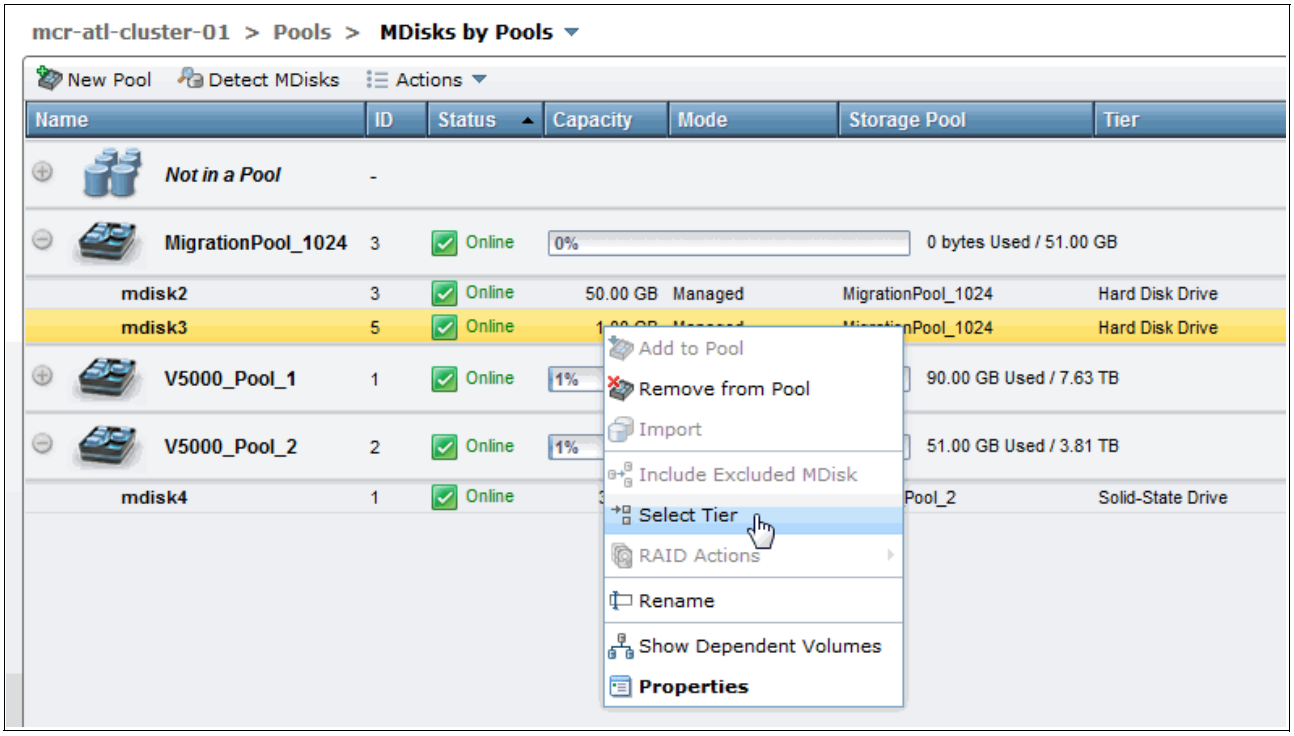

Figure 7-68 Select Tier for an MDisk

For demonstration purposes, we assign the tier SSD to `mdisk3`, as shown in Figure 7-69. This MDisk is a LUN made of SAS HDDs in an external storage system. The tier that was assigned by default is Hard Disk Drive.

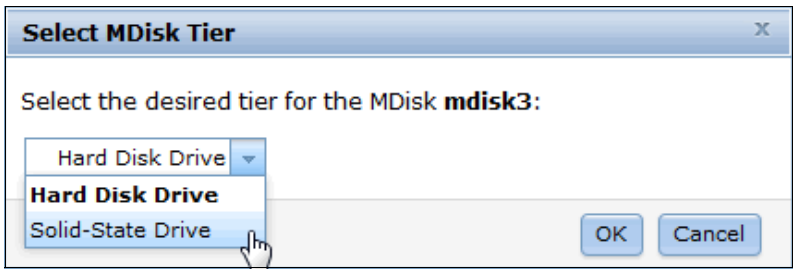

Figure 7-69 Assign wanted tier to an MDisk

After the action completes successfully, the MDisk can be found in the SSD tier, as shown in Figure 7-70.

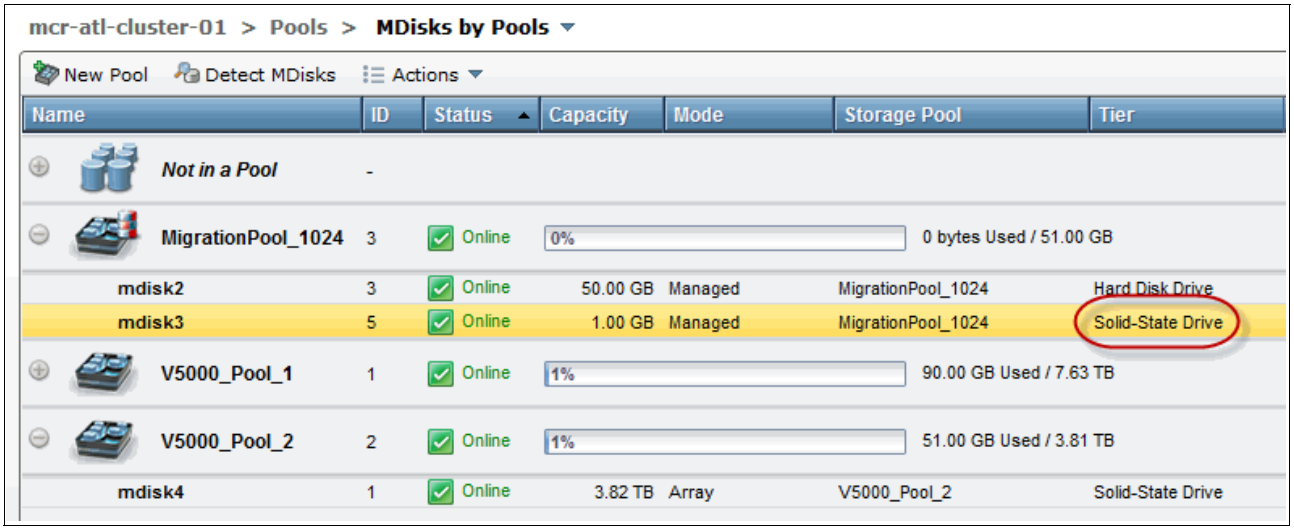

Figure 7-70 Wanted tier that is assigned to the MDisk

7.3.6 More actions on MDisks

The following actions can be performed on MDisks:

- Detect MDisks

  The Detect MDisks button at the upper left of the MDisks by Pools window is useful if you have external storage controllers in your environment (for more information, see Chapter 11, “External storage virtualization” on page 547). The Detect MDisk action starts a rescan of the Fibre Channel network. It discovers any new MDisks that were mapped to the IBM Storwize V5000 storage system and rebalances MDisk access across the available controller device ports. This action also detects any loss of controller port availability and updates the IBM Storwize V5000 configuration to reflect any changes.

  When external storage controllers are added to the IBM Storwize V5000 environment, the IBM Storwize V5000 automatically discovers the controllers and the LUNs that are presented by those controllers are listed as unmanaged MDisks. However, if you attached new storage and the IBM Storwize V5000 did not detect it, you might need to use the Detect MDisk button before the system detects the new LUNs. If the configuration of the external controllers is modified afterward, the IBM Storwize V5000 might be unaware of these configuration changes. Use the Detect MDisk button to rescan the Fibre Channel network and update the list of unmanaged MDisks.

  Figure 7-71 on page 340 shows the Detect MDisks button.

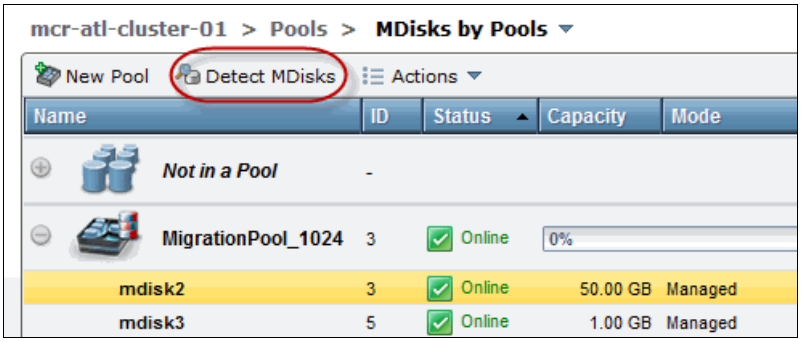

Figure 7-71 Detect MDisks

MDisks detection: The Detect MDisks action is asynchronous. Although the task appears to be finished, it still might be running in the background.

- Include Excluded MDisks

  An MDisk can be excluded from the IBM Storwize V5000 because of multiple I/O failures. These failures might be caused, for example, by link errors. After a fabric-related problem is fixed, the excluded disk can be added back into the IBM Storwize V5000 by selecting the MDisks and clicking **Include Excluded MDisk** from the Actions drop-down menu.

Some of the other actions are available by clicking **MDisk by Pool** → **Actions**, as shown in Figure 7-72.

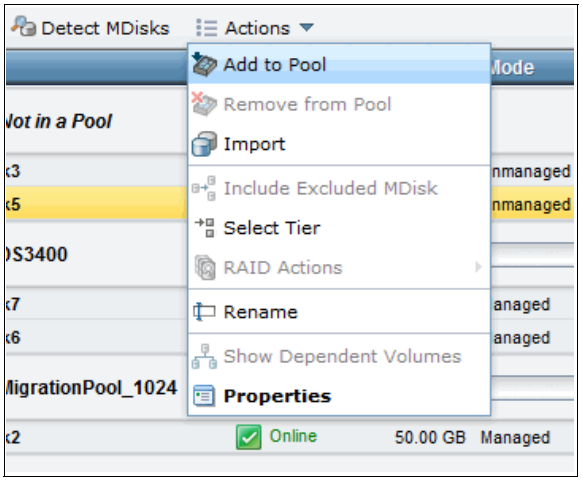

Figure 7-72 MDisk actions on externally virtualized storage

Rename

MDisks can be renamed by selecting the MDisk and clicking **Rename** from the Actions menu. Enter the new name of your MDisk (as shown in Figure 7-73) and click **Rename**.

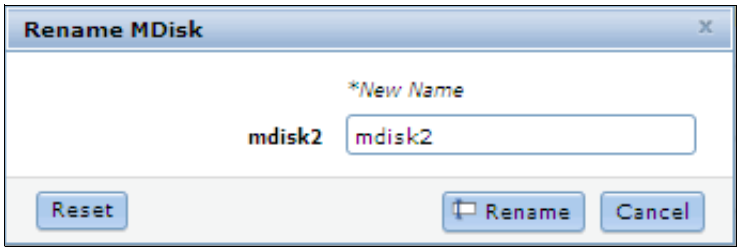

Figure 7-73 Rename MDisk

Show Dependent Volumes

Figure 7-74 shows the volumes that are dependent on an MDisk. The volumes can be displayed by selecting the MDisk and clicking **Show Dependent Volumes** from the Actions menu. The volumes are listed with general information.

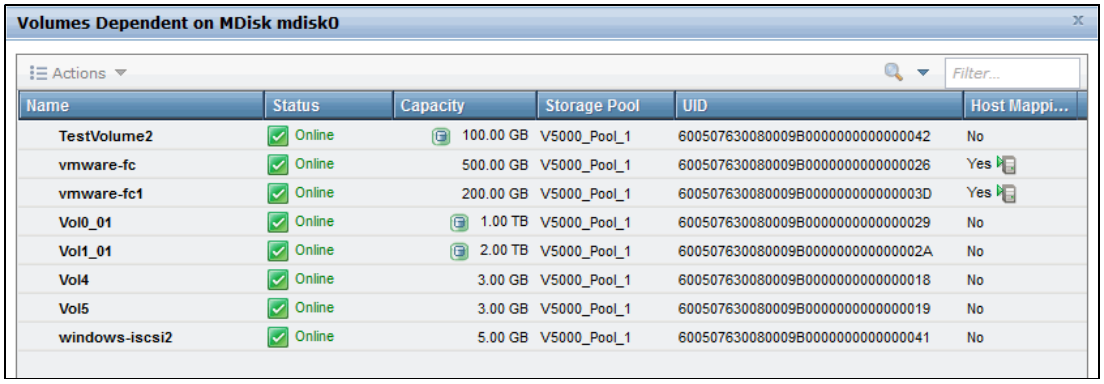

Figure 7-74 Show dependent volumes

Properties

The Properties action for an MDisk shows the information that you need to identify it. In the MDisks by Pools window, select the MDisk and click **Properties** from the Actions menu. The following tabs are available in this information window:

- The Overview tab (as shown in Figure on page 342) contains information about the MDisk. To show more details, click **Show Details**.

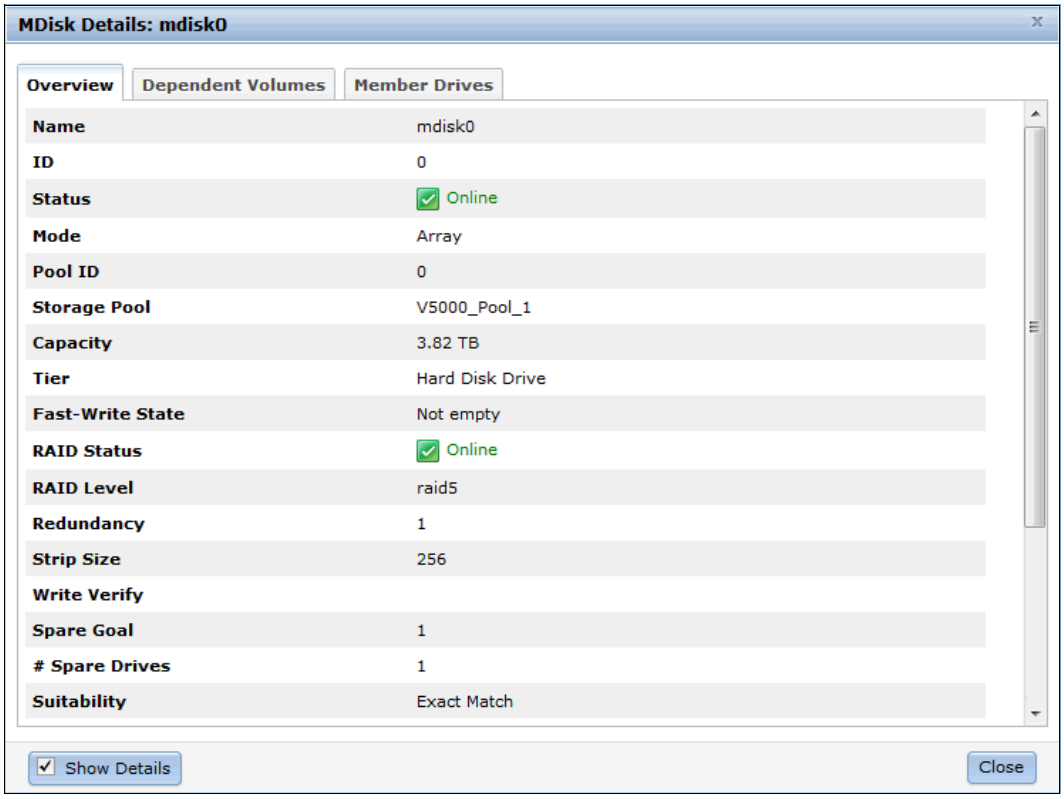

Figure 7-75 MDisk properties overview

- The Dependent Volumes tab (as shown in Figure 7-76) lists all of volumes that use extents on this MDisk.

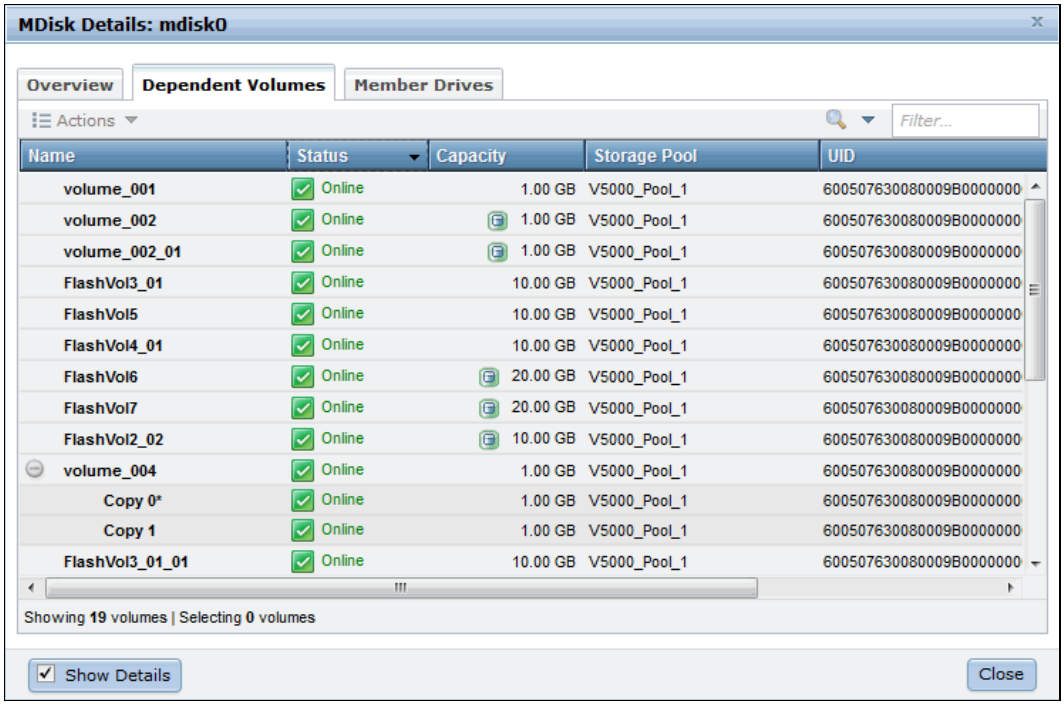

Figure 7-76 MDisk dependent volumes

- In the Member Drives tab (as shown in Figure 7-77), you find all of the member drives of this MDisk. Also, all actions that are described in 7.1.2, “Actions on internal drives” on page 299 can be performed on the drives that are listed here.

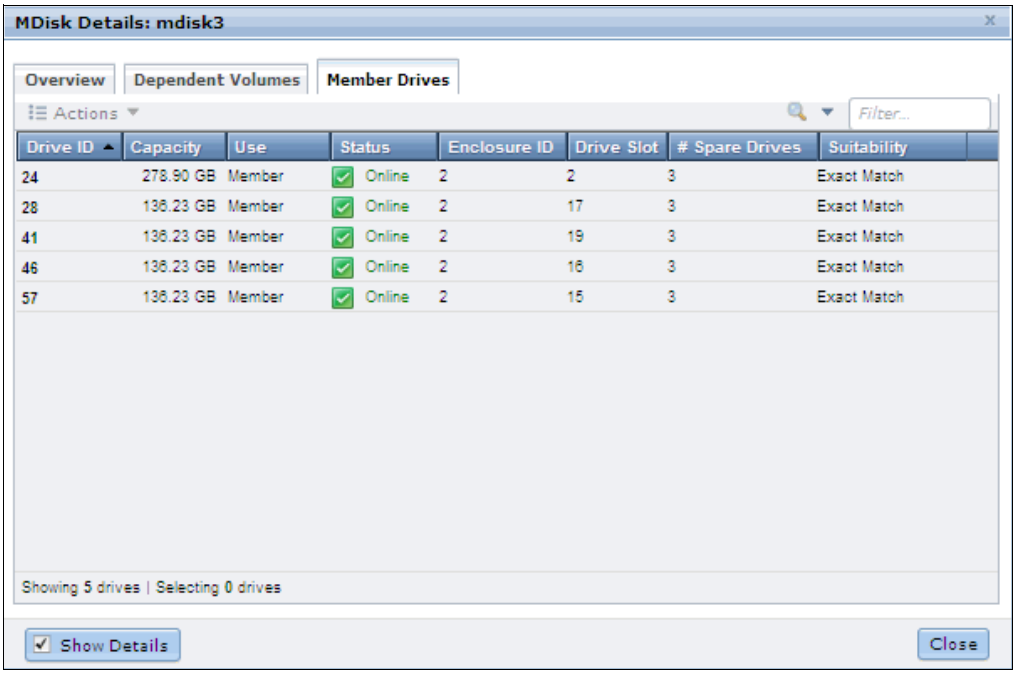

Figure 7-77 MDisk properties member

7.4 Working with storage pools

Storage pools act as a container for MDisks and provision the capacity to volumes. IBM Storwize V5000 organizes storage in storage pools to ease storage management and make it more efficient. Storage pools and MDisks are managed via the MDisks by Pools window. You can access the MDisks by Pools window by clicking **Home** → **Overview** and then clicking the **Pools** icon. Extended help information for storage pools is displayed. If you click **Visit Pools**, the MDisks by Pools window opens, as shown in Figure 7-78 on page 344.

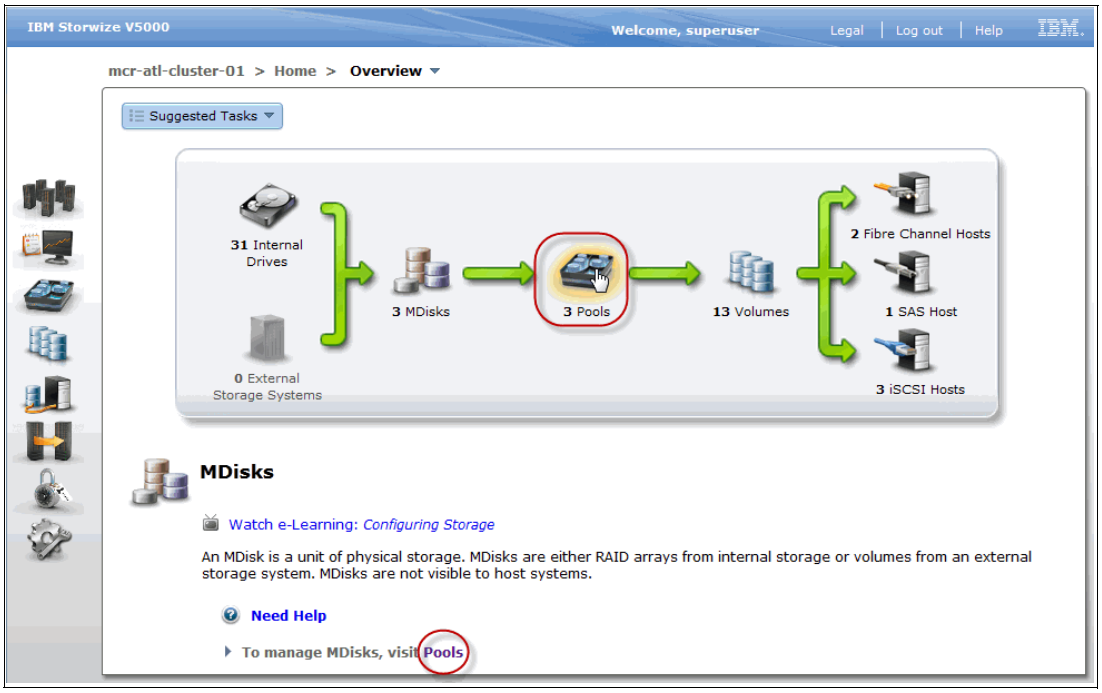

Figure 7-78 Pools from the overview window

An alternative path to the Pools window is to click **Pools** → **MDisks by Pools**, as shown in Figure 7-79.

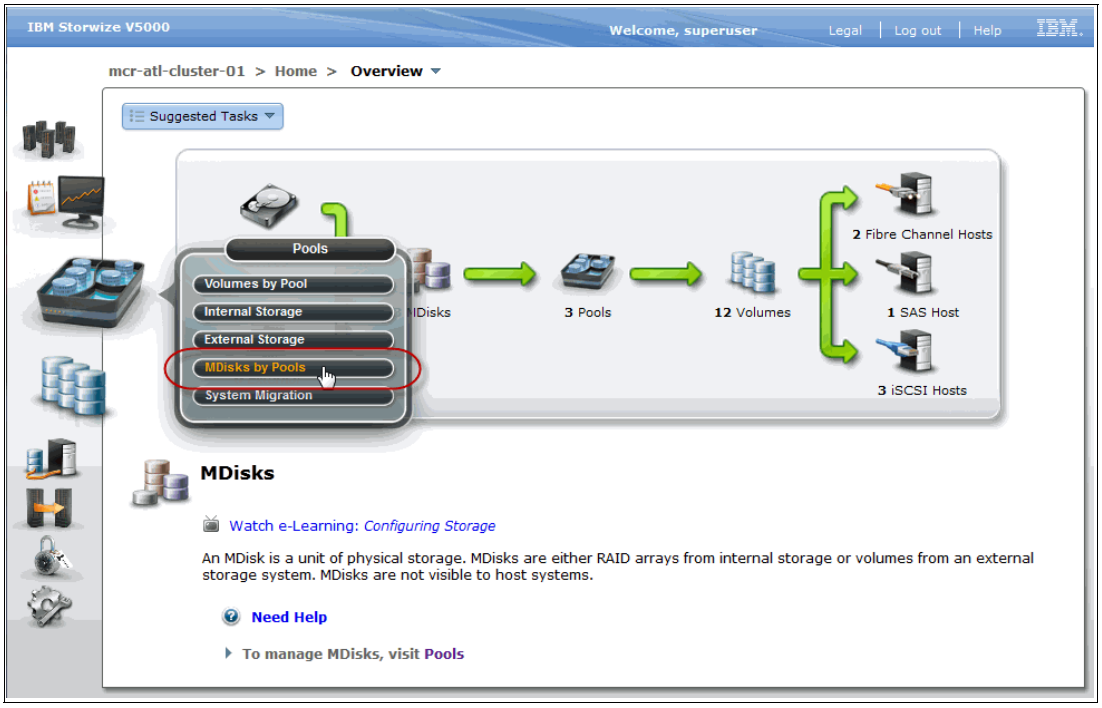

Figure 7-79 Pools from MDisk by Pools window

By using the MDisk by Pools window (as shown in Figure 7-80), you can manage internal and external storage pools. All existing storage pools are displayed row-by-row. The first row features the item Not in a Pool, which contains all unmanaged MDisks, if any exist. Each defined storage pool is displayed with its assigned icon and name, numerical ID, status, and a graphical indicator that shows that the ratio the pool's capacity that is allocated to volumes.

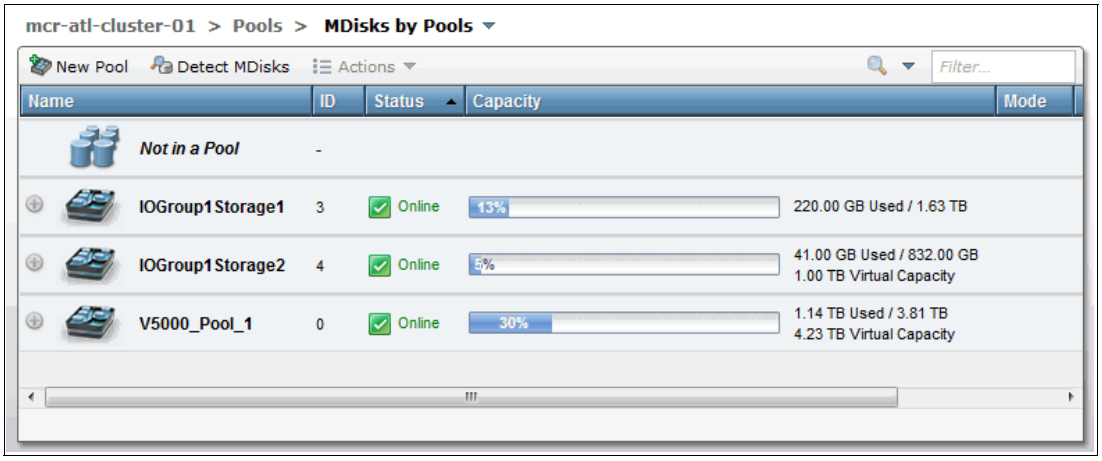

Figure 7-80 Pool window

When you expand a pool's entry by clicking the plus sign (+) to the left of the pool's icon, you can access the MDisks that are associated with this pool. You can perform all actions on them, as described in 7.3, "Working with MDisks on internal and external storage" on page 320.

7.4.1 Create Pool option

New storage pools are built when an MDisk is created if this MDisk is not attached to an existing pool. To create an empty pool, click the **New Pool** option in the pool window.

The only required parameter for the pool is the pool name, as shown in Figure 7-81.

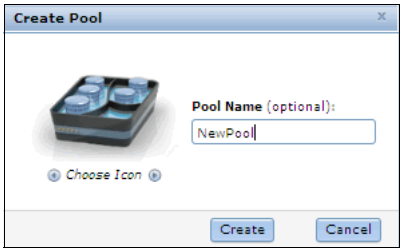

Figure 7-81 Create pool name input

The new pool is included in the pool list with 0 bytes, as shown in Figure 7-82.

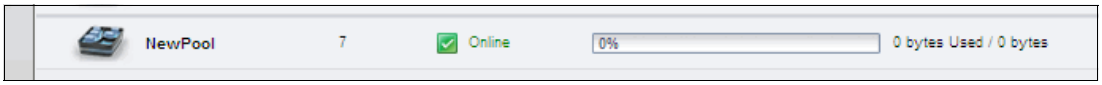

Figure 7-82 Empty pool that is created

7.4.2 Actions on storage pools

A few actions can be performed on storage pools by using the Actions menu, as shown in Figure 7-83. A pool can be renamed or deleted and its icon can be changed.

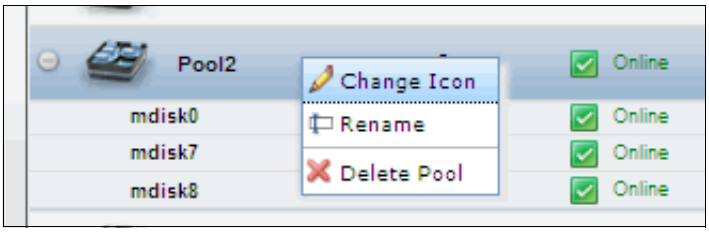

Figure 7-83 Pool action overview

Change Storage Pool icon

There are different storage pool icons available that can be selected, as shown in Figure 7-84. These icons can be used to differentiate between different storage tiers or types of drives.

Figure 7-84 Change storage pool icon

Rename storage pool

The storage pool can be renamed at any time, as shown in Figure 7-85.

Figure 7-85 Rename storage pool

Deleting a storage pool

Pools can be deleted only if there are no MDisks or volumes that are assigned to it. A confirmation panel appears to confirm that all associated MDisk and volumes can be deleted with the pool, as shown in Figure 7-86.

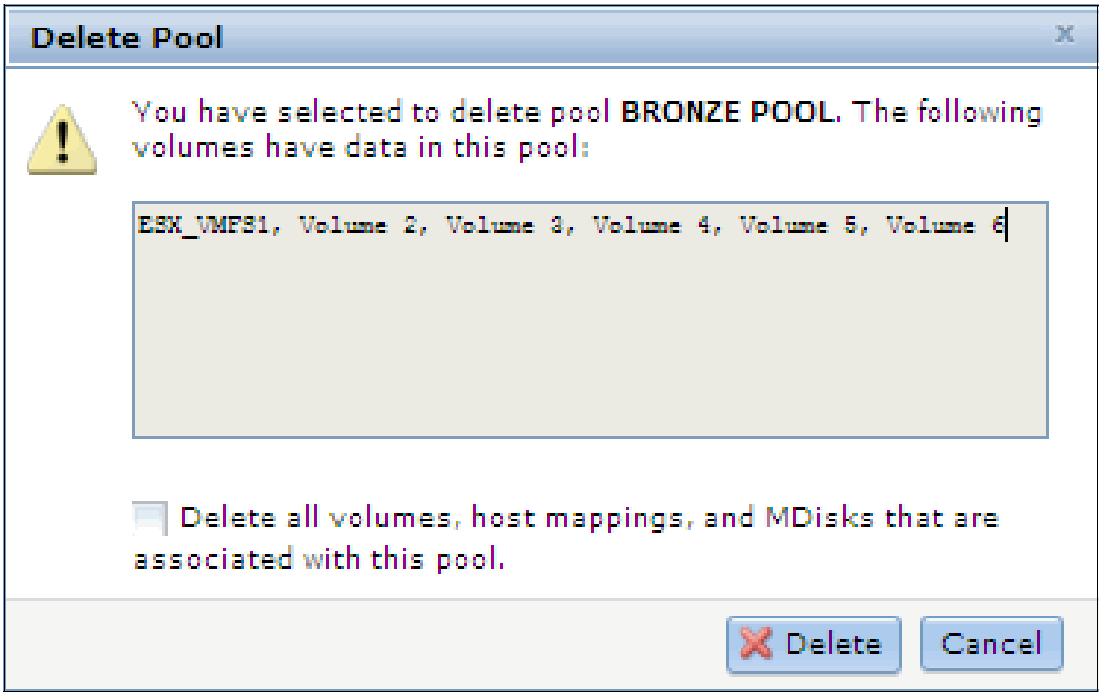

Figure 7-86 Confirmation to delete the storage pool

If it is safe to delete the pool, the option must be selected.

Important: After you delete the pool, all data that is stored in the pool is lost except for the image mode MDisks; their volume definition is deleted, but the data on the imported MDisk remains untouched.

After you delete the pool, all the associated volumes and their host mappings are removed. All the array mode MDisks in the pool are removed and all the member drives return to candidate status. All the managed or image mode MDisks in the pool return to a status of unmanaged after the pool is deleted.

8

Advanced host and volume administration

The IBM Storwize V5000 offers many functions for volume and host configuration. The basic host and volume features of IBM Storwize V5000 are described in Chapter 4, “Host configuration” on page 153 and Chapter 5, “I/O Group basic volume configuration” on page 161. Those chapters also describe how to create hosts and volumes and how to map them to a host.

This chapter includes the following topics:

- Advanced host administration
- Adding and deleting host ports
- Host mappings overview
- Advanced volume administration
- Volume properties
- Advanced volume copy functions
- Volumes by Storage Pool
- Volumes by host

8.1 Advanced host administration

This section describes host administration, including host modification, host mappings, and deleting hosts. Basic host creation and mapping are described in Chapter 4, “Host configuration” on page 153. It is assumed that you created some hosts and that some volumes are mapped to them.

The following topics are covered in this section:

- All Hosts, as described in 8.1.1, “Modifying Mappings menu” on page 352.
- Ports by Host, as described in 8.2, “Adding and deleting host ports” on page 367.
- Host Mappings, as described in 8.3, “Host mappings overview” on page 373.

The IBM Storwize V5000 GUI for hosts menu is shown in Figure 8-1.

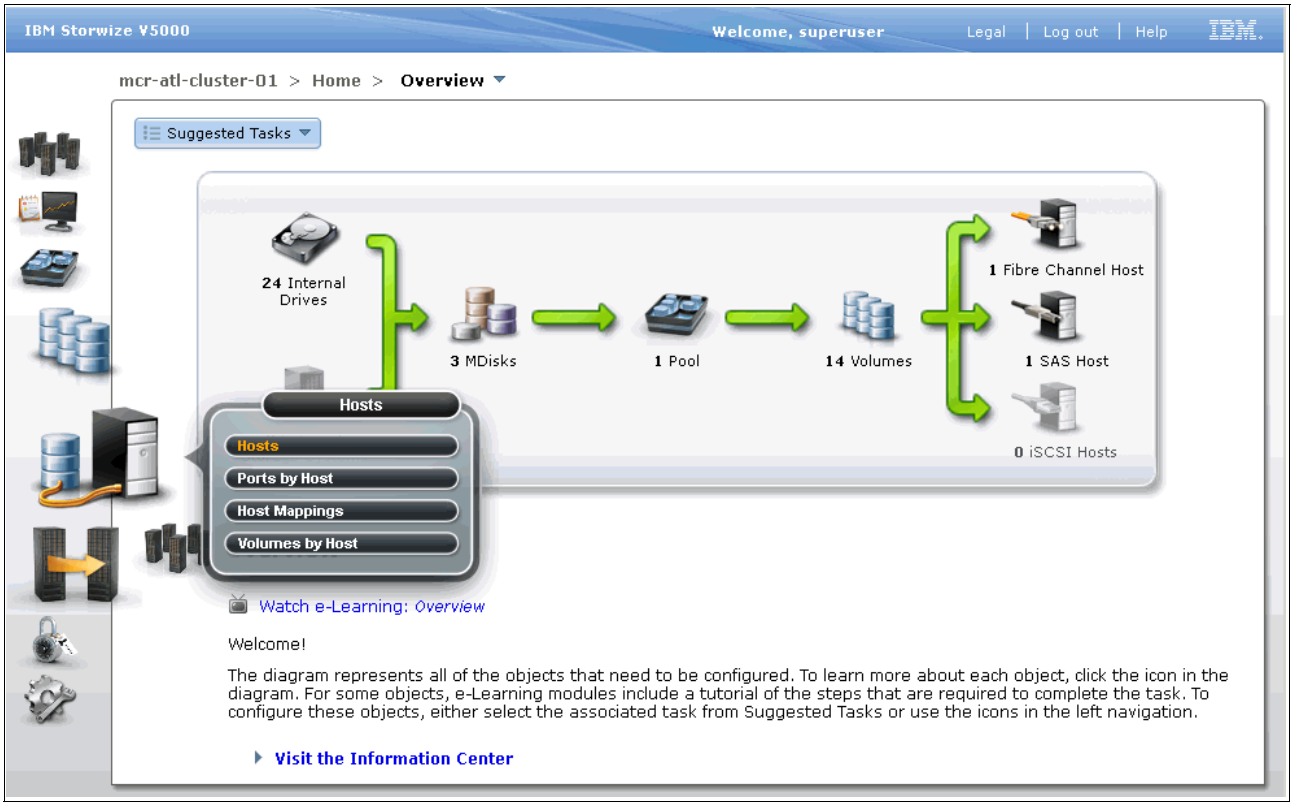

Figure 8-1 Host menu

If you click **Hosts**, the Hosts window opens, as shown in Figure 8-2.

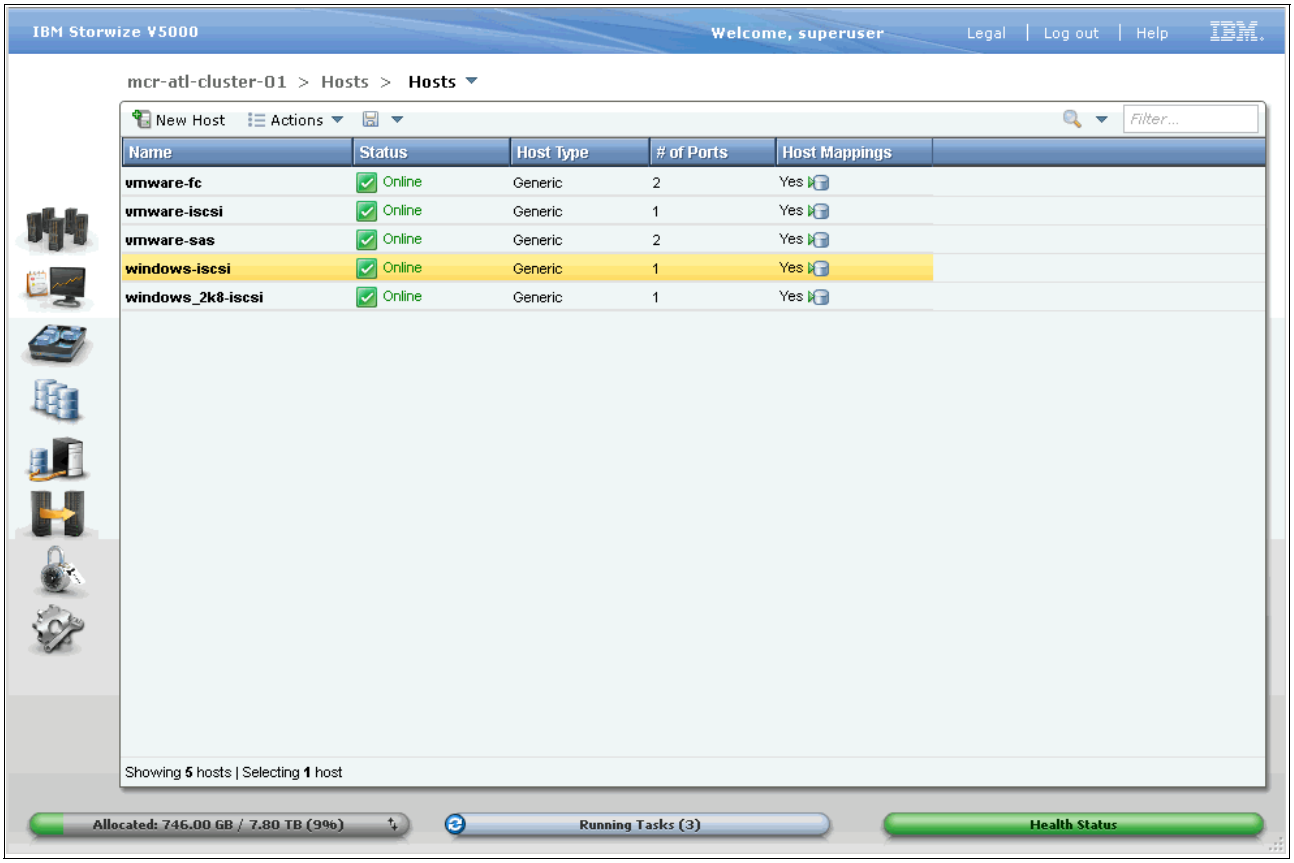

Figure 8-2 Hosts

As you can see in Figure 8-2, a few hosts are created and there are volumes that are mapped to all of them. These hosts are used to show all the possible modifications.

If you highlight a host, you can click **Action** (as shown in Figure 8-3 on page 352) or right-click the host to see all of the available tasks.

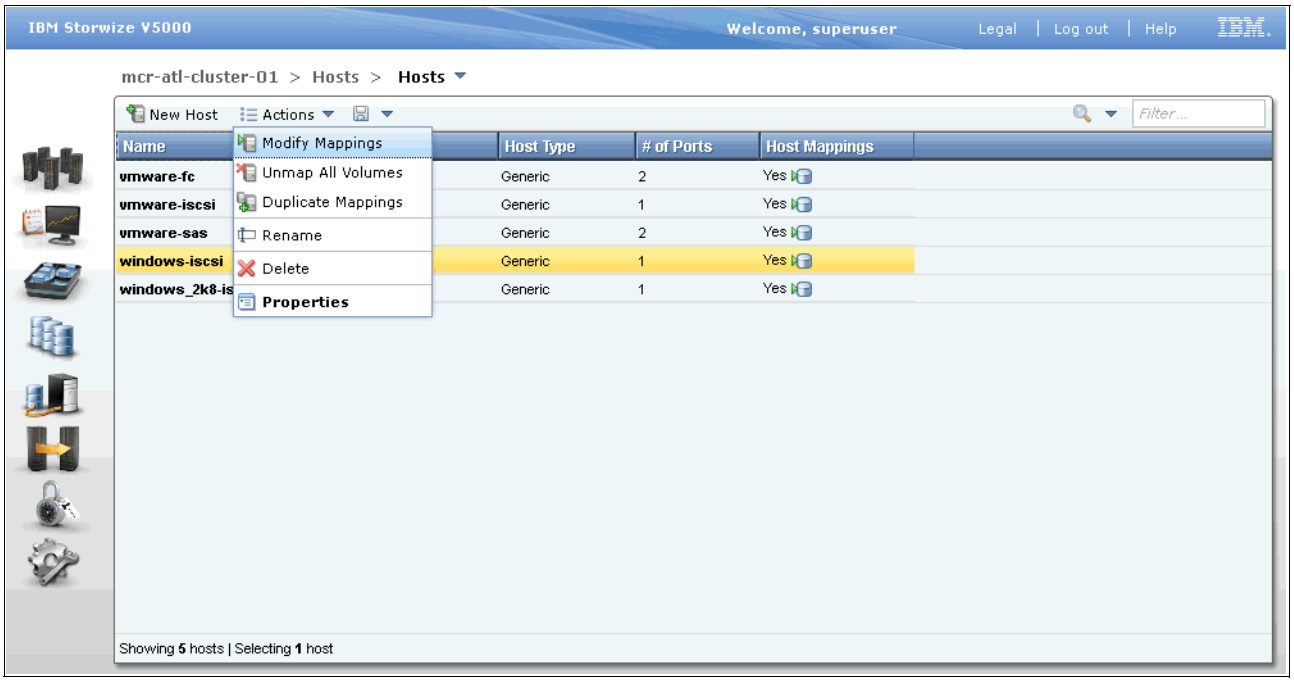

Figure 8-3 Host menu options

As figure Figure 8-3 shows, there are a number of tasks that are related to host mapping. For more information, see 8.1.1, “Modifying Mappings menu” on page 352 and 8.1.2, “Unmapping volumes from a host” on page 356.

8.1.1 Modifying Mappings menu

From the host window, highlight a host and select **Modify Mappings**, as shown in Figure 8-3. The Modify Host Mappings window opens, as shown in Figure 8-4 on page 353.

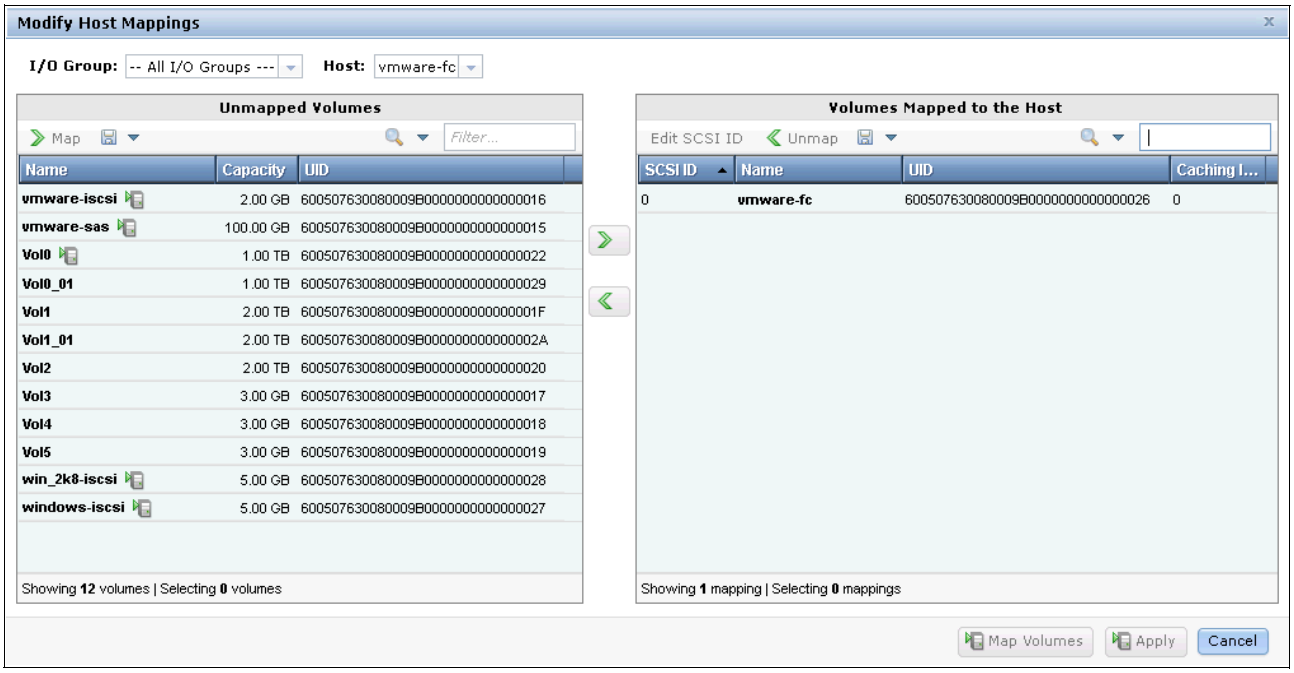

Figure 8-4 Host mappings window

At the upper left, there is a drop-down menu that shows the I/O Group selection. By selecting individual I/O Groups, the IBM Storwize V5000 GUI lists only the volumes that correspond to that I/O Group. The next drop-down menu lists the host that is attached to the IBM Storwize V5000.

Important: Before you change host mappings, always ensure that the host can access volumes from the correct I/O group.

The two panels show all of the available unmapped and mapped volumes for a particular host. The left pane shows the volumes that are available for mapping to the chosen host. The right pane shows the volumes that are already mapped. In our example, one volume with `SCSI ID 0` is mapped to the host `vmware-fc1`, and 12 more volumes are available. In our example, we selected **I/O groups B** and **vmware-fc1** as host and **Vol3** from Volume panel, as shown in Figure 8-5 on page 354.

Important: The unmapped volumes panel refers to volumes that are not mapped to the chosen host.

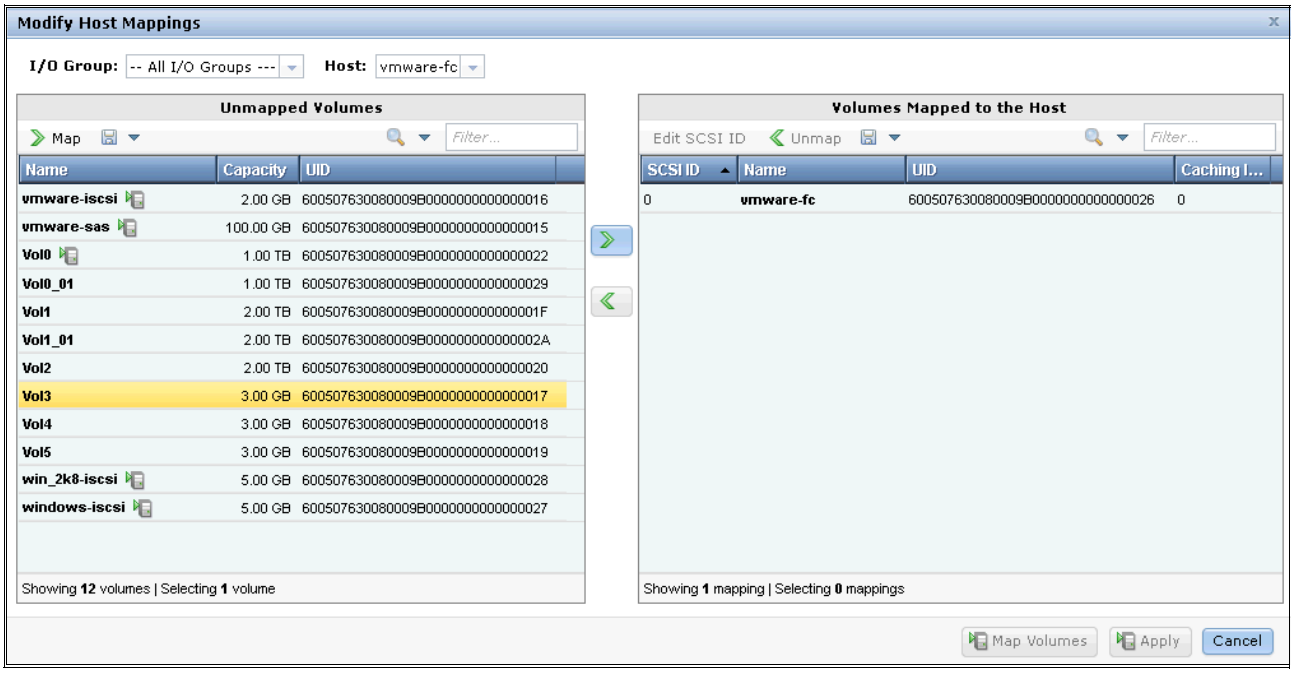

Figure 8-5 Modify Host Mappings

To map a volume, highlight the volume in the left pane and select the right-pointing arrow to move the volume from pane to pane. The changes are marked in yellow and now the Map Volumes and Apply buttons are enabled, as shown in Figure 8-6.

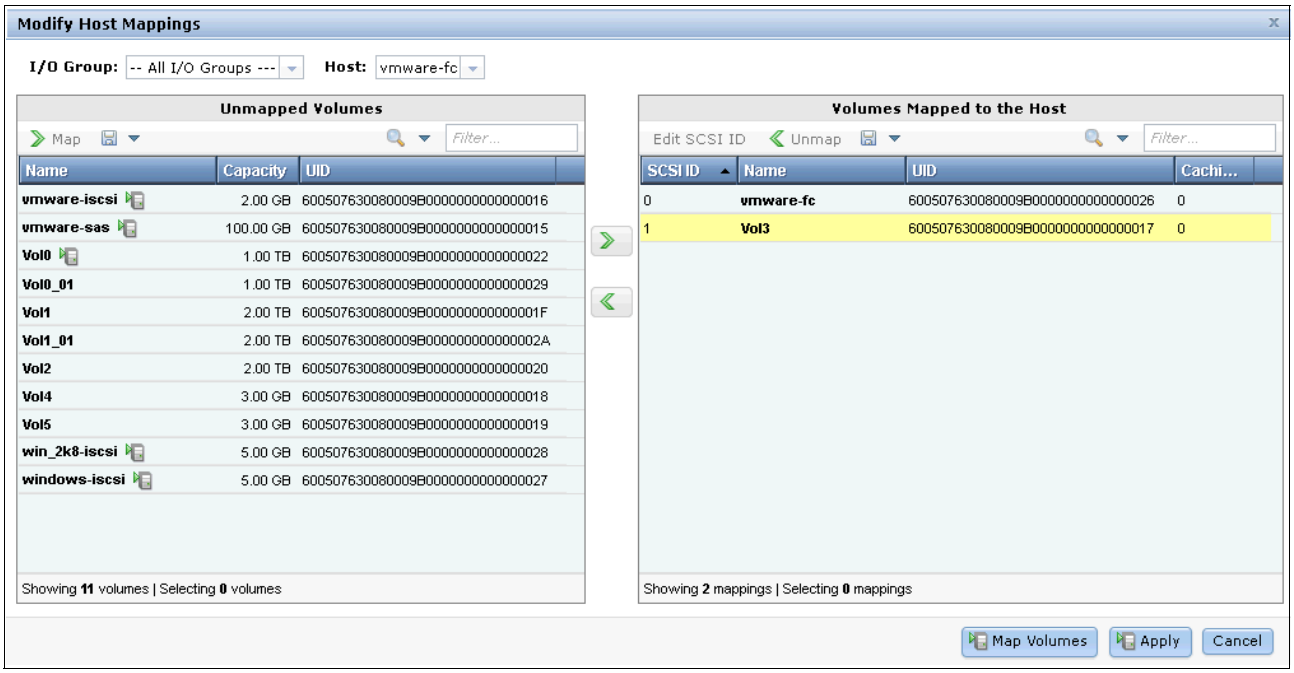

Figure 8-6 Modify Host Mappings

If you click **Map Volumes**, the changes are applied and the Modify Mappings window shows that the task completed successfully, as shown in Figure 8-7.

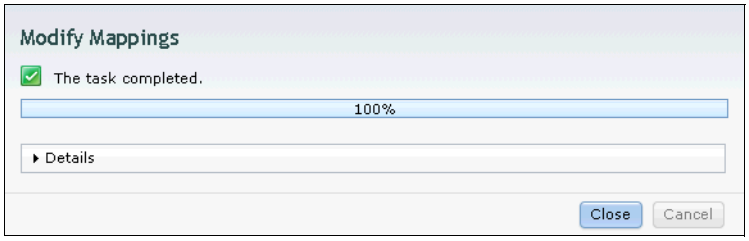

Figure 8-7 Modify Mappings task completed

After you click **Close**, the Modify Host Window closes. If you clicked **Apply**, the changes are submitted to the system, but the Modify Host window remains open for further changes.

You can now choose to modify another host by selecting it from the Hosts drop-down menu or continue working with the host that is already selected, as shown in Figure 8-8.

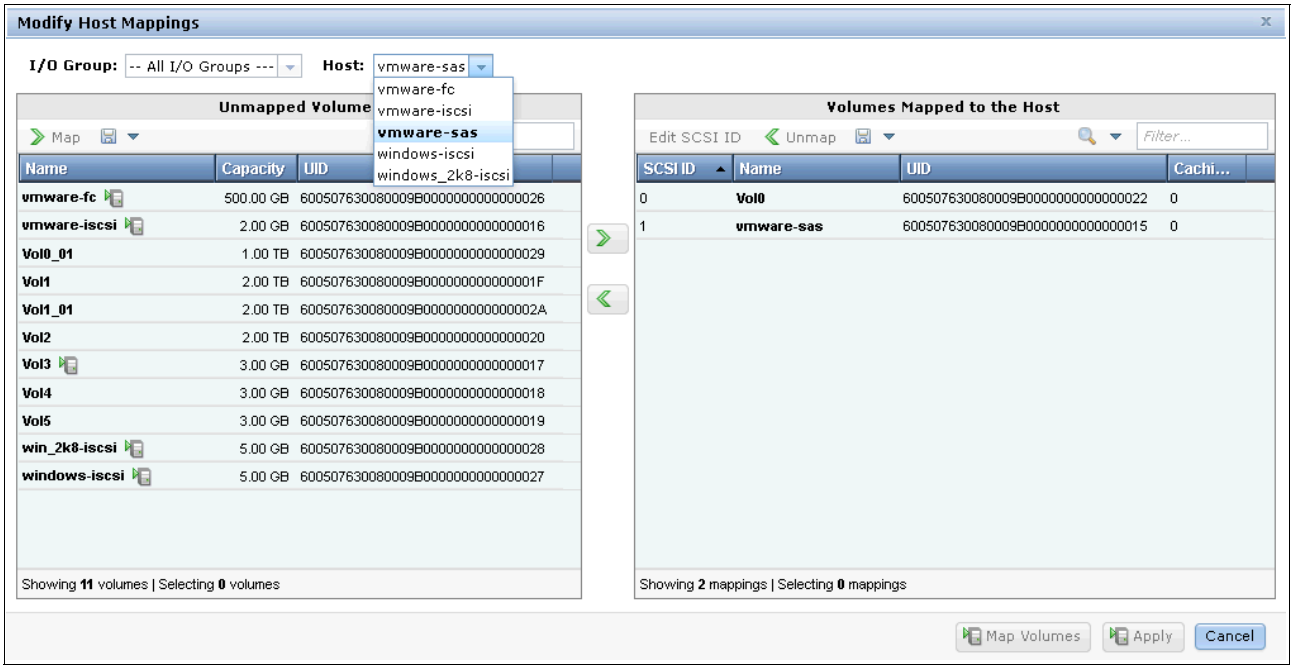

Figure 8-8 Selecting another host to modify

Highlight the volume that you want to modify again and click the right-pointing arrow to move it to the right side pane. The changes are shown in yellow in Figure 8-9 on page 356.

If you right-click the yellow unmapped volume, you can change the SCSI ID, which is used for the host mapping, as shown in Figure 8-9 on page 356.

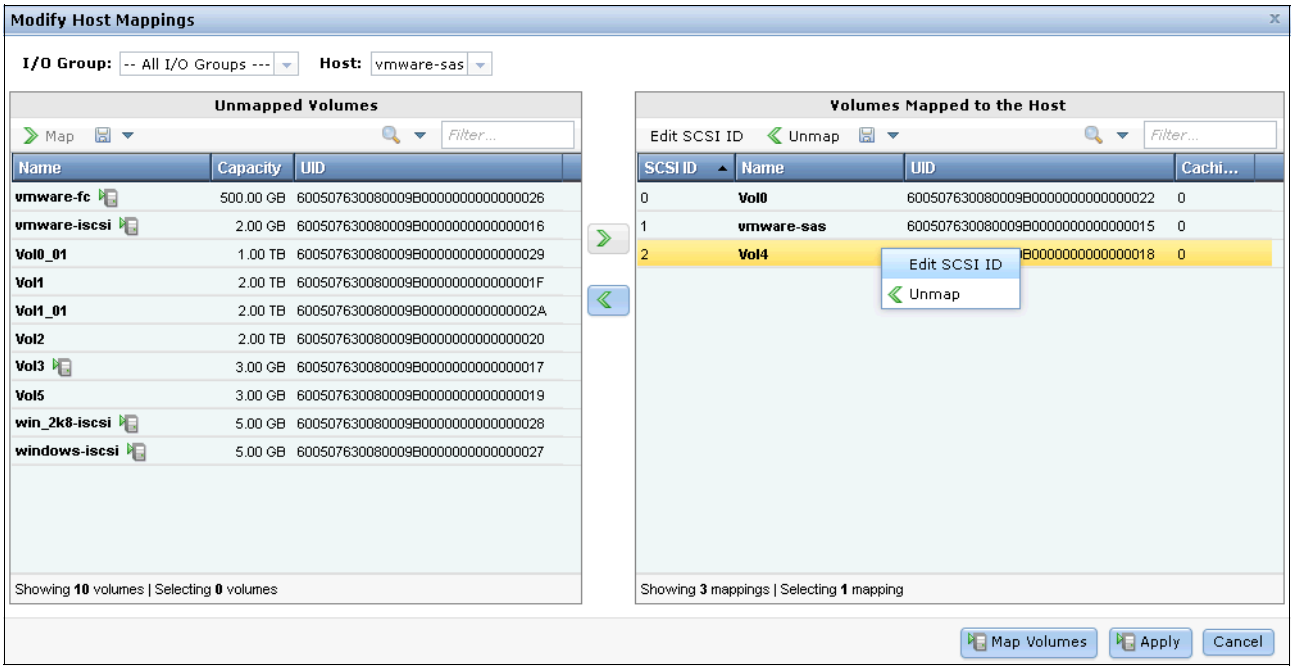

Figure 8-9 Editing iSCSI ID

Click **Edit SCSI ID** and then click **OK** to change the SCSI D. Click **Apply** to submit the changes and complete the host volume mapping.

Important: IBM Storwize V5000 automatically assigns the lowest available SCSI ID if none is specified. However, you can set an SCSI ID for the volume. The SCSI ID cannot be change while volume is assigned to host.

If you want to remove a host mapping, the required steps are the same. For more information about unmapping volumes, see 8.1.2, “Unmapping volumes from a host” on page 356.

8.1.2 Unmapping volumes from a host

If you want to remove host access to certain volumes on your IBM Storwize V5000, you select the volumes by holding the Ctrl key and highlighting the volumes, as shown in Figure 8-10 on page 357.

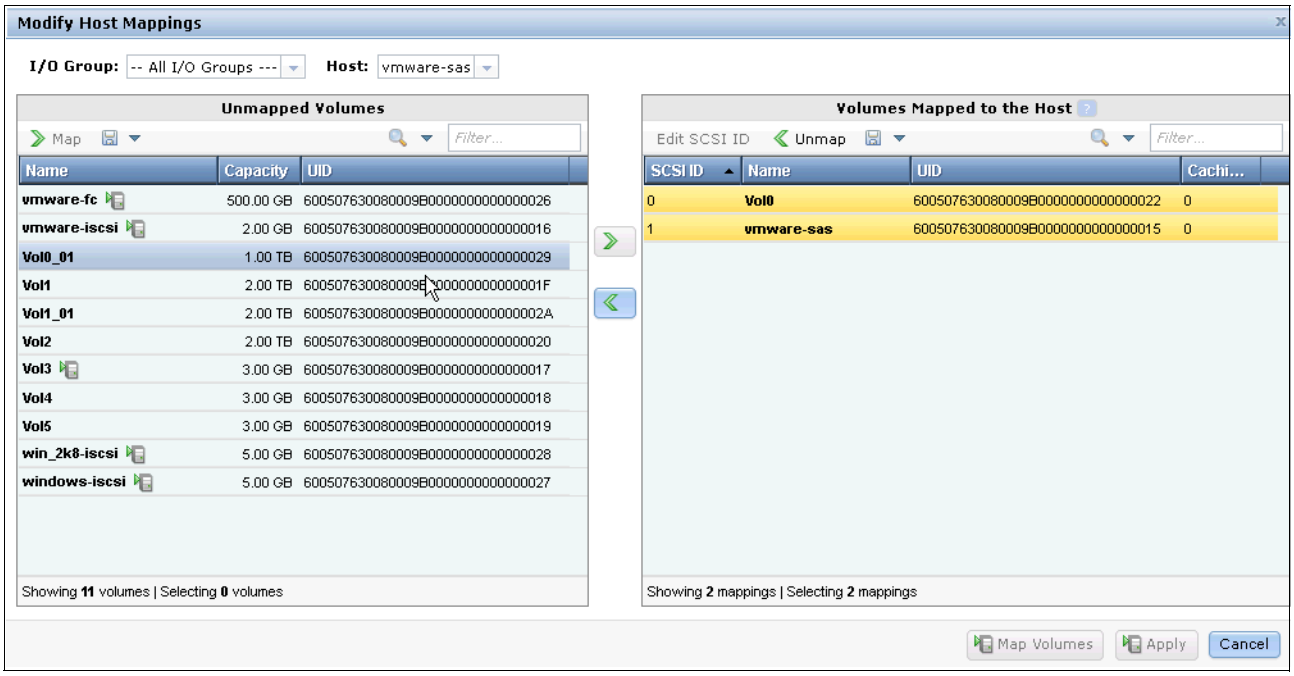

Figure 8-10 Unmapping certain volumes

You can remove access to all volumes in your IBM Storwize V5000 from a host by highlighting the host from the Hosts window and clicking **Unmap all Volumes**, as shown in Figure 8-11.

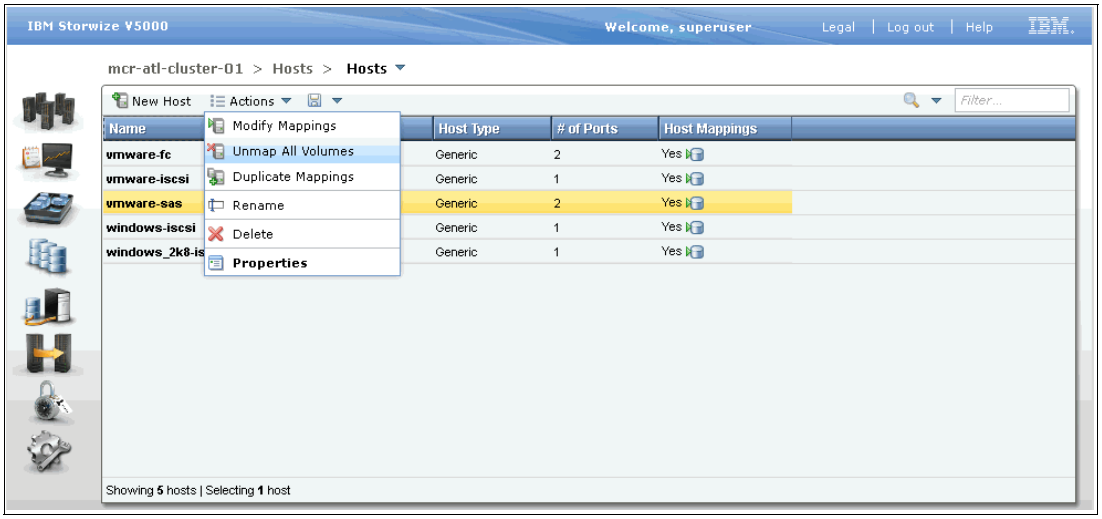

Figure 8-11 Unmap all volumes

You are prompted to confirm the number of mappings you want to remove. Enter the number of mappings and click **Unmap**. In our example, we remove three mappings, as shown in Figure 8-12.

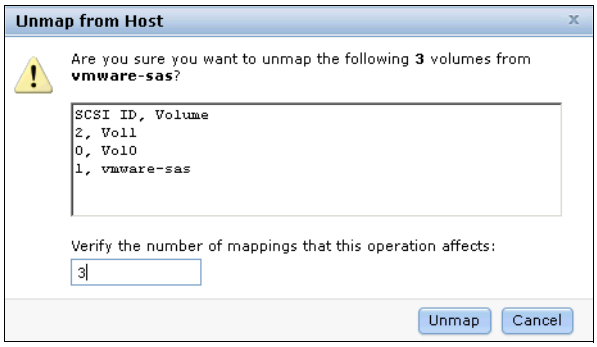

Figure 8-12 Enter the number of mappings to be removed

Unmapping: By clicking **Unmap**, all access for this host to volumes that are controlled by IBM Storwize V5000 system is removed. Ensure that you run the required procedures in your host operating system before the unmapping procedure is done.

The changes are applied to the system, as shown in Figure 8-13. Click **Close** after you review the output.

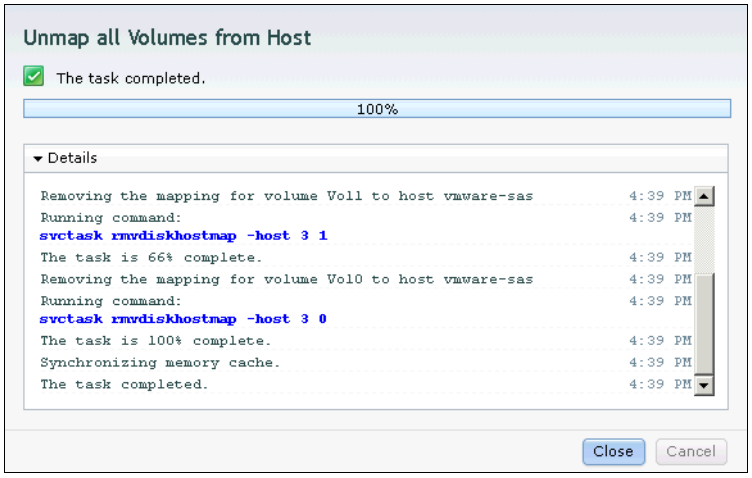

Figure 8-13 Unmapping all volumes from host

Figure 8-14 shows that the selected host no longer has any volume mappings.

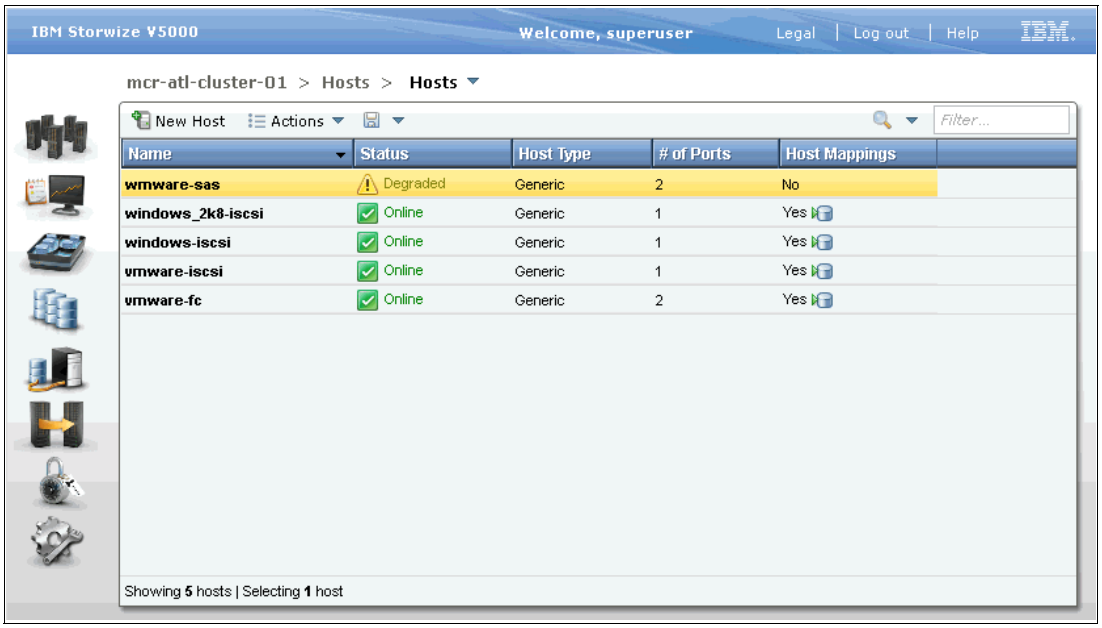

Figure 8-14 Host mapping

8.1.3 Renaming a host

To rename a host object in the IBM Storwize V5000, highlight the host from Host window and click **Rename**, as shown in Figure 8-15.

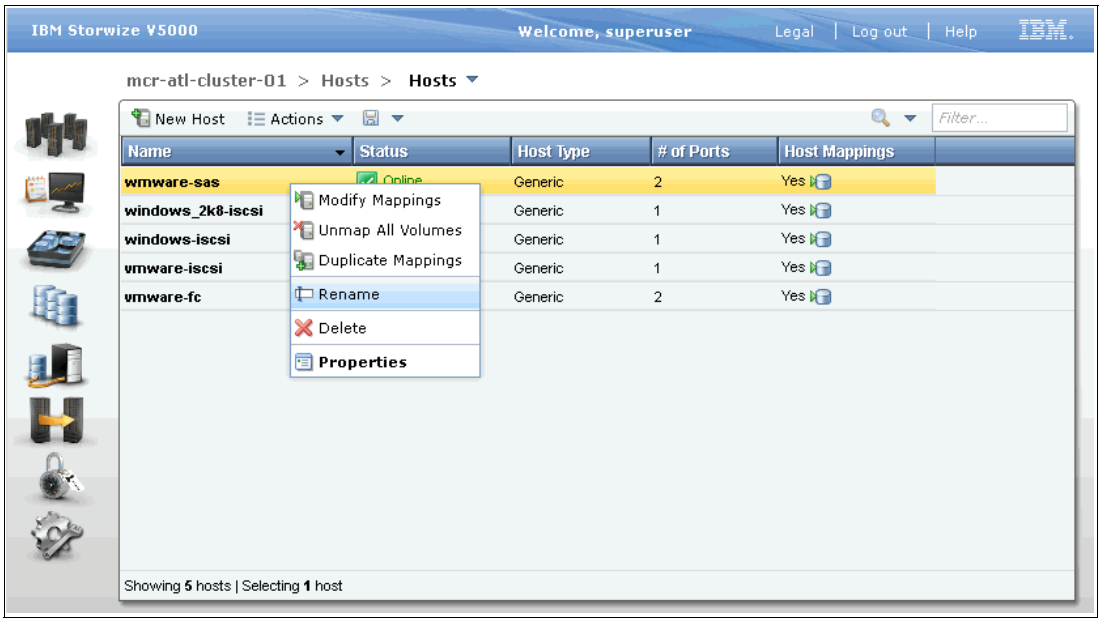

Figure 8-15 Renaming a host

Enter a new name and click **Rename**, as shown in Figure 8-16. If you click **Reset**, your changes are not saved and the host retains its original name.

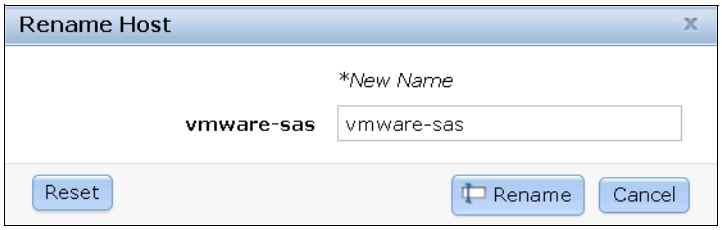

Figure 8-16 Renaming a host window

After the changes are applied to the system, click **Close**, as shown in Figure 8-17.

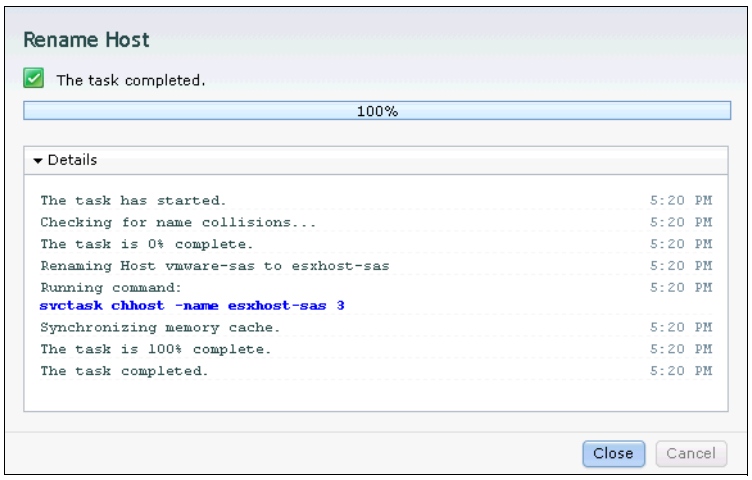

Figure 8-17 Rename a host task completed

8.1.4 Deleting a host

To delete a host, go to the Host window, highlight the host, then click **Delete**, as shown in Figure 8-18.

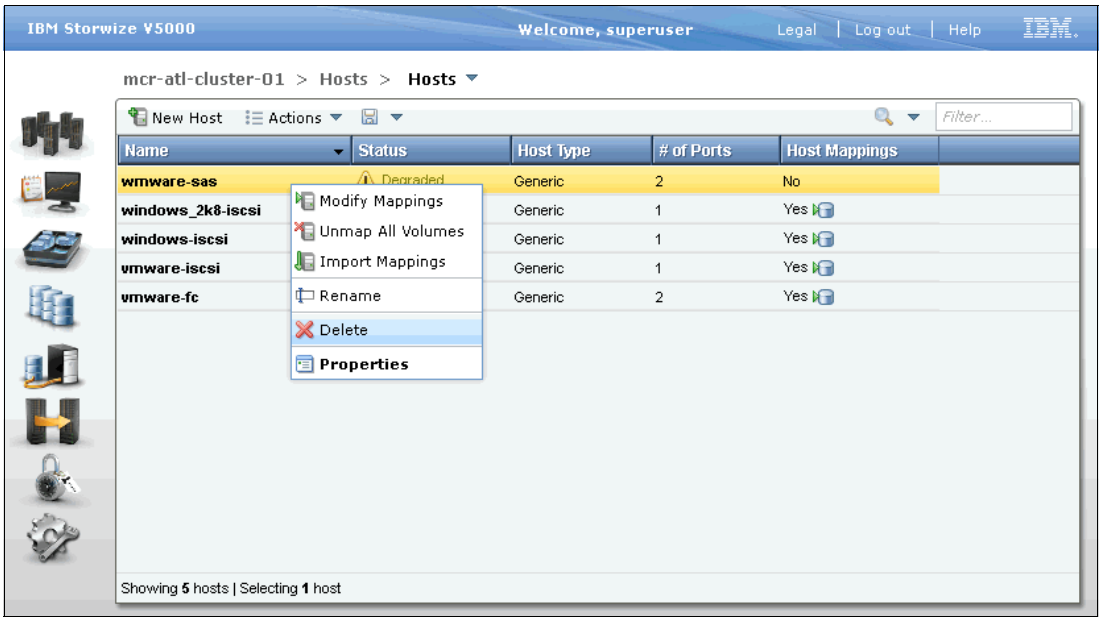

Figure 8-18 Deleting a host

You are prompted to confirm the number of hosts you want to delete. Click **Delete**, as shown in Figure 8-19.

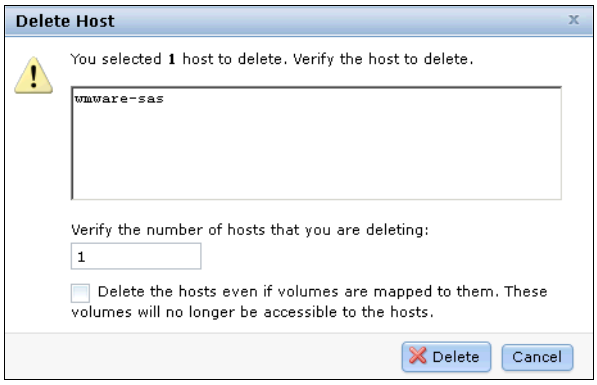

Figure 8-19 Deleting a host

If you want to delete a host with volumes assigned, you must force the deletion by selecting the option in the lower part of the window (see Figure 8-19). If you select this option, the host is removed from the IBM Storwize V5000.

After the task is complete, click **Close** to return to the mappings window, as shown in Figure 8-20.

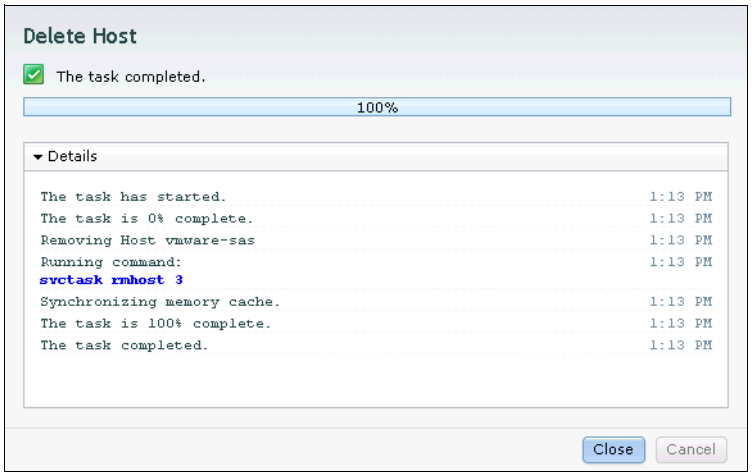

Figure 8-20 Delete host task completed

8.1.5 Host properties

This section describes the host properties. Relevant host information can be found through the next steps. The Host Details window gives you an overview of your host from the following tabs:

- Overview
- Mapped Volumes
- Port Definitions

To open the Host Properties window, highlight the host. From the Action drop-down menu, click **Properties.** You also can highlight the host and right-click it, as shown in Figure 8-21.

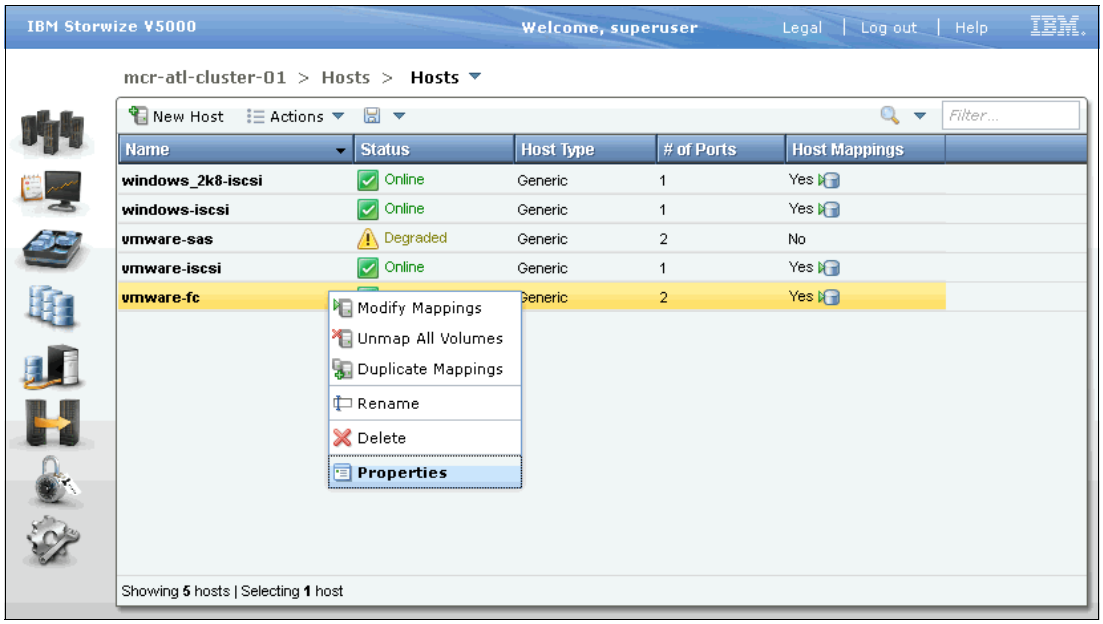

Figure 8-21 Opening host properties

In the next example, we selected host `vmware-fc` to show the host properties information.

As the Overview tab opens, select **Show Details** in the lower left to see more information about the host, as shown in Figure 8-22.

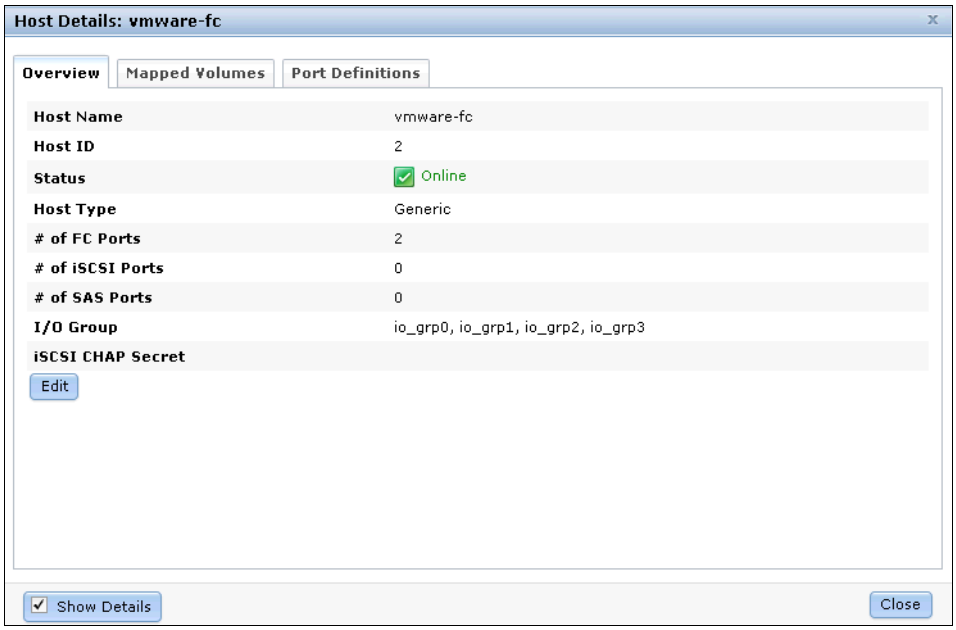

Figure 8-22 Host detail information

This tab provides the following information:

- Host Name: Host object name.
- Host ID: Host object identification number.
- Status: The current host object status; it can be Online, Offline, or Degraded.
- # of FC: The number of host Fibre Channel ports that IBM Storwize V5000 can see.
- # of iSCSI Ports: The number of host iSCSI names or host IQN ID.
- # of SAS Ports: The number of host SAS ports that are connected to IBM Storwize V5000.
- I/O Group: The I/O Group from which the host can access a volume (or volumes).
- iSCSI CHAP Secret: The Challenge Handshake Authentication Protocol information if it exists or is configured.

To change the host properties, click **Edit** and several fields can be edited, as shown in Figure 8-23.

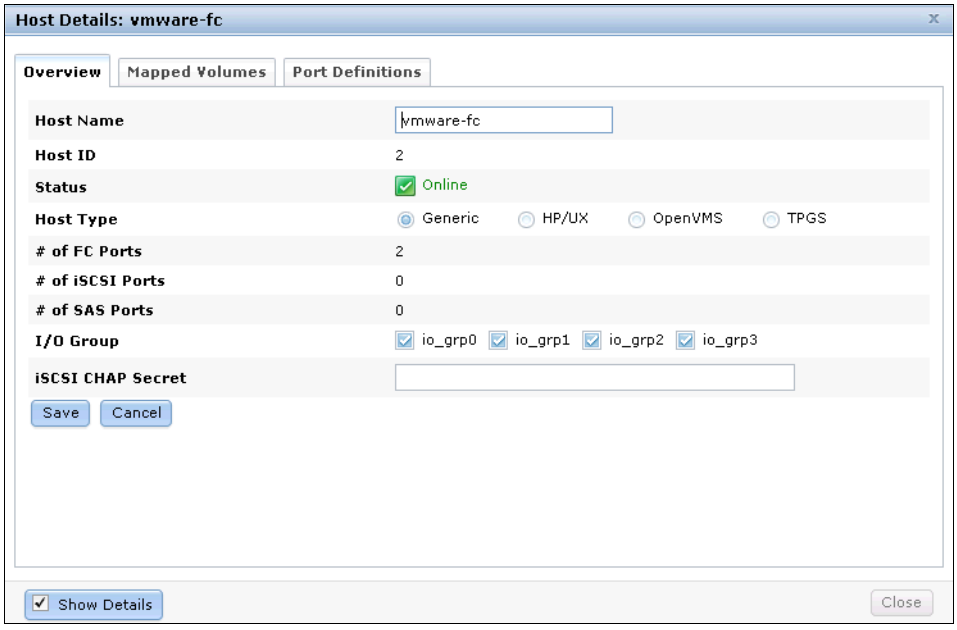

Figure 8-23 Host properties: Editing host information

The following changes can be made:

- Host Name: Change the host name.
- Host Type: Change this setting if you intend to change host type to HP/UX, OpenVMS, or TPGS hosts.
- I/O Group: Change the I/O Group from which the host can access volumes.
- iSCSI CHAP Secret: Enter or change the iSCSI CHAP secret for this host.

I/O Group: You can use I/O Group options to control the number of I/O Groups the host can access volumes through.

Make any necessary changes and click **Save** to apply them. Figure 8-24 on page 365 shows the progress bar of the changes that were made.

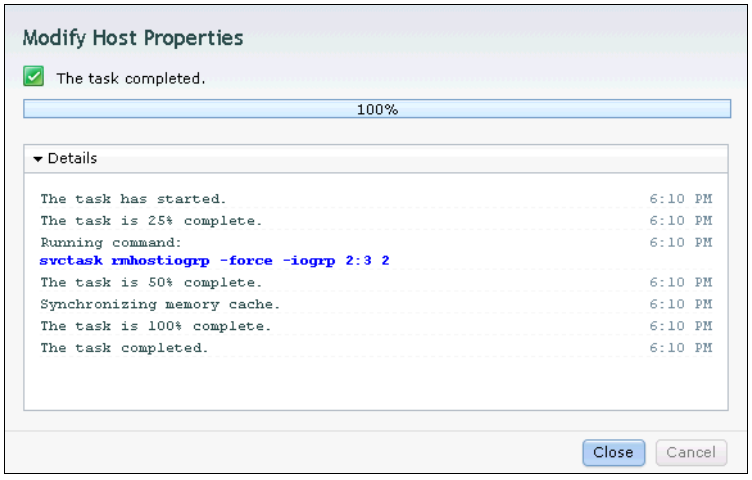

Figure 8-24 Editing host properties task completed

Click **Close** to return to the Host Details window.

The Mapped Volume tab (as shown in Figure 8-25) gives you an overview of which volumes are mapped to this host. The details that are shown are SCSI ID, volume name, UID, (volume ID) and the caching I/O Group per volume. Clicking the **Show Details** option does not show any detailed information.

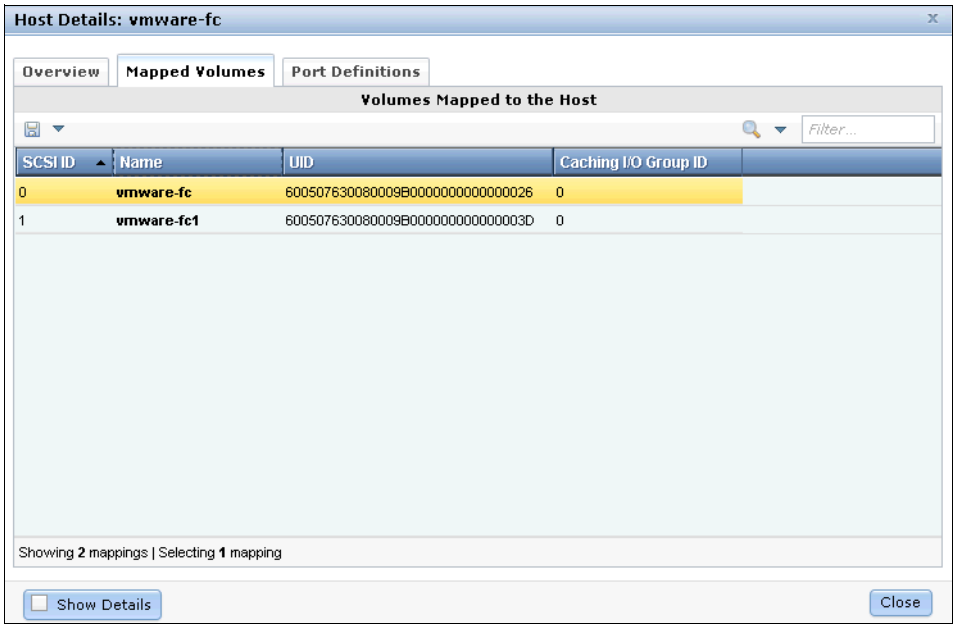

Figure 8-25 Host Details: Mapped volumes information

The Port Definitions tab (as shown in Figure 8-26) shows the following information:

- Configured host ports and their status
- The worldwide port names (WWPNs) (for SAS and FC hosts)
- iSCSI Qualified Name (IQN) for iSCSI hosts
- Type column: Shows the port type information.
- # Nodes Logged In column: Lists the number of IBM Storwize V5000 node canisters that each port (initiator port) logged on to.

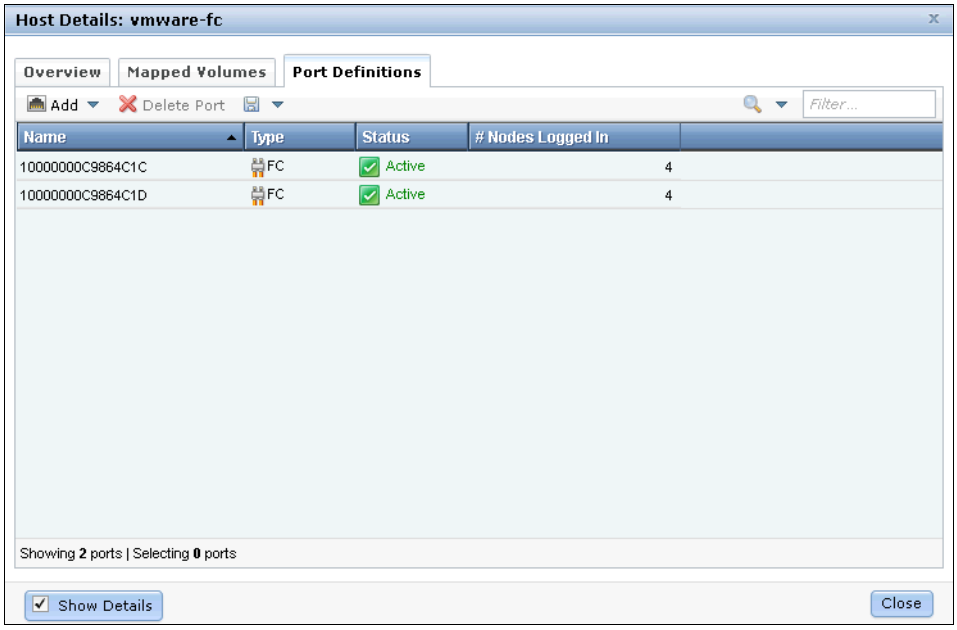

Figure 8-26 Host port details

By using this window, you can also Add and Delete Host Port (or ports), as described in 8.2, “Adding and deleting host ports” on page 367. Selecting the **Show Details** option does not show any further information.

Click **Close** to close the Host Details section.

8.2 Adding and deleting host ports

To configure host ports, use IBM Storwize V5000 GUI by clicking **Host** → **Ports by Host** to open the Ports by Host window, as shown in Figure 8-27.

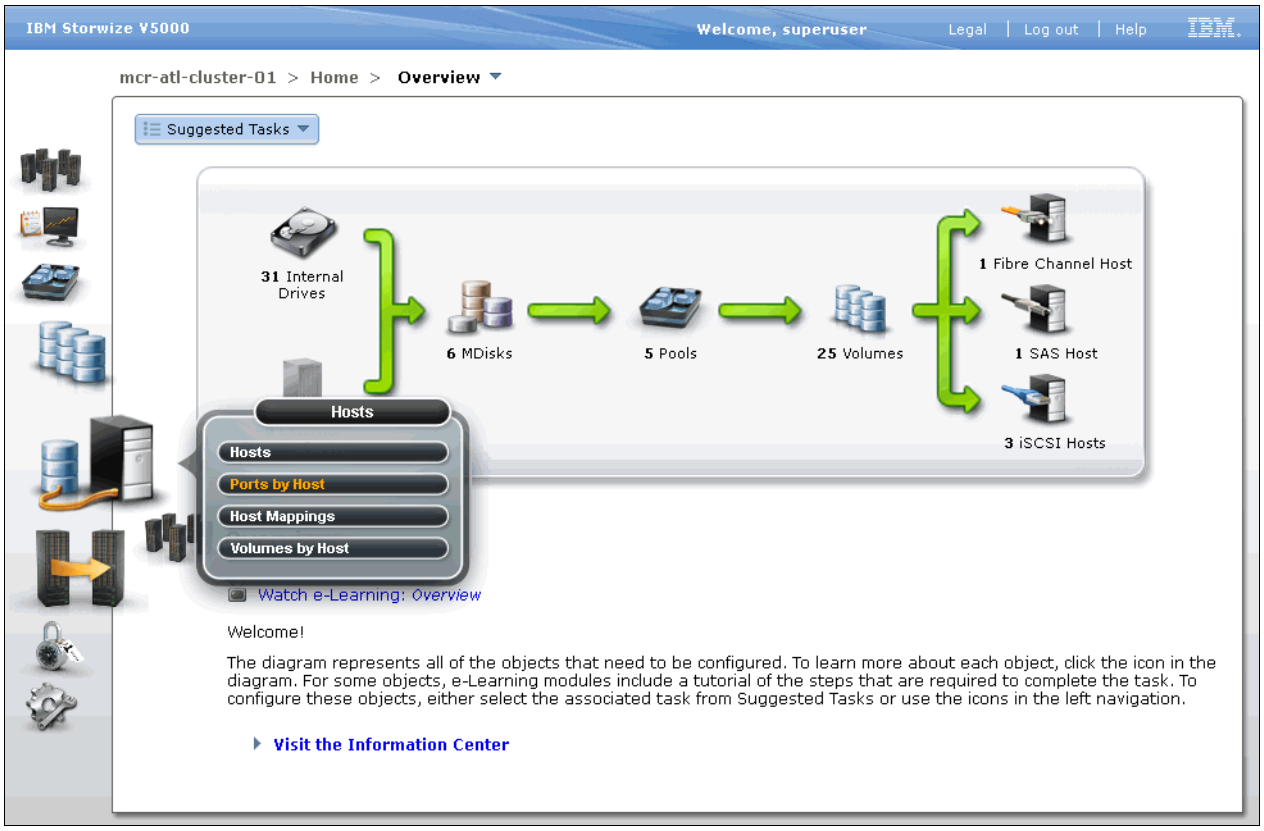

Figure 8-27 Ports by Host window

Hosts are listed in the pane on the left side of the window. The Function Icons show an orange cable for Fibre Channel host, a black cable for SAS host, and a blue cable for an iSCSI host.

The properties of the highlighted host are shown in the right side pane. If you click **New Host**, the wizard that is described in Chapter 4, “Host configuration” on page 153 starts.

If you click the **Action** drop-down menu (as shown in Figure 8-28 on page 368), the tasks that are described in the previous sections can be started from this location.

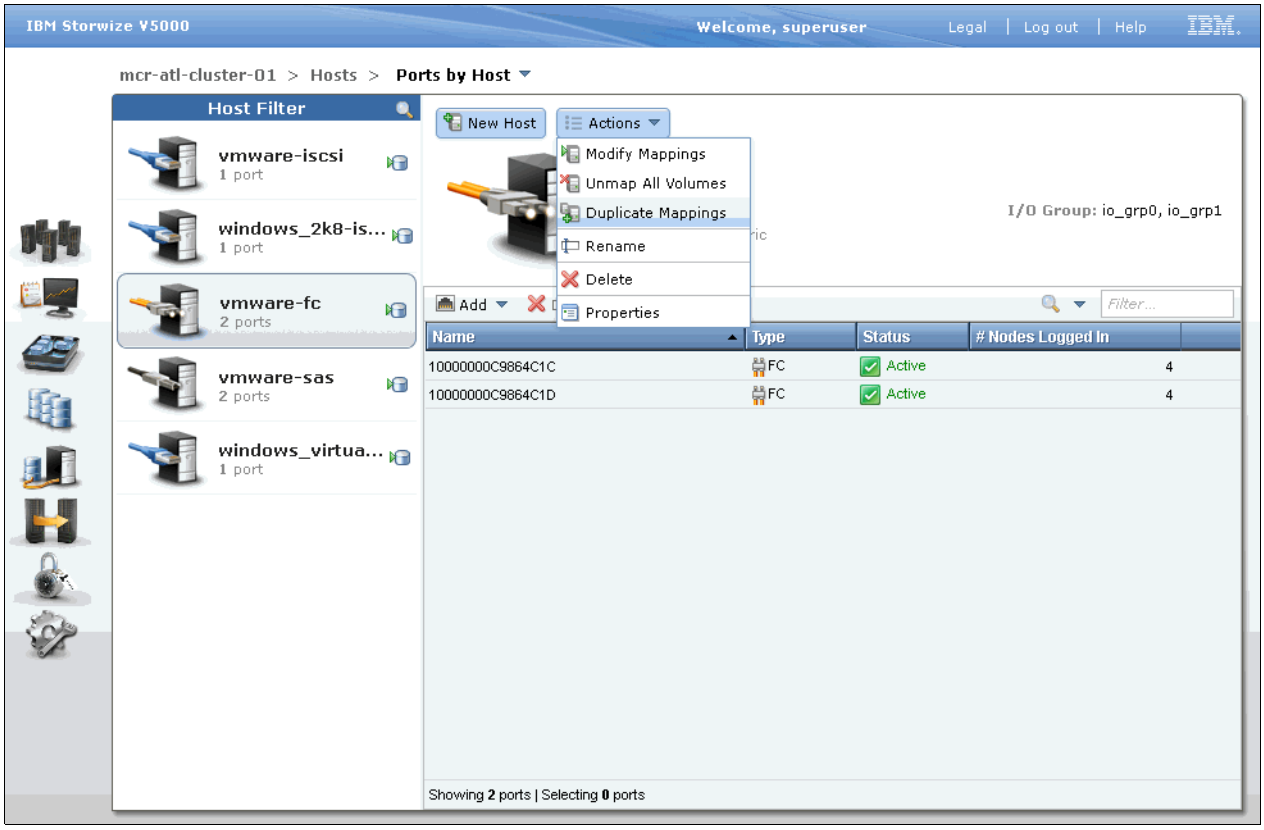

Figure 8-28 Host Action menu

8.2.1 Adding a host port

To add a host port, highlight the host from left side panel, click **Add**, and then choose a Fibre Channel, SAS, or an iSCSI port, as shown in Figure 8-29.

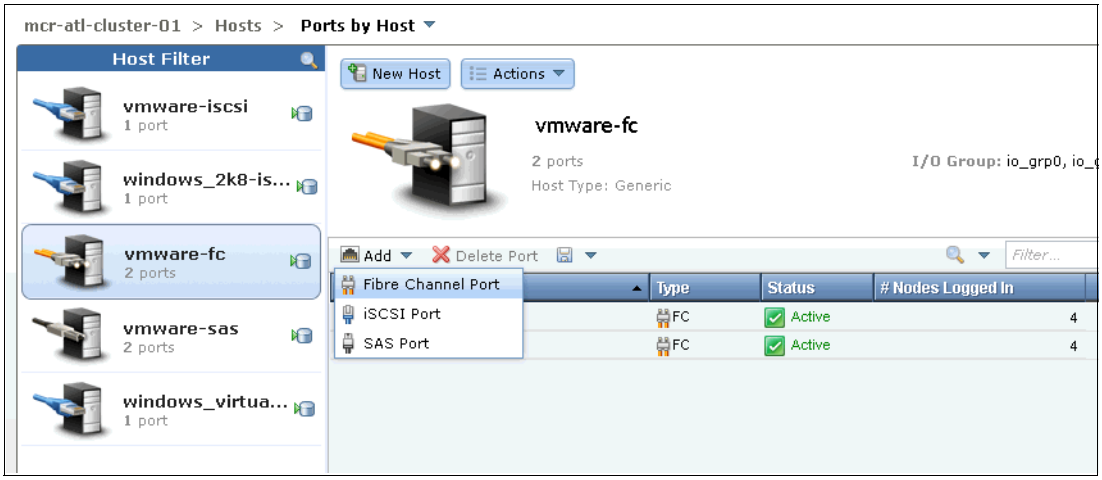

Figure 8-29 Adding a host port

Important: A host system can have a mix of Fibre Channel, iSCSI, and SAS connections. If a configuration requires you to mix protocols, check the capabilities of your operating system and plan carefully to avoid miscommunication or data loss.

8.2.2 Adding a Fibre Channel port

As shown in Figure 8-29 on page 368, click **Fibre Channel Port** and the Add Fibre Channel Ports window opens.

If you click the **Fibre Channel Ports** drop-down menu, you see a list of all available Fibre Channel host ports. If the WWPN of your host is not available in the menu, check your SAN zoning and rescan the SAN from the host. You might also try to rescan by clicking **Rescan**.

Select the WWPN to add and click **Add Port to List**, which shows the new port is added to the list.

Repeat this step to add more ports to a host. If you want to add an offline port, manually enter the WWPN of the port into the Fibre Channel Ports field and click **Add Port to List**, as shown in Figure 8-30.

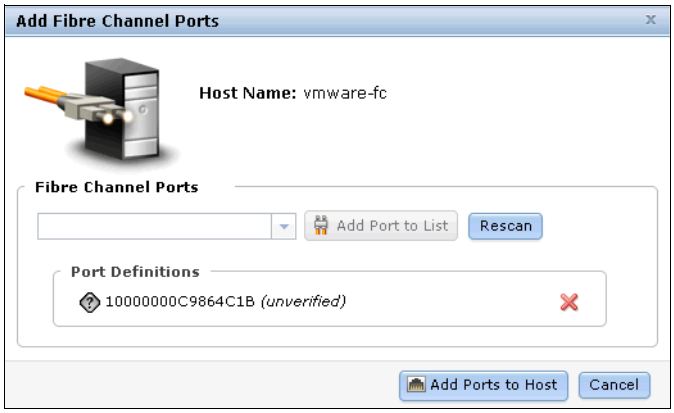

Figure 8-30 Adding offline port

As shown in Figure 8-30, the port appears as unverified because it is not logged on to the IBM Storwize V5000. The first time the port logs on, the state automatically changes to online and the mapping is applied to this port.

To remove one of the ports from the list, click the red X next to it. In Figure 8-30, we manually added an FC port.

Important: If you are removing online or offline ports, IBM Storwize V5000 prompts you to add the number of ports you want to delete but does not warn about mappings. Disk mapping is associated to the host object and Logical Unit Number (LUN) access is lost if all ports are deleted.

Click **Add Ports to Host** and the changes are applied. Figure 8-31 shows the output after ports are added to the host. Even if it is an offline port, the IBM Storwize V5000 still adds it.

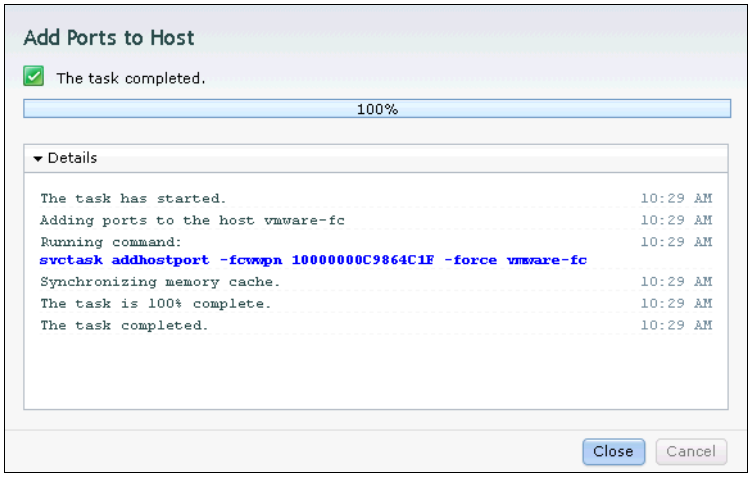

Figure 8-31 Adding a host port

8.2.3 Adding a SAS host port

As shown in Figure 8-29 on page 368, from the IBM Storwize V5000 GUI, click **Host** → **Port by Host** and then click **Add** → **SAS Port** to add an SAS host port to an existing host.

The **Add SAS Host Port** window opens. If you click the **SAS Ports** drop-down menu, you see a list of all known SAS Ports that are connected to IBM Storwize V5000. If SAS WWPNs are not available, try the Rescan option or check the physical connection (or connections).

Important: IBM Storwize V5000 allows the addition of an offline SAS port. Enter the SAS WWPN in SAS Port field and then click **Add Port to List**.

Select the SAS WWPN you want to add to the existing host and click **Add Port to List**, as shown in Figure 8-32.

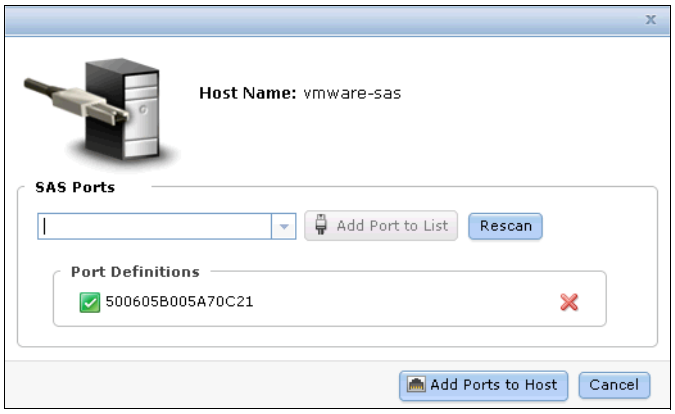

Figure 8-32 Adding an online SAS port

The Add Port to Host task completes successfully.

8.2.4 Adding an iSCSI host port

To add an iSCSI host port, click **iSCSI Port** (as shown in Figure 8-29 on page 368) and the Add iSCSI Ports window opens, as shown in Figure 8-33.

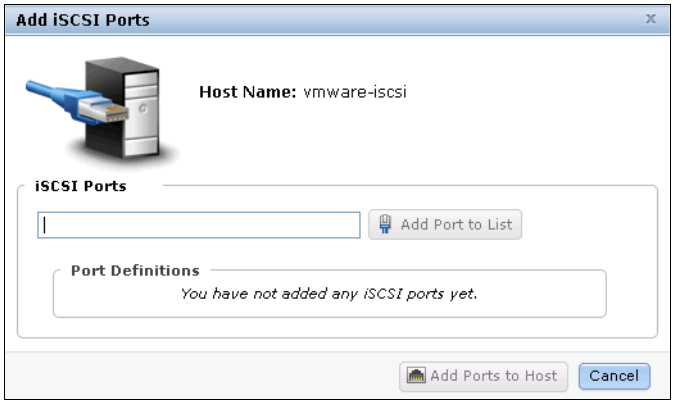

Figure 8-33 Adding iSCSI Host Port

Enter the initiator name of your host and click **Add Port to List**. After you add the iSCSI Port, click **Add Ports to Host** to complete the tasks and apply the changes to the system. The iSCSI port status remains unknown until it is added to the host and a host rescan process is completed. Figure 8-34 shows the output after an iSCSI port is added.

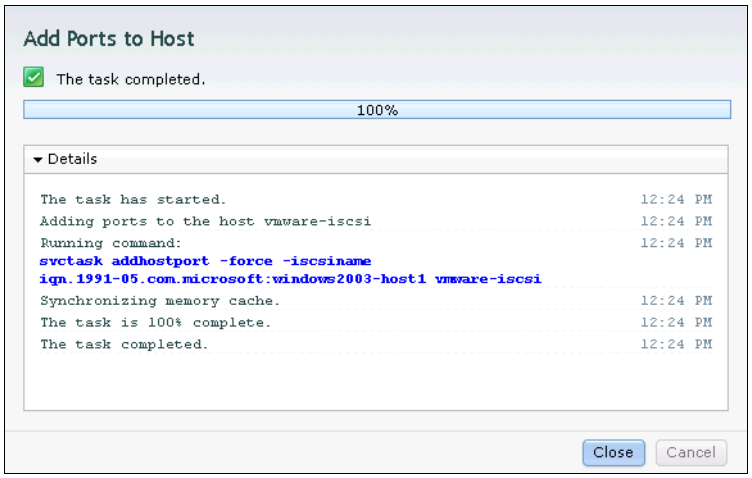

Figure 8-34 Successful iSCSI port addition

Click **Close** to return to the Ports by Host window.

Important: An error message with code `CMMVC6581E` is shown if one of the following conditions occurs:

- The IQNs exceed the maximum number that is allowed.
- There is a duplicated IQN.
- The IQN contains a comma or leading or trailing spaces.
- The IQN is invalid in some other way.

8.2.5 Deleting a host port

To delete host ports, click **Host** → **Ports by Host** to open the Ports by Host window, as shown in Figure 8-27 on page 367.

Select the host in left pane, highlight the host port that you want to delete and the Delete Port button becomes available, as shown in Figure 8-35.

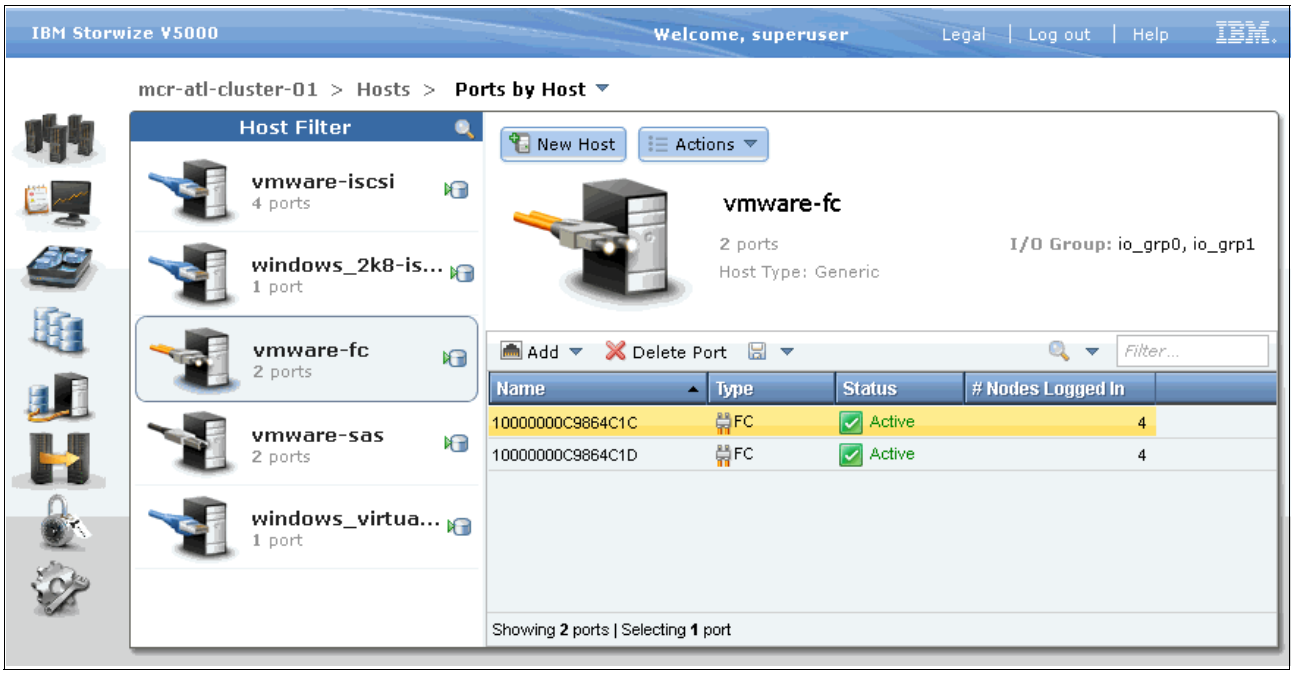

Figure 8-35 Delete host port

If you press and hold the Ctrl key, you can also select several host ports to delete.

Click **Delete** and you are prompted to enter the number of host ports that you want to delete, as shown in Figure 8-36.

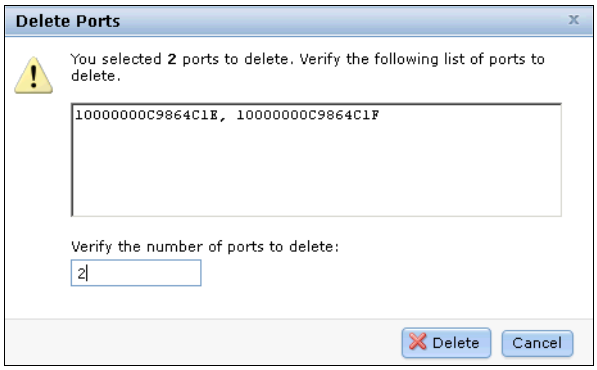

Figure 8-36 Deleting host port

Click **Delete** to apply the changes to the system. A task window opens that shows the results. Click **Close** to return to the Ports by Host window.

8.3 Host mappings overview

From IBM Storwize V5000 GUI, select **Host** → **Host Mappings** to open the Host Mappings overview window, as shown in Figure 8-37.

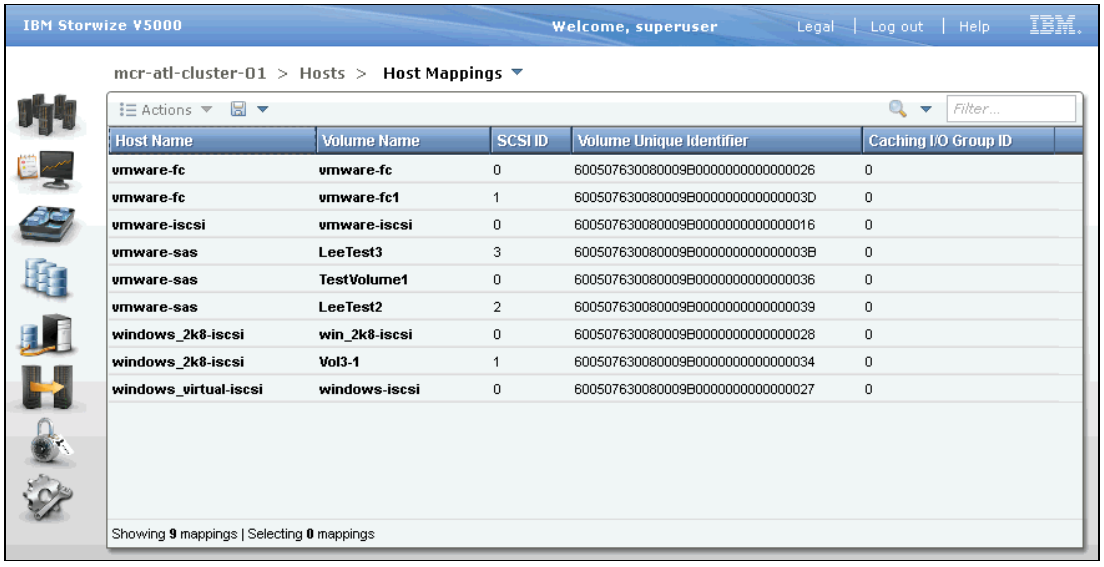

Figure 8-37 Host volume mappings

The window shows a list of all the hosts and volumes and the respective SCSI ID and Volume Unique Identifier (UID). In our example in Figure 8-37, the host `vmware-fc` has two mapped volumes (`vmware-fc` and `vmware-fc1`), and the associated SCSI ID (0 and 1), Volume Name, UID, and Caching I/O Group ID.

If you highlight one line and click **Actions** (as shown in Figure 8-38), the following options are available:

- Unmap Volumes
- Properties (Host)
- Properties (Volume)

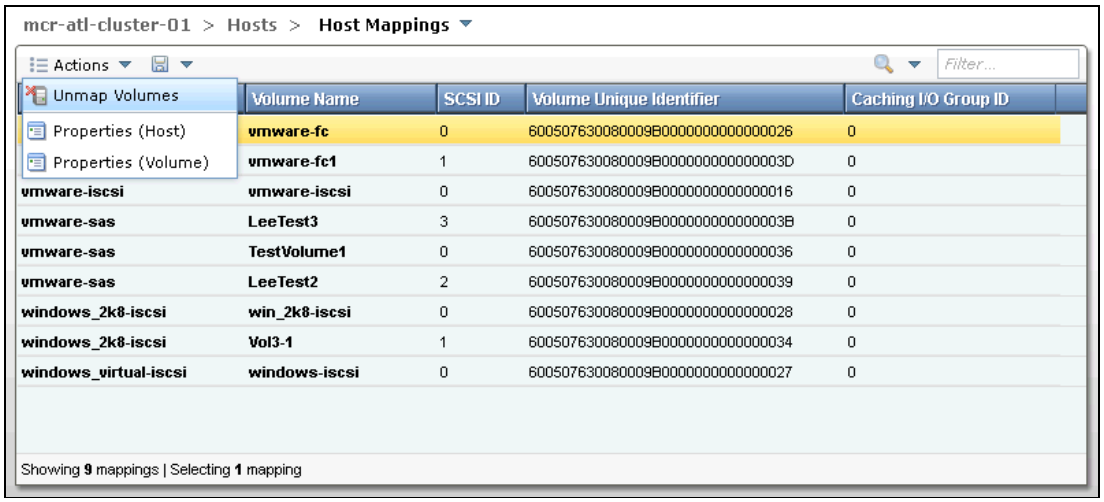

Figure 8-38 Host mapping options

If multiple lines are highlighted (by holding the Ctrl key), only the Unmap Volumes option is available.

8.3.1 Unmap Volumes

Highlight one or more lines and click **Unmap Volumes**, enter the number of volumes to remove (as shown in Figure 8-39), and click **Unmap**. The mappings for all selected entries are removed.

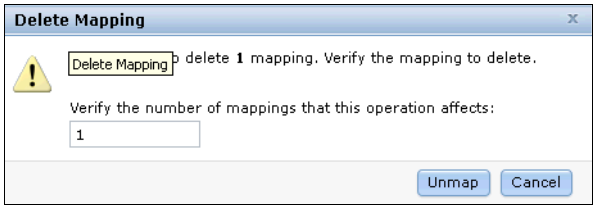

Figure 8-39 Unmapping a volume from host

A window opens that shows the status and completion of volume unmapping. Figure 8-40 shows volume `windows2k8-s` is unmapped from host `windows2k8-sas`.

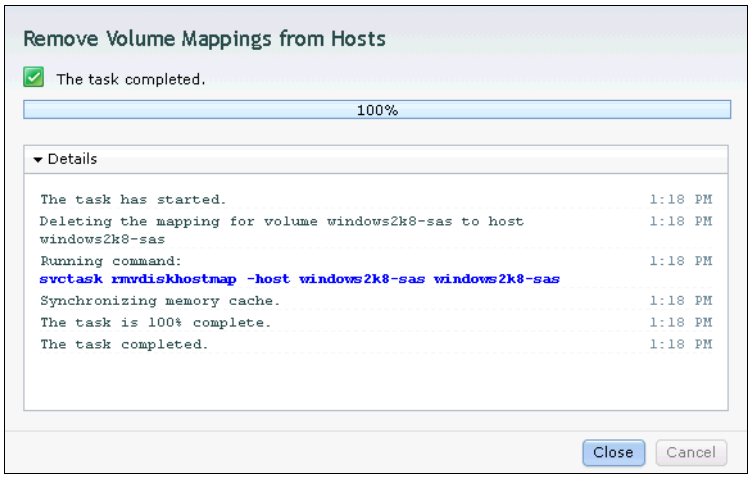

Figure 8-40 Unmapping a volume from host

Warning: Always ensure that you run the required procedures in your host operating system before you unmap volumes in the IBM Storwize V5000 GUI.

8.3.2 Properties (Host)

Selecting an entry and clicking **Properties (Host)** (as shown in Figure 8-38 on page 373) opens the Host Properties window. For more information, see 8.1.5, “Host properties” on page 362.

8.3.3 Properties (Volume)

Selecting an entry and clicking **Properties (Volume)** (as shown in Figure 8-38 on page 373) opens the Volume Properties view. For more information about volume properties, see 8.5, “Volume properties” on page 388.

8.4 Advanced volume administration

This section describes volume administration tasks, such as, volume modification, volume migration, and creation of volume copies. We assume that volumes were created on your IBM Storwize V5000 and you are familiar with generic, thin-provision, mirror, and thin-mirror volumes.

For more information about basic volume configuration, see Chapter 5, “I/O Group basic volume configuration” on page 161.

Figure 8-41 shows the following options that are available within the Volumes menu for advanced features administration:

- Volumes
- Volumes by Pool
- Volumes by Host

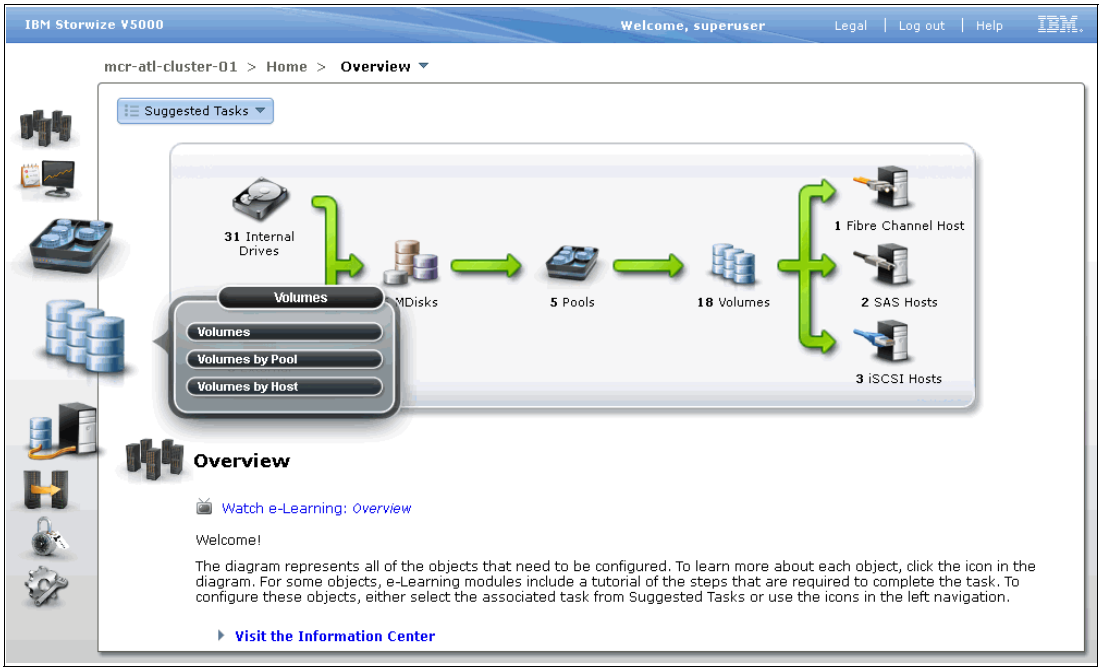

Figure 8-41 Volume options menu

8.4.1 Advanced volume functions

Click **Volumes** (as shown in Figure 8-41 on page 375) and the Volumes window opens, as shown in Figure 8-42.

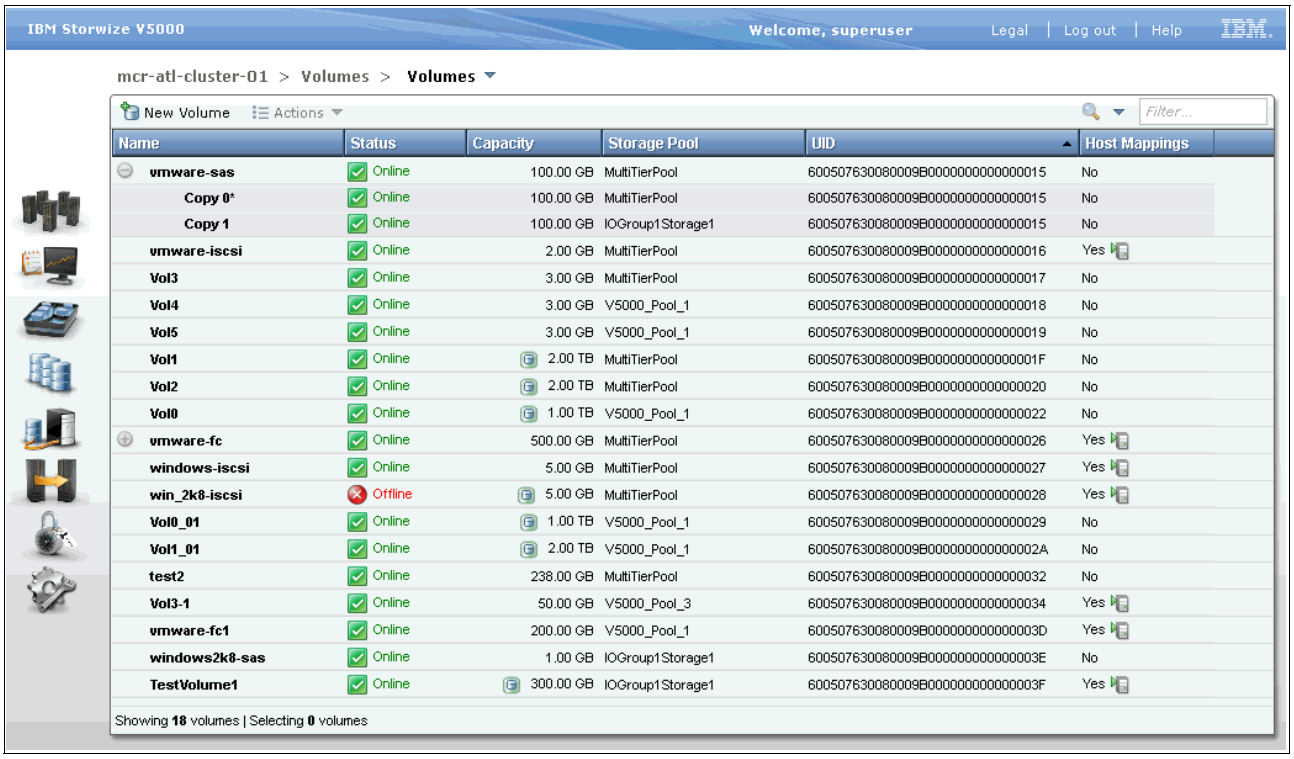

Figure 8-42 Volume window

By default, this window lists all configured volumes on the system and provides the following information:

- Name: Shows the name of the volume. If there is a + sign next to the name, this sign means that there are two copies of this volume. Click the + sign to expand the view and list the copies, as shown in Figure 8-43 on page 377.
- Status: Provides the status information about the volume, which can be online, offline, or degraded.
- Capacity: The disk capacity that is presented to the host. If a blue volume is listed next to the capacity, this means that this volume is a thin-provisioned volume. Therefore, the listed capacity is the virtual capacity, which might be more than the real capacity on the system.
- Storage Pool: Shows in which Storage Pool the volume is stored. The primary copy is shown unless you expand the volume copies.
- UID: The volume unique identifier.
- Host Mappings: Shows if a volume has host mapping: `Yes` when host mapping exists (along with small server icon) and `No` when there are no hosting mappings.

Important: If you right-click anywhere in blue title bar, you can customize the volume attributes that are displayed. You might want to add some useful information, such as, Caching I/O Group and Real Capacity.

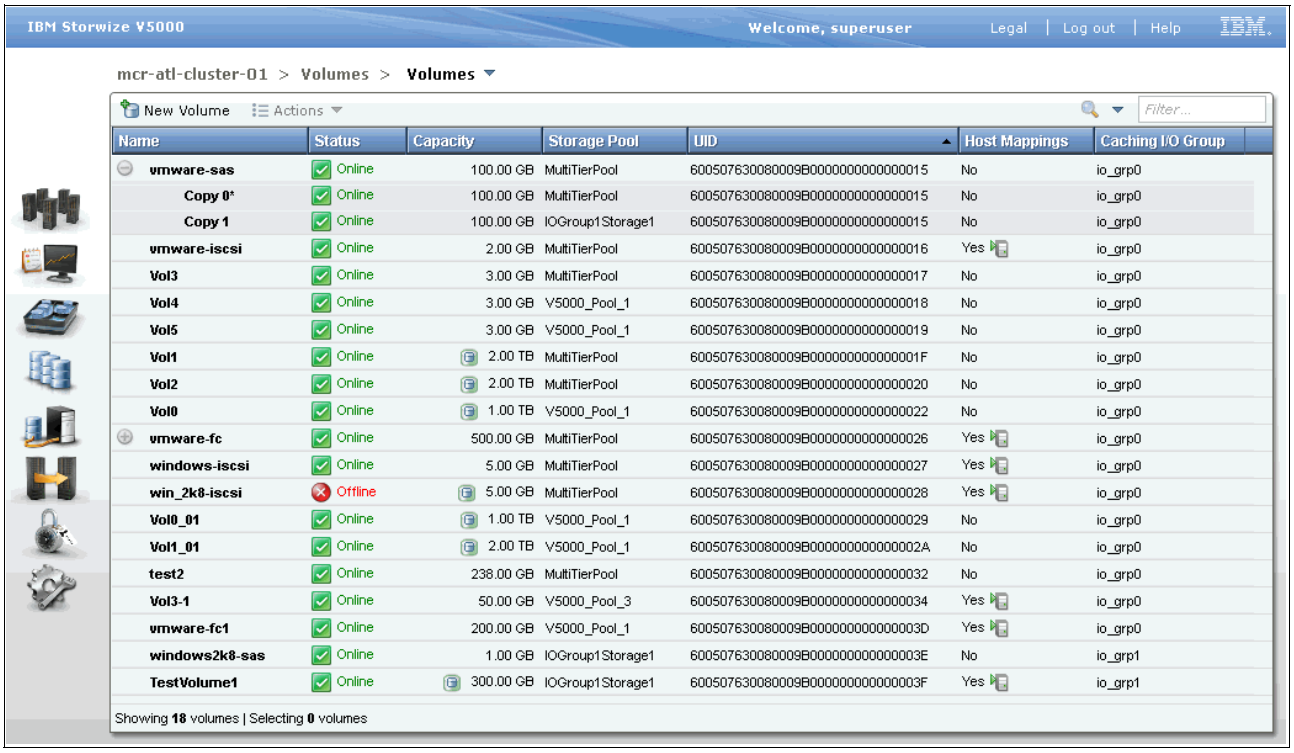

Figure 8-43 Expand volume copies

To create a volume, click **New Volume** and complete the steps that are described in 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162.

You can right-click or highlight a volume and select **Actions** to see the available actions for a volume, as shown in Figure 8-44 on page 378.

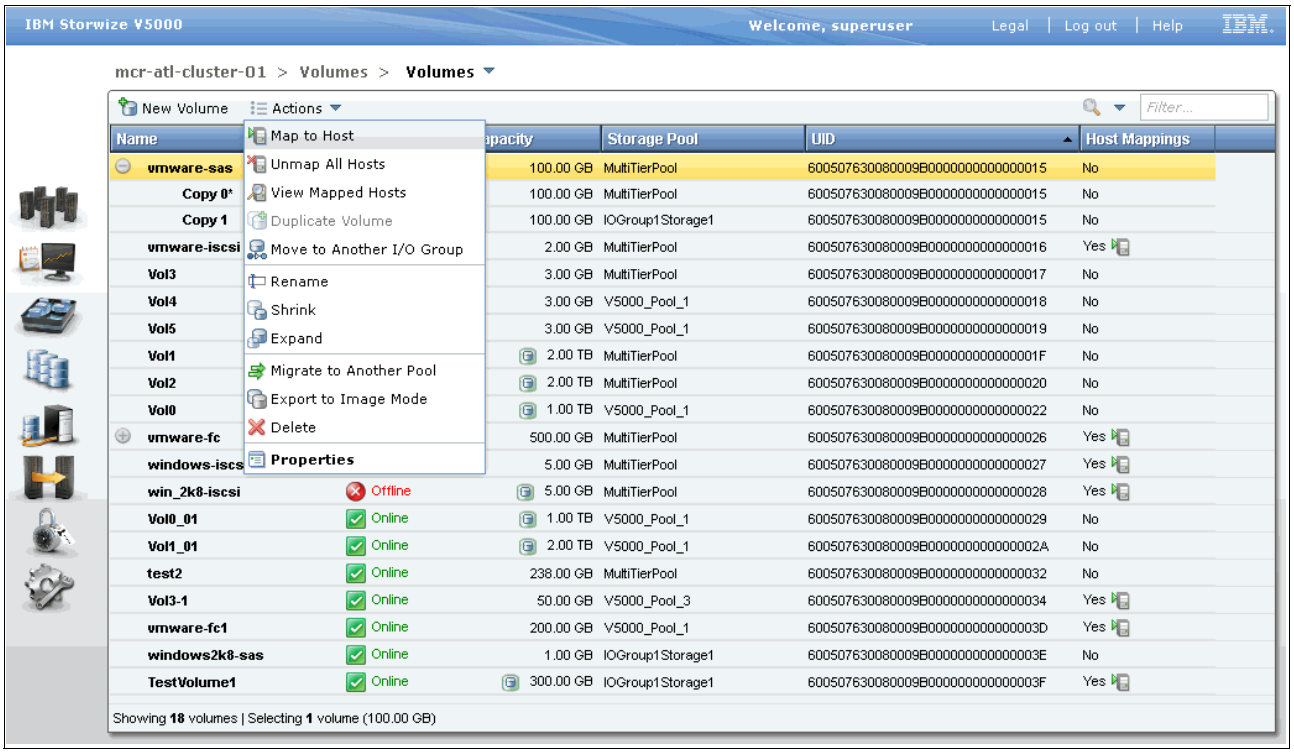

Figure 8-44 Listing the action options for volume.

Depending on which volume you highlighted, the following Volume Copy options are available:

- Map to Host
- Unmap All Hosts
- View Mapped Host
- Duplicate Volume
- Move to Another I/O Group
- Rename
- Shrink
- Expand
- Migration to Another Pool
- Export to Image Mode
- Delete
- Properties

For Thin-Provisioned with single copy, the following options are available:

- Add Mirror Copy: Only available for generic volumes.
- Thin Provisioned: Only available for the following thin-provisioned volumes:
  - Shrink
  - Expand
  - Properties

These options are described in the next sections.

8.4.2 Mapping a volume to a host

If you want to map a volume to a host, select **Map to Host** from the menu that is shown in Figure 8-44 on page 378. Select the I/O Group and Host to which you want to map the volume and click **Next**. Figure 8-45 shows the Modify Host Mappings menu.

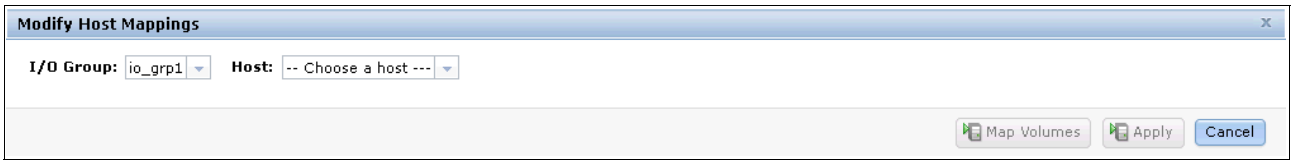

Figure 8-45 Modify Host Mappings menu

Important: You cannot change the caching I/O Group by using the I/O Group drop-down menu. Instead, the menu it is used to list hosts that have access to the specified I/O Group.

After you select a host, the Modify Mappings window opens. In the upper left, you see your I/O Group and selected host. The yellow volume is the selected volume that is ready to be mapped, as shown in Figure 8-46. Click **Map Volumes** to apply the changes to the system.

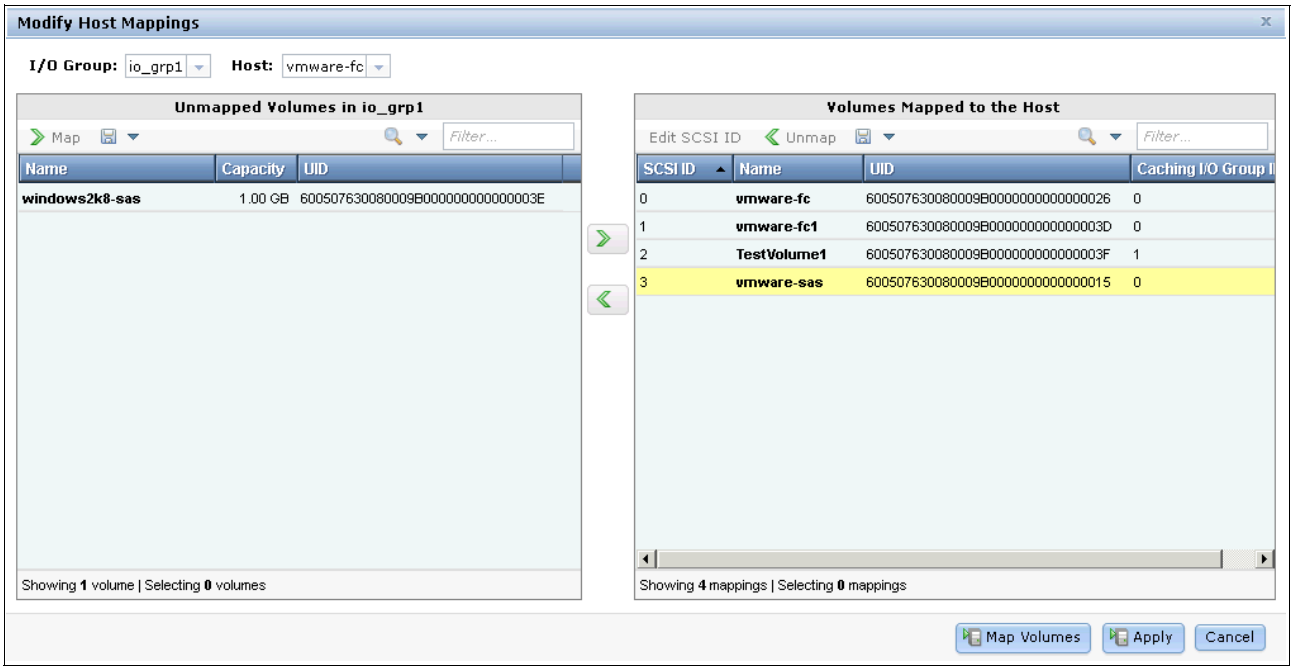

Figure 8-46 Modify Host Mappings

After the changes are made, click **Close** to return to the All Volumes window.

Modify Mappings window: For more information about the Modify Mappings window, see 8.1.1, “Modifying Mappings menu” on page 352.

8.4.3 Unmapping volumes from all hosts

If you want to remove all host mappings from a volume, click **Unmap All Hosts** (as shown in Figure 8-44 on page 378). This action removes all host mappings, which means that no hosts can access this volume. Enter the number of mappings that are affected and click **Unmap**, as show in Figure 8-47.

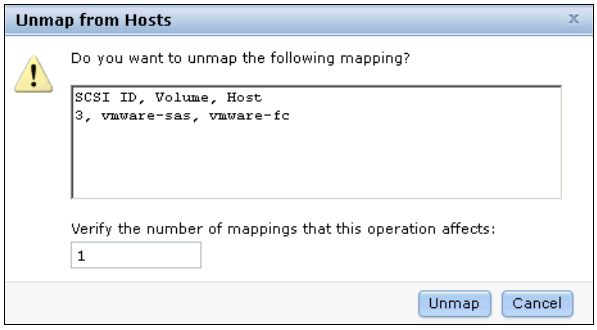

Figure 8-47 Unmapping from host (or hosts)

After the task completes, click **Close** to return to the All Volumes window.

Important Always ensure that you run the required procedures in your host operating system before the unmapping procedure.

8.4.4 Viewing a host that is mapped to a volume

If you want to know which host mappings are configured, highlight a volume and click **View Mapped Host** (as shown in Figure 8-44 on page 378). The Host Maps tab of the Volume Details window opens, as shown in Figure 8-48 on page 381. In this example, you see that there is one existing host mapping to the `vmware-sas` volume.

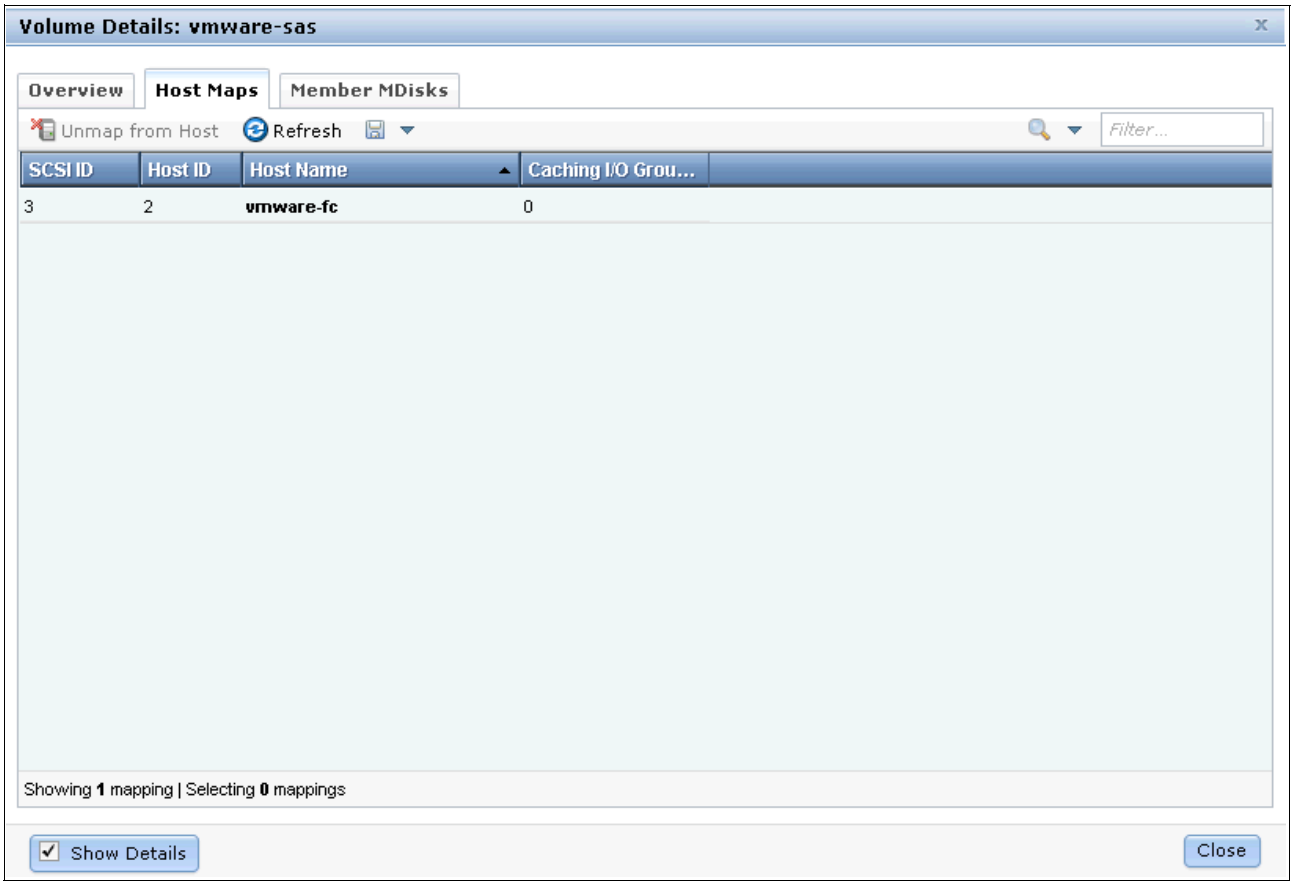

Figure 8-48 Volume to host mapping

If you want to remove a mapping, highlight the host and click **Unmap from Host**, which removes the access for the selected host after you confirm it. If several hosts are mapped to this volume (for example, in a cluster), only the highlighted host is removed.

8.4.5 Renaming a volume

To rename a volume, select **Rename** (as shown in Figure 8-44 on page 378). The Rename Volume window opens. Enter the new name, as shown in Figure 8-49.

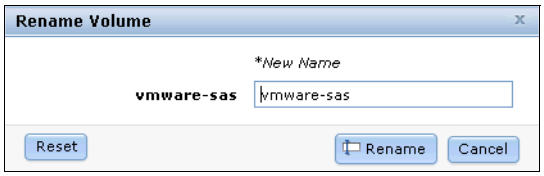

Figure 8-49 Renaming a volume

If you click **Reset**, the name field is reset to the active name of the volume. Click **Rename** to apply the changes and click **Close** after task window completes.

8.4.6 Shrinking a volume

The IBM Storwize V5000 can shrink volumes. This feature should be used only if your host operating system supports it. This capability reduces the capacity that is allocated to the particular volume by the amount that you specify. To shrink a volume, click **Shrink**, as shown in Figure 8-44 on page 378. You can enter the new size or by how much the volume should shrink. If you enter a value, the other line updates automatically, as shown in Figure 8-50.

Important: Before you shrink a volume, ensure that the volume is not mapped to any host object and does not contain data. If both conditions are ignored, it is likely that your operating system logs disk errors or data corruption occurs.

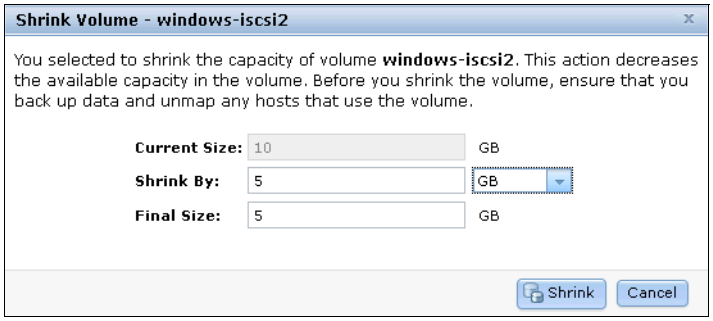

Figure 8-50 Shrink Volume window

Click **Shrink** to start the process and then click **Close** when task window completes and to return to the All Volumes window.

Run the required procedures on your host after the shrinking process.

Important: For volumes that contain more than one copy, you might receive a CMMVC6354E error; run the **`lsvdisksyncprogress`** command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by running the **`chvdisk`** command. When the copy is synchronized, resubmit the shrink process.

8.4.7 Expanding a volume

If you want to expand a volume, click **Expand** (as shown in Figure 8-44 on page 378) and the Expand Volume window opens. Before you continue, check if your operating system supports online volume expansion. Enter the new volume size and click **Expand**, as shown in Figure 8-51 on page 383.

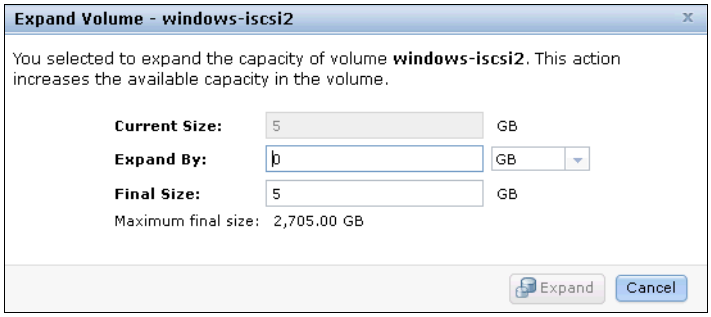

Figure 8-51 Expand Volume window

After the tasks complete, click **Close** to return to the All Volumes window.

Run the required procedures in your operating system to use the available space.

8.4.8 Migrating a volume to another storage pool

The IBM Storwize V5000 supports online volume migration while applications are running. By using volume migration, you can move volumes between storage pools, whether the pools are internal pools or on an external storage system. The migration process is a low priority and one extent is moved at a time and has a slight effect on the performance of the IBM Storwize V5000.

Important: For the migration to be acceptable, the source and target storage pool must have the same extent size. For more information about extent size, see Chapter 1, “Overview of the IBM Storwize V5000 system” on page 1.

To migrate a volume to another storage pool, click **Migrate to Another Pool** (as shown in Figure 8-44 on page 378). The Migrate Volume Copy window opens. If your volume consists of more than one copy, you are asked which copy you want to migrate to another storage pool, as shown in Figure 8-52. If the selected volume consists of one copy, this option does not appear. Notice that the `vmware-sas` volume has two copies stored in two different storage pools. The storage pools to which they belong are shown in parentheses.

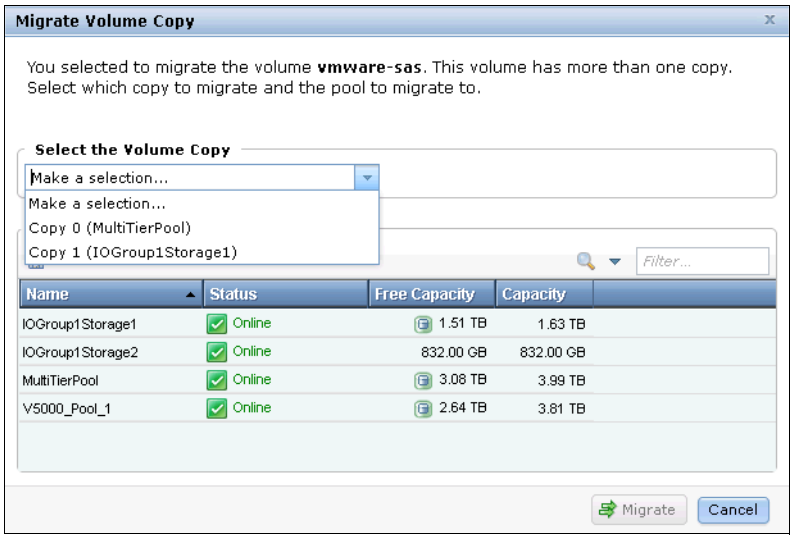

Figure 8-52 Migrate Volume

Select the new target storage pool and click **Migrate**, as shown in Figure 8-52 on page 383.

The volume copy migration starts, as shown in Figure 8-53. Click **Close** to return to the All Volumes window.

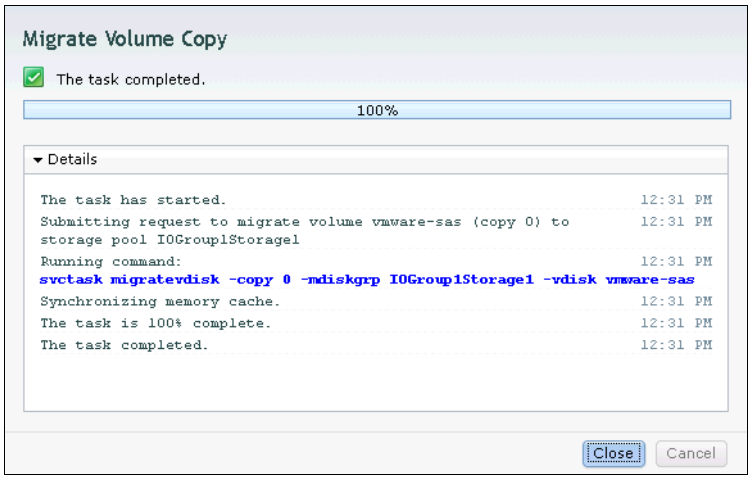

Figure 8-53 Volume Copy Migration starts

Depending on the size of the volume, the migration process can take some time. You can monitor the status of the migration in the running tasks bar at the bottom of the window. Volume migration tasks cannot be interrupted.

After the migration completes, the “copy 0” from the `vmware-sas` volume is shown in the new storage pool, as shown in Figure 8-54.

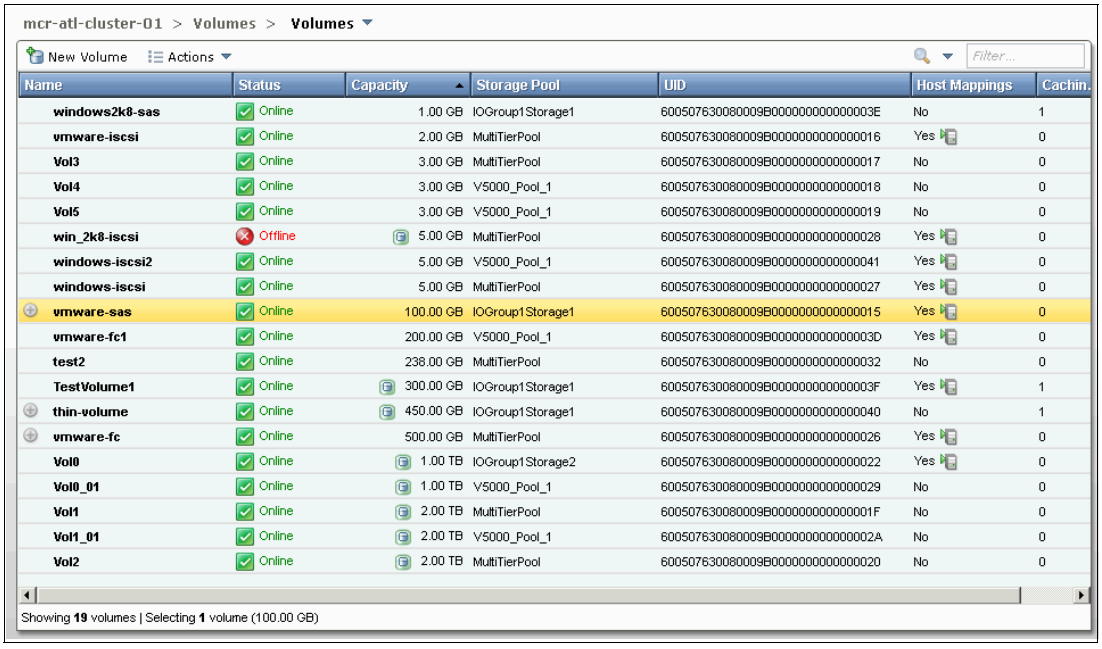

Figure 8-54 Volume showing at new storage pool

The volume copy was migrated without any downtime to the new storage pool. It is also possible to migrate both volume copies to other storage pools.

The volume copy feature also can be used to migrate volumes to a different pool, as described in 8.6.5, “Migrating volumes by using the volume copy features” on page 404.

8.4.9 Exporting to an image mode volume

Image mode provides a direct block-for-block translation from MDisk to a Volume with no virtualization. An image mode MDisk is associated with exactly one volume. This feature can be used to export a volume to a non-virtualized disk and to remove the volume from storage virtualization.

To export a volume to an image volume, browse to IBM Storwize V5000 GUI and click **Volumes** → **Volumes**, as shown in Figure 8-55.

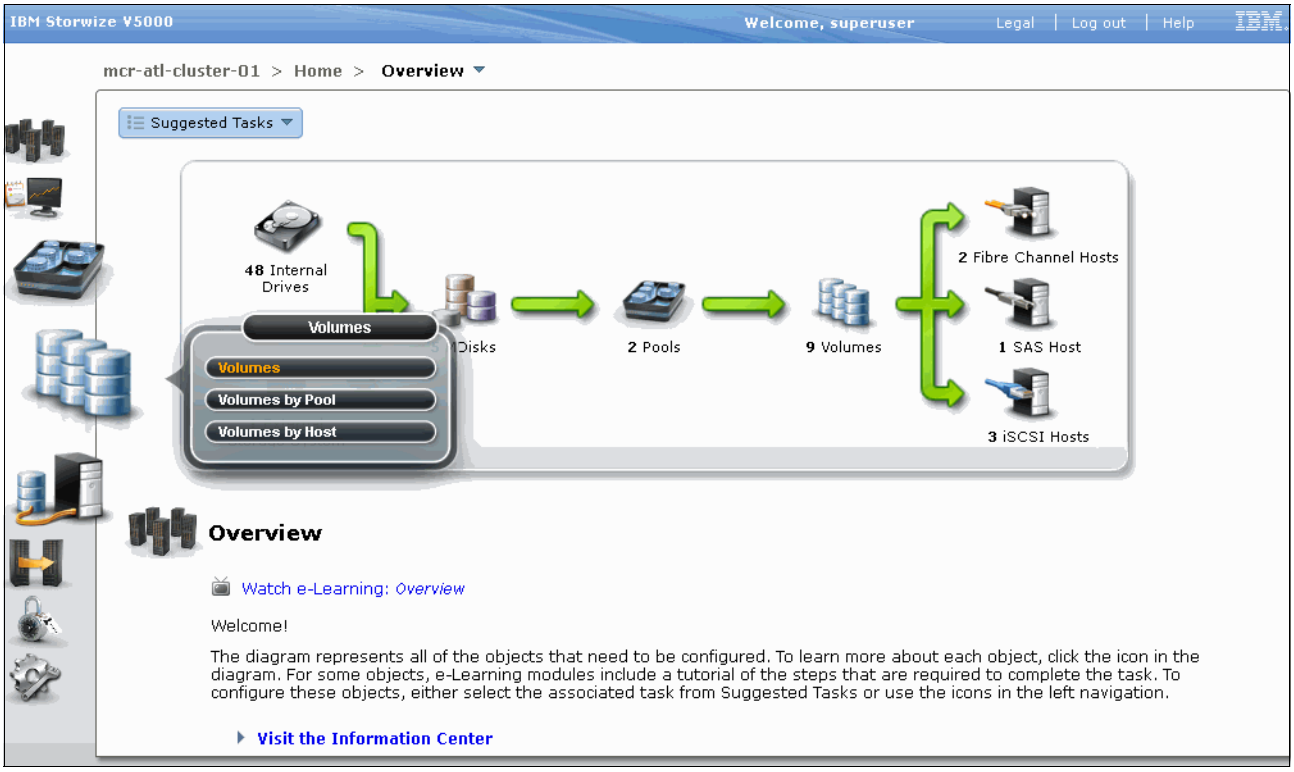

Figure 8-55 Exporting a volume to an image mode

Highlight the volume that you want to export to an image mode and, from the Actions menu, select **Export to Image Mode**, as shown in Figure 8-56.

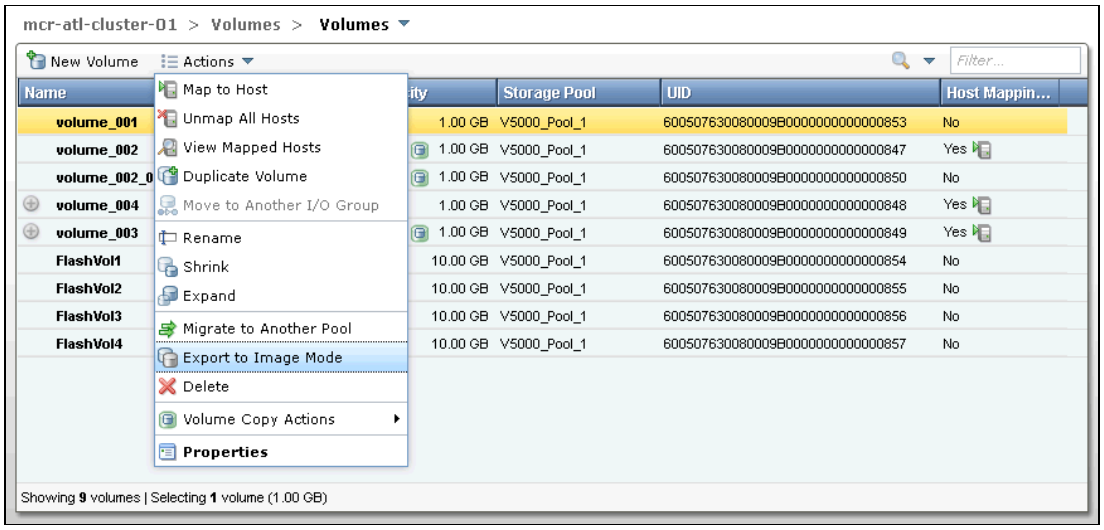

Figure 8-56 Exporting a volume to an image mode

The Export to Image Mode wizard opens that shows all available MDisk. Select the MDisk you want to export and click **Next.** In our example, we are exporting the volume `volume_001` to an image mode MDisk named `mdisk5`, as shown in Figure 8-57.

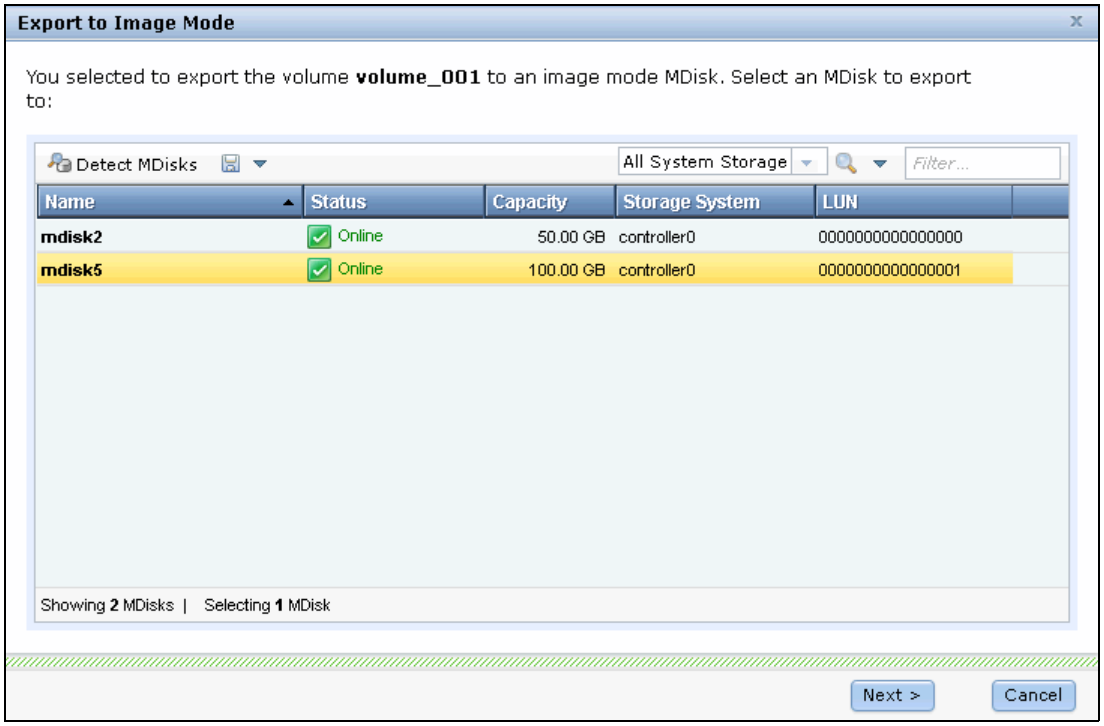

Figure 8-57 Selecting the Manage Disk to export the volume.

By clicking **Next**, you must select the storage pool into which the image-mode volume is placed after migration is completed, as shown in Figure 8-58.

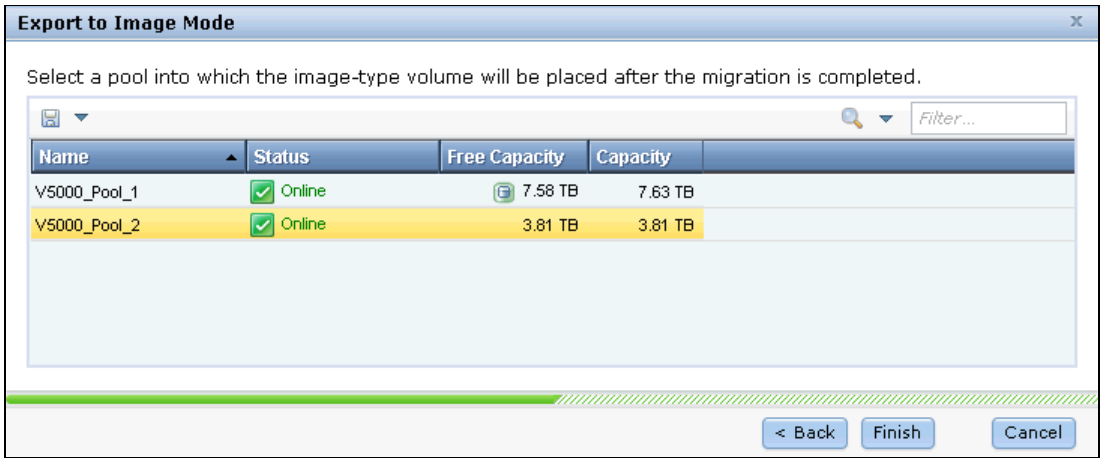

Figure 8-58 Select the Storage Pool

Click **Finish** to start the migration. After the task is complete, click **Close** to return to Volumes window.

Important: Use image mode to import or export existing data into or out of the IBM Storwize V5000. Migrate such data from image mode MDisks to other storage pools to benefit from storage virtualization.

For more information about importing volumes from external storage, see Chapter 6, “Storage migration wizard” on page 237 and Chapter 7, “Storage pools” on page 295.

8.4.10 Deleting a volume

To delete a volume, select **Delete**, as shown in Figure 8-44 on page 378. Enter the number of volumes that you want to delete and select the option if you want to force the deletion. Figure 8-59 shows the Delete Volume window.

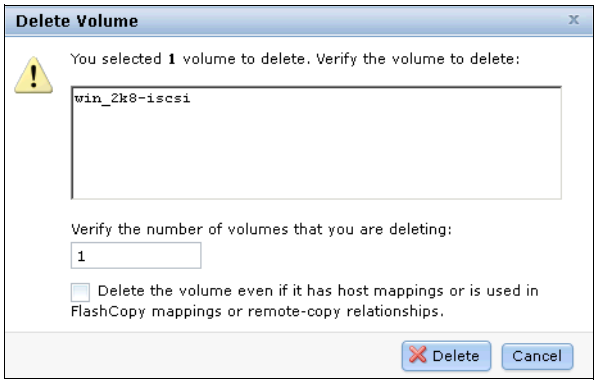

Figure 8-59 Delete Volume window

Click **Delete** and the volume is removed from the system. If you must force a volume removal, select the option.

Click **Close** to return to Volumes window**.**

Important: You must force the deletion if the volume has host mappings or is used in FlashCopy mappings. To be cautious, always ensure that the volume has no association before you delete it.

8.5 Volume properties

This section provides an overview of all available information that is related to IBM Storwize V5000 volumes.

To open the advanced view of a volume, select **Properties** (as shown in Figure 8-44 on page 378), and the Volume Details window opens, as shown in Figure 8-60. The following tabs are available:

- Overview
- Host Maps
- Member MDisk

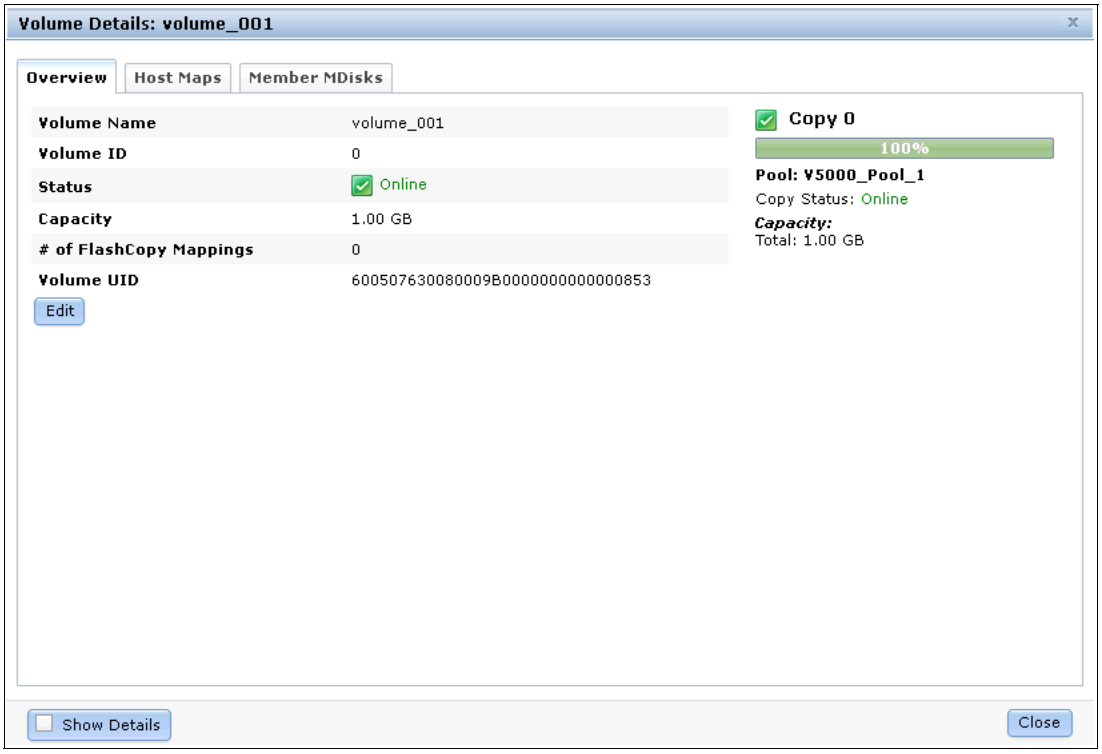

Figure 8-60 Volume Details: Overview tab

8.5.1 Overview tab

The Overview tab that is shown in Figure 8-61 on page 389 gives you a complete overview of the volume properties. In the left part of the window, you find common volume properties. In the right part of the window, you see information about the volume copies. The detailed view was chosen by clicking the **Show Details** option in the lower left.

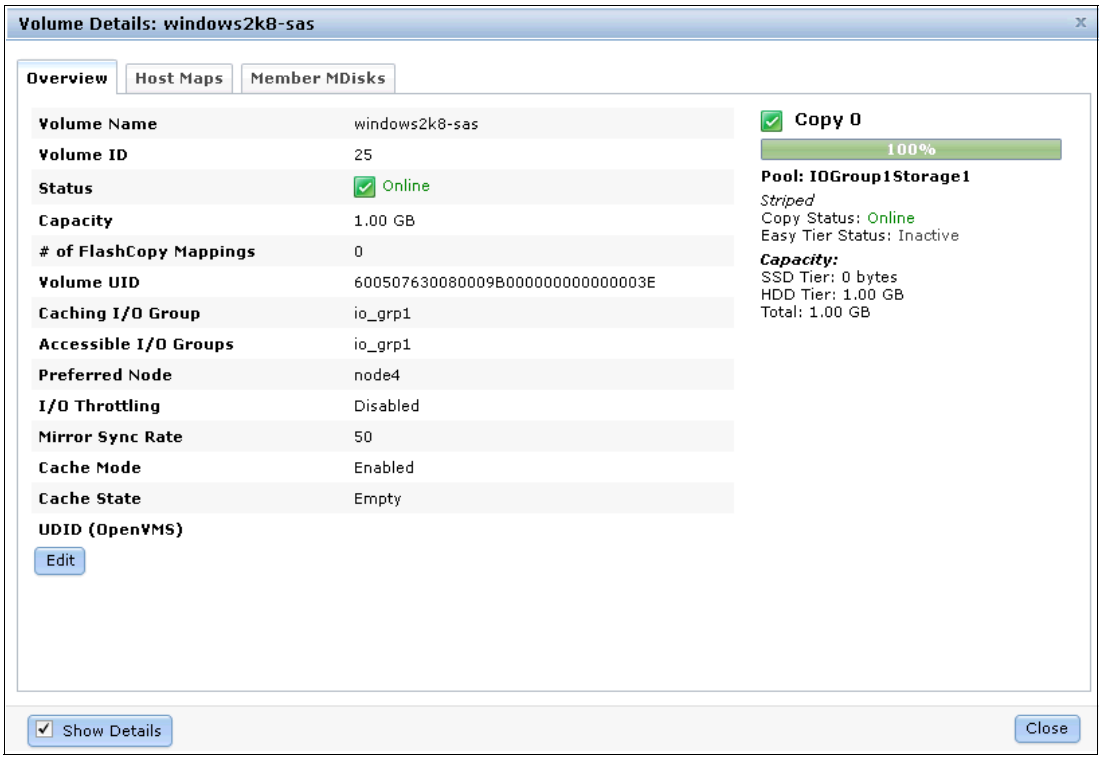

Figure 8-61 Volume properties window

The following details are available:

- Volume Properties:
  - Volume Name: Shows the name of the volume.
  - Volume ID: Shows the ID of the volume. Every volume has a system-wide unique ID.
  - Status: Gives status information about the volume, which can be online, offline, or degraded.
  - Capacity: Shows the capacity of the volume. If the volume is thin-provisioned, this number is the virtual capacity; the real capacity is displayed for each copy.
  - # of FlashCopy Mappings: The number of existing FlashCopy relationships. For more information, see Chapter 10, “Copy services” on page 449.
  - Volume UID: The volume unique identifier.
  - Caching I/O Group: Specifies the volume Caching I/O Group.
  - Accessible I/O Group: Shows the I/O Group the host can use to access the volume.
  - Preferred Node: Specifies the ID of the preferred node for the volume.
  - I/O Throttling: It is possible to set a maximum rate at which the volume processes I/O requests. The limit can be set in I/Os to MBps. This feature is an advanced feature and it is possible to enable it only through the CLI, as described in Appendix A, “Command-line interface setup and SAN Boot” on page 609.
  - Mirror Sync Rate: After creation, or if a volume copy is offline, the mirror sync rate weights the synchronization process. Volumes with a high sync rate (100%) complete the synchronization faster than volumes with a lower priority. By default, the rate is set to 50% for all volumes.
  - Cache Mode: Shows if the cache is enabled or disabled for this volume.

- Cache State: Provides feedback if open I/O requests are inside the cache that is not destaged to the disks.
- UDID (OpenVMS): The unit device identifiers are used by OpenVMS hosts to access the volume.

- Copy Properties:
  - Storage Pool: Provides information about which pool the copy is in, what type of copy it is (generic or thin-provisioned), the status of the copy, and Easy Tier status.
  - Capacity: Shows the allocated (used) and the virtual (Real) capacity from both Tiers (SSD and HDD) and the warning threshold, and the grain size for Thin-Provisioned volumes.

If you want to modify any of these settings, click **Edit** and the window changes to modify mode. Figure 8-62 shows the Volume Details Overview tab in modify mode.

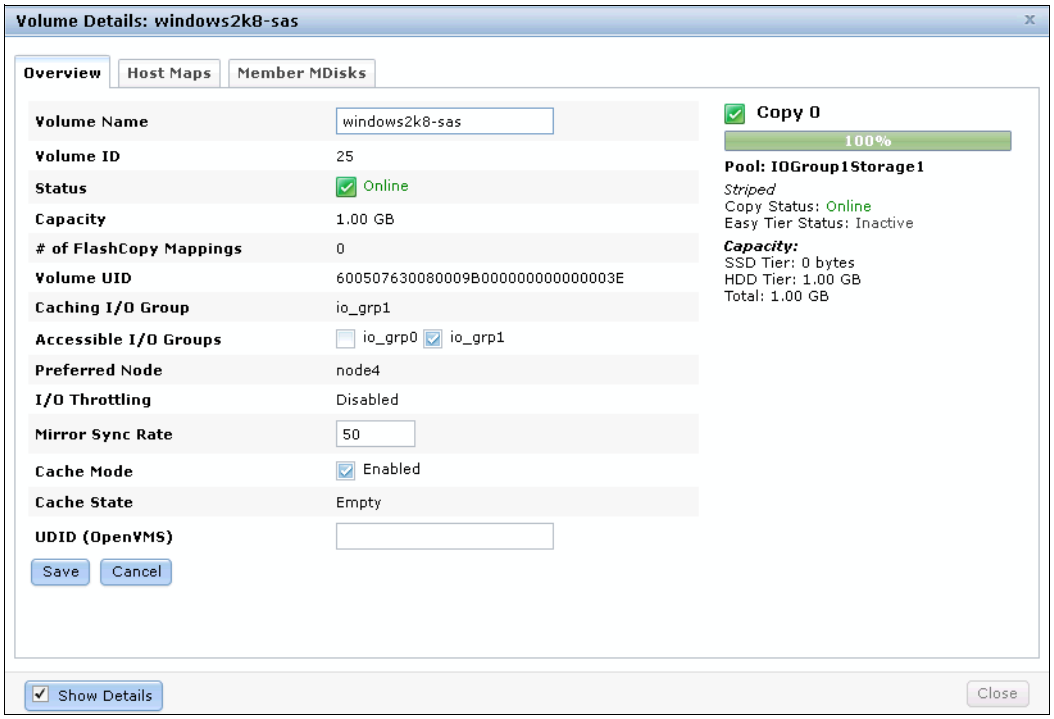

Figure 8-62 Modify Volume Details window

In the Modify Volume Details window, the following properties can be changed:

- Volume Name
- I/O Group
- Mirror Sync Rate
- Cache Mode
- UDID

Make any required changes and click **Save**.

Important: Changing the I/O Group can cause loss of access because of cache reload and host-I/O Group access. Also, setting the Mirror Sync Rate to 0% disables synchronization.

8.5.2 Host Maps tab

The second tab of the Volume Properties window is Host Maps, as shown in Figure 8-63. All hosts that are mapped to the selected volume are listed in this view.

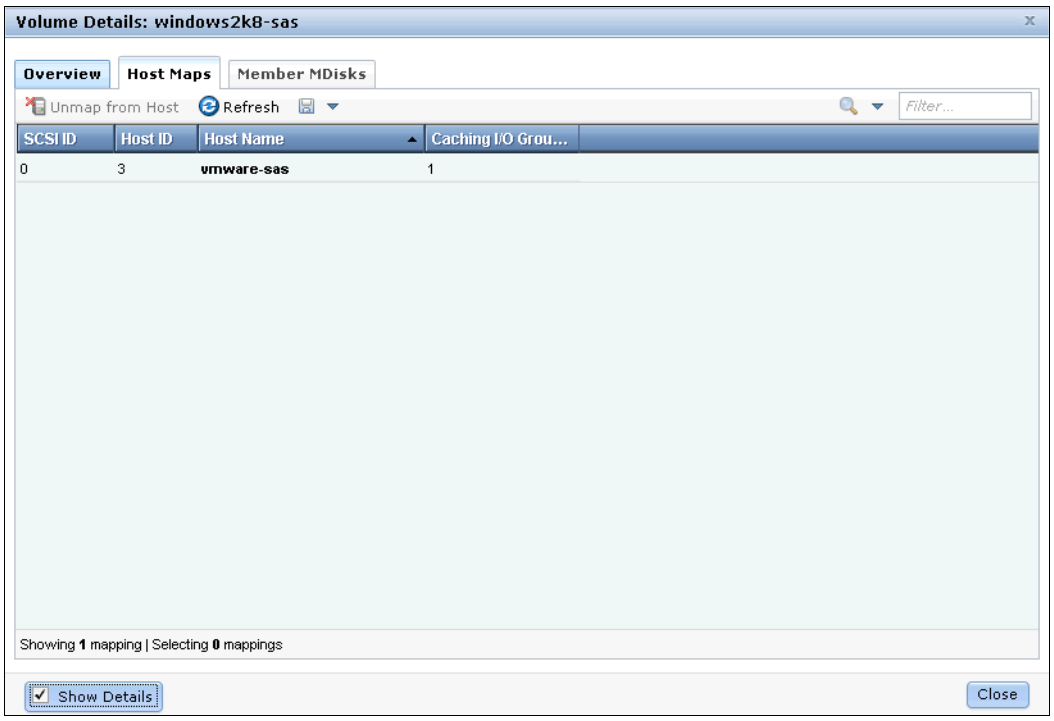

Figure 8-63 Host Maps

To unmap a host from the volume, highlight it and click **Unmap from Host**. Confirm the number of mappings to remove and click **Unmap**. Figure 8-64 shows the Unmap Host window.

Figure 8-64 Unmap Host window

The changes are applied to the system. The selected host no longer has access to this volume. Click **Close** to return to the Host Maps window. For more information about host mappings, see 8.3, “Host mappings overview” on page 373.

8.5.3 Member MDisk tab

The third tab is the Member MDisk tab, which lists all MDisks on which the volume is located. Select a copy and the associated MDisks is shown in the window, as shown in Figure 8-65.

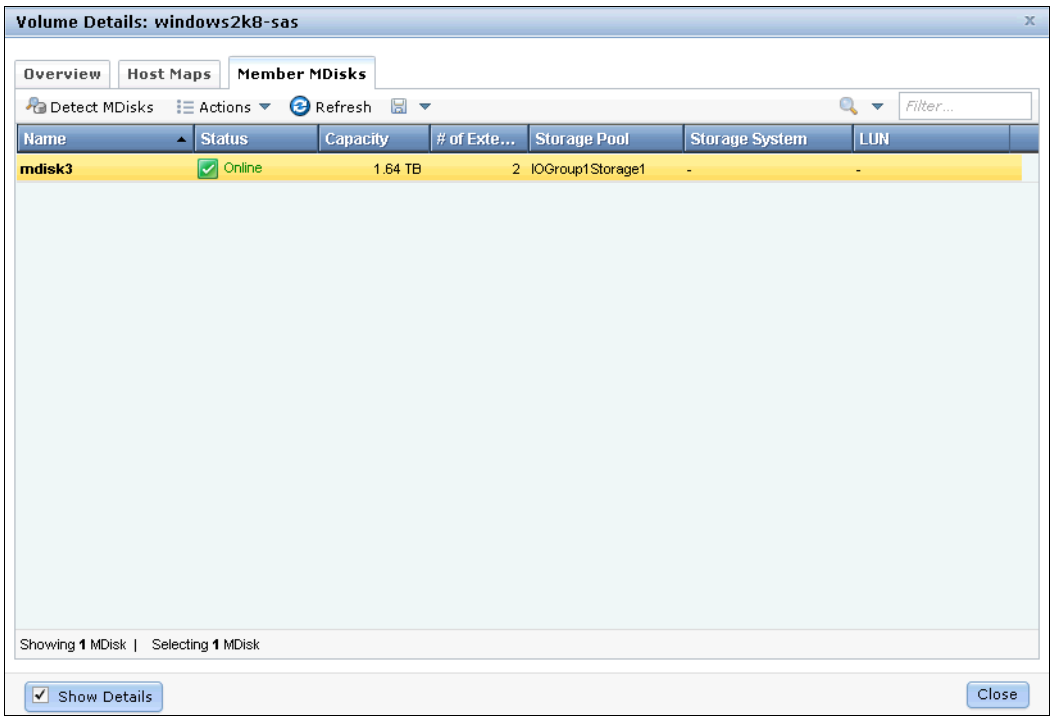

Figure 8-65 Member MDisk tab

When an image mode volume is using external storage, you should see the Storage Subsystem name and the external LUN ID, as shown in Figure 8-66 on page 393.

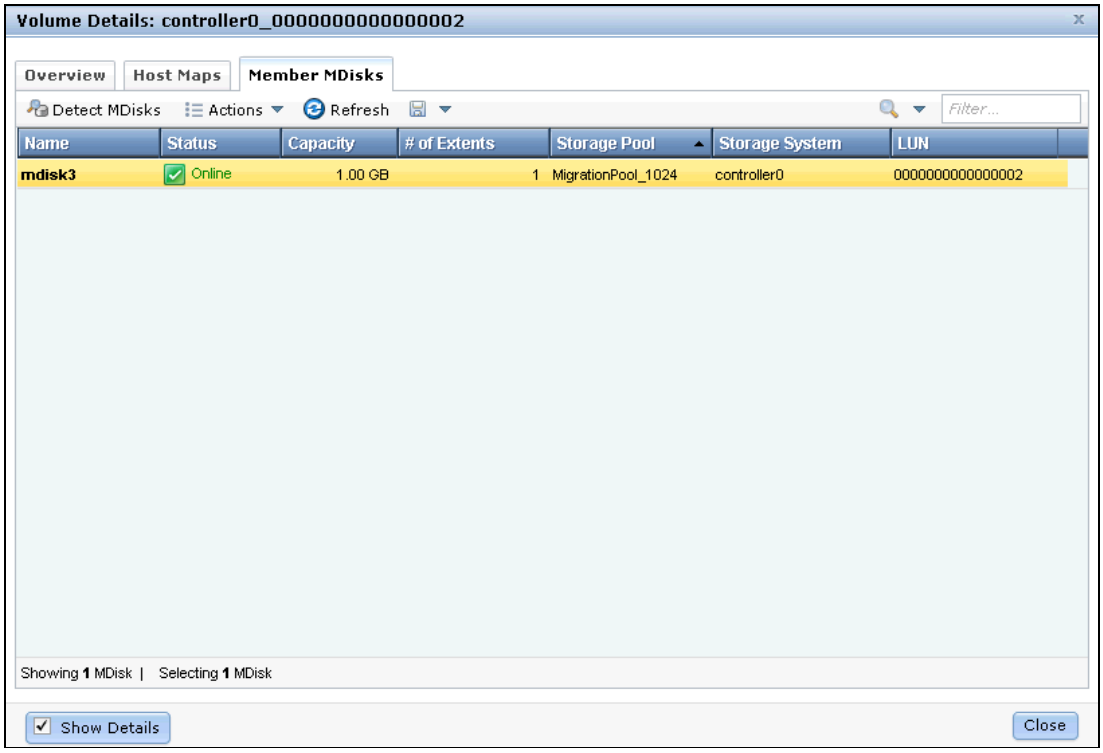

Figure 8-66 Image mode volume details

Highlight an MDisk and click **Actions** to see the available tasks, as shown in Figure 8-67. The Show Details option in the lower left does not provide more information. For more information about the available tasks, see Chapter 7, “Storage pools” on page 295.

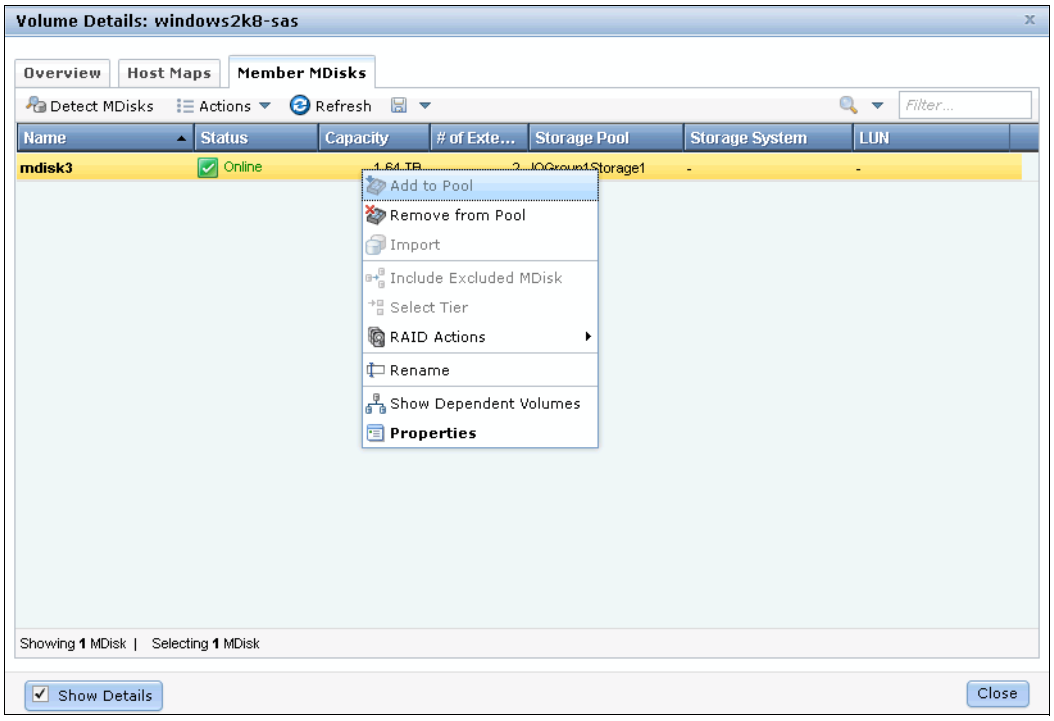

Figure 8-67 MDisk action menu

Click **Close** to return to the All Volumes window.

8.5.4 Adding a mirrored volume copy

If you have a volume that consists of only one copy, you can add a second mirrored copy to the volume. This action creates a second online copy of your volume. This second copy can be generic or thin-provisioned.

You also can use this method to migrate data across storage pools with different extent size.

To add a second copy, highlight the volume and click **Actions** → **Volume Copy Actions** → **Add Mirrored Copy**, as shown in Figure 8-68.

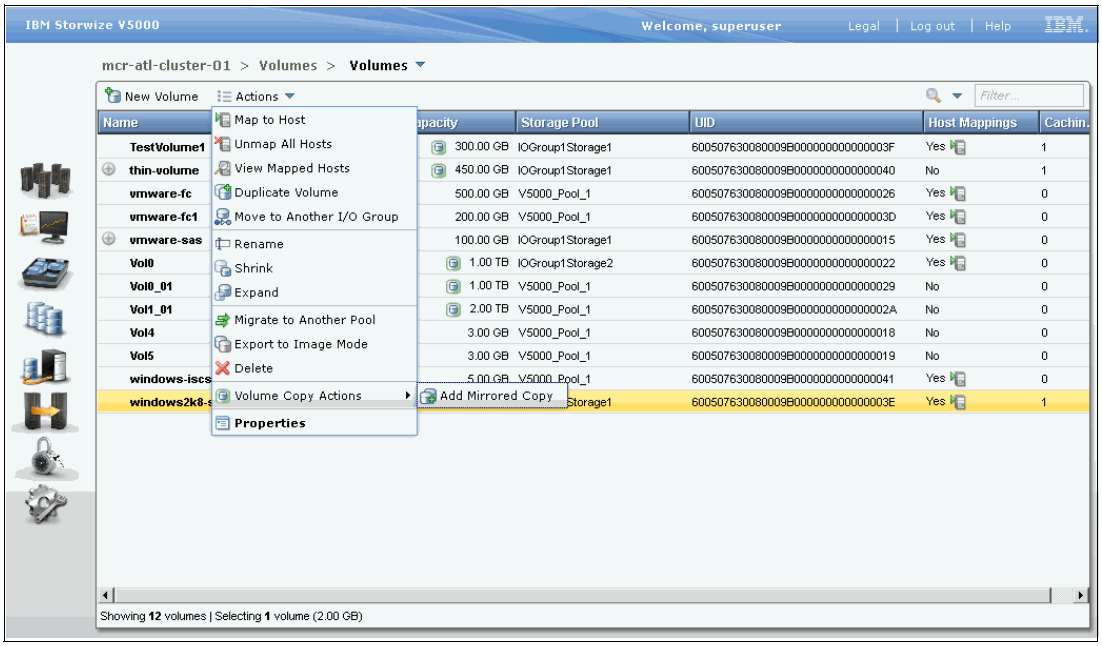

Figure 8-68 Add mirrored copy

Select the storage pool to which the new copy should be created, as shown in Figure 8-69. If the new copy should be thin-provisioned, select the Thin-Provisioned option and click **Add Copy**.

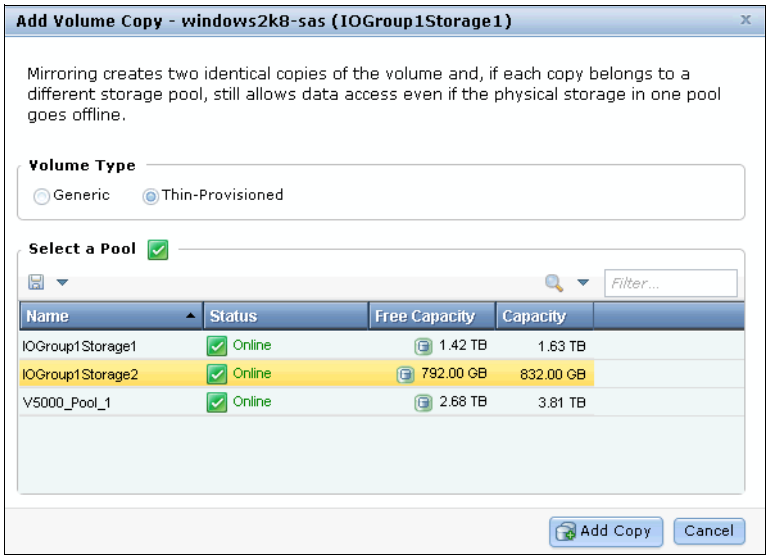

Figure 8-69 Select storage pool

The copy is created after you click **Add Copy** and data starts to synchronize as a background task. Figure 8-70 shows you that the volume named `windows2k8-sas` now has two volume copies that are stored in two different storage pools.

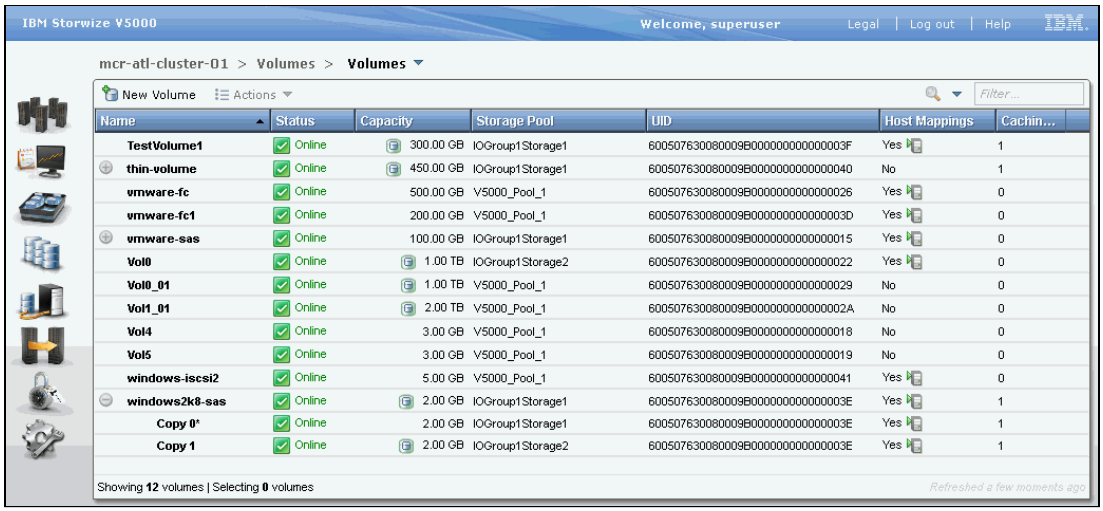

Figure 8-70 Volume containing two copies

8.5.5 Editing thin-provisioned volume properties

The processes that are used to modify the volume size that is presented to a host are described in 8.4.6, “Shrinking a volume” on page 382 and 8.4.7, “Expanding a volume” on page 382. However, if you have a thin-provisioned volume, you can also edit the allocated size and the warning thresholds. To edit these settings, select the volume copy, then select **Actions** → **Thin-Provisioned** or highlight and right-click **Thin-Provisioned** → **Shrink**, as shown in Figure 8-71 on page 396.

The following options are available as shown in Figure 8-71:

- Shrink
- Expand
- Edit Properties

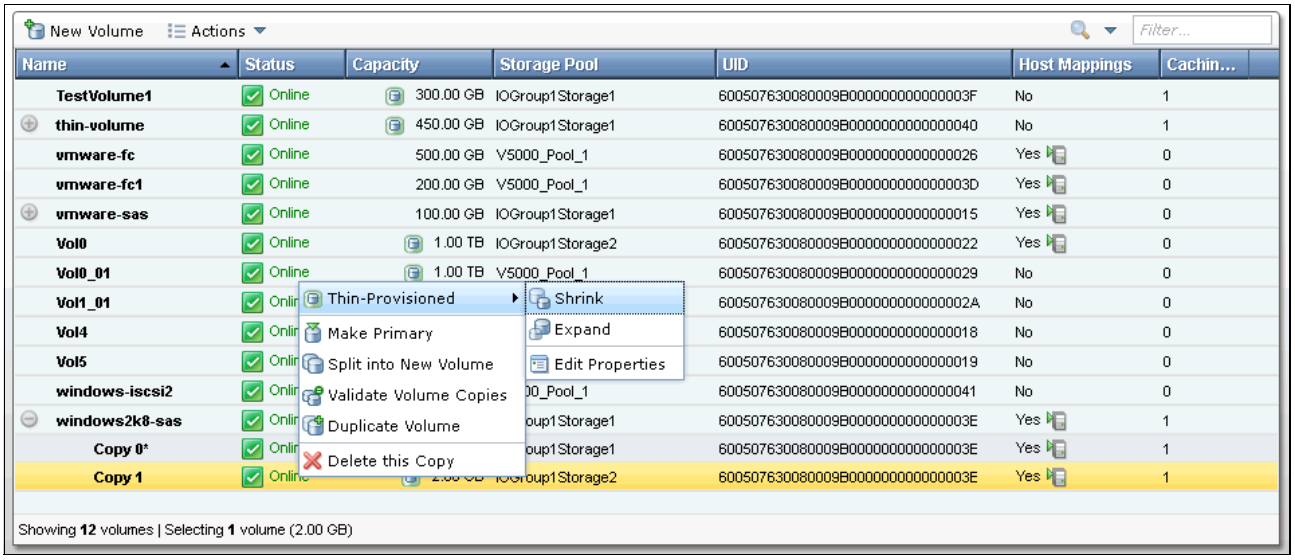

Figure 8-71 Working with thin-provisioned volumes

These changes are made only to the internal storage; no changes to your host are necessary.

Shrinking thin-provisioned space

Select **Shrink** (as shown in Figure 8-71) to reduce the allocated space of a thin-provisioned volume. Enter the amount by which the volume should shrink or the new final size and click **Shrink**.

Deallocating extents: You can deallocate only extents that do not include stored data on them. If the space is allocated because there is data on them, you cannot shrink the allocated space and an out-of-range warning message appears.

Figure 8-72 shows the Shrink Volume window.

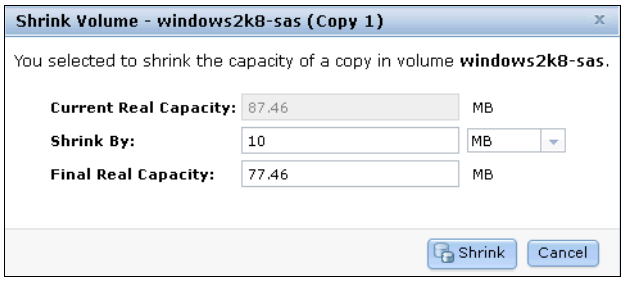

Figure 8-72 Shrink Volume window

After the task completes, click **Close**. The allocated space of the thin-provisioned volume is reduced.

Expanding thin-provisioned space

To expand the allocated space of a thin-provisioned volume, select **Expand**, as shown in Figure 8-71 on page 396. Enter the amount by which space should be allocated or the new final size and click **Expand**. In our example that is shown in Figure 8-73, we are expanding the thin-provisioned space by 10 MB.

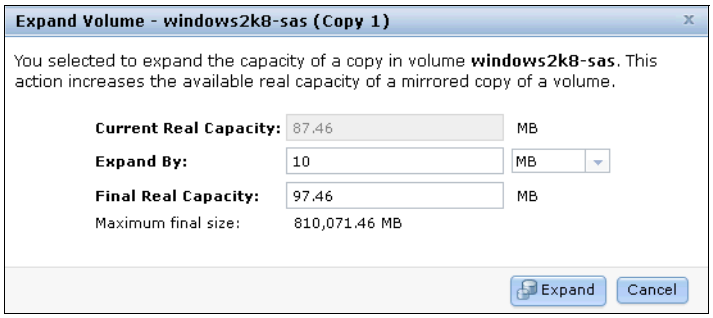

Figure 8-73 Expand Volume window

The new space is now allocated. Click **Close** after task is complete.

Editing thin-provisioned properties

To edit thin-provisioned properties, select **Edit Properties**, as shown in Figure 8-71 on page 396. Edit the settings (if required) and click **OK** to apply the changes.

Figure 8-74 shows the Edit Properties window.

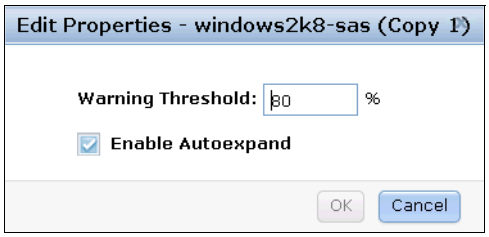

Figure 8-74 Edit Properties window

After the task completes, click **Close** to return to the All Volumes window.

8.6 Advanced volume copy functions

In 8.4.1, “Advanced volume functions” on page 376, we described all of the available actions at a volume level and how to create a second volume copy. In this section, we focus on volumes that consist of two volume copies and how to apply the concept of two copies for business continuity and data migration.

If you expand the volume and highlight a copy, the following volume copy actions are available, as shown in Figure 8-75:

- Thin-provisioned (for Thin volumes)
- Make Primary (for non-primary copy)
- Split into New Volume
- Validate Volume Copies
- Delete Copy option

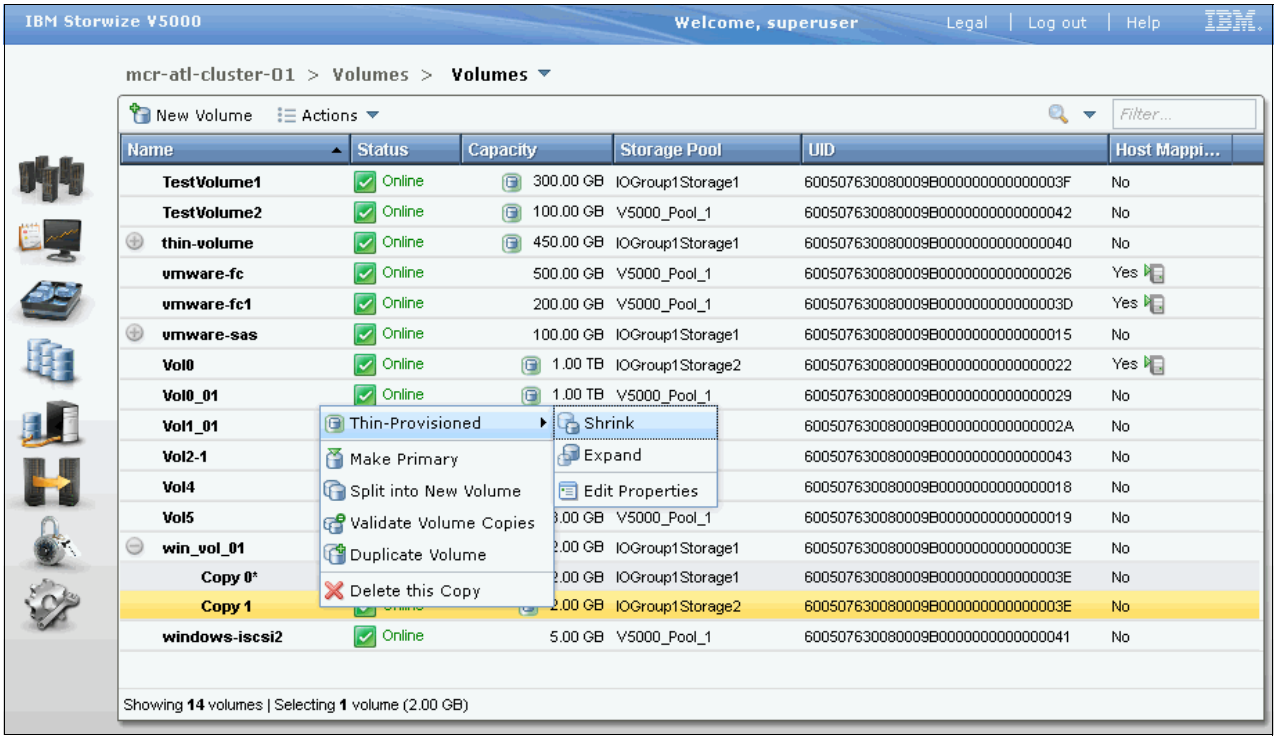

Figure 8-75 Volume copy actions

If you look at the volume copies that are shown in Figure 8-75, you see that one of the copies has a star displayed next to its name, as shown in Figure 8-76.

Figure 8-76 Volume copy names

Each volume has a primary and a secondary copy, and the star indicates the primary copy. The two copies are always synchronized, which means that all writes are destaged to both copies, but all reads are always done from the primary copy. Two copies per volume is the maximum number configurable and you can change the roles of your copies.

To accomplish this task, highlight the secondary copy and then click **Actions** → **Make Primary**. Usually, it is a best practice to place the volume copies on storage pools with similar performance because the write performance is constrained if one copy is on a lower performance pool.

Figure 8-77 shows the secondary copy Actions menu.

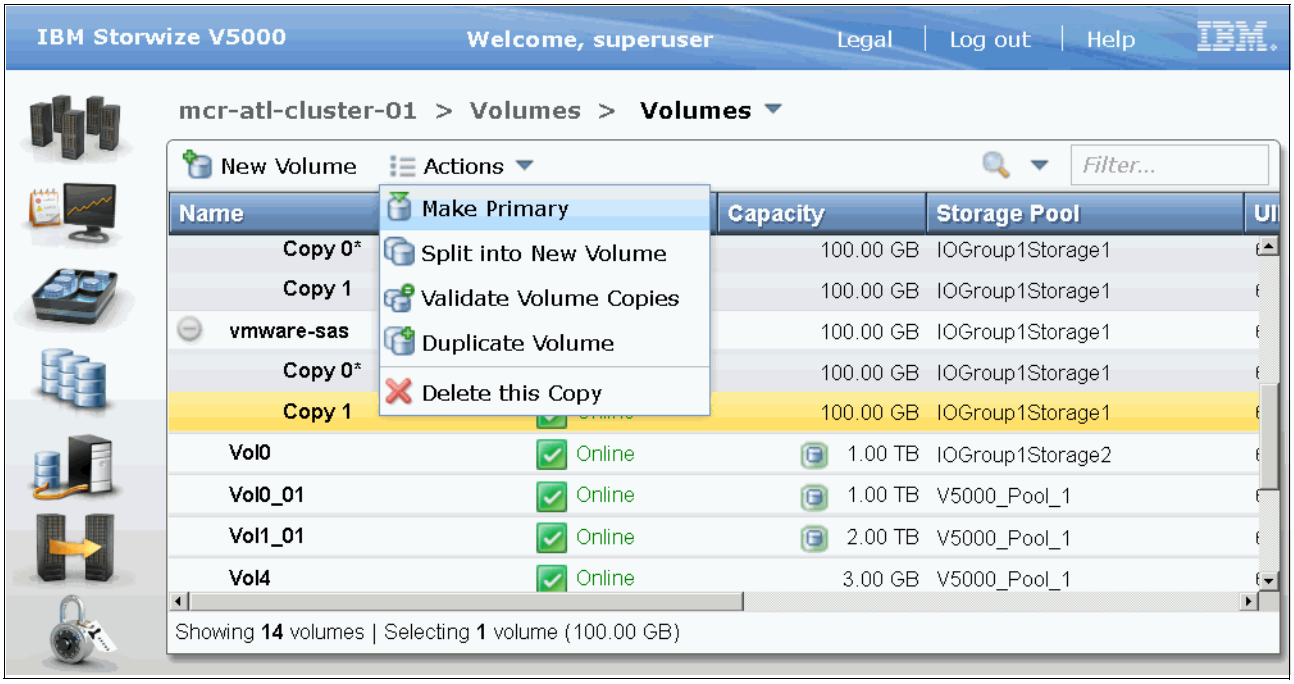

Figure 8-77 Make primary

If you demand high read performance only, another possibility is to place the primary copy in an SSD pool and the secondary copy in a normal disk storage pool. This action maximizes the read performance of the volume and makes sure that you have a synchronized second copy in your less expensive disk pool. It is possible to migrate online copies between storage pools. For more information about how to select which copy you want to migrate, see 8.4.8, “Migrating a volume to another storage pool” on page 383.

Click **Make Primary** and the role of the copy is changed to online. Click **Close** when the task completes.

The volume copy feature also is a powerful option for migrating volumes, as described in 8.6.5, “Migrating volumes by using the volume copy features” on page 404.

8.6.1 Thin-provisioned menu

This menu item includes the same functions that are described in “Shrinking thin-provisioned space” on page 396, “Expanding thin-provisioned space” on page 397, and “Editing thin-provisioned properties” on page 397. You can specify the same settings for each volume copy.

Figure 8-78 shows the Thin-Provisioned menu item.

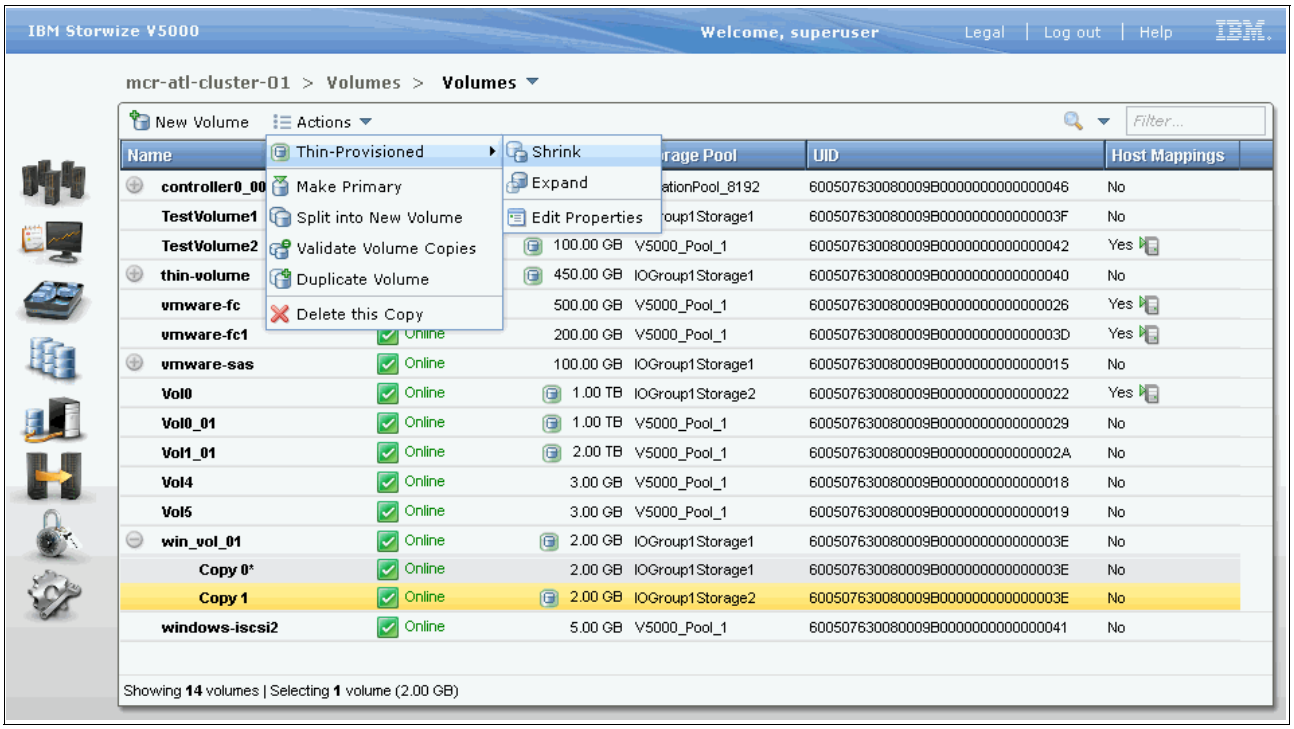

Figure 8-78 Thin-Provisioned menu item

8.6.2 Splitting into a new volume

If your two-volume copies are synchronized, you can split one of the copies to a new volume and map this new volume to another host. From a storage point of view, this procedure can be performed online, which means you can split one copy from the volume and create a copy from the remaining volume without any host impact. However, if you want to use the split copy for testing or backup purposes, you must make sure that the data inside the volume is consistent. Therefore, you must flush the data to storage to make the copies consistent.

For more information about flushing the data, see your operating system documentation. The easiest way to flush the data is to shut down the hosts or application before a copy is split.

In our example, volume `win_vol_01` has two copies: Copy 0 as primary and Copy 1 as secondary. To split a copy, click **Split into New Volume** (as shown in Figure 8-75 on page 398) on any copy and the remaining secondary copy automatically becomes the primary for the source volume.

Optionally, enter a name for the new volume and click **Split Volume Copy**.

Figure 8-79 shows the Split Volume Copy window.

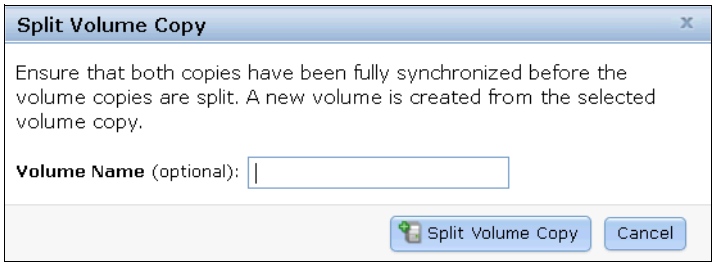

Figure 8-79 Split Volume Copy window

After the task completes, click **Close** to return to the All Volumes window, where the copy appears as a new volume named `vdisk0` that can be mapped to a host, as shown in Figure 8-80.

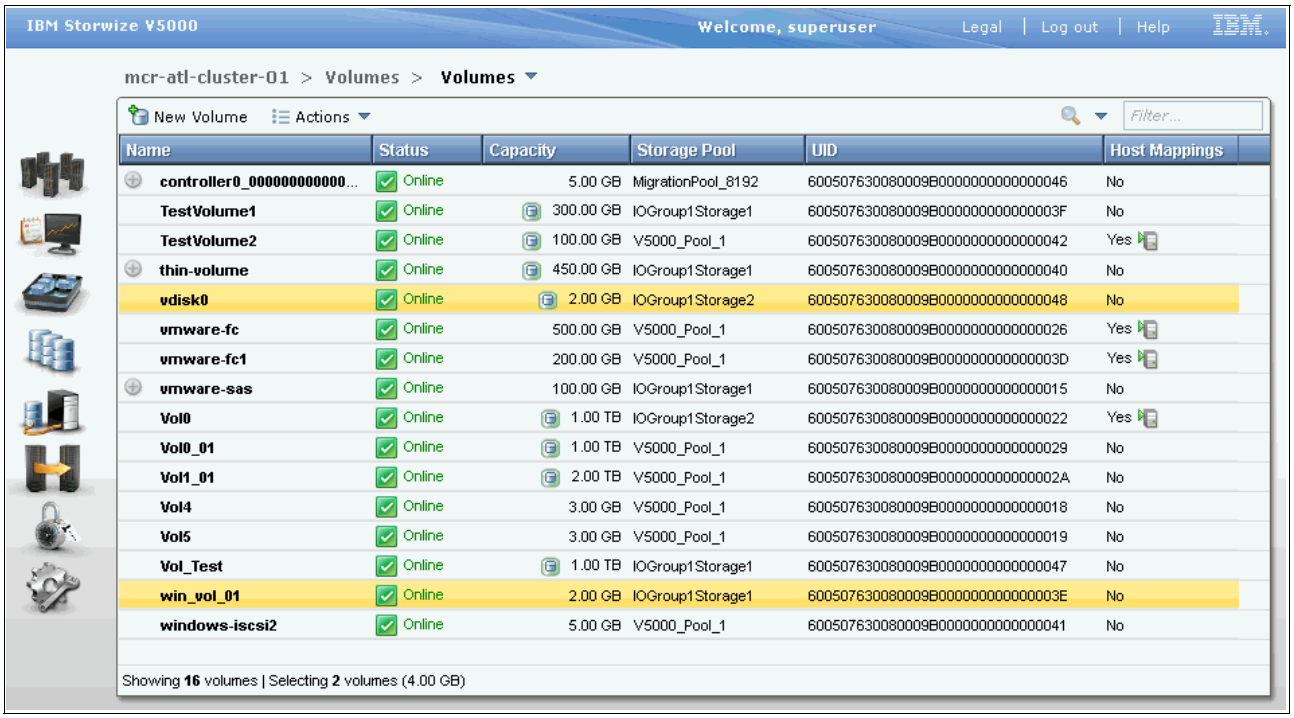

Figure 8-80 All Volumes: New volume from split copy

Important: If you receive error message code `CMMVC6357E` while you are splitting volume copy, use the **`lsvdisksyncprogress`** command to view the synchronization status or wait for the copy to synchronize. Example 8-1 on page 402 shows an output of **`lsvdisksyncprogress`** command.

Example 8-1 Output of lsvdisksyncprogress command

```
IBM_Storwize:mcr-atl-cluster-01:superuser>lsvdisksyncprogress
vdisk_id vdisk_name  copy_id progress estimated_completion_time
3        vmware-sas  1       3        130605014819
14       thin-volume 1       38       130606032210
25       win_vol_01  1       55       130604121159
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

8.6.3 Validate Volume Copies option

By using the IBM Storwize V5000 GUI, you can check volume copies that are identical or process the differences between them.

To validate the copies of a mirrored volume, complete the following steps:

1. Select **Validate Volume Copies**, as shown in Figure 8-75 on page 398. The Validate Volume Copies window opens, as shown in Figure 8-81.

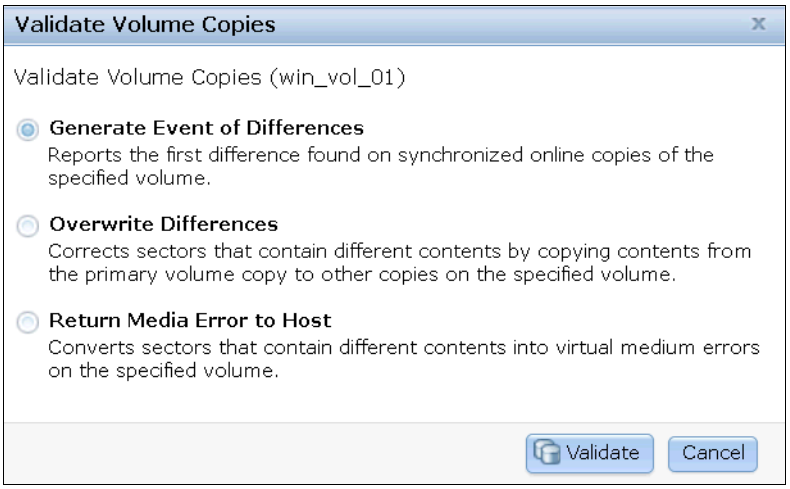

Figure 8-81 Validate Volume Copies window

The following options are available:

– Generate Event of Differences

  Use this option if you want to verify only that the mirrored volume copies are identical. If any difference is found, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this option, starting at a different LBA each time, to count the number of differences on a volume.

– Overwrite Differences

  Use this option to overwrite contents from the primary volume copy to the other volume copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies that are compared. Upon completion, the command process logs an event that indicates the number of differences that were corrected. Use this option if you are sure that the primary volume copy data is correct or that your host applications can handle incorrect data.

– Return Media Error to Host

  Use this option to convert sectors on all volume copies that contain different contents into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is and you do not want an incorrect version of the data to be used.

2. Select which action to perform and click **Validate** to start the task. The volume is now checked. Click **Close**.

   Figure 8-82 shows the output when a volume copy **Generate Event of Differences** option is chosen.

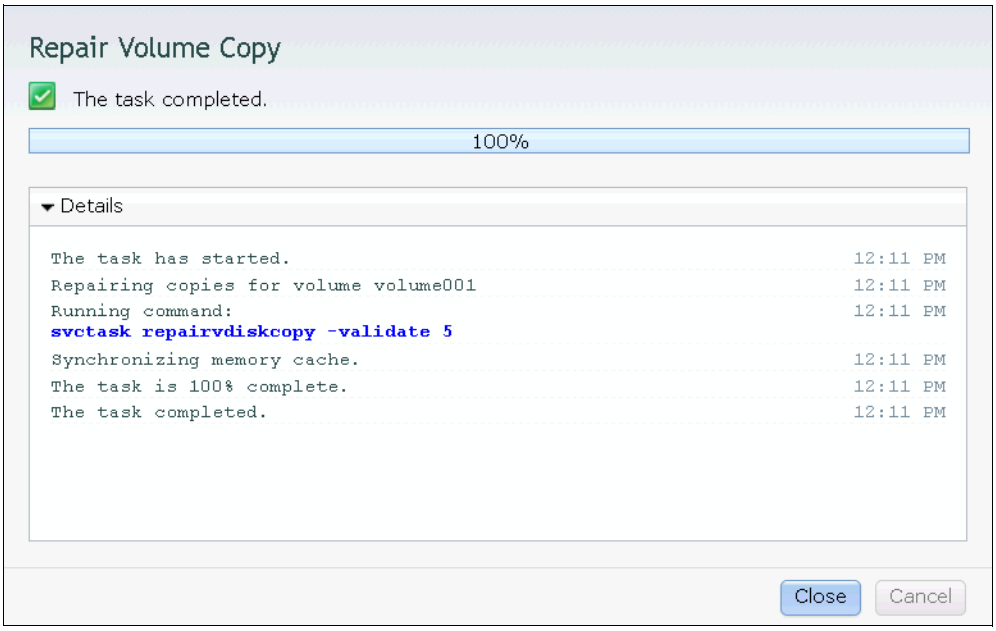

Figure 8-82 Volume copy validation output

The validation process runs as a background process and might take some time, depending on the volume size. You can check the status in the Running Tasks window, as shown in Figure 8-83 on page 404.

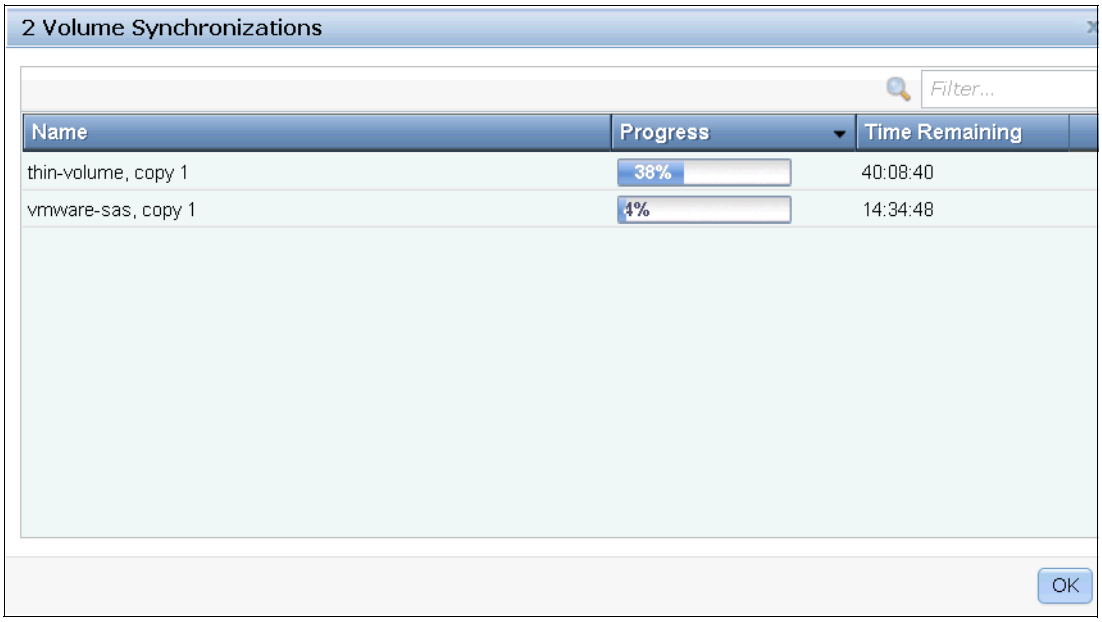

Figure 8-83 Validate Volume Copies: Running Tasks

8.6.4 Delete Volume Copy option

Click **Delete** (as shown in Figure 8-75 on page 398) to delete a volume copy. The copy is deleted, but the volume remains online by using the remaining copy. Confirm the deletion process by clicking **Yes**. Figure 8-84 shows the copy deletion warning window.

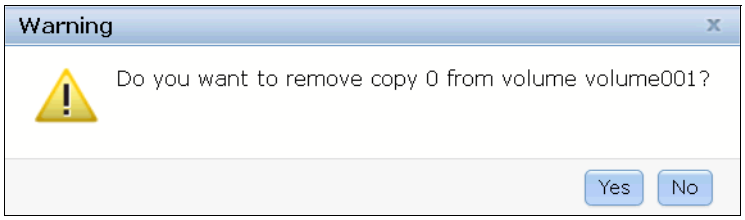

Figure 8-84 Delete a copy

After the copy is deleted, click **Close** to return to the All Volumes window.

8.6.5 Migrating volumes by using the volume copy features

In the previous sections, we showed that it is possible to create, synchronize, split, and delete volume copies. A combination of these tasks can be used to migrate volumes to other storage pools.

The easiest way to migrate volume copies is to use the migration feature that is described in 8.4.8, “Migrating a volume to another storage pool” on page 383. If you use this feature, one extent after another is migrated to the new storage pool. However, the use of volume copies provides another way to migrate volumes if you have different storage pool characteristics in terms of extent size.

To migrate a volume, complete the following steps:

1. Create a second copy of your volume in the target storage pool. For more information, see 8.5.4, “Adding a mirrored volume copy” on page 394.
2. Wait until the copies are synchronized.
3. Change the role of the copies and make the new copy the primary copy. For more information, see 8.6, “Advanced volume copy functions” on page 398.
4. Split or delete the old copy from the volume. For more information, see 8.6.2, “Splitting into a new volume” on page 400 or 8.6.4, “Delete Volume Copy option” on page 404.

This migration process requires more user interaction with the IBM Storwize V5000 GUI, but it offers some benefits.

As an example, we look at migrating a volume from a tier 1 storage pool to a lower performance tier 2 storage pool.

In step 1, you create the copy on the tier 2 pool, while all reads are still performed in the tier 1 pool to the primary copy. After the synchronization, all writes are destaged to both pools, but the reads are still done only from the primary copy.

Because the copies are fully synchronized, you can switch their role online (see step 3), and analyze the performance of the new pool. When you are done testing your lower performance pool, you can split or delete the old copy in tier 1 or switch back to tier 1 in seconds if the tier 2 storage pool did not meet your requirements.

8.7 Volumes by Storage Pool

To see an overview of which volumes are on which storage pool, click **Volumes by Pool**, as shown in Figure 8-85.

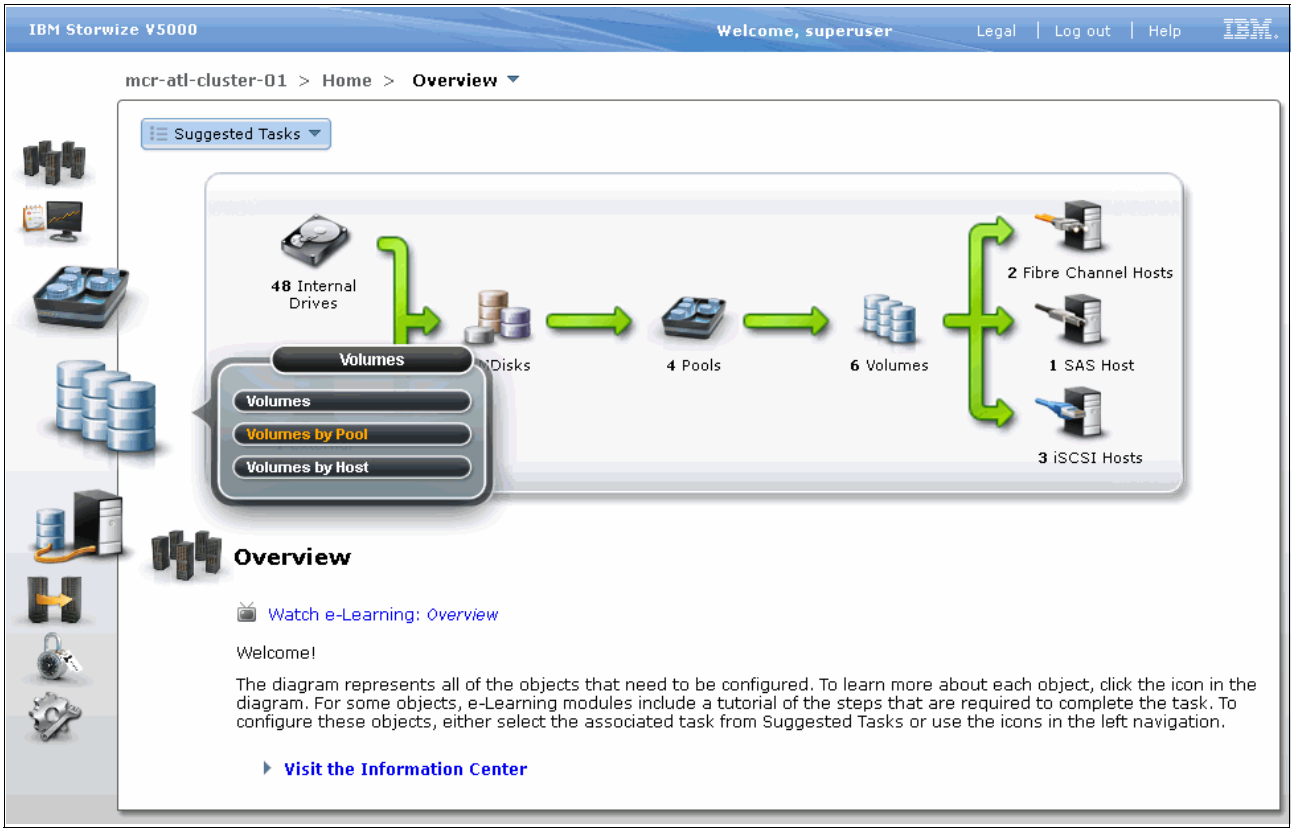

Figure 8-85 Volumes by Pool

The Volumes by Pool window opens, as shown in Figure 8-86.

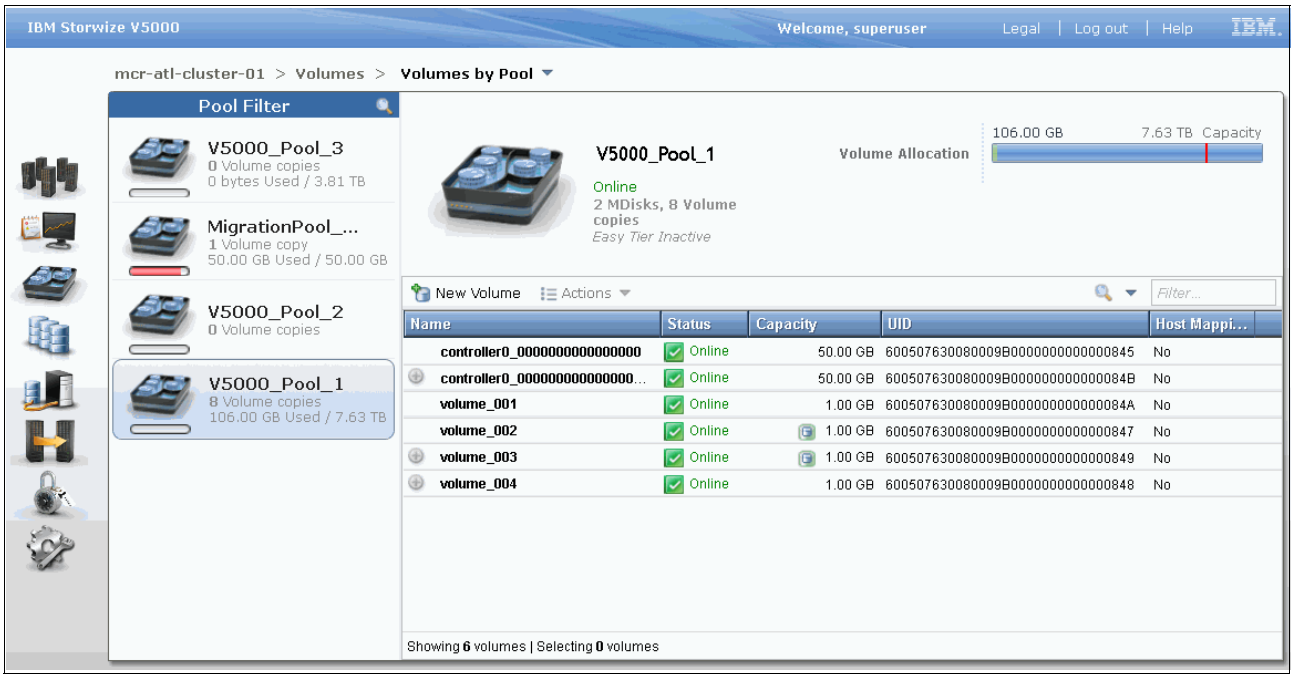

Figure 8-86 Volumes by Pool window

The left pane is called Pool Filter and all of your existing storage pools are displayed there. For more information about storage pools, see Chapter 7, “Storage pools” on page 295.

In the upper right, you see information about the pool that you selected in the pool filter. The following information is also shown:

- Pool icon: Because storage pools can have different characteristics, you can change the storage pool icon. For more information, see 7.4, “Working with storage pools” on page 343.
- Pool Name: The name that is given during the creation of the storage pool. For more information about changing the storage pool name, see “Rename” on page 341.
- Pool Details: Shows you the information about the storage pools, such as, status, the number of managed disks, and Easy Tier status.
- Volume allocation: Shows you the amount of capacity that is allocated to volumes from this storage pool.

The lower right section (as shown in Figure 8-87 on page 408) lists all volumes that have at least one copy in the selected storage pool. The following information is provided:

- Name: Shows the name of the volume.
- Status: Shows the status of the volume.
- Capacity: Shows the capacity that is presented to host.
- UID: Shows the volume unique identifier.
- Host Mappings: Shows if host mapping exists.

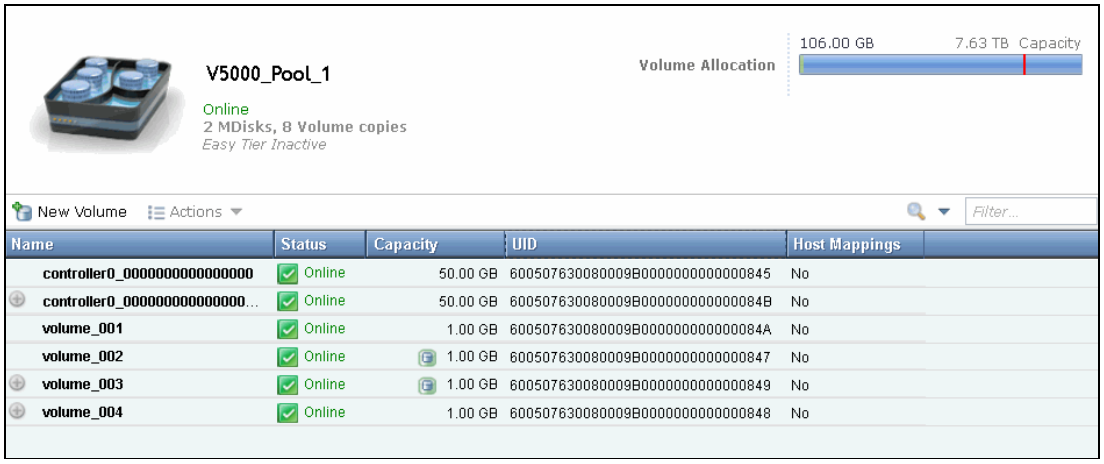

Figure 8-87 Volumes by Storage Pool

It is also possible to create volumes from this window. Click **Create Volume** to start the Volume Creation window. The steps are the same as the steps that are described in Chapter 5, “I/O Group basic volume configuration” on page 161.

If you highlight a volume and select **Actions** or right-click the volume, the same options are shown as described in 8.4, “Advanced volume administration” on page 375.

8.8 Volumes by Host

To see an overview about which volume a host can access, click **Volumes by Host** (as shown in Figure 8-85 on page 406) and the Volumes by Host window opens, as shown in Figure 8-88.

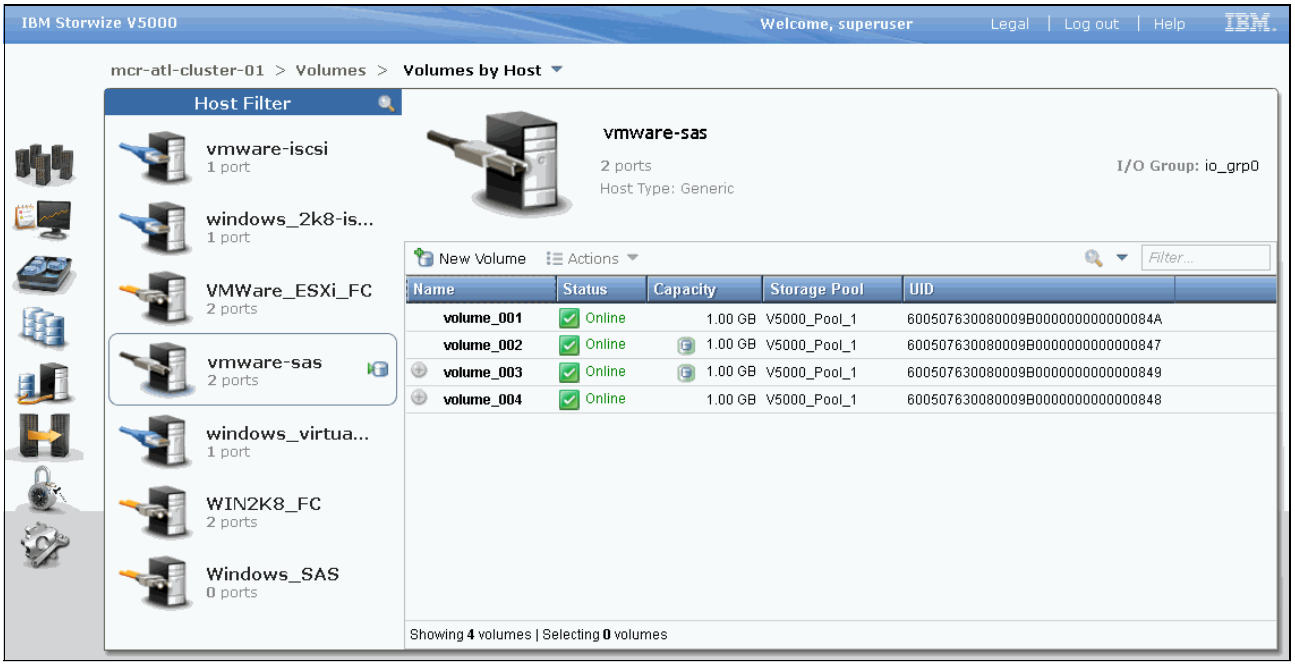

Figure 8-88 Volumes by host

In the left pane of the view is the Host Filter. If you select a host, its properties appear in the right pane, such as, the host name, number of ports, host type, and the I/O Group to which it has access.

The hosts with the orange cable represent the Fibre Channel host. The black cable represents the SAS hosts and the blue cable represents the iSCSI hosts.

The volumes that are mapped to this host are listed, as shown in Figure 8-89.

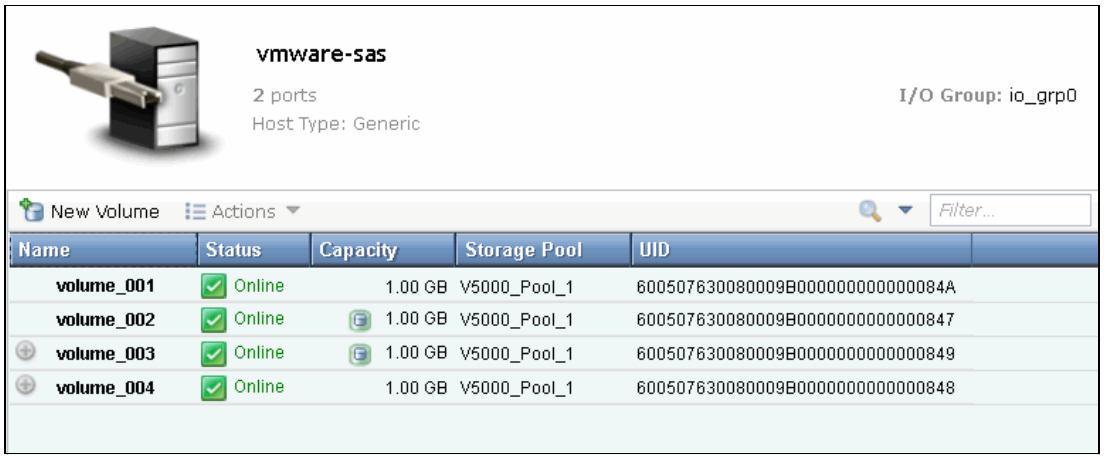

Figure 8-89 Volumes by Host

It is also possible to create a volume from this window. If you click **New Volume**, the same wizard opens that is described in 5.1, “Provisioning storage from IBM Storwize V5000 and making it available to the host” on page 162.

If you highlight the volume, the Actions button becomes available and the options are the same as those actions that are described in 8.4, “Advanced volume administration” on page 375.

9

Easy Tier

In today's storage market, solid-state drives (SSDs) are emerging as an attractive alternative to hard disk drives (HDDs). Because of their low response times, high throughput, and IOPS-energy-efficient characteristics, SSDs have the potential to allow your storage infrastructure to achieve significant savings in operational costs. However, the current acquisition cost per GB for SSDs is much higher than for HDDs. SSD performance depends greatly on workload characteristics, so SSDs must be used with HDDs. It is critical to choose the right mix of drives and the right data placement to achieve optimal performance at low cost. Maximum value can be derived by placing "hot" data with high I/O density and low response time requirements on SSDs, while HDDs are targeted for "cooler" data that is accessed more sequentially and at lower rates.

Easy Tier automates the placement of data among different storage tiers, and can be enabled for internal and external storage. This IBM Storwize V5000 licensable feature boosts your storage infrastructure performance to achieve optimal performance through a software, server, and storage solution.

This chapter describes the function that is provided by the Easy Tier disk performance optimization feature of the IBM Storwize V5000. It also describes how to activate the Easy Tier process for both evaluation purposes and for automatic extent migration. We included Storage Tier Advisor Tool (STAT) and Tivoli Storage Productivity Center for performance monitoring.

This chapter includes the following topics:

- Easy Tier overview
- Easy Tier for IBM Storwize V5000
- Easy Tier process
- Easy Tier configuration by using the GUI
- Easy Tier configuration by using the command-line interface
- IBM Storage Tier Advisor Tool
- Tivoli Storage Productivity Center
- Administering and reporting an IBM Storwize V5000 system through Tivoli Storage Productivity Center